

POLYHEDRAL AND COMBINATORIAL ASPECTS IN OPTIMIZATION

Habilitationsschrift

zur Erlangung des akademischen Grades

doctor rerum naturalium habitatus (Dr. rer. nat. habil.)

genehmigt

durch die Fakultät für **Mathematik**

der Otto-von-Guericke-Universität Magdeburg

von Dr. rer. nat. Dirk Oliver Theis
geb. am 28.07.1972 in Schwalmstadt

Gutachter:

Prof. Dr. rer. nat. habil. Amin Coja-Oghlan (University of Warwick)

Prof. Dr. rer. nat. habil. Michael Jünger (Universität zu Köln)

Prof. Dr. rer. nat. habil. Volker Kaibel (Otto-von-Guericke-Universität Magdeburg)

Magdeburg, den 02.05.2012
(Beschlussdatum)

Deutsche Zusammenfassung

Diese Schrift enthält Manuskripte und Nachdrucke aus den Gebieten der Optimierung und Kombinatorik. Zu Beginn beschäftigt sie sich mit den Polyedern des bekannten Problems des Handlungsreisenden. Zu diesem sind zwei Polyeder assoziiert (bekannt als graphisches bzw. symmetrisches), deren Verhältnis zwei Artikel beleuchten. Zum Einen stellt sich heraus, dass das graphische Polyeder sich allein geometrisch aus dem symmetrischen ergibt, unter Zuhilfenahme des Metrischen Kegels. Zum Anderen kann man den Chamber Komplex der Projektion der Polare der Polyeder elementar charakterisieren.

Weiterhin enthält die Schrift eine Arbeit über Kanten eines Polyeders zu Graph Labeling Problemen, sowie eine Arbeit zu Network Design mit nicht-linearen Kosten.

Stärker kombinatorisch ausgerichtet sind Arbeiten über kleine Minoren in Graphen mit großem, konstanten Durchschnittsgrad, sowie über Edge-Labelings in Graphen ohne kurze Kreise. Die Beziehungen zur Optimierung werden in der Einführung erläutert. Ein Artikel widmet sich der Färbungszahl von zufälligen Überlagerungen gewisser Graphen, mit endlicher, nach unendlich gehender Faser; ein weiterer der Erfüllbarkeit zufälliger logischer Formeln mit Bedingungen der Form $x \in I$ für Intervalle $I \subset [0, 1]$. Die Arbeit schließt mit Resultaten zum sogenannten Cops-&-Robber Spiel auf Graphen.

Contents

Part 1. Introduction	1
Chapter 1. Optimization and Polyhedra	3
Chapter 2. Optimization and Combinatorics	11
Part 2. Reprints & Manuscripts	19
Chapter 3. The relationship between the GTSP, STSP, and Metric Cone	21
Chapter 4. Facial structure of Symmetric and Graphical Traveling Salesman polyhedra	25
Chapter 5. On a class of metrics related to graph layout problems	45
Chapter 6. The VPN problem with concave costs	63
Chapter 7. Small minors	73
Chapter 8. Good edge labelings and graphs with girth at least five	93
Chapter 9. Random lifts	113
Chapter 10. Random 3-SAT with interval constraints	121
Chapter 11. Cops & Robber on non-orientable surfaces	151
Chapter 12. Cops & Robber on graphs with forbidden (induced) subgraphs	155
Bibliography	163

Part 1

Introduction

Optimization and Polyhedra

In 1954, Dantzig, Fulkerson and Johnson [DFJ54] initiated what is currently the most successful practical method for solving large-scale NP-hard discrete optimization problems. The idea was to use Dantzig's Simplex method for Linear Programming to try to solve the *Traveling Salesman Problem (TSP)*: Given a complete graph K_n together with lengths for its edges, find a *tour* visiting each vertex exactly once (or Hamiltonian cycle) with minimum total cost of its edges. The approach of Dantzig, Fulkerson and Johnson was iterative. They first decided on a Linear Programming formulation whose optimal solution would provide a lower bound to the length of the optimal tour. Due to the exponential size of the formulation, its solution would not be computationally feasible. Hence, only a considerably smaller Linear Program, containing a subset of the constraints, would actually be solved by the Simplex method. If the solution to the LP were found to violate some of the constraints which had been omitted, those constraints would be added to the Linear Program, and thus, an iterative procedure would generate successively better lower bounds on the length of the optimal tour.

This iterative *cutting plane method* is usually combined with branch-and-bound techniques. The resulting *Branch-and-Cut* method has proved tremendously successful in solving a great variety of NP-hard combinatorial optimization problems. Even though it is also at the heart of state-of-the-art Integer Programming solvers, a key to its current success in solving problems like the TSP lies in large parts with the understanding of polyhedra associated with the problems. For problems which are defined combinatorially, like the TSP, this results in fruitful interactions of polyhedral-geometric and combinatorial techniques.

1.1. Understanding the relationship between the Symmetric and Graphical TSP

The *Symmetric Traveling Salesman Polytope* is the convex hull of all characteristic vectors of edge sets of cycles (i.e., circuits) on the vertex set $V_n := \{1, \dots, n\}$ (in other words, Hamiltonian cycles in the complete graph with vertex set V_n). For the formal definition, denote by E_n the set of all two-element subsets of V_n . This is the set of all possible edges of a graph with vertex set V_n . The Symmetric Traveling Salesman Polytope is then the following set:

$$S_n := \text{conv} \left\{ \chi^C \mid C \text{ is the edge set of a Hamiltonian cycle with vertex set } V_n \right\} \subset \mathbb{R}^{E_n}.$$

Here, for an edge set F , χ^F is the characteristic vector in \mathbb{R}^{E_n} with $\chi_e^F = 1$ if $e \in F$, and zero otherwise.

In the mid nineteen-fifties, the first theoretical research about Symmetric Traveling Salesman Polytopes appeared in a series of short communications and papers [Hel55a, Hel55b, Hel56, Kuh55, Nor55]. With few exceptions (for example [FN92, Nor55] for the case $n \leq 5$; [BC91] for $n = 6, 7$; [CJR91, CR96, CR01, ORT07] for $n = 8, 9$), no complete characterization of the facets of S_n are known. In fact, since the TSP is NP-hard, there cannot exist a polynomial time algorithm producing, for every n and every point $x \in \mathbb{R}^{E_n} \setminus S_n$, a hyperplane separating x from S_n , unless $P=NP$ (I have omitted some technical conditions here). Another noteworthy argument for the complexity of these polytopes is a result of Billera & Sarangarajan [BS96]: For every 0/1-polytope P , there exists an n such that P is affinely isomorphic to a face of S_n .

Since the seminal work of Naddef & Rinaldi [NR91, NR93], a second polyhedron also has been used: the *Graphical Traveling Salesman polyhedron*. It is the convex hull of all characteristic vectors of edge multi-sets of connected (loopless) Eulerian multi-graphs on the vertex set V_n . (Recall that *Eulerian* means that there exists a walk containing all edges.) A (loopless) multi-graph with vertex set V_n has as its edge set a sub-multi-set of E_n . By defining, for any multi-set F of edges of K_n , its characteristic vector $\chi^F \in \mathbb{R}^{E_n}$ so that χ_e^F counts the number of occurrences of e in F , the Graphical Traveling Salesman Polyhedron is formally defined as

$$P_n := \text{conv} \left\{ \chi^F \mid F \text{ is the edge multi-set of a connected Eulerian multi-graph} \right. \\ \left. \text{with vertex set } V_n \right\} \subset \mathbb{R}^{E_n}.$$

Ever since Naddef & Rinaldi's papers on the Graphical and Symmetric Traveling Salesman Polyhedra [NR91, NR93], P_n is considered to be an important tool for investigating the facets of S_n . Moreover, the Graphical Traveling Salesman Polyhedron is also occasionally more convenient to work with computationally: In works of Carr [Car04] and Applegate, Bixby, Chvátal & Cook [ABCC01], P_n is (proposed to be) used algorithmically within Branch-and-Cut frameworks solving the TSP.

Numerous authors have expressed how close the connection between Graphical and Symmetric Traveling Salesman Polyhedra is. The most basic justification for this opinion is the fact that S_n is a face of P_n — consisting of all points x whose “degree” is two at every vertex —, but the links are far deeper (see [Nad02] or [ORT07] and the references therein).

The connections known before the publication of the short communication [The10], which is reprinted as Chapter 3 of this thesis, were established combinatorially by comparing Hamiltonian cycles with spanning Eulerian multi-graphs. Surprisingly, though, the relationship of the two polyhedra can be understood entirely geometrically.

Theorem 1.1. P_n is the intersection of the positive orthant with the Minkowski sum of S_n and the polar C_n^Δ of the metric cone C_n :

$$P_n = (S_n + C_n^\Delta) \cap \mathbb{R}_+^{E_n}$$

The metric cone consists of all $a \in \mathbb{R}^{E_n}$ which satisfy the triangle inequality

$$a_{uv} \leq a_{uw} + a_{wv}$$

for all pairwise distinct vertices $u, v, w \in V_n$. Consequently, its polar is generated as a cone by the vectors

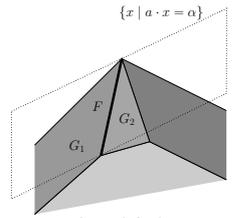
$$\chi^{\{uw\}} + \chi^{\{wv\}} - \chi^{\{uv\}}.$$

While the importance of the triangle inequality was realized already by Naddef & Rinaldi [NR91, NR93], the depth of this link has not been noticed for 20 years.

1.2. On the facial structure of Symmetric and Graphical TSP

Although surprising, Theorem 1.1 scratches only on the surface of the connection between Symmetric and Graphical Traveling Salesmen.

As mentioned above, S_n is a face of P_n . This means that every inequality valid for S_n can be “rotated” to make it valid for P_n . By “rotation” we mean modifying left and right hand sides of an inequality $a \bullet x \geq \alpha$ in such a way that the set of points in the affine hull of S_n which satisfy the inequality with equation remains the same, yet the hyperplane the inequality defines in the ambient space changes. Technically, this amounts to adding equations valid for S_n to $a \bullet x \geq \alpha$. Once the inequality is rotated so that it is valid for P_n , one may ask which face of P_n is defined by the rotated inequality. Since $S_n \neq P_n$, there are always several such



faces, but even when we aim for inclusion-wise maximal faces of P_n defined by some rotated version of $a \bullet x \geq \alpha$, in general, these are not unique. In the picture above, by properly tilting the hyperplane defined by $a \bullet x = \alpha$, we can obtain the faces F , G_1 and G_2 .

Extending results obtained in [The05], the manuscript [The], which forms Chapter 4 of this thesis, deals with the following question:

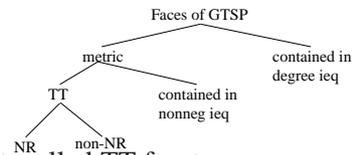
Question 1.2. *Given a valid inequality $a \bullet x \geq \alpha$ for S_n , what is the largest possible dimension of a face of P_n defined by a rotated version of that inequality?*

It turns out that to answer this question, next to knowing what face of S_n is defined by the inequality, one only needs to look at the $n(n - 1)(n - 2)/2$ numbers $a_{uv} + a_{vw} - a_{uw}$, for all triples u, v, w of distinct vertices in V_n . More accurately, only the ordering relations between these numbers are important.

Rotation is a standard tool in Discrete Optimization. The most prominent example is sequential lifting, which is a constrained form of rotation. In the setting of sequential lifting, P is a polyhedron for which the non-negativity inequality $x_j \geq 0$ for a coordinate j is valid, defining a non-empty face $S := P \cap \{x \mid x_j = 0\}$. Then, an inequality valid for S is rotated by adding scalar multiples of the equation $x_j = 0$ to it in such a way that it becomes valid for P and the face defined by the rotated inequality is strictly greater than the face of S defined by it. By iterating this procedure, one may “sequentially” lift inequalities which are valid for a smaller face S . The face of P defined by the sequentially lifted inequality may in general depend on the order in which the coordinates j are processed. The same procedure works when generic inequalities $c \bullet x \geq \gamma$ are used instead of the non-negativity inequalities.

Sequential lifting or other rotation-based tools are applied manually to find facets of polyhedra which contain faces which are better understood. Moreover, mechanisms of this kind are used computationally in cutting-plane algorithms where some separation procedure first works on a face and then lifts the obtained inequalities.

Naddef & Rinaldi [NR91, NR93] proved a theorem saying that, if an inequality defines a facet of S_n , then there is a unique maximal face of P_n which can be obtained by rotating the inequality, and this maximal obtainable face is a facet of P_n . Naddef & Rinaldi classified the facets of P_n into three types — non-negativity facets, degree facets, and the rest, called TT-facets — based on properties of the coefficients. While the degree facets and non-negativity facets are both small in number and easily understood, the interesting class both for understanding the polyhedron and for applications is the huge set of TT-facets. By the theorem just mentioned, once one knows that the degree facets of P_n are precisely those which contain S_n — also proved in Naddef & Rinaldi’s paper —, this also classifies the facets of S_n into two types: non-negativity and TT-facets.



In an earlier paper [ORT05, ORT07] we have refined the classification by splitting the TT-facets of P_n into two subclasses: NR-facets and non-NR-facets, depending on whether the intersection of the facet with S_n is a facet of S_n (these P_n -facets are called NR-facets) or a face of S_n of smaller dimension. Our main result was the fact that the non-NR class is not empty.

In terms of rotation, this shows that there are (“TT-type”) valid inequalities for S_n which do not define facets of S_n , but which can be rotated to define facets of P_n .

The answer to Question 1.2 is formulated using the terminology of polar polyhedra, polyhedral complexes, and polyhedral subdivisions. A polar polyhedron S^Δ of a polyhedron S has the property that the points of S^Δ are in bijection with the linear inequalities (up to scaling) for S . Moreover, a point a is contained in a face of dimension k of S^Δ , if, and only if, the corresponding inequality defines a face of dimension at least $\dim S + 1 - k$ of S . In particular, the vertices of S^Δ are in bijection with the facets of S . A polyhedral complex is a finite set of polyhedra,

closed under taking faces, such that the intersection of any two polyhedra in the set is a face of both.

For the reasons explained above, the results in [The] pertain to the “important” part of the polar of S_n , namely the part which remains if we delete the vertices corresponding to non-negativity facets. This corresponds to taking only the “TT-type” valid inequalities for S_n . This subset of faces of the polar of S_n is a polyhedral complex, which we denote by \mathcal{A} .

Now, take a point in \mathcal{A} , consider the corresponding valid inequality for S_n , and rotate it in all possible ways yielding inequalities valid for P_n . A certain set of faces of P_n can be defined by the rotated versions of this inequality. We partition the points contained in \mathcal{A} in the following way: two points are in the same cell of the partition, if, by rotating the corresponding valid inequalities, the two sets of faces of P_n which can be defined, coincide.

In fact, the partition whose definition I have just outlined, gives a polyhedral subdivision \mathcal{S} of \mathcal{A} , i.e., the set of closures of the cells is a polyhedral complex, and every face of \mathcal{A} is a disjoint union of cells. This is true in the general situation when a polytope S is a face of another polytope P . Indeed, in the general situation, \mathcal{S} is known as the *chamber complex* of the canonical projection of the polar of P onto the polar of S . I call such a polyhedral subdivision a *rotation complex*. In [The], the following results are established about the rotation complex in the TSP situation:

- (A) The decomposition of \mathcal{A} into cells can be described in an elementary way that does not refer to rotation; moreover, it does not refer to any Graphical Traveling Salesman concepts whatsoever. Indeed, to describe the subdivision, for a point a contained in \mathcal{A} , it suffices to check the order relations of the expressions $a_{uv} - a_{uw} - a_{vw}$, with u, v, w three distinct vertices in V .
- (A') The rotation complex \mathcal{S} is the common refinement of \mathcal{A} with a projection of a natural sub-complex of the boundary complex of the metric cone. (The common refinement of two polyhedral complexes is the set of all intersections of polyhedra in the two complexes, see Fig. 4.1, left, on page 28.)
- (B) The points in \mathcal{A} are in *bijection(!)* with the “important” part of the polar of P_n , and this bijection maps faces of the polar of P_n onto faces of the rotation complex \mathcal{S} . In other words, the polar of P_n can be “flattened” onto the polar of S_n , see Fig. 4.1, right.

Again, “important” is meant to be understood in the sense that it corresponds to considering TT-type inequalities only. Item (B) is not a consequence of known facts about the chamber complex (injectivity fails to hold in general). The picture in Fig. 4.1, left, on page 28, illustrates Items (A) and (A'). It shows a hypothetical drawing of \mathcal{A} (solid lines) with two points a, a' . To decide whether these two points, when viewed as valid inequalities for S_n , yield the same faces of P_n when rotated, one has to check the expressions $a_{uv} - a_{uw} - a_{vw}$. This amounts to checking if they are contained in the same cone in the picture on the left (dotted lines). Due to the Theorem 1.1 discussed in the previous section, occurrence of the metric cone here is no surprise, of course.

The rigorous formulation of the two theorems corresponding to (A), (A') and (B) requires a larger technical apparatus, and is omitted at this point. Germs of Items (A) and (B) had already been proved in [The05], albeit with a considerably more complicated proof.

An outlook on polyhedral STSP/GTSP problems

Earlier versions of methods developed in Chapter 4 here helped resolve two open questions regarding Symmetric vs. Graphical TSP: The existence of non-NR facets and the complete description of P_9 (together with a computer search) [ORT07]. A number of open problems remain, which I would like to address here.

0-Node lifting. The polyhedron P_n has the pleasant property that a very simple lifting operation called *0-node lifting* [NR91, NR93] preserves the facet defining property. In other words,

if an inequality defines a TT-type facet of P_n , then duplicating a vertex in V_n and joining the two twins by an edge with coefficient 0 yields a facet-defining inequality for P_{n+1} . Annoyingly, to this date, it is not known whether if one starts with an inequality defining an NR-facet (i.e., the inequality also defines a facet of S_n), the facet defined by the 0-node lifted inequality can be non-NR (meaning, it does not define a facet of S_{n+1}).

Conjecture 1.3. *There exist NR-facet defining inequalities which, after 0-node lifting, define non-NR facets.*

It is known that when, starting with an NR-facet, 0-node lifting is applied twice at the same vertex, then the resulting facet is NR [QW93]. Hence, in terms of rotation complexes, an example as in the conjecture would behave like this: By 0-node lifting, a vertex of the polar of S_n jumps into the interior of a face of S_{n+1} which is a vertex of the rotation complex, and by 0-node lifting again at the same vertex, it jumps to a vertex of the polar of S_{n+2} . Consequently it appears as if the rotation complex theory ought to be able to prove wrong Conjecture 1.3. However, I believe that it is actually true.

Computationally checking millions of NR-facets (with $n = 10, 11, 12, 13$) has not unearthed such an inequality. The problem with computational methods in searching such an example is the following. For $n \leq 8$, none of the 24 TT-type facet classes (i.e., facets modulo permutation of vertices) of P_n are non-NR. For $n = 9$ there is exactly one non-NR facet class among the 192 TT-type facet classes of P_9 (0.52%). Among the (conjectured) 15621 TT-type facet classes of P_{10} , there are (conjectured) 243 non-NR facet classes (1.56%). While the ratio seems to be increasing with n , for those values of n in which computation can be done in any significant scale (up to 15), the non-NR facets seem to appear to be statistically scarce. On the other side, there is the observation that the 0-node lifted facet classes of P_n are also statistically very scarce in the TT-type facet classes of P_{n+1} (13% for $n = 8$, 1.23% for $n = 9$, appears to be decreasing with n). This makes it appear unlikely to hit, “by chance”, an example of a 0-node lifted NR-facet which is also a non-NR facet.

Parsimonious relaxations. In Theorem 4.5 (see also [The05, ORT07]), a necessary condition is given for a certain subgraph of the ridge-graph of GTSP to be connected. I believe that this condition actually characterizes so-called *parsimonious relaxations* (see Section 4.2.4 for the definition).

Recall that the ridge-graph has as its vertices the facets of a polyhedron, with two facets being adjacent if their intersection has maximal possible dimension.

Conjecture. *Suppose a system of inequalities defining NR-facets has the following property: If the corresponding vertices are deleted from the ridge graph, then every connected component contains an NR-facet. Then the relaxation given by this system of inequalities has the parsimonious property.*

The formulation here is not exact, see Conj. 4.24 for the exact formulation. The conjecture holds for the known relaxations of S_n consisting of NR-inequalities described in [ORT07] which fail the parsimonious property.

1.3. Embeddable Metrics and the Linear Arrangement Problem

The study of polyhedra consisting of metrics or semimetrics on some fixed finite space has a long tradition in Polyhedral Combinatorics and Convex Geometry; see, e.g. [DL97], for a starting point. (A *semimetric* satisfies all the requirements of a metric, except the distance of distinct points may be zero; it is customary in this area to use the term “metric” also for semimetrics.) Without additional constraints on the semimetrics, one has the metric cone, which we have already encountered above. Often, embeddability constraints are required: One studies convex/polyhedral-geometric properties of the set of (semi-)metrics embeddable in a fixed normed space, possibly with additional conditions.

It turns out that these constrained sets of semimetrics are related to combinatorial optimization problems. As a famous example, the set of all semimetrics on an n -point set which are embeddable in $\mathbb{R}^{\binom{n}{2}}$ with the 1-norm coincides with the cut cone.

In the paper [LRST10], which is reprinted as Chapter 5 of this thesis, we study metrics which are embeddable in the real line in such a way that every two points are at least some fixed $\varepsilon > 0$ apart. These metrics are related to the Linear Arrangement and other graph layout problems.

Let me first note that changing the ε amounts to a dilation of the set of metrics, so that fixing $\varepsilon := 1$ is no loss of generality. We call those metrics “ \mathbb{R} -embeddable 1-separated”. Secondly, the cardinality of the point set will be denoted by n .

The following classes of semimetrics, which are closely related to the \mathbb{R} -embeddable 1-separated metrics, have been extensively studied in the literature:

- As indicated above, the cut cone, CUT_n , coincides with the set of all semimetrics which can be embedded into a space with 1-norm, i.e., semimetrics d (on a fixed n -point set) for which there exists an m and points $x_1, \dots, x_n \in \mathbb{R}^m$ such that $d(i, j) = \|x_i - x_j\|_1$ for all pairs of points i, j . These semimetrics are customarily referred to as ℓ_1 -semimetrics.
- The ℓ_2 -embeddable semimetrics, which are defined as in the ℓ_1 case, except that $d(i, j) = \|x_i - x_j\|_2$. Denote the set of these metrics by M_n^{L2} .
- The \mathbb{R} -embeddable semimetrics, which are the special case of ℓ_1 - (or ℓ_2 -) embeddable semimetrics obtained when m is fixed to 1. The set of these metrics is denoted by M_n^R .

In general, the sets M_n^{L2} and M_n^R are not convex, but the convex hull of both M_n^{L2} and M_n^R is CUT_n . Similarly, the set M_n^{R1} of \mathbb{R} -embeddable 1-separated metrics is not convex (it is the union of $n!/2$ disjoint simplicial cones, see Prop. 5.3), but its convex hull occurs in the context of graph layout problems.

Connection to Graph Layout problems. Given a graph G a *layout* is a bijection $\pi: V \rightarrow \{1, \dots, |V(G)|\}$. Several important combinatorial optimization problems, collectively known as *graph layout problems*, call for a layout minimizing a function of the distances $|\pi(u) - \pi(v)|$, $uv \in E(G)$ (see the survey [DPS02]). In the Linear Arrangement Problem, the objective is to minimize $\sum_{uv \in E} |\pi(u) - \pi(v)|$. In the Bandwidth Problem, the objective is to minimize $\max_{uv \in E} |\pi(u) - \pi(v)|$.

Now let $d(u, v)$ for $\{u, v\} \subset V(G)$ be an integral variable representing the quantity $|\pi(u) - \pi(v)|$. It has been observed by several authors that interesting relaxations of graph layout problems can be formed by deriving valid linear inequalities that are satisfied by all feasible symmetric functions d . Some of these inequalities have been used to derive approximation algorithms for various graph layout problems (e.g. [ENRS00, RR05]). It is thus natural to study the following *permutation metrics polytope*: [AL09]

$$P_n = \text{conv} \left\{ d: V \times V \rightarrow \mathbb{R} \mid \exists \pi \in S(n) \text{ s.t. } d(i, j) = |\pi(i) - \pi(j)| \forall i, j \in V \right\},$$

where $S(n)$ stands for the set of all permutations of $V = \{1, \dots, n\}$. In [AL09], it is shown that P_n is of dimension $\binom{n}{2} - 1$ and that its affine hull is defined by the equation $\sum_{\{i, j\} \subset V} d(i, j) = \binom{n+1}{3}$.

The connection with \mathbb{R} -embeddable 1-separated metrics lies in the fact that the closure of the convex hull of these metrics equals the Minkowski sum of P_n and the cut cone (Prop. 5.6).

Denote the convex hull of M_n^{R1} by Q_n (this is not a polyhedron), and its closure by $\overline{Q_n}$.

Unbounded edges of Q_n and $\overline{Q_n}$. From this starting point, in the paper [LRST10], we move to study the unbounded edges of the convex set Q_n (Prop. 5.20) and the polyhedron $\overline{Q_n}$ (Theorem 5.22). Since $\overline{Q_n} = P_n + \text{CUT}_n$ (see above), the directions of the unbounded edges

correspond to cuts, whereas the vertices in which these edges originate correspond to permutations. Characterizing, for a given permutation/vertex, the set of cuts corresponding to unbounded edges originating in the vertex requires combinatorial/geometric investigations which are done with some ease in the case of Q_n , but, in the case of $\overline{Q_n}$, they are quite subtle. I cannot restrain myself from pointing the reader to the beautiful connections between geometry and combinatorics in Lemmas 5.32 (“reduction”) and 5.34 (“induction”), as well as in the examples of that section.

The resulting relationship between permutations and cuts given by the incidence of extreme rays of CUT_n on vertices of P_n is the following: for a permutation π , a cut $(U : V \setminus U)$ corresponds to an unbounded edge ending in the vertex π if there is no $k \leq n - 1$ such that either U or $V \setminus U$ equals $\pi^{-1}(\{1, \dots, k - 1, k + 1\})$.

(The paper [LRST10] also contains studies of facets of $\overline{Q_n}$. These were mainly done by A. Letchford, who also proposed the study of unbounded edges of Q_n and $\overline{Q_n}$.)

Outlook. Our work raises some further questions, most prominently the following:

Question 1.4. *Do the bounded edges of $\overline{Q_n}$ have a simple combinatorial interpretation?*

The bounded edges of $\overline{Q_n}$ are of course those of P_n .

Generally speaking, bounded edges have received more attention than their longer brothers. This fact roots in hopes to adopt the simplex algorithm to make use of edges in a more “direct” way, without requiring a complete description of the polyhedron by inequalities, and has given rise to a number of conjectures and questions about (bounded) edges of combinatorial optimization polyhedra, or even polytopes in general, most famously, of course, to the Hirsch conjecture. Given that for other classes of unbounded combinatorial optimization polyhedra, no characterization of the extreme rays or even vertex / unbounded edge incidences is known, our result might stir hopes that for P_n , the adjacency relation of vertices admits a combinatorial characterization.

1.4. Virtual Private Network design with non-linear costs

The fourth paper reprinted in this thesis has a somewhat lesser connection to polyhedra. In the *symmetric Virtual Private Network design* (sVPN) problem, vertices want to communicate with each other. The exact amount of traffic between pairs of vertices is not known in advance, but for each vertex v the cumulative amount of traffic that it may send or receive is bounded from above by a given number b_v . The aim is to install minimum cost capacities on the edges of the graph supporting *any* possible communication scenario subject to these bounds. The cost for installing one unit of capacity on an edge e is c_e .

Goyal, Olver and Shepherd [GOS08b] proved that the symmetric Virtual Private Network Design (sVPN) problem has the so-called *tree routing property*, namely, that there always exists an optimal solution to the problem which installs non-zero capacities only on edges which do not contain a cycle. Earlier, Fingerhut, Suri and Turner [FST97] and Gupta, Kleinberg, Kumar, Rastogi and Yener [GKK⁺01] had shown that such a tree-shaped solution can be found in polynomial time. Thus, sVPN can be solved in polynomial time.

In the paper [FOST10] which is reprinted as Chapter 6, we consider an APX-hard generalization of sVPN , where the contribution of each edge to the total cost is proportional to some non-negative, concave, non-decreasing function f of the capacity reservation (f does not depend on the edge and is given by an oracle).

The polyhedral part of that paper is the proof for the fact that the non-linear version has the tree routing property, too. For this, we associate polyhedra with instances of a related problem (the sCR , see Fig. 6.1 on page 65) in such a way that the tree routing property for an instance can be expressed as a property of the extreme points of the associated polyhedron. We then

show how the transition from linear to concave amounts to a coordinate-wise concave mapping of the corresponding polyhedra, which preserves the property of the extreme points.

Building on this, we study approximation algorithms for the concave version of $sVPN$. For a general concave function, using known results on the so-called Single Source Buy at Bulk (SSBB) problem, we give a 24.92-approximation algorithm. For a more restricted class of concave functions, by reducing to the so-called Single Source Rent or Buy (SSRB) problem, we are able to obtain a 2.92-approximation.

Optimization and Combinatorics

2.1. Small Minors in graphs

I would like to motivate the paper [FJTW], which forms chapter 7 through its connection to optimization. Many optimization problems are of the following form. Given a graph G (possibly with costs on vertices or edges) find a subgraph of G which belongs to a target graph class, by deleting as few vertices and edges as possible (i.e., by incurring smallest possible cost). A trivial example is the Maximum Spanning Forest problem: Delete as few edges as possible such that the resulting graph is a disjoint union of trees. Fiorini, Joret, and Pietropaoli [FJP10] considered what they called the “Diamond Hitting Set” problem: Delete as few vertices as possible such that the resulting graph is a disjoint union of cactus graphs. Informally, a cactus graph is obtained from a tree by replacing some of the edges with cycles. Formally, a graph H is a disjoint union of cactus graphs if and only if H does not contain a diamond $K_4 \setminus e$ as a minor.

To obtain an approximation algorithm for their Diamond Hitting Set problem, Fiorini et al. used a lemma saying that every graph G with average degree at least three contains, as a subgraph, a subdivision of a diamond of size $O(\log n)$, with n being the number of vertices in G . Moreover, such a subgraph can be found in polynomial time. This lemma allows for a very simple $O(\log n)$ approximation algorithm (for the unweighted problem), one of whose core ingredients is iteratively finding and deleting $O(\log n)$ -sized diamond-subdivisions as long as there are any.

Their result left the following obvious question. If the target graph class is defined by forbidding another minor H instead of the diamond, does there still exist a $O(\log n)$ approximation algorithm? Since Fiorini et al.’s algorithm relies heavily on the existence of a small diamond-subdivision given large enough average degree, a necessary (though not sufficient) condition to successfully apply their techniques would be the existence, in every graph G with large enough average degree, of an H minor supported on a small subgraph of G .

Our manuscript [FJTW] deals with this problem.

The case $H = K_3$ asks for a short cycle in a graph. An easy and well-known theorem states that, if a graph has minimum degree larger than two, then it contains a cycle of size $O(\log n)$. Alon, Hoory, and Linial [AHL02] extended this to graphs with average degree larger than two (but see Lemma 7.4 for a different proof; Alon et al. prove considerably more than this statement).

For general K_t minors, Kostochka and Pyber [KP88] proved that, given $t, \varepsilon > 0$, every graph with at least $t4^{t(t-1)}n^{1+\varepsilon}$ edges contains a K_t -subdivision with at most $(7t^2 \log t)/\varepsilon$ vertices. Taking $\varepsilon := 1/\log n$, $t = 4$, for example, and a conjectured improvement of $7(t^2 \log t)/\varepsilon$ to $O(t^2/\varepsilon)$, this gives a bound of 4^{13} for the average degree.

Our first result is the following.

Theorem (Theorem 7.5 on page 76). *Every graph with average degree at least $4 + \varepsilon$ contains a K_4 -subdivision of size $O(\log n)$.*

There are simple examples of (even planar) graphs with average degree four whose only K_4 -minors have size $\Omega(n)$ (see Fig. 7.1 on page 74). Using an inductive approach, this theorem can be extended to yield the following.

Theorem (Theorem 7.6 on page 77). *Every graph with average degree at least $2^t + \varepsilon$ contains a K_t minor supported on $O(\log n)$ vertices.*

Planar graphs. It turns out that, if G is planar, one can do better than Theorem 7.5. Using the discharging method (see Lemma 7.16), we could prove the following.

Theorem (Theorem 7.19 on page 83). *Every planar graph with average degree at least $4 + \varepsilon$ contains a K_4 -subdivision of size $O(1)$.*

Outlook. While the average degree bound in Theorem 7.5 is optimal, the one in Theorem 7.6 is not (even though it is a big improvement on Kostochka and Pyber’s $t4^{t(t-1)}$). Thus, the most obvious question is the following.

Question 2.1. *Is there a sub-exponential function f such that every graph G with average degree at least $f(t)$ contains a K_t minor supported on $O(\log n)$ vertices?*

The following, I find somewhat more intriguing. Let $f(t)$ be the infimum over all numbers d for which the following holds: Every graph G with average degree at least d contains a K_t minor (regardless how much of the graph it covers).

Question 2.2. *Does every graph with average degree at least $f(t) + \varepsilon$ have a K_t minor supported on $O(\log n)$ vertices?*

For $t = 4$, our theorem proves just that. For $t = 5$, we have $f(5) = 6$, and the question is open.

2.2. Good edge labelings

Our manuscript [BFT], which is included as Chapter 8 in this thesis, deals with a theoretical problem arising in the context of so-called Wavelength Division Multiplexing problems. Given a network, the *Routing and Wavelength Assignment Problem* asks for finding routes and associated wavelengths, such that a set of traffic requests is satisfied, while minimizing the number of used wavelengths [BCCP06]. In a recent paper, Bermond, Cosnard, and Pérennes [JCBP09] establish a relationship with good edge-labelings.

A *good edge-labeling* of a graph G is a labeling of its edges $\phi: E(G) \rightarrow \mathbb{R}$ such that, for any ordered pair of distinct vertices u and v , there is at most one nondecreasing path from u to v . Equivalently:

An edge-labeling is good, if, and only if, every cycle has at least two local minima.

For simplicity, let us say that a local minimum is an edge e whose label is strictly less than the labels of the two edges incident to e on the cycle (this simplification requires to assume, wlog, that all labels are distinct).

Araujo, Cohen, Giroire, and Havet [ACGH09, ACGH12] have studied good edge-labelings in more depth. They call a graph with no good edge-labeling *bad*, and say that a *critical* graph is a minimal bad graph, that is, every proper subgraph has a good edge-labeling. It is easy to see that C_3 and $K_{2,3}$ are critical. Araujo et al.’s [ACGH12] paper comprises an infinite family of critical graphs; results that graphs in some classes always have a good edge-labelings (planar graphs with girth at least 6, $(C_3, K_{2,3})$ -free outerplanar graphs, $(C_3, K_{2,3})$ -free sub-cubic graphs); the algorithmic complexity of recognizing bad graphs; and a connection to matching-cuts. (A *matching-cut*, also known as “simple cut” [Gra70], is a set of independent edges which is an edge-cut.)

In fact, all their arguments for proving non-criticality rely on the existence of matching-cuts. One of the central contributions of our paper [BFT] is that we move beyond using matching-cuts.

Araujo et al. also pose a number of problems and conjectures. In particular, they have the following conjecture, which is one of the two motivations behind our paper.

Conjecture (Araujo et al. [ACGH12]). *There is no critical graph with average degree less than three, with the exception of C_3 and $K_{2,3}$.*

Araujo et al. [ACGH12] prove a weaker version of this conjecture, relying in part on a theorem by Farley and Proskurowski [FP84, BFP11] stating that a graph with sufficiently few edges always has a matching cut. They also use a characterization of extremal graphs with no matching-cut by Bonsma [Bon05, BFP11]. From the proofs in Araujo et al. [ACGH12], it appears that the depths of the arguments increases rapidly as the upper bound 3 is approached.

In our paper, we show that there is no critical graph with average degree less than three and girth at least five. Put differently, we prove Conjecture 8.1 in the case when the graph has girth at least five.

Theorem (Theorem 8.2 on page 94). *There is no critical graph with average degree less than three and girth at least five.*

The second motivation behind our paper is the fact that no bad graph with girth larger than four is known. In particular, the bad graphs in Araujo et al.’s construction contain many 4-cycles. It is quite natural to ask whether there exists a number g such that every graph with girth at least g has a good edge-labeling. As mentioned above, Araujo et al. [ACGH12] have shown that with the additional restriction that the graphs be planar, $g := 6$ does the trick.

We prove a structural theorem on critical graphs with girth at least five (Theorem 8.42). Roughly speaking, it says that a critical graph with girth at least five cannot contain a subgraph which is a “windmill”. A windmill essentially consists of a number of shortest paths meeting in an “axis”, with the paths originating from vertices of degree two and having in their interior only vertices of degree three.

Of this Theorem 8.42, the above state Theorem 8.2 is a corollary, which is proved using an approach inspired by the discharging method from topological graph theory.

For our proof of Theorem 8.42, we define a class of graphs which we call “decent”, which have the property that they cannot be contained in a critical graph. More importantly, we give a quite general *gluing* operation which preserves “decency”. Starting from a small family of basic “decent” graphs, by gluing inductively, this approach allows us to show that certain more complicated configurations cannot be contained in critical graphs, which leads to the proof of Theorem 8.42.

Outlook. I believe that the following question is the most fundamental one concerning good edge-labelings.

Question. *Is there a constant g such that every graph with girth at least g has a good edge-labeling?*

Araujo et al. propose the following conjecture.

Conjecture ([ACGH12]). *For every $c < 4$, the number of (isomorphism classes of) critical graphs with average degree at most c is finite.*

In view of our work on good edge-labelings and girth, I think that the following conjecture, if true, might be considerably easier to answer in the affirmative.

Conjecture. *For every $c < 4$, the number of (isomorphism classes of) critical graphs with girth at least five and average degree at most c is finite.*

2.3. Coloring random lifts

Let G be a graph, and h a positive integer. An h -lift of G is a graph \tilde{G} which is an h -fold covering of G in the sense of CW-complexes. Put differently, there is a graph homomorphism

$\phi: \tilde{G} \rightarrow G$ which maps the neighborhood of any vertex v in \tilde{G} one-to-one onto the neighborhood of the vertex $\phi(v)$ of G . The graph G is called the *base graph* of the lift.

More concretely, we may say that an h -lift of G has vertex set $V(G) \times [h]$, with $[h] := \{1, \dots, h\}$ as usual. The set $\{v\} \times [h]$ is called the *fiber over v* . Fixing an arbitrary orientation of the edges of G , the edge set of an h -lift is of the following form: There exist permutations σ_e of $[h]$, $e \in E(G)$, such that for every edge uv of G , oriented $u \rightarrow v$, the edges between the fibers $\{v\} \times [h]$ and $\{u\} \times [h]$ are $(u, j)(v, \sigma_{uv}(j))$, $j \in [h]$. Changing the orientation of the edges in the graph obviously does not change the lift — just replace each affected permutation by its inverse.

A *random h -lift* of G is a graph drawn uniformly at random from the graphs just described, which amounts to choosing a permutation, uniformly at random, independently for every edge of G .

Amit, Linial, Matoušek, and Rozenman [ALMR01], proposed to study properties of random lifts in the limit $h \rightarrow \infty$. Their conference paper sketched results on connectivity, independence number, chromatic number, perfect matchings, and expansion of random lifts, and was followed by a series of articles containing broader and more detailed results [AL02, AL06, ALM02, LR05], and e.g. [BL06, DL06, LP10, BCCF06, GJR10], to name a few.

In [ALM02] Amit, Linial, and Matoušek focused on independence and chromatic numbers of random lifts of graphs. They asked the following question.

Question 2.3. *Is there a zero-one law for the chromatic number of random lifts?*

By zero-one law, they mean that the chromatic number of a random lift (of a fixed base graph) is asymptotically almost surely (aas) for $h \rightarrow \infty$ equal to a fixed number, depending only on the base graph.

For the base graph is K_n , Amit et al. prove that $\chi(\tilde{G}) = \Theta(n/\log n)$ aas (with absolute constants in the Θ). The smallest value for n , for which this is not trivial, is $n = 5$. Amit et al. ask the following:

Question 2.4. *Is the chromatic number of a random lift of K_5 aas equal to a single number?*

It is easy to see that the only two numbers which might occur with positive probability are 3 and 4.

In our paper [FT], which is reprinted as Chapter 9 we give an algorithm which 3-colors random lifts of $K_5 \setminus e$, the graph obtained by deleting one edge from K_5 , and prove that it succeeds aas:

Theorem 2.5 (Theorem 9.1 on page 114). *A random lift of $K_5 \setminus e$ is aas 3-colorable.*

The theorem can easily be extended to a larger graph of base graphs consisting of a cycle joined to an independent set.

Shi and Wormald [SW07] proved that the chromatic number of random 4-regular graphs (with uniform distribution) is three, and random lifts of K_{d+1} have some similarity to random d -regular graphs. However, the cycle structure of random lifts is more delicate than that of uniformly random regular graphs (it is related to the distribution of fixed points of words of random permutations, which is understood [Nic94, LP10]), and the Shi-Wormald proof makes explicit use of the cycle structure. Still, I believe that adapting the Shi-Wormald algorithm and proof ought to be possible to settle Question 2.4. On the other hand, I do not think there is an answer to the question which is simpler than the corresponding question for uniformly random 4-regular graphs.

2.4. Random 3-SAT with interval constraints

The result of Shi and Wormald [SW07] just mentioned uses an ODE-based technique by Wormald for proving concentration for random processes [Wor95]. In algorithmic settings,

these variables often observe parameters within an algorithm on a random object, with the parameters changing from one iteration to the next.

Another situation to which this method has been successfully applied, is the famous k -SAT problem. Here, it has been used to analyze algorithms which, given a uniformly drawn k -SAT formula on n variables and $m = m(n)$ clauses, attempt to find an assignment of values to the variables (*interpretation*) satisfying the formula. Indeed, ODE-based techniques have a tradition in random k -SAT beyond Wormald’s method (see, e.g. [CF86, CF90, FS96, Ach00, AS00]). Historically, at that time, one was not so much interested in algorithms which succeed aas (for $n \rightarrow \infty$), because, by invoking a strong theorem of Friedgut’s [Fri99], proving success with positive probability already allowed to infer aas satisfiability of the formula. (In more recent work, however, algorithms succeeding aas have come into focus again, e.g., [CO10, COF].)

In our manuscript [BT], which forms Chapter 10 of this thesis, we deal with a variant of k -SAT which arises in applications.

Let M be a (usually finite) set, \mathcal{S} a set of subsets of M , and k a positive integer. For the *signed k -satisfiability problem*, or *signed k -SAT*, one is given as input a finite set of variables X and a formula in *signed conjunctive normal form (CNF)*. This means that there is a list of clauses, each of which is a disjunction of *signed literals* of the form $x \in \mathcal{S}$ where x is a variable in X and the “sign” \mathcal{S} is a set in \mathcal{S} . As in classical k -SAT, the question is then whether there exists a satisfying *interpretation*, i.e., an assignment of values to the variables such that each of the clauses is satisfied. This setting includes as a special case the classical SAT problem: choose for M the 2-element set $\{0, 1\}$ and $\mathcal{S} = \{\{0\}, \{1\}\}$.

In case M is a totally ordered set and the set \mathcal{S} is the set of all intervals in M , we speak of *Interval SAT*, or *iSAT*. In our manuscript, we study the case when $M = [0, 1]$.

Our interest in this particular version of signed SAT arises from applications in computational systems biology. There, iSAT yields a generalization of modeling with Boolean networks, where biological systems are represented by logical formulas with variables correspond to biological components like proteins. Reactions are modeled as logical conditions which have to hold simultaneously, and then transferred into CNF.

Although the model is widely used by practitioners, often, this binary approach is not sufficient to model real life behavior or even accommodate all known data. A typical situation is that an experiment yields several “activation levels” of a component. Thus, one wants to make statements of the form: If the quantity of component A reaches a certain threshold but does not exceed another, and component B occurs in sufficient quantity, then another component C is in a certain frame of activation levels. The collection of such rules accurately models the global behavior of the system.

On the theoretical side, signed SAT originated in the area of so-called multi-valued logic, where variables can take a (usually finite) number of so-called *truth values*, not just TRUE or FALSE. The motivation for studying signed formulas was to be able to better cover practical applications. Most applications and a great deal of the earlier complexity results focus on *regular* signed SAT, where M is a totally ordered set, and the signs may only be of the form $S = \{j \mid j \geq i\}$ or $S = \{j \mid j \leq i\}$. For regular signed SAT, random formulas have been investigated computationally. Manyà et al. [MBE198] study uniformly generated random regular 3-SAT instances, and observe a phase transition similar to that observed in classical SAT. Moreover, in [BM99a, BMC⁺07] a bound on the ratio m/n is given, beyond which a random formula is aas unsatisfiable. To my knowledge, however, ours is the first rigorous analysis of an algorithm for random signed SAT.

In our paper, we present and analyze an algorithm which solves uniformly random 3-iSAT instances with high probability, provided that the ratio between the number m of clauses and the number n of variables is at most 2.3. Our algorithm is an adaption of the well-known Unit

Clause algorithm from classical SAT, where, in an inner loop, 1-clauses are treated if any exist, and in an outer loop, a variable is chosen freely and assigned some value. This Unit Clause approach is enhanced with a “backtracking” subroutine, which is not completely unlike the one used in [FS96] for classical 3-SAT, see Algorithm 5 on page 127.

For the analysis of the outer loop, we use the Wormald’s ODE-method mentioned above, and the value 2.3 arises from the numerical solution to an initial value problem. The analysis of the inner loop requires to study the first busy period of a certain stable server system, or, in our case, more accurately, the total number of individuals in a type of branching process.

Discussion. The way I see it, the manuscript [BT] highlights some of the specific problems of random iSAT. To understand the algorithm, first of all, it is important to realize that the bottleneck lies in the rate at which 2-clauses become 1-clauses (by deleting literals whose variable has been set in such a way that the literal is not satisfied) during the run of the inner loop. In the branching process terminology, this amounts to the number of offspring of one individual.

If the corresponding algorithm is analyzed for classical 3-SAT, whenever a variable is set to some value in the inner loop, the probability that a fixed literal containing this variable is not satisfied by the chosen value, is $1/2$, regardless of the chosen value. In 3-iSAT, this probability depends on the value. Thus, for choosing a value for a variable in the inner loop, there are two possibilities.

- (1) By looking only at the 1-clause, choose the best value possible.
- (2) By looking at the 1-clause and all literals in 2- or even 3-clauses containing the variable, choose a value which satisfies a large fraction of them.

Possibility (2) requires to find, for a Poisson random variable R the expectation of the random variable $X(R)$, where $X(r)$ is the (cardinality of the) largest subset K of $\{1, \dots, r+1\}$ such that $I_0 \cap \bigcap_{k \in K} I_k \neq \emptyset$, for random intervals I_0, \dots, I_{r+1} . This expectation, as a function of the mean of R , then forms one key term in the system of ODEs, with the mean of R being a quotient of two parameters.

Asymptotically, for $r \rightarrow \infty$, deciding only based on I_0 $X(r)$ is optimal. However, for small values of r , the difference between (1) and (2) can be large, e.g., for $r = 1$, the mean for (1) is $1^{11}/24$, that for (2) is $12/3$.

In our manuscript, we have decided for the much simpler but also much worse possibility (1), until we knew how to deal with computing the mean in (2). An alternative would have been to use “cheap tricks”: For, say $r \leq 3$, the computation of $\mathbf{E} X(r)$ can be done by hand. Since $\mathbf{E} R < 3$ most of the time, this would have recovered a significant part of the gap, at the expense of adding some lengthy computations and making the ODE more complicated.

Apart from this central issue, it would also be interesting to find a bound for the ratio above which random 3-iSAT formulas are as not satisfiable. Moreover, our analysis of 2-iSAT (to which our 3-iSAT algorithm reduces its instances) is quite superficial, and could be much improved.

2.5. Cops & Robber

The game of *Cops and Robber* is played on a connected graph by two players — the cops and the robber. The cop player has at his disposal k pieces (the “cops”), for some integer $k \geq 1$, and the robber player has one piece (the “robber”). The pieces will always be on vertices of the graph. We will usually speak informally of “the cops” instead of the “cop player”, and “the robber” instead of the “robber player”.

The game begins with the cops positioning themselves (i.e., placing their k pieces) on (not necessarily distinct) starting vertices. Next, the robber chooses his starting vertex. Now, starting with the cop player, the two players move their pieces alternately. In the cops’ move, they decides for each of them whether he stands still or moves to an adjacent vertex. In the

robber's move, he can choose to move to an adjacent vertex, or to pass. The game ends when a cop and the robber are on the same vertex (that is, the cops catch the robber); in this case the cops win. The robber wins if he is never caught by the cops, i.e., the game continues forever. Both players have complete information, i.e., they know the graph and the positions of all the pieces.

A *winning strategy* for a player is one by following which the player wins, regardless of the moves of the other player. It follows from standard arguments in Game Theory that one of the two players always has a winning strategy (cf. [BI93]).

The key problem in this game is to know how many cops are needed to catch a robber on a given graph. For a connected graph G , the smallest integer k such that with k cops, the cops have a winning strategy, is called the *cop number* of G and is denoted by $\text{cop}(G)$. The cop number of a non-connected graph is the maximum cop number of its connected components.

Nowakowski and Winkler [NW83] and Quilliot [Qui78] characterized the class of graphs with cop number 1. Finding a combinatorial characterization of graphs with cop number k (for $k \geq 2$) is a major open problem in the field, to which Clarke and MacGillivray [CM11] have recently made an important contribution. On the other hand, algorithmic characterizations of such graphs, which are polynomial in the size of the graph but not in k , do exist [BI93, GR95, HM06]. However, determining the cop number of a graph is a computationally hard problem [FGK08].

I would like to make the reader aware of the new book by Bonato and Nowakowski [BN11] on Cops & Robber on graphs. The Cops & Robber game belongs to a larger class of search problems on graphs (cf. [FT08])¹.

2.5.1. Cops & Robber on non-orientable surfaces. By surface, we mean a closed surface, i.e. a compact two dimensional topological manifold without boundary. For any non-negative integer g , we denote by $\text{cop}(g)$ the supremum over all $\text{cop}(G)$, with G ranging over all graphs embeddable in an orientable surface of genus g , and we call this the cop number of the surface. Similarly, we define the cop number $\widetilde{\text{cop}}(g)$ of a non-orientable surface of genus g to be the supremum over all $\text{cop}(G)$, with G ranging over all graphs embeddable in this surface.

Aigner and Fromme [AF84] proved that the cop number of the sphere is equal to three: $\text{cop}(0) = 3$. Quilliot [Qui85] gave an inductive argument to the effect that the cop number of an orientable surface of genus g is at most $2g + 3$. Schröder [Sch01] was able to sharpen this result to $\text{cop}(g) \leq \frac{3}{2}g + 3$. He also proved that the cop number of the double torus is at most 5.

Generalizing the work of Aigner and Fromme, Andreae [And86] proved that, for any graph H satisfying a mild connectivity assumption, the class of graphs which do not contain H as a minor has cop number bounded by a constant depending on H . Using this, and the well-known formula for the non-orientable genus of a complete graph, he obtained an upper bound for the cop number of a non-orientable surface of genus g , namely

$$\widetilde{\text{cop}}(g) \leq \left(\frac{\lfloor 7/2 + \sqrt{6g + 1/4} \rfloor}{2} \right).$$

Nowakowski and Schröder [NS] use a series of technically challenging arguments to prove a much stronger bound: $\widetilde{\text{cop}}(g) \leq 2g + 1$.

In our note [CFJT], which forms Chapter 11 of this thesis, we prove the following.

Theorem (Theorem 11.1 on page 152). *For every positive integer g , $\text{cop}(\lfloor g/2 \rfloor) \leq \widetilde{\text{cop}}(g) \leq \text{cop}(g - 1)$.*

The proof uses of the following tool: If \tilde{G} is a lift of G then $\text{cop}(G) \leq \text{cop}(\tilde{G})$. We have made considerable effort to use quite sophisticated generalization of this tool to obtain

¹Some of them actually have real applications.

better bounds for the cop number of orientable surfaces, but were unable to overcome a lattice-geometric question concerning the homology classes of the cycles of the graph.

2.5.2. Cops & Robber and forbidden (induced) subgraphs. As mentioned above, excluding a minor forces bounded cop number. In our paper [JKT10], which is reprinted as Chapter 12 of this thesis, we studied the corresponding question for the subgraph and induced subgraph relations. The results we obtained are the following.

Theorem 2.6 (Theorem 12.1 on page 156). *The class of H -free graphs has bounded cop number if, and only if, every connected component of H is a path.*

Here, a graph is H -free, if it contains no induced subgraph isomorphic to H . The cop number of a graph not containing an induced path of length $\ell \geq 2$ is at most $\ell - 1$ (Prop. 12.2). Similarly, every graph with no induced cycle of length at least $\ell \geq 3$ has cop number at most $\ell - 2$ (Prop. 12.3).

Let us say that a graph is H -subgraph-free, if it contains no subgraph isomorphic to H .

Theorem 2.7 (Theorem 12.4 on page 156). *The class of H -subgraph-free graphs has bounded cop number if, and only if, every connected component of H is a tree with at most three leaves.*

As an intermediate step towards Theorem 12.4, we study how the cop number of a graph G is related to its tree-width, and obtain that the cop number of a graph G is at most one plus half its tree-width (Prop. 12.5).

Purportedly, people fall in two groups depending on whether, when they first learn about the Cops & Robber game, they identify with the Cops, or with the Robber. I must admit that I am in the ‘Cops’ group. That may be the reason why I find the proof for upper bounds for the cop number in Chapter 12 especially appealing. Thus, I would like to point the reader to cops strategies used in the proofs of Prop. 12.5, and, particularly, Prop. 12.2. For the non-boundedness statements of the two theorems, robber strategies are given.

Outlook on Cops & Robber problems

There are several open problems in the area of Cops & Robber. I would like to mention my three favorite ones. The first two are about graphs on surfaces, the first one is directly related to Theorem 11.1 (see above):

Question 2.8. *Is the cop number of a non-orientable surface of genus $2g$ equal to that of the orientable surface of genus g ? In other words, is it true that, for every non-negative integer g , we have $\text{cop}(g) = \text{cop}(\lfloor g/2 \rfloor)$?*

The second question reflects the fact that the lower bound for the cop number of an orientable surface of genus g is $\Theta(g^{1/4})$, which is far away from the $O(g)$ upper bounds.

Conjecture 2.9. *The cop number of orientable surfaces is $o(g)$, where g is the genus.*

Finally, a more structural question. It has been observed [BI93] that the class of k -copwin graphs, i.e., the graphs with cop number at most k , are closed under taking retracts. (For the definition of a retract, one assumes that every vertex has a tiny loop attached to it. A *retraction* is then a homomorphism $r: G \rightarrow G$ with $r^2 = r$; we say that $r(G)$ is a *retract* of G .) Thus, the k -copwin graphs can be characterized by giving a set of forbidden retracts. However, one graph being a retract of another is a very strong condition (considerably stronger than induced subgraph), so for small values of k , this set is likely to be enormous. In fact, it is possible that the only set, for which proving bounded cop number is feasible, consists of essentially all (isomorphism classes of) not-copwin- k graphs (possibly after applying some simple reduction operations to discard some redundant ones).

Question 2.10. *Can $O(1)$ -copwin graphs be defined by forbidden retracts in a meaningful way?*

Part 2

Reprints & Manuscripts

The relationship between the GTSP, STSP, and Metric Cone

Abstract. In this short communication, we observe that the Graphical Traveling Salesman Polyhedron is the intersection of the positive orthant with the Minkowski sum of the Symmetric Traveling Salesman Polytope and the polar of the metric cone. This follows almost trivially from known facts. There are two reasons why we find this observation worth communicating none-the-less: It is very surprising; it helps to understand the relationship between these two important families of polyhedra.

3.1. Introduction

The *Symmetric Traveling Salesman Polytope* is the convex hull of all characteristic vectors of edge sets of cycles (i.e., circuits) on the vertex set $V_n := \{1, \dots, n\}$ (in other words, Hamiltonian cycles in the complete graph with vertex set V_n). For the formal definition, denote by E the set of all two-element subsets of V_n . This is the set of all possible edges of a graph with vertex set V_n . The Symmetric Traveling Salesman Polytope is then the following set:

$$S_n := \text{conv} \left\{ \chi^C \mid C \text{ is the edge set of a Hamiltonian cycle with vertex set } V_n \right\} \subset \mathbb{R}^E.$$

Here, for an edge set F , χ^F is the characteristic vector in \mathbb{R}^E with $\chi_e^F = 1$ if $e \in F$, and zero otherwise. The importance of the Symmetric Traveling Salesman Polytope comes mainly, but not exclusively, from its use in the solution of the so-called Symmetric Traveling Salesman Problem, which consists in finding a Hamiltonian cycle of minimum cost.

The *Graphical Traveling Salesman Polyhedron* is the convex hull of all characteristic vectors of edge multi-sets of connected Eulerian multi-graphs on the vertex set V_n . A multi-graph with vertex set V_n has as its edge set a sub-multi-set of E , which is to say that our multi-graphs can have parallel edges but no loops. By defining, for any multi-set F of edges of K_n , its characteristic vector $\chi^F \in \mathbb{R}^E$ in such a way that χ_e^F counts the number of occurrences of e in F , the Graphical Traveling Salesman Polyhedron is formally defined as

$$P_n := \text{conv} \left\{ \chi^F \mid F \text{ is the edge multi-set of a connected Eulerian multi-graph} \right. \\ \left. \text{with vertex set } V_n \right\} \subset \mathbb{R}^E.$$

Ever since the seminal work of Naddef & Rinaldi [NR91, NR93] on the two polyhedra, P_n is considered to be an important tool for investigating the facets of S_n . Moreover, in works of Carr [Car04] and Applegate, Bixby, Chvátal & Cook [ABCC01], P_n has been used algorithmically in contributing to solution schemes for the Symmetric Traveling Salesman Problem.

Numerous authors have expressed how close the connection between Graphical and Symmetric Traveling Salesman Polyhedra is. The most basic justification for this opinion is the fact that S_n is a face of P_n — consisting of all points x whose “degree” is two at every vertex: $\sum_{v \neq u} x_{uv} = 2$ for all $u \in V_n$. However, the connections are far deeper (see [Nad02] or [ORT07] and the references therein). In this short communication, we contribute the following surprising geometric observation to the issue of the relationship between these two polyhedra:

Theorem 3.1. P_n is the intersection of the positive orthant with the Minkowski sum of S_n and the polar C_n^Δ of the metric cone C_n :

$$(1) \quad P_n = (S_n + C_n^\Delta) \cap \mathbb{R}_+^E$$

The metric cone consists of all $a \in \mathbb{R}^E$ which satisfy the triangle inequality:

$$(2) \quad a_{uv} \leq a_{uw} + a_{wv}$$

for all pairwise distinct vertices $u, v, w \in V_n$. Consequently, its polar is generated as a cone by the vectors (we abbreviate $\chi^{\{e\}}$ to χ^e)

$$(3) \quad \chi^{uw} + \chi^{wv} - \chi^{uv}.$$

The proof of this theorem is an application of three or four known facts or techniques in the area of Symmetric and Graphical Traveling Salesman polyhedra.

3.2. Proof

We start with showing that $P_n \subset (S_n + C_n^\Delta) \cap \mathbb{R}_+^E$. While $P_n \subset \mathbb{R}_+^E$ holds trivially, $P_n \subset S_n + C_n^\Delta$ follows from an argument of [NR93], which we reproduce here for the sake of completeness.

Let $x \in \mathbb{Z}_+^E$ be the characteristic vector of the edge multi-set of a connected Eulerian multi-graph G with vertex set V_n . We prove by induction on the number m of edges of G , that x can be written as a sum of a cycle and a number of vectors (3). If $m = n$, then there is nothing to prove. Let $m \geq n+1$. There exists a vertex w of degree at least four in G . We distinguish two cases. The easy case occurs when $G \setminus w$ is still connected. Here, we let u and v be two arbitrary (possibly identical) neighbors of w . By either replacing the edges uw and wv of G with the new edge uv , if $u \neq v$, or deleting uw and wv , if $u = v$, one obtains a connected Eulerian multi-graph G' with fewer edges than G . The change in the vector x amounts to subtracting the expression (3): $x' = x - (\chi^{uw} + \chi^{wv} - \chi^{uv})$, if $u \neq v$, and $x' = x - (\chi^{uw} + \chi^{wv})$, if $u = v$. In the slightly more difficult case when the graph $G \setminus w$ has at least two connected components, we can let u and v be two neighbors of w in distinct components of $G \setminus w$. This makes sure that the graph G' is still connected. We conclude by induction that x' , and hence x , can be written as a sum of a cycle and a number of vectors (3).

We now prove $P_n \supset (S_n + C_n^\Delta) \cap \mathbb{R}_+^E$. For this, we show that any inequality which is facet-defining for P_n is valid for $(S_n + C_n^\Delta) \cap \mathbb{R}_+^E$.

We again invoke an argument from [NR93]: Naddef & Rinaldi have shown¹ that the inequalities defining facets of P_n fall into one of two categories: the non-negativity inequalities $x_e \geq 0$, with $e \in E$ (or positive scalar multiples thereof), or inequalities whose coefficient vectors satisfy the triangle inequality (2). We reproduce the proof of this statement.

First recall that an inequality $a \bullet x \geq \alpha$ is said to be *dominated* by another inequality $b \bullet x \geq \beta$, if the face defined by the first inequality is contained in the face defined by the second inequality.

Suppose that $a \bullet x \geq \alpha$ is not dominated by a non-negativity inequality (it need not be define a facet, though), and let u, v, w be three distinct vertices in V_n . Then there exists an $x \in \mathbb{Z}_+^E$ defining the edge multi-set of a connected Eulerian multi-graph G which has an edge between u and v , such that $a \bullet x = \alpha$. If we replace the edge uv of G by the two edges uw and wv , then we obtain a connected Eulerian multi-graph, whose edge multi-set is given, in terms of its characteristic vector, by $x' := x + \chi^{uw} + \chi^{wv} - \chi^{uv}$. Now $a \bullet x' \geq \alpha$, implies $a_{uw} + a_{wv} - a_{uv} \geq 0$, i.e., the triangle inequality.

¹In fact, Proposition 2.2 of [NR93] states that the facet-defining inequalities for P_n fall into three classes — one of which is the class of non-negativity inequalities and the other two satisfy the triangle inequality.

We now conclude the proof of the inclusion $P_n \supset (S_n + C_n^\Delta) \cap \mathbb{R}_+^E$. Let $a \bullet x \geq \alpha$ be an inequality which is facet-defining for P_n . First note that the non-negativity inequalities are clearly satisfied by the right hand side of (1). Hence, using what we have just discussed, let us assume that a satisfies the triangle inequality. This means that a is a member of the metric cone C_n . Consequently, the inequality $a \bullet x \geq 0$ is valid for C_n^Δ . Further, since $S_n \subset P_n$, the inequality $a \bullet x \geq \alpha$ is clearly valid for S_n . Hence the inequality is valid for $S_n + C_n^\Delta$.

This concludes the proof of the theorem.

Note that, en passant, we have proved the following. If we define P'_n to be the set of all $y \in \mathbb{R}^E$ which satisfy $a \bullet y \geq \alpha$ for every inequality $a \bullet x \geq \alpha$ defining a facet of P_n but not being a scalar multiple of a non-negativity inequality, then we have $S_n + C_n^\Delta \subset P'_n$.

Acknowledgments

Thanks are extended to the *Deutsche Forschungsgemeinschaft*, DFG, for funding this research, and to the *Communauté française de Belgique – Actions de Recherche Concertées* for supporting the author during the time the paper was written.

Facial structure of Symmetric and Graphical Traveling Salesman polyhedra

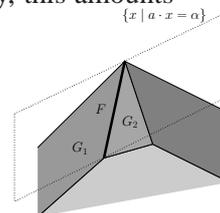
Abstract. The Symmetric Traveling Salesman Polytope S for a fixed number n of cities is a face of the corresponding Graphical Traveling Salesman Polyhedron P . This has been used to study facets of S using P as a tool. In this paper, we study the operation of “rotating” (or “lifting”) valid inequalities for S to obtain a valid inequalities for P .

As an application, we describe a surprising relationship between (a) the parsimonious property of relaxations of the Symmetric Traveling Salesman Polytope and (b) a connectivity property of the ridge graph of the Graphical Traveling Salesman Polyhedron.

4.1. Introduction

Suppose that S and P are polyhedra, and that S is a proper face of P . If $a \bullet x \geq \alpha$ is a valid inequality for S , it can be “rotated” so that it becomes also valid for P . By “rotation” we mean modifying left and right hand sides of the inequality in such a way that the set of points in the affine hull of S which satisfy the inequality with equation remains the same, yet the hyperplane the inequality defines in the ambient space changes. Technically, this amounts to adding equations valid for S to $a \bullet x \geq \alpha$.

Once the inequality is rotated so that it is valid for P , one may ask which face of P is defined by the rotated inequality. Since $S \neq P$, there are always several such faces, but even when we aim for inclusion-wise maximal faces of P defined by some rotated version of $a \bullet x \geq \alpha$, in general, these are not unique either. In the picture to the right, by properly tilting the hyperplane defined by $a \bullet x = \alpha$, we can obtain the inequalities F , F_1 and F_2 .



A prominent example is of course sequential lifting, where S is an intersection of faces defined by non-negativity inequalities $x_j \geq 0$.

Sequential lifting or other rotation-based tools are applied manually to find facets of polyhedra which contain faces which are better understood. Moreover, mechanisms of this kind are used computationally in cutting-plane algorithms where some cutting-plane generation procedure first works on a face and then lifts the obtained inequalities.

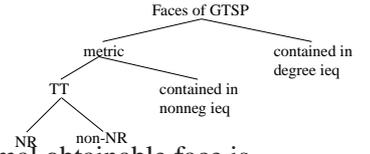
In this paper, we study what rotating inequalities does for the Symmetric Traveling Salesman Polytope S and the Graphical Traveling Salesman Polyhedron P . Let $n \geq 3$ be an integer. Let $V := \{1, \dots, n\}$ and E be the set of all unordered pairs $\{u, v\} \in V$, i.e., the set of edges of the complete graph with vertex set V . The two polyhedra are subsets of the space \mathbb{R}^E of vectors indexed by the elements of E . The Symmetric Traveling Salesman Polytope S is the convex hull of all incidence vectors of edge sets of cycles with vertex set V (or of tours; or of Hamilton cycles of the complete graph K_n). The Graphical Traveling Salesman Polyhedron P is the convex hull of all vectors corresponding to connected Eulerian multi-graphs with vertex set V . (The precise definitions will be given in (5) on page 28 below.)

Ever since the groundbreaking work of Dantzig, Fulkerson, and Johnson on the computational solution of the Traveling Salesman Problem (TSP), the facet-structure of these polytopes has received much attention (e.g., [ABCC06, DFJ54, GP85, JRR95, Nad02, Sch03]).

Their combinatorial and linear-algebraic properties have also been researched. For example, questions about properties of the graph (e.g., its diameter) have been addressed [GP85, Ris98, RC98, Sie98, STT95, ST92].

With few exceptions (for example [FN92, Nor55] for the case $n \leq 5$; [BC91] for $n = 6, 7$; [CJR91, CR96, CR01] for $n = 8, 9$), no complete characterization of the facets of S or P are known. In fact, since the TSP is NP-hard, there cannot exist a polynomial time algorithm producing, for every n and every point $x \in \mathbb{R}^E \setminus S$, a hyperplane separating x from S , unless $P=NP$. Another noteworthy argument for the complexity of these polytopes is a result of Billera & Sarangarajan [BS96]: For every 0/1-polytope P , there exists an n such that P is affinely isomorphic to a face of S .

Since the seminal work of Naddef & Rinaldi [NR91, NR93] on the Graphical and Symmetric Traveling Salesman polyhedra, it is known that S is a face of P . Moreover, Naddef & Rinaldi proved a theorem saying that, if an inequality defines a facet of S , then there is a unique maximal face of P which can be obtained by rotating the inequality, and this maximal obtainable face is a facet of P .



Naddef & Rinaldi classified the facets of P into three types — non-negativity facets, degree facets, and the rest, called TT-facets — based on properties of the coefficients. While the degree facets and non-negativity facets are both small in number and easily understood, the interesting class both for understanding the polyhedron and for applications is the huge set of TT-facets. By the theorem just mentioned, once one knows that the degree facets of P are precisely those which contain S — also proved in Naddef & Rinaldi’s paper —, this also classifies the facets of S into two types: non-negativity and TT-facets.

Oswald, Reinelt and Theis [ORT05, ORT07] have refined the classification by splitting the TT-facets of P into two subclasses: NR-facets and non-NR-facets, depending on whether the intersection of the facet with S is a facet of S (these P -facets are called NR-facets) or a face of S of smaller dimension, the main result being the fact that the non-NR class is not empty. The existence of non-NR-facets has unpleasant consequences both for theoretical research and practical computational approaches to solving TSP instances. On the theoretical side, it is much easier to prove facet-defining property of inequalities for P than for S . Moreover, P pleasantly preserves facet-defining property when a certain important lifting operation for facet-defining inequalities (which replaces vertices by sets of vertices) is performed. For S , this is not known to be true. On the computational side, in the context of cutting-plane methods for S , certain generic separation algorithms produce inequalities which are facet-defining for P , but sometimes it is not clear whether these inequalities must be strengthened if they are to define facets of S . Examples of such separation algorithms include the local cuts method of Applegate, Bixby, Chvátal & Cook [ABCC01, ABCC03, ABCC06] (see the discussion in [ORT07]) or the path-lifting method of Carr [Car04].

In terms of rotation, the result in [ORT05, ORT07] shows that there are valid inequalities for S which do not define facets of S , but which can be rotated to define facets of P . The starting point of the present paper is the question what properties these valid inequalities for S might have. The results we propose are most easily formulated using the terminology of polar polyhedra. A polar polyhedron S^Δ of a polyhedron S has the property that the points of S^Δ are in bijection with the linear inequalities (up to scaling) for S . Moreover, a point a is contained in a face of dimension k of S^Δ , if, and only if, the corresponding inequality defines a face of dimension at least $\dim S + 1 - k$ of S . In particular, the vertices of S^Δ are in bijection with the facets of S . Also recall the concept of a polyhedral complex: a finite set of polyhedra, closed under taking faces, such that the intersection of any two polyhedra in the set is a face of both.

We give results about the “important” part of the polar of S , namely the part which remains if we delete the vertices corresponding a non-negativity facets. This corresponds to taking only the “TT-class” of valid inequalities for S ; the details are made precise below (Section 4.2).

This subset of faces of the polar of S is a polyhedral complex, which we denote by \mathcal{A} . Take a point in \mathcal{A} , consider the corresponding valid inequality for S , and rotate it in all possible ways yielding inequalities valid for P . A certain set of faces of P can be defined by the rotated versions of this inequality. Now we partition the points contained in \mathcal{A} in the following way: two points are in the same cell of the partition, if, by rotating the corresponding valid inequalities, the two sets of faces of P which can be defined coincide.

In fact, the partition whose definition we have just outlined, gives a polyhedral subdivision \mathcal{S} of \mathcal{A} , i.e., the set of closures of the cells is a polyhedral complex, and every face of \mathcal{A} is a disjoint union of cells. This is true in the general situation when a polytope S is a face of another polytope P . Indeed, \mathcal{S} is known as the *chamber complex* of the canonical projection of the polar of P onto the polar of S . We call such a polyhedral subdivision a *rotation complex*. We give the following results about the rotation complex in the TSP situation:

- (A) The decomposition of \mathcal{A} into cells can be described in a natural way that does not refer to rotation; moreover, it does not refer to any Graphical Traveling Salesman concepts whatsoever. Indeed, to describe the subdivision, for a point a contained in \mathcal{A} , it suffices to check the order relations of the expressions $a_{uv} - a_{uw} - a_{vw}$, with u, v, w three distinct vertices in V . (As customary, we use the abbreviated notation $uv := \{u, v\}$.)
- (B) The points in \mathcal{A} are in *bijection(!)* with the “important” part of the polar of P (the definition of polar here is not canonical and will be made precise), and this bijection maps faces of the polar of P onto faces of the rotation complex \mathcal{S} . In other words, the polar of P can be “flattened” onto the polar of S , see Fig. 4.1, right.

Again, “important” is meant to be understood in the sense that it corresponds to considering TT-type inequalities only. Item (B) is not a consequence of known facts about the chamber complex (injectivity fails to hold in general).

The picture in Fig. 4.1, left, illustrates Item (A). It shows a hypothetical drawing of \mathcal{A} (solid lines) with two points a, a' . To decide whether these two points, when viewed as valid inequalities for S , yield the same faces of P when rotated, one has to check the expressions $a_{uv} - a_{uw} - a_{vw}$. This amounts to checking if they are contained in the same cone in the picture (dotted lines). Indeed, Item (A) can be restated as saying that the rotation complex \mathcal{S} is the common refinement of \mathcal{A} with a projection of a natural sub-complex of the boundary complex of the metric cone. (The common refinement of two polyhedral complexes is the set of all intersections of polyhedra in the two complexes, and the metric cone consists of all functions $E \rightarrow \mathbb{R}_+$ satisfying the triangle inequality). The occurrence of the metric cone in the context of the two polyhedra S and P is no surprise: It is known [The10] that the metric cone plays a rôle in the relationship between the polyhedra S and P . One can construct P by gluing together S and the dual of the metric cone, and then cutting off the waste (see [The10] for a rigorous statement). Item (B) addresses the uniqueness question for faces defined by rotated inequalities addressed above (second paragraph of the introduction). Note, though, that having a point-wise bijection is a stronger statement than saying that the maximal faces obtainable by rotation are unique.

We apply these results to a problem concerning the ridge graph of P . The ridge graph has as its vertices the facets, and two facets are linked by an edge if and only if their intersection is a ridge, i.e., a face of dimension $\dim P - 2$. The ridge graph is of importance for the problem of computing a complete system of facet-defining inequalities, when the points and extreme rays are given. A common solution here is to search in the ridge graph, i.e., once a facet is found, its neighbors are computed. A problem which may occur is that, for some facets, computing the neighbors is not computationally feasible. Due to the connectivity of the ridge graph, some of its vertices can be skipped in the search, and still all vertices are reached. For example, when the facets of a d -dimensional polytope are computed in this way, by Balinski’s Theorem, one

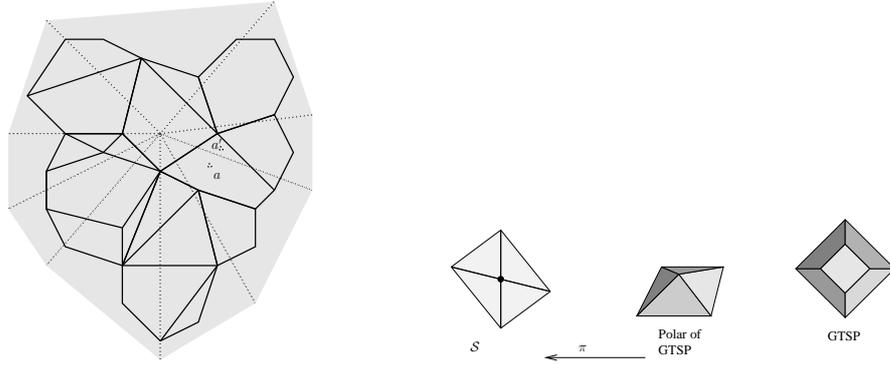


FIGURE 4.1. Left: S as a common refinement of \mathcal{A} and \mathcal{T} , a projection of a sub-complex of the boundary complex of the metric cone; Right: “Flattening” of parts of the polar of P onto the polar of S .

may skip $d - 1$ arbitrarily selected facets. Very often, however, the number of facets whose neighbors cannot be computed is too large (exponential in the dimension). Thus, one would like to prove connectivity properties of the ridge graph which allow for these vertices to be dead ends in the search.

We prove the following. If a system of NR-facet-defining inequalities satisfies the so-called parsimonious property [GB93, Goe95], the removal of the corresponding vertices from the ridge graph leaves connected components, each of which contains a vertex corresponding to an NR-facet. The proof of this makes use of (B) above in an essential way. In deed, by pressing the boundary of P “flat” onto a lower dimension, one can use linear algebra arguments, which cannot be used when the boundary is molded around P in the higher dimension.

The statement has been used in a computation proof of the completeness of an outer description for the Graphical Traveling Salesman Polyhedron P in the case $n = 9$ in [ORT07] in the scenario sketched above.

This paper is organized as follows. In the second section, we define some basic concepts from polyhedral theory. Section 4.2 provides rigorous formulations of our results. Section 4.3 contains the proofs of the results about the rotation complex, while the results about the ridge graph are proved in Section 4.4.

4.2. Exposition of results

We refer the reader to [Grü03] and [Zie98] for background material on polyhedra, polarity, projective transformations, and polyhedral complexes. For a polyhedron P , let $\mathcal{C}(P)$ be the set of all of its faces. This is a polyhedral complex with underlying point set P .

Fix an integer $n \geq 3$. The *Symmetric Traveling Salesman Polytope* is defined as the convex hull in \mathbb{R}^E of all edge sets of cycles with vertex set V (or Hamiltonian cycles in the complete graph K_n):

$$(4) \quad S := \text{conv}\{\chi^{E(C)} \mid C \text{ is the cycle with } V(C) = V\},$$

where χ^F denotes the characteristic vector of a set F , i.e., $\chi_e^F = 1$, if $e \in F$, and 0 otherwise.

The second polyhedron which we will consider is defined to be the convex hull of all edge multi-sets of connected Eulerian multi-graphs on the vertex set V :

$$(5) \quad P := \text{conv}\{x \in \mathbb{Z}_+^E \mid x \text{ defines a connected Eulerian multi-graph with vertex set } V\},$$

where we identify sub-multi-sets of E with vectors in \mathbb{Z}_+^E (i.e., there are x_e copies of edge e present in the multi-graph). This polyhedron was introduced in [CFN85] under the name of

Graphical Traveling Salesman Polyhedron and has since frequently occurred in the literature on Traveling Salesman Polyhedra. It is particularly important in the study of properties, mainly facets, of Symmetric Traveling Salesman Polytopes (e.g., [Goe95, NP01, NR91, NR93, NR07], see [ABCC06, Nad02] for further references).

The polyhedron P has been called the *Graphical Relaxation* of S by Naddef & Rinaldi [NR91, NR93] who discovered and made use of the fact that S is a face of P : While the latter is a full-dimensional unbounded polyhedron in \mathbb{R}^E [CFN85], the former is a polytope of dimension $\binom{n}{2} - n$ [Nor55], and the inequality $\sum_{e \in E} x_e \geq n$ is valid for P and satisfied with equality only by cycles, thus attesting to the face relation.

4.2.1. Definitions of the polars. We denote by $x \bullet y$ the standard scalar product in Euclidean space.

The set of facets of P containing S is known. For $u \in V$, let δ_u be the point in \mathbb{R}^E which is $1/2$ on all edges incident to u and zero otherwise. It is proven in [CFN85] that the inequalities $\delta_u \bullet x \geq 1$, $u \in V$, define facets of P , the so-called *degree facets*. Clearly, S is the intersection of all the degree facets (because summing all degree inequalities gives $\sum_e x_e \geq n$).

It is customary to write inequalities valid for P in the form $a \bullet x \geq \alpha$, and we define the polars accordingly. Define the linear space L to be the set of solutions to the n linear equations $\delta_u \bullet x = 0$, $u \in V$. Note that the δ_u are linearly independent, $\dim S = \dim L$, and the affine hull of S is a translated copy of L . Whenever z is a relative interior point of S , the polar of S may be defined as the following set:

$$(6) \quad S^\Delta := \{a \in L \mid (-a) \bullet (x - z) \leq 1 \forall x \in S\}.$$

(The *relative interior* $\text{relint } P$ of a polyhedron P is the interior (in the topological sense) of P in the affine space spanned by P , in other words, $\text{relint } P = P \setminus \bigcup_{F \subsetneq P} F$, where the union runs over all faces of P .) So a point $a \in S^\Delta$ corresponds to a valid inequality $a \bullet x \geq a \bullet z - 1$ of S . Changing z amounts to submitting S^Δ to a projective transformation. Although it can be seen that our results do not depend on the choice of z , it makes things easier to define

$$(7) \quad z := \frac{2}{n-1} \mathbf{1} = \frac{1}{(n-1)!/2} \sum_C \chi^{E(C)} = \frac{2}{n-1} \sum_{u=1}^n \delta_u,$$

where the first sum extends over all cycles with vertex set V . So z is at the same time the average of the vertices $\chi^{E(C)}$ of S and a weighted sum of the left-hand sides δ_u of the equations.

Next, we construct a kind of polar for P . For this, we will use the blocking polyhedron construction, which is well-known in polyhedral combinatorics. Goemans [Goe95] has been observed that P is of so-called *blocking type*, i.e., it is the Minkowski sum of \mathbb{R}_+^E with the convex hull of a finite set of points in \mathbb{R}_+^E . Thus we define

$$P^\Delta := \{a \in \mathbb{R}^m \mid a \bullet x \geq 1 \forall x \in P\}.$$

This set is sometimes called the *blocking polyhedron* of P . Note that $P^\Delta \subset \mathbb{R}_+^E$ (see [CFN85]; this is well-known to be true for the blocking polyhedron of any blocking type polyhedra). Other known facts about blocking type polyhedra and their blocking polyhedra includes the fact that P^Δ is also of blocking type. In particular, the extreme rays of both P and P^Δ are the positive coordinate directions.

Calling P^Δ the *polar (polyhedron)* of P is justified by that fact that, essentially, it has the defining properties of a polar polytope. Let us elaborate. For any polytope Q containing 0 as an interior point, there is a mapping assigning to every face the face of its polar consisting of all points corresponding to inequalities which are satisfied with equality by all points of F . This mapping is an inclusion reversing bijection. In the case of blocking type polyhedra and their blocking polyhedra, something similar holds. The following definitions and lemma will make this clear. For a face F of P , define its *conjugate face* F^\diamond to be the set of points $a \in P^\Delta$

satisfying $a \bullet x = 1$ for every $x \in F$. The following lemma will establish that the mapping $F \mapsto F^\diamond$ has the properties. For brevity, we say that a face F of P is *good* if it is not contained in a *non-negativity facet*, i.e., a facet defined by $x_e \geq 0$ (these inequalities do define facets of P [CFN85]). The non-negativity inequalities are valid for P^Δ , and hence P^Δ has (possibly empty) non-negativity faces, too. As is customary, the *co-dimension* of a face F of a polyhedron $Q \subset \mathbb{R}^m$ is $\dim Q - \dim F$.

Lemma 4.1. *The polar P^Δ of P has the following properties.*

- (a) *Let $a \in \mathbb{R}^E \setminus \{0\}$ and $d \geq -1$. Then a is a relative interior point of a non-trivial face of P^Δ with co-dimension $d + 1$ if and only if the inequality $a \bullet x \geq 1$ is valid for P and defines a face of dimension d of P .*
- (b) *Let $\mathcal{N} \subset \mathcal{C}(P)$ be the set of intersections of non-negativity facets of P (with $\emptyset, P \in \mathcal{N}$), and similarly $\mathcal{N}' \subset \mathcal{C}(P^\Delta)$ be the set of all intersections of non-negativity faces of P^Δ . Then conjugation of faces $\mathcal{C}(P) \setminus \mathcal{N} \rightarrow \mathcal{C}(P^\Delta) \setminus \mathcal{N}'$, $F \mapsto F^\diamond := \{a \in P^\Delta \mid a \bullet x = 1 \forall x \in F\}$ defines an inclusion reversing bijection.*
- (c) *A face F of P is good if and only if F^\diamond is bounded. □*

We leave the proof of this lemma to the reader.

Thus we see that blocking polyhedra behave like polar polytopes, except that the non-negativity faces are set apart.

Since we will construct projective mappings between parts of the polar polyhedra, for our results, the realization of the polar as a concrete polyhedron plays a great role, not just the properties of its face lattice. Thus, a word is in order why we chose the blocking polyhedron “without loss of generality”. Clearly, the most natural definition of a polar would be to intersect the polar cone $\{(\alpha, a) \in \mathbb{R} \times \mathbb{R}^E \mid a \bullet x \geq \alpha \forall x \in P\}$ with the hyperplane $\alpha + \sum_e a_e = 1$. (From the above mentioned fact that P is the Minkowski sum of \mathbb{R}_+^E with the convex hull of a finite set of points in \mathbb{R}_+^E , we see that this hyperplane intersects all extreme rays of the polar cone except for $\mathbb{R}_+(\alpha, 0)$ which does not correspond to a facet of P .) However, a moments thought will convince the reader that the result is projectively isomorphic to the blocking polyhedron.

The points δ_u defined above are vertices of P^Δ , more precisely, they are the vertices of the face S^\diamond of P^Δ .

4.2.2. Definitions of the polyhedral complexes. A *polyhedral complex* is a finite set of polyhedra \mathcal{C} with the properties that (a) if $F \in \mathcal{C}$ and G is a face of F , then $G \in \mathcal{C}$; and (b) if $F, G \in \mathcal{C}$, then $F \cap G$ is a face of both F and G . The polyhedra in \mathcal{C} are called the faces of \mathcal{C} , and faces of a \mathcal{C} having dimension 0 (or 1, respectively) are called vertices (or edges, respectively) of \mathcal{C} . A *sub-complex* of a polyhedral complex \mathcal{C} is a polyhedral complex \mathcal{D} with $\mathcal{D} \subset \mathcal{C}$.

We consider the set of faces of S^Δ which do not contain a vertex corresponding to a non-negativity facet of S (as for P , a non-negativity facet of S is one defined by an inequality $x_e \geq 0$ for some $e \in E$). In symbols, if N denotes the set of vertices of S^Δ corresponding to a non-negativity facet of S , we deal with the polyhedral complex

$$(8) \quad \mathcal{A} := \text{dl}(N, S^\Delta) := \{F \text{ face of } S^\Delta \mid F \cap N = \emptyset\} = \text{dl}(\{\{x\} \mid x \in N\}, \mathcal{C}(S^\Delta))$$

where $\text{dl}(N, S^\Delta)$ is a slight abuse of notation: For a polyhedral complex \mathcal{C} and a set of faces $\mathcal{D} \subset \mathcal{C}$, we define the *deletion of \mathcal{D} in \mathcal{C}* to be the polyhedral sub-complex of \mathcal{C} consisting of all faces $F \in \mathcal{C}$ whose intersection with all faces in \mathcal{D} is empty, i.e.,

$$\text{dl}(\mathcal{D}, \mathcal{C}) := \{F \in \mathcal{C} \mid \forall G \in \mathcal{D}: F \cap G = \emptyset\}$$

4.2.2.1. Tight triangularity. Let $a \in \mathbb{R}^E$. We say that a is *metric*,¹ if it satisfies the triangle inequality, i.e., $a_{vu} + a_{uw} - a_{vw} \geq 0$ for all three distinct u, v, w . (As is customary for graphs,

¹Note that this implies $a_e \geq 0$ for all e .

we abbreviate $\{v, w\}$ with $v \neq w$ to vw .) We follow [NR93] in calling a *tight triangular (TT)*, if it is metric and for each $u \in V$ there exists vw such that the triangle inequality is satisfied with equation: $a_{vu} + a_{uw} - a_{vw} = 0$. Abusively, we say that a linear inequality is metric, or TT, if the left hand side vector has the property.

4.2.2.2. *Metric cone, TT-fan and flat TT-fan.* A polyhedral complex is a (pointed) *fan* if it contains precisely one vertex, and each face which is not a vertex is empty or a pointed cone.

The *metric cone*, C , consists of all (semi-)metrics on V . In our context, a (semi-)metric is a metric point $d \in \mathbb{R}^E$, i.e.,

$$(9) \quad d_{vu} + d_{uw} - d_{vw} \geq 0$$

holds for all distinct $u, v, w \in V$.

For a polyhedral complex \mathcal{C} , we denote by $|\mathcal{C}| := \bigcup_{F \in \mathcal{C}} F$ its *underlying point set*, and, informally, we say that a point x is in \mathcal{C} , if $x \in |\mathcal{C}|$.

We now define the *TT-fan* \mathcal{T}' , which is a sub-fan of the fan of all faces of the metric cone. Heuristically, the elements of $|\mathcal{T}'|$ are metrics on V satisfying the following: for every point $u \in V$, there exist two other points $v, w \in V$ such that u is the “middle point” of the “line segment” between v and w . More accurately, letting $F_{u,vw}$ denote the face of C defined by inequality (9) the TT-fan is defined as follows:

$$(10) \quad \mathcal{T}' := \bigcap_{u \in V} \bigcup_{v, w \neq u} \mathcal{C}(F_{u,vw}) \subset \mathcal{C}(C).$$

\mathcal{T}' is indeed a fan. “TT” stands for “tight triangular”, a term coined by Naddef & Rinaldi [NR93] for a point’s property of being in $|\mathcal{T}'|$. However, we are not aware of any reference to this fan in the literature. Denote by $p: \mathbb{R}^E \rightarrow L$ the orthogonal projection. We will prove in the next section (Lemma 4.10) that applying p to \mathcal{T}' produces a fan \mathcal{T} isomorphic to \mathcal{T}' :

$$(11) \quad \mathcal{T} := \{p(F) \mid F \in \mathcal{T}'\}.$$

We call \mathcal{T} the *flat TT-fan*.

4.2.2.3. *Definition of the edge sets $E^u(a)$.* Let $a \in S^\Delta$. For every $u \in V$, we let $E^u(a)$ be the set of edges on which the slack of the triangle inequality (9) is minimized:

$$(12) \quad E^u(a) := \left\{ vw \in E \mid u \neq v, w, \text{ and} \right. \\ \left. a_{vu} + a_{vw} - a_{vw} = \min_{v', w' \neq u} a_{v'u} + a_{uw'} - a_{v'w'} \right\}.$$

4.2.2.4. *The TT-sub-complex of P^Δ .* We let $\hat{\mathcal{C}}(P)$ be the polyhedral complex of all bounded faces of P .

We define a sub-complex \mathcal{B} of $\mathcal{C}(P^\Delta)$: \mathcal{B} is what remains of the complex $\hat{\mathcal{C}}(P^\Delta)$ of bounded faces of P^Δ after deleting the conjugate face of S in P^Δ , in symbols

$$\mathcal{B} := \text{dl}(S^\diamond, \hat{\mathcal{C}}(P^\Delta)).$$

It will become clear in the next section (see Remark 4.9) that the points of the complex \mathcal{B} are precisely the points in $|\hat{\mathcal{C}}(P^\Delta)|$ which are tight triangular.

4.2.3. Rotation and statements of the results. Let \mathcal{C} and \mathcal{D} be two polyhedral complexes. \mathcal{D} is called a *subdivision* of \mathcal{C} , if, (a) every face of \mathcal{D} is contained in some face of \mathcal{C} ; and (b) every face of \mathcal{C} is a union of faces of \mathcal{D} .

We now give the rigorous definition of “rotation” and of the rotation complex, as outlined in the introduction. More accurately, we define a “rotation partition” of $|\mathcal{A}|$, which will turn out to be a polyhedral complex subdividing \mathcal{A} .

A point $a \in S^\Delta$ corresponds to an inequality $a \bullet x \geq a \bullet z - 1$ valid for S . Rotating this inequality amounts to adding an equation valid for S . The left-hand side vector q of such an

equation is a linear combination of the left-hand side vectors of the equations $\delta_u \bullet x = 1$, and the right-hand side coincides with $q \bullet z$. Hence, for a fixed q , rotating the inequality $a \bullet x \geq a \bullet z - 1$ by q gives the following

$$(13) \quad (a + q) \bullet x \geq a \bullet z - 1 + q \bullet z.$$

For $a \in |\mathcal{A}|$, let $\mathfrak{F}(a) \in \mathcal{C}(P)$ be the set of faces of P which can be defined by the rotated version of the inequality corresponding to a . More precisely, a set $F \subset \mathbb{R}^E$ is in $\mathfrak{F}(a)$ if, and only if, there exists a q as above, such that the rotated inequality (13) is valid for P , and F is the set of points in P satisfying it with equality: $F = \{x \in P \mid (a + q) \bullet x = a \bullet z - 1 + q \bullet z\}$.

Now we define a partition \mathcal{S}° of $|\mathcal{A}|$, by letting two points a, b be in the same cell of \mathcal{S}° if and only if $\mathfrak{F}(a) = \mathfrak{F}(b)$. Moreover, let \mathcal{S} be the set of all closures of cells of \mathcal{S}° :

$$\mathcal{S} := \{\overline{X} \mid X \in \mathcal{S}^\circ\},$$

where, for $X \subset \mathbb{R}^m$, we denote by \overline{X} the closure of X in the topological sense. We call \mathcal{S} the *rotation complex* (the word ‘‘complex’’ is justified by Theorem 4.2).

Let \mathcal{C} and \mathcal{D} be two polyhedral complexes. The *common refinement* of \mathcal{C} and \mathcal{D} is the polyhedral complex whose faces are all the intersections of faces of \mathcal{C} and \mathcal{D} : $\mathcal{C} \vee \mathcal{D} := \{F \cap G \mid F \in \mathcal{C}, G \in \mathcal{D}\}$. The common refinement $\mathcal{C} \vee \mathcal{D}$ is a subdivision of both \mathcal{C} and \mathcal{D} .

Theorem 4.2. *\mathcal{S} is a polyhedral complex. Moreover, $X \mapsto \overline{X}$ and $F \mapsto \text{relint } F$ are inverse bijections between \mathcal{S}° and \mathcal{S} . The following is true.*

- (a) *The rotation complex \mathcal{S} is the common refinement of \mathcal{A} and the flat TT-fan \mathcal{T} .*
- (b) *Two points a, b in $|\mathcal{A}|$ are in the relative interior of the same face of the rotation complex \mathcal{S} if, and only if, they are in the relative interior of same face of S^Δ and $E^u(a) = E^u(b)$ for all $u \in V$.*

This corresponds to item (A) on page 27 in the introduction, while the next theorem corresponds to item (B).

Two polyhedral complexes \mathcal{C} and \mathcal{D} are called *combinatorially equivalent*, if there exists a bijection $\Phi: \mathcal{C} \rightarrow \mathcal{D}$, which preserves the inclusion relation of faces, i.e., if $F \subset F'$ are two faces of \mathcal{C} , then $\Phi(F) \subset \Phi(F')$. We say that a mapping $f: |\mathcal{C}| \rightarrow |\mathcal{D}|$ *induces a combinatorial equivalence*, if f induces the polyhedral complex \mathcal{D} , and the mapping $F \mapsto f(F)$ is one-to-one. In this case, \mathcal{C} and \mathcal{D} are combinatorially equivalent via $F \mapsto f(F)$.

Theorem 4.3. *There is a projective homeomorphism $\pi: |\mathcal{B}| \rightarrow |\mathcal{A}|$, such that the mapping $F \mapsto \pi(F)$ is a combinatorial equivalence between the polyhedral complex \mathcal{B} and the rotation complex \mathcal{S} .*

Remark 4.4. Let us say that a vertex of P^Δ is a TT-vertex, if, as a point, it is TT in the above sense, or, equivalently, if the vertex corresponds to a TT-facet of P . Similarly, let us call a TT-vertex of P^Δ an NR-vertex (non-NR-vertex), if the corresponding facet of P is an NR-facet (non-NR-facet, resp.). Theorems 4.2 and 4.3 imply that the NR-vertices of P^Δ are in bijection with the vertices of \mathcal{A} via π , while the non-NR vertices of P^Δ are mapped to non-vertex points by π .

4.2.4. Parsimonious property of relaxations and the ridge graph. Given a system $Bx \geq b$ of linear inequalities which are valid for S , one may ask how the minimum value of a linear function $x \mapsto c^\top x$ changes if either degree inequalities or degree equations are present, in other words, whether the following inequality is strict:

$$(14a) \quad \min\{c^\top x \mid Bx \geq b, \delta_v \bullet x \geq 1 \forall v, x \geq 0\}$$

$$\leq$$

$$(14b) \quad \min\{c^\top x \mid Bx \geq b, \delta_v \bullet x = 1 \forall v, x \geq 0\}$$

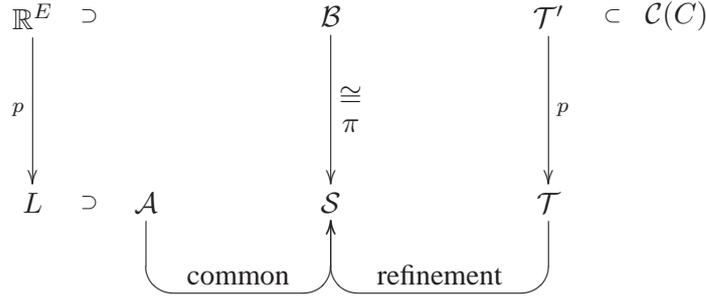


FIGURE 4.2. Mappings and sets

We say that the system of linear inequalities and equations in (14a),

$$(15) \quad \begin{aligned} Bx &\geq b \\ \delta_v \bullet x &\geq 1 \quad \forall v \in V \\ x &\geq 0 \end{aligned}$$

is a *relaxation of S* . Such a relaxation is said to have the *parsimonious property* [GB93] if equality holds in (14) for all c satisfying the triangle inequality.

Goemans [Goe95] raised the question whether all relaxations of S consisting of inequalities defining NR-facets of P (in other words, they are facet-defining for P and for S) have the parsimonious property.

The parsimonious property had earlier been proved to be satisfied for the relaxation consisting of all inequalities defining facets of P by Naddef & Rinaldi [NR91], in other words: optimizing an objective function satisfying the triangle inequality over P yields the same value as optimizing over S . The parsimonious property has been verified by Goemans and Bertsimas [GB93] for the relaxation consisting of all non-negativity inequalities $x_e \geq 0$, $e \in E$, and all so-called subtour elimination inequalities. For every $U \subsetneq V$ with $|U| \geq 2$, the corresponding *subtour elimination inequality*

$$(16) \quad \sum_{\substack{uv \in E \\ |\{u,v\} \cap U| = 1}} x_{uv} \geq 2,$$

is valid and facet-defining for S (whenever $n \geq 5$) [GP79a, GP79b].

To our knowledge, the first example of a relaxation of S which does not have the parsimonious property is due to Letchford [Let05]. While the operative inequalities in his relaxation do not define facets of S or of P , in [ORT05, ORT07], a relaxation consisting of inequalities defining facets of P was derived, which does not have the parsimonious property.

As an application of Theorems 4.2 and 4.3, we give a necessary condition for a relaxation of S consisting of inequalities defining NR-facets of P to have the parsimonious property. The condition is based on connectivity properties of the ridge graph of P . Recall that the *ridge graph* \mathcal{G} of P is the graph whose vertex set consists of all facets of P where two facets are adjacent if their intersection has dimension $\dim P - 2$, i.e., it is a *ridge*. We will relate a given relaxation to the induced subgraph \mathcal{G}_B of the ridge graph of P which is obtained if all vertices corresponding to the facets defined by inequalities in (15) are deleted.

Theorem 4.5. *Suppose $Bx \geq b$ consists of inequalities defining NR-facets of P . If the relaxation (15) of S has the parsimonious property, then every connected component of \mathcal{G}_B contains vertices corresponding to NR-facets of P .*

Thus, we link the optimization view given by the parsimonious property question with combinatorial properties of the polyhedral complex $\mathcal{C}(P)$, or, more precisely, of \mathcal{B} . In the proof,

Theorem 4.3 is used to “flatten” the latter complex, which then allows us to use a separating-hyperplane argument for constructing a path in the ridge graph.

4.3. Proofs for Theorems 4.2 and 4.3

In Subsection 4.3.1, we will need to discuss some properties of Symmetric and Graphical Traveling Salesman polyhedra. Most of them are generalizations of facts in the seminal papers by Naddef & Rinaldi [NR91, NR93]. The proof of Theorems 4.2 and 4.3 then takes up Subsections 4.3.2 and 4.3.3.

As said before, we assume in the whole section that $n \geq 5$, because we require the technical fact that non-negativity inequalities $x_e \geq 0$, for an $e \in E$, define facets of S , which is true if and only if $n \geq 5$, see [GP79a, GP79b].

4.3.1. Preliminaries on P . Naddef & Rinaldi [NR93] proved that every facet of S is contained in precisely $n + 1$ facets of P : the n degree facets and one additional facet. This fact and its generalizations are useful for our purposes. For the sake of completeness, we will sketch its proof, and introduce some of the tools for the proofs of our main theorems along the way.

First we set up some notations. For a linear subspace $L \subset \mathbb{R}^m$, denote by $L^\perp := \{q \in \mathbb{R}^m \mid q \bullet x = 0 \ \forall x \in L\}$ the orthogonal complement of L . Let D be the $V \times E$ -matrix whose rows are the δ_u^\top , $u \in V$. Recall from Section 4.2.2.2 that p is the orthogonal projection from \mathbb{R}^E onto $L = \ker D$. Note that the orthogonal complement $L^\perp = \ker p$ of L is equal to $\text{im } D^\top = \{D^\top \xi \mid \xi \in \mathbb{R}^V\}$, the space of all linear combinations of the δ_u .

In the following lemma, we summarize basic facts about tight triangularity. Recall from page 4.2.1 that a face of P is good, if it is not contained in a non-negativity facet.

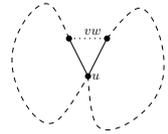
Lemma 4.6.

- (i) A metric inequality which is valid for S is also valid for P .
- (ii) An inequality defining a good face of P is metric.
- (iii) An inequality defining a good face F of P is TT if and only if F is not contained in a degree facet.
- (iv) If F is a good face of P , then $S \cap F$ is also a good face of P .
- (v) Let the TT inequality $a \bullet x \geq 1$ be valid for P . If it defines a face of co-dimension c of S , then it defines a face of co-dimension at most c of P .
- (vi) For every $a \in \mathbb{R}^E$ there is a unique TT representative in the co-set $a + L^\perp = \{a + D^\top \xi \mid \xi \in \mathbb{R}^V\}$. More precisely, we can obtain a unique $\lambda(a) \in \mathbb{R}^V$ for which $a - D^\top \lambda(a)$ is TT by letting

$$(17) \quad \lambda_u(a) := \min_{v,w \neq u} (a_{vu} + a_{uw} - a_{vw})$$

The mapping $\lambda: \mathbb{R}^E \rightarrow \mathbb{R}^V$ is defined as $a \mapsto (\lambda_u(a))_{u \in V}$.

Given a vertex u and an edge vw not incident to u , a *shortcut* is a vector $s_{u,vw} := \chi^{vw} - \chi^{vu} - \chi^{uw} \in \mathbb{R}^E$. Here, we abbreviate $\chi^{\{e\}}$ to χ^e . Note that $-a \bullet s_{u,vw} = a_{vu} + a_{uw} - a_{vw}$ is the slack of the corresponding triangle inequality.



PROOFS FOR LEMMA 4.6 (SKETCHES). Please note that there is not a shred of an argument in the proofs for the statements of Lemma 4.6, which is not present in the [NR93] paper, only that the arguments are applied to faces instead of facets.

The key ingredient in (a–c) is the *shortcut argument* which Naddef & Rinaldi pioneered in [NR93]. Let $x \in \mathbb{Z}_+^E$ represent the edge multi-set of a connected Eulerian multi-graph H with vertex set V . If H is not a cycle, i.e., if H has a vertex u of degree four or more, then one can find an edge vw such that vu and vw are in H , and $H' := H + vw - \{vu, uw\}$ is still a connected Eulerian multi-graph; cf. the picture on the right. If y represents its edge multi-set,

then $y = x + s_{u,vw}$. This gives (a), the implication “ \Rightarrow ” in (c), and by carefully selecting the edge vw , (d). Similarly, one can subtract a shortcut from an x , which gives (b), the other direction in (c), and, by taking for each vertex u a shortcut $s_{u,vw}$, implies (e).

The proof can be found in [NR93], but we present the basic computation which is used here and in some other arguments in the present paper. Let a as in Item (f), and suppose $q = \sum_{j=1}^n \mu_j \delta_j$ for some real numbers μ_1, \dots, μ_n . For every selection of u and disjoint $vw \in E$, the slack of the corresponding triangle inequality for $a + q$ can be computed as follows

$$(18) \quad (a + q)_{vu} + (a + q)_{uw} - (a + q)_{vw} = a_{vu} + a_{uw} - a_{vw} + \mu_u.$$

Thus, $a + q$ is TT if and only if the μ_u are equal to the λ_u in (17). \square

The proof of (f) gives the following.

Remark 4.7. If a is as in (f) and $q = \sum_{j=1}^n \mu_j \delta_j$ then for every $u \in V$ $\mu_u = \lambda_u(a)$ implies $a \bullet s_{u,e} = 0$ for all $e \in E^u(a)$.

We now prove the important theorem of Naddef & Rinaldi.

Theorem 4.8 ([NR93]).

- (i) If a facet G of P contains S , then G is a degree facet.
- (ii) Let F be a good facet of S (i.e., a facet of S which is a good face of P). There exists a unique facet G of P with $F = G \cap S$.

PROOF. (a). If $G \supset S$, then G is trivially good, because S is not contained in a non-negativity facet. If G is not equal to a degree facet, then, by Lemma 4.6(c), it is defined by a TT inequality, which contradicts Lemma 4.6(e).

(b). Clearly, G exists because S is a face of P . Let G be defined by an inequality $a \bullet x \geq \alpha$. Then a is TT by Lemma 4.6(c), hence, by Lemma 4.6(f), unique in the set $a + L^\perp$ of all left hand sides of inequalities defining the facet F of S . \square

4.3.1.1. *Related aspects of the polar polyhedra.* By passing to the polar, Theorem 4.8(b) is equivalent to the following. If a is a vertex of P^Δ such that the inequality $a \bullet x \geq 1$ defines a facet (called F in Theorem 4.8(b)) of S , then a and $\delta_u, u \in V$, are the vertices of an n -simplex which is a face of P^Δ .

Remark 4.9. By Lemma 4.1(b) and Lemma 4.6(c), the points of the complex $\mathcal{B} = \text{dl}(S^\diamond, \hat{\mathcal{C}}(P^\Delta))$ are precisely the points in $|\hat{\mathcal{C}}(P^\Delta)|$ which are tight triangular.

4.3.2. Descriptions of the rotation complex. We will now prove Theorem 4.2. We start by proving that the two refinements of $\mathcal{A} = \text{dl}(N, S^\Delta)$ defined in (a) and (b) respectively of Theorem 4.2 are identical: the one using the flat TT-fan defined in (11) and the one using the sets $E^u(a)$ defined in (12).

Let us first verify that the orthogonal projection p maps the TT-fan $|\mathcal{T}'|$ bijectively onto L . For this, we define some mappings, based on (17):

$$(19) \quad \begin{aligned} \lambda_u: \mathbb{R}^E &\rightarrow \mathbb{R}: & a &\mapsto \min_{v,w \neq u} a_{vu} + a_{uw} - a_{vw}, \\ \lambda: \mathbb{R}^E &\rightarrow \mathbb{R}^V: & a &\mapsto (\lambda_1(a), \dots, \lambda_n(a))^\top, \\ \vartheta: \mathbb{R}^E &\rightarrow \mathbb{R}^E: & a &\mapsto a - D^\top \lambda(a), \\ \tilde{\vartheta}: \mathbb{R} \times \mathbb{R}^E &\rightarrow \mathbb{R} \times \mathbb{R}^E: & (\alpha, a) &\mapsto (\alpha - \mathbf{1} \bullet \lambda(a), \vartheta(a)). \end{aligned}$$

Note that $\tilde{\vartheta}$ is essentially the same as ϑ except that the former takes the “right hand side” α into account.

Lemma 4.10. The mappings $p: |\mathcal{T}'| \rightarrow L$ and $\vartheta_L: L \rightarrow |\mathcal{T}'|$ are inverses of each other.

(The restriction of a mapping $f: X \rightarrow Y$ to a set $Z \subset X$ is denoted by $f|_Z$.)

PROOF. By Lemma 4.6(vi), every co-set $a + L^\perp$ of L^\perp contains a unique TT point, namely $\vartheta(a)$. The co-set also contains a unique point of L , namely the orthogonal projection $p(a)$ of a onto L . Hence, the two mappings are inverses of each other. \square

In view of Lemma 4.10, p transports the fan \mathcal{T}' into a fan $\mathcal{T} := p(\mathcal{T}')$ in L , the flat TT-fan defined in Section 4.2. It is a complete fan in its ambient space L . (A fan \mathcal{C} is *complete*, if $|\mathcal{C}|$ is equal to the ambient space.) The next lemma states that the refinements of \mathcal{A} used in Theorem 4.2 are identical. The proof is a direct verification based on the definitions of $E^u(\cdot)$ and ϑ , using Lemma 4.10.

Lemma 4.11. *For two points $a, b \in L$, the following are equivalent:*

- (i) $E^u(a) = E^u(b)$ for all $u \in V$
- (ii) a and b are in the relative interior of the same face of the flat TT-fan \mathcal{T} . \square

For easy reference, let \mathcal{D} denote the common refinement of \mathcal{A} and the flat TT-fan \mathcal{T} . This is certainly a polyhedral complex, and the previous lemma implies that two points in the relative interior of a face of \mathcal{A} are in the relative interior of the same face of \mathcal{D} , if, and only if, (i) holds and they are in the relative interior of same face of S^Δ .

This shows that items (a) and (b) of Theorem 4.2 are just reformulations of each other, one using the formulation involving the sets $E^u(\cdot)$, the other using the common refinement with \mathcal{T} . Moreover, to establish Theorem 4.2, it remains to prove that the partition of $|\mathcal{A}|$ into open faces of \mathcal{D} coincides with the partition S° : Once this is established, both the statement about the closures and relative interiors in Theorem 4.2, and items (a) and (b) follow.

To prove that these two partitions coincide, we need to descend deeper into the properties of P . For $X \subset \mathbb{R}^m$, we denote by $\text{aff } X$ the affine hull of X , i.e., the smallest affine subspace of \mathbb{R}^m containing X . We let $\text{dir } X$ denote the “space of directions” in X , i.e., the linear space generated by the points $y - x$, $x, y \in X$. Hence, $\text{aff } X = x + \text{dir } X$ holds for every $x \in \text{aff } X$. If F is a face of P , then a shortcut is said to be *feasible* for F , if it is contained in the space $\text{dir } F$. We note the following for easy reference.

Lemma 4.12. *If F is a good face of P , then a shortcut $s_{u,vw}$ is feasible for F if and only if $a \bullet s_{u,vw} = 0$ for one (and hence for all) $a \in \text{relint } F^\diamond$.*

PROOF. If F is a good face, then the polarity relations of Lemma 4.1 hold between F and F^\diamond . The details are left to the reader. \square

The following lemma highlights the importance of shortcuts in the relationship between S and P .

Lemma 4.13. *A good face F of P is uniquely determined by*

- *the set of cycles whose characteristic vectors are contained in F , plus*
- *the set of its feasible shortcuts.*

PROOF. A face is uniquely determined by the vertices it contains the extreme rays of its characteristic cone. By the shortcut argument, every vertex of F is either itself a cycle, or it can be constructed from a cycle by successively subtracting feasible shortcuts. As for the rays, $\mathbb{R}_+ \chi^{uv}$ is a ray of F if and only if, for any $a \in \text{relint } F^\diamond$, we have $a_{uv} = 0$ (by Lemma 4.6(b)). By Lemma 4.12, this is equivalent to the property that for every $w \neq u, v$, both $s_{u,vw}$ and $s_{v,uw}$ are feasible shortcuts. \square

We can now finish the proof of Theorem 4.2.

PROOF OF THEOREM 4.2(B). Let $a \in |\mathcal{A}|$. The inequalities of the form (13) all define good faces of P , because a defines a face of S not contained in a non-negativity facet of S .

Moreover, since every inequality of the form (13) defines the same face of S , Lemma 4.13 implies that every member of the set $\mathfrak{F}(a)$ of faces of P defined by inequalities of the form (13) is uniquely determined by its set of feasible shortcuts.

We claim that the set $\mathfrak{F}(a)$ is in bijection with the set of all subsets of V , where the bijection is accomplished in the following way: To a subset $I \subset V$, there is a face in $\mathfrak{F}(a)$ whose set of feasible shortcuts is precisely

$$(*) \quad \bigcup_{u \in I} \{s_{u,e} \mid e \in E^u(a)\}.$$

The faces obtainable in this way are clearly pairwise distinct by what we have just said (note that $E^u(a) \neq \emptyset$). We have to construct a corresponding inequality for every set I , and we have to show that all faces in $\mathfrak{F}(a)$ can be reached in this way.

For the former issue, for $I \subset V$ we define $q := \sum_{u \notin I} \delta_u$, and consider the inequality

$$(\vartheta(a) + q) \bullet x \geq -1 + a \bullet z - \mathbf{1} \bullet \lambda(a) + q \bullet z,$$

which is of the form (13) because $\mathbf{1} = Dz$, and defines a good face of P . The set of feasible shortcuts of this inequality is easily verified to be (*): Compute the slacks of the triangle inequalities as in (18) and then use Lemma 4.12.

To see that every face F in $\mathfrak{F}(a)$ can be obtained in this way we argue that if there exists an edge vw such that $s_{u,vw}$ is feasible for F , then $vw \in E^u(a)$ and for every $e \in E^u(a)$ the shortcut $s_{u,e}$ is feasible for F . But this is an immediate consequence of Remark 4.7 following Lemma 4.6.

This completes the proof of Theorem 4.2. \square

4.3.3. Projective equivalence of the two complexes. We now proceed to prove Theorem 4.3. We want to define a mapping π by letting

$$(20a) \quad \pi(a) := \frac{1}{a \bullet z - 1} p(a),$$

for $a \in P^\Delta$. The denominator will be zero, if, and only if, $a \bullet x \geq 1$ is satisfied by equality for all $x \in S$, in other words, $\pi(a)$ is well-defined for all $a \in P^\Delta \setminus S^\diamond$.

By Lemma 4.6, a point a in the complex $\hat{C}(P^\Delta)$ of bounded faces of P^Δ defines a good face of S , so we have $\pi(a) \in |\mathcal{A}|$, whenever $a \notin S^\diamond$. Hence, we have the mapping

$$(20b) \quad \pi: |\mathcal{B}| \rightarrow |\mathcal{A}|.$$

In this subsection, we will prove that π as given in (20) is a homeomorphism, and show that it induces a combinatorial equivalence between \mathcal{B} and the rotation complex \mathcal{S} ; i.e., we prove Theorem 4.3. We will explicitly construct the inverse mapping π^{-1} , which, essentially, transforms a point into its TT-representative in the sense of Lemma 4.6(vi).

When we write the projective mapping π as a linear mapping from $\mathbb{R} \times \mathbb{R}^E \rightarrow \mathbb{R} \times L$ by homogenization,² it has the following form:

$$\tilde{\pi} := \begin{pmatrix} -1 & z \bullet \square \\ 0 & p \end{pmatrix},$$

where \square replaces the variable, i.e., $X \square Z$ (X and Y are parts of a formula) is short for $y \mapsto XyZ$.

As a technical intermediate step in the construction of π^{-1} , we define a linear mapping $I: \mathbb{R} \times \mathbb{R}^E \rightarrow \mathbb{R} \times \mathbb{R}^E$ taking points in $\mathbb{R} \times L$ to points in $\mathbb{R} \times \mathbb{R}^E$ by the matrix

$$I := \begin{pmatrix} -1 & z \bullet \square \\ 0 & \text{id} \end{pmatrix},$$

²Recall that a projective mapping $\mathbb{R}^m \rightarrow \mathbb{R}^m$ can be understood as a linear mapping $\mathbb{R} \times \mathbb{R}^m \rightarrow \mathbb{R} \times \mathbb{R}^m$: If $f = \begin{pmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \end{pmatrix}$ for linear mappings $f_{k,\ell}$, then $\Pi(f^0)(x) = \frac{f_{11}(1)+f_{12}(x)}{f_{21}(1)+f_{22}(x)}$ is a projective mapping. This commutes with concatenation of mappings: $\Pi(f \circ g) = \Pi(f) \circ \Pi(g)$.

Now we let $(\gamma, c) := \tilde{\vartheta} \circ I(1, \square)$. Recalling the definition of $\tilde{\vartheta}$ from (19), in long, this reads

$$(21) \quad (\gamma, c): a \mapsto (\gamma(a), c(a)) := \tilde{\vartheta}(I(1, a)) = (-1 + a \bullet z - \mathbf{1} \bullet \lambda(a), a - D^T \lambda(a)).$$

Clearly, for all $a \in L$, the point $c(a)$ is TT. If $a \in S^\Delta$, i.e., if the inequality $a \bullet x \geq -1 + a \bullet z$ is valid for S , then the inequality $c(a) \geq \gamma(a)$ is of the form (13) (cf. the corresponding statement in the proof of Theorem 4.2 above). We note the following fact as a lemma for the sake of easy reference.

Lemma 4.14. *If $a \in S^\Delta$, the two inequalities $a \bullet x \geq -1 + a \bullet z$ and $c(a) \bullet x \geq \gamma(a)$ define the same face of S . \square*

Finally, we define

$$(22) \quad \varphi: |\mathcal{A}| \rightarrow |\mathcal{B}|: a \mapsto \frac{1}{\gamma(a)} c(a).$$

Recall that, if \mathcal{D}' is a polyhedral complex, a homeomorphism $f: |\mathcal{C}| \rightarrow |\mathcal{D}'|$ is called *refinement map*, if the image of \mathcal{C} under the mapping $F \mapsto f(F)$ is a polyhedral complex \mathcal{D} which is a subdivision of \mathcal{D}' . In this case, we say that f *induces* \mathcal{D} .

PROOF OF THEOREM 4.3. In the remainder of this section, we will discuss the following issues:

- (a) φ is well-defined (in 4.3.3.1)
- (b) φ is a left-inverse of $\pi: |\mathcal{B}| \rightarrow |\mathcal{A}|$ (in 4.3.3.2)
- (c) $\pi: |\mathcal{B}| \rightarrow |\mathcal{A}|$ is onto (in 4.3.3.3)

Items (b) and (c) imply that

$$(23) \quad \varphi \circ \pi = \text{id}_{|\mathcal{B}|} \quad \text{and} \quad \pi \circ \varphi = \text{id}_{|\mathcal{A}|},$$

so that φ is a homeomorphism $|\mathcal{A}| \rightarrow |\mathcal{B}|$.

- (d) $\pi: |\mathcal{B}| \rightarrow |\mathcal{A}|$ is a refinement map inducing the rotation complex \mathcal{S} (in 4.3.3.4).

From this and (d), Theorem 4.3 follows. \square

4.3.3.1. *We show: φ is well-defined.* We start by showing that the quotient in (22) is well-defined. The key ingredient here is the fact that we are only considering good faces.

Lemma 4.15. *For all $a \in |\mathcal{A}|$ we have $\gamma(a) > 0$.*

PROOF. Assume to the contrary that $\gamma(a) = 0$. Since $c(a)$ is metric, $c(a) \geq 0$ holds. We distinguish two cases: $c(a) = 0$ and $c(a) \succeq 0$. In the first case, the hyperplane defined by $c(a) \bullet x = \gamma(a)$ contains S , while $a \bullet x \geq -1 + a \bullet z$ defines a proper face of S , a contradiction to Lemma 4.14. On the other hand, if $c(a) \succeq 0$, then the inequality $c(a) \bullet x \geq \gamma(a)$ is a non-negative linear combination of non-negativity inequalities, and hence the face defined by $c(a) \bullet x = \gamma(a)$ is contained in a non-negativity facet of P . But since $a \in |\mathcal{A}|$, i.e., a is not a relative interior point of a face of S^Δ which contains a vertex of S^Δ corresponding to a non-negativity facet of S , the face of S defined by $a \bullet x \geq -1 + a \bullet z$ is not contained in a non-negativity facet of S . Thus Lemma 4.14 yields a contradiction. \square

It remains to be shown that the image of $|\mathcal{A}|$ under φ is really contained in the target space given in (22): For all $a \in |\mathcal{A}|$ we have $\varphi(a) \in |\mathcal{B}|$. This also follows from Lemma 4.14: The inequality $\varphi(a) \bullet x \geq 1$ is valid for P , and the face it defines is good. Since $\varphi(a)$ is TT, the conclusion follows from Remark 4.9.

4.3.3.2. We show: φ is a left-inverse of π , i.e., for all $a \in |\mathcal{B}|$ the identity $\varphi(\pi((a))) = a$ holds.

Lemma 4.16. For all $a \in |\mathcal{B}|$ we have $(\gamma, c)(\tilde{\pi}(1, a)) = (1, a)$. In particular, we have that $\varphi \circ \pi$ restricted to $|\mathcal{B}|$ is equal to the identity mapping on this set.

PROOF. To see this we compute

$$I(\tilde{\pi}(1, a)) = I(-1 + a \bullet z, p(a)) = (1 - a \bullet z - z \bullet p(a), p(a)) = ((p(a) - a) \bullet z + 1, p(a))$$

Using that a is TT (Remark 4.9), we conclude

$$\tilde{\vartheta}(I(\tilde{\pi}(1, a))) = \left((p(a) - a) \bullet z + 1 - \lambda(p(a)) \bullet \mathbf{1}, a \right).$$

Since a is TT, by Lemma 4.6(vi), $\lambda(p(a))$ is a solution to $p(a) - a = D^\top \lambda$. Thus, using $\mathbf{1} = Dz$, it follows that

$$(p(a) - a) \bullet z + 1 - \lambda(p(a)) \bullet \mathbf{1} = (p(a) - a) \bullet z + 1 - D^\top \lambda(p(a)) \bullet z = 1.$$

From the statement about $(\tilde{\vartheta} \circ I) \circ \tilde{\pi}$, the statement about the projective mappings $\varphi \circ \pi$ follows by a slight generalization of the well-known fact that concatenation of projective mappings commutes with homogenization. We omit the computation, and only note that it makes use of the fact that the two mappings $h_1: a \mapsto a - D^\top \lambda(a)$ and $h_2: a \mapsto a \bullet z + \lambda(a) \bullet \mathbf{1}$ are positive homogeneous (i.e., $h_i(\eta a) = \eta h_i(a)$ for $\eta \geq 0$, $i = 1, 2$, which follows directly from the definition of λ). \square

4.3.3.3. We show: φ is one-to-one. Since we already know that $\varphi \circ \pi = \text{id}$, surjectivity of π is equivalent to injectivity of φ . It is actually easier to prove the following slightly stronger statement.

Lemma 4.17. Let $a, b \in L$. If there exists an $\eta \in \mathbb{R}_+$ such that $(\gamma(a), c(a)) = \eta(\gamma(b), c(b))$ then $\eta = 1$ and $a = b$. In particular, φ is injective.

PROOF. Let such a, b, η be given. We have

$$0 = c(a) - \eta c(b) = a - D^\top \lambda(a) - \eta \left[b - D^\top \lambda(b) \right] = a - \eta b - D^\top \left[\lambda(a) - \eta \lambda(b) \right].$$

Since $a, b \in L$ and $D^\top [\lambda(a) - \eta \lambda(b)] \in L^\perp$ we have

$$(*) \quad a - \eta b = 0 = D^\top \lambda(a) - \eta D^\top \lambda(b)$$

Applying $z \bullet \square$ to the second equation, we obtain

$$0 = \mathbf{1} \bullet \lambda(a) - \eta \mathbf{1} \bullet \lambda(b)$$

Now we use $\gamma(a) = \eta \gamma(b)$ and compute

$$\begin{aligned} 0 = \gamma(a) - \eta \gamma(b) &= -1 + a \bullet z - \mathbf{1} \bullet \lambda(a) - \eta \left[-1 + b \bullet z - \mathbf{1} \bullet \lambda(b) \right] \\ &= -1 + \eta + (a - \eta b) \bullet z. \end{aligned}$$

Since $z \in L^\perp$ we have $(a - \eta b) \bullet z = 0$, whence $\eta = 1$. Now $a = b$ follows from (*). \square

4.3.3.4. π induces the rotation complex. We are finally ready to prove that π is a refinement map inducing the rotation complex.

Let \mathcal{C} be a polyhedral complex, and $f: |\mathcal{C}| \rightarrow \mathbb{R}^k$ a mapping. We say that f induces the polyhedral complex \mathcal{D} , if, for every $F \in \mathcal{C}$, its image $f(F)$ under f is a polyhedron, and the set of all these polyhedra is equal to \mathcal{D} .

Lemma 4.18. For every face F of \mathcal{S} there exists a unique face $\Phi(F)$ of \mathcal{B} with $\varphi(\text{relint } F) \subset \text{relint } \Phi(F)$. Moreover, if $F_1 \neq F_2$ are faces of \mathcal{S} , then $\Phi(F_1) \neq \Phi(F_2)$.

PROOF. Let F' be the face of \mathcal{A} with $\text{relint } F \subset \text{relint } F'$. Now, let $a \in \text{relint } F$ and $G^\#$ be the face of P defined by the inequality $\varphi(a) \bullet x \geq 1$. Since this inequality defines the same face of P as the inequality $c(a) \bullet x \geq \gamma(a)$ which is of the form (13), the set of cycles whose characteristic vectors are in $G^\#$ coincides with the set of cycles contained in the face F'^\diamond of S , where the conjugate face is taken in S vs. S^Δ (not in P vs. P^Δ), and thus does not depend on the choice of $a \in \text{relint } F'$. Moreover, the set of feasible shortcuts for $G^\#$ is in bijection with $E^u(a)$, $u \in V$, and hence, by Theorem 4.2, depends only on F not on the choice of $a \in \text{relint } F$. Thus, by Lemma 4.13, $G^\#$ does not depend on the choice of $a \in \text{relint } F$. Hence, with $\Phi(F) := (G^\#)^\diamond$, we have $\varphi(a) \in \text{relint } \Phi(F)$ for all $a \in \text{relint } F$.

The injectivity follows from Lemma 4.13 because, as we have just argued, $a \in \text{relint } F$ uniquely determines the set of cycles and shortcuts. \square

Lemma 4.18 provides us with a mapping $\Phi: \mathcal{S} \rightarrow \mathcal{B}$: $\Phi(F)$ is the unique face of \mathcal{B} whose relative interior contains $\varphi(a)$ for some a in the relative interior of F . Since φ is surjective by what we have proved in Subsection 4.3.3.2, this immediately implies that Φ is, too: For $G \in \mathcal{B}$, let $b \in \text{relint } G$, choose $a \in \mathcal{S}$ with $\varphi(a) = b$ and let F be the unique face of \mathcal{S} containing a as a relative interior point. Then $\Phi(F) = G$.

Hence, we obtain the following.

Lemma 4.19. *There is a bijection $\Phi: \mathcal{S} \rightarrow \mathcal{B}$ with $\Phi(F) = \varphi(F)$.*

Recall that *boundary* of a polyhedron F in the relative topology of $\text{aff } F$ is $\partial F := F \setminus \text{relint } F = \bigcup_{G \subsetneq F} G$ where the union runs over all faces of F .

PROOF. What remains to be shown is the final statement: $\Phi(F) = \varphi(F)$. We already know that $\varphi(\text{relint } F) \subset \text{relint } \Phi(F)$, by the definition of Φ . By the surjectivity of φ , we have, in fact equality in this relation: For every $b \in \text{relint } \Phi(F)$, there is a a such that $\varphi(a) = b$, but by the injectivity of Φ , we must have $a \in \text{relint } F$.

Moreover, $\varphi(\text{relint } F) = \text{relint } \Phi(F)$ implies $\varphi(F) \subset \Phi(F)$ by continuity of φ .

Standard Euclidean topology arguments show that φ maps the boundary ∂F of F into the boundary of $\varphi(F)$. (This is most easily seen by noting that φ is the inverse of a projective mapping; see equations (23).) The boundary of F is the union of its facets, and we can apply Lemma 4.18 to those. In particular, we obtain $\Phi(F') \cap \text{relint } \Phi(F) = \emptyset$ by the injectivity of Φ . Thus, we have $\partial\varphi(F) \subset \mathbb{C}(\text{relint } \Phi(F))$, but $\varphi(\partial F) \subset \Phi(F)$. From this, we conclude that $\varphi(\partial F) \subset \partial\Phi(F)$. Now $\text{relint } F$ and $\text{relint } \Phi(F)$ have the same dimension, because φ is a homeomorphism; see (23). In such a case, the Borsuk-Ulam theorem states that if a continuous mapping φ maps a topological sphere ∂F into another topological sphere $\partial\Phi(F)$ of the same dimension, but leaves out a point, cannot be injective, a contradiction to what we have proved in Subsection 4.3.3.3. Hence, $\varphi(\partial F) = \partial\Phi(F)$, and we conclude $\varphi(F) = \Phi(F)$. \square

Remark 4.20. The topological arguments contained in the proof of Lemma 4.19 can be replaced by more technical but more elementary ones from linear algebra. In any case, they reflect basic topological facts.

4.4. Proof of Theorem 4.5

We will apply Theorem 4.3 to prove Theorem 4.5. The following lemma is the link between parsimonious property and geometry.

Lemma 4.21. *Let $Bx \geq 1$ be a system of inequalities defining NR-facets of P such that the relaxation (15) has the parsimonious property. If $c^\top x \geq \gamma$ defines a non-NR facet of P , then c, γ cannot be written in the form*

$$(24) \quad \begin{aligned} c &= b - \sum_{v \in V} \mu_v \delta_v \\ \gamma &= \beta - \sum_{v \in V} \mu_v \end{aligned}$$

with $b^\top = \sum_j t_j b_j$ a non-negative linear combination of rows b_j of B , $\beta = \sum_j t_j$, and $\mu_v \in \mathbb{R}$ for all $v \in V$.

PROOF. Suppose that c, γ can be written as in (24). Then minimizing the cost function c over the relaxation consisting of

- all non-negativity inequalities
- all degree equations(!) $\delta_v \bullet x = 1, v \in V$;
- all inequalities in the system $Bx \geq \mathbf{1}$.

yields γ as the minimum. If the degree equations are relaxed to inequalities, then, by the parsimonious property of (15), the minimum is still γ . By Farkas's Lemma (or LP-duality), this implies that the inequality $c \bullet x \geq \gamma$ is dominated by non-negativity inequalities, degree inequalities, and inequalities in $Bx \geq \mathbf{1}$. This is impossible since (c, γ) defines a non-NR facet of P and all facets in $Bx \geq \mathbf{1}$ are NR. \square

We are now ready to prove Theorem 4.5.

PROOF OF THEOREM 4.5. Let $a_o \bullet x \geq 1$ be an inequality defining a non-NR facet of P which is not in the system $Bx \geq \mathbf{1}$. By Lemma 4.1, the paths in the ridge graph of P not touching non-negativity facets are precisely the paths in the 1-skeleton of P^Δ . (The 1-skeleton or graph of a polyhedral complex \mathcal{C} is the graph G whose vertices are the vertices of \mathcal{C} , with two vertices of G being adjacent if and only if there exists an edge of \mathcal{C} containing them both.)

Thus, we have to find a path in the graph of P^Δ which starts from a_o , ends in an NR-vertex, and does not use any degree vertices or vertices corresponding to rows of B .

By Theorem 4.3, we know that there exists a projective homeomorphism $\pi: |\mathcal{B}| \rightarrow |\mathcal{A}|$ transporting the polyhedral complex \mathcal{B} onto the rotation complex. We let $\phi := \pi^{-1}$.

Let $a := \phi^{-1}(a_o)$. This point is contained in the relative interior of a unique face F of S^Δ containing no non-negativity vertex. Let \mathcal{D}_F denote the set of all faces of the rotation complex \mathcal{D} which are contained in F , and let B_F denote the set of vertices b of F for which $\varphi(b)^\top$ is a row of B . We will prove the following:

Claim 4.22. Let F be a face of \mathcal{A} , and let a be a relative interior point of F which is a vertex of \mathcal{D}_F such that $\varphi(a)^\top$ is not a row of B (see Fig. 4.3). Then there is a path in the 1-skeleton of \mathcal{D}_F starting at a , ending in a vertex of F , and not touching any of the vertices in B_F .

By Theorem 4.3, the paths in \mathcal{D}_F correspond to paths in \mathcal{B} . Moreover, the vertices corresponding to rows of B are avoided in the path in \mathcal{D}_F . Thus, we have in \mathcal{B} a path from a to an NR-vertex not touching any vertices of the parsimonious formulation, which concludes the proof of Theorem 4.5. \square

PROOF OF CLAIM 4.22. The proof of the claim is by induction on $\dim F$. For $\dim F = 0$, we are done, because then a is a vertex of F . Let $\dim F \geq 1$, and assume the claim holds for relative interior points a' of faces F' with dimension $\dim F' < \dim F$.

If $B_F = \emptyset$, we are done. Otherwise let $Q := \text{conv } B_F$. This is a non-empty polytope which is contained in F . Using Lemma 4.21 we will show the following:

Claim 4.23. Let c be a vertex of \mathcal{D}_F which is not a member of B_F . Then c cannot be contained in Q .

The proof of Claim 4.23 is technical, and we postpone it till the proof of Claim 4.22 is finished. If Claim 4.23 is true, however, then we know that a is not in Q . Let p, π define a hyperplane separating a from Q , i.e., $q \bullet p < \pi$ for all $q \in Q$, and $a \bullet p > \pi$. See Fig. 4.3 for an illustration. It assumes the face F is an 8-gon.

By a standard general position argument, we can assume that p is not parallel to any face with co-dimension at least one in \mathcal{D}_F . Hence, there exists an $\varepsilon > 0$ such that the line segment $a +]0, \varepsilon[\cdot p$ is contained in the relative interior of a $\dim F$ -dimensional face G of \mathcal{D}_F , of which a

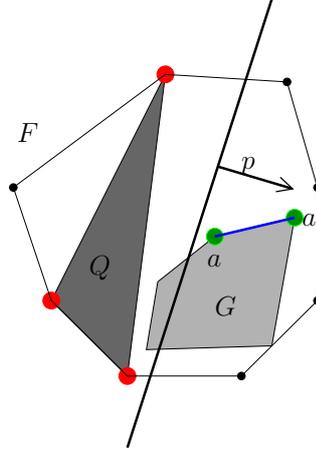


FIGURE 4.3. One step of the path

is a vertex. By elementary polytope theory (the edges of a polyhedron incident to a fixed vertex span a cone of the same dimension as the polyhedron), G must have a vertex a_1 adjacent to a with $a \bullet p < a_1 \bullet p$. Clearly $a_1 \notin B_F$.

If a_1 is in the boundary of F , then the induction hypotheses implies the existence of a path from a_1 to a vertex of F not using any vertex in B_F . If that is not the case, we apply the argument in the previous paragraph inductively to obtain a path a, a_1, \dots, a_k in the 1-skeleton of \mathcal{D}_F with $a \bullet p < a_1 \bullet p < \dots < a_j \bullet p < a_{j+1} \bullet p < \dots < a_k \bullet p$. Since the 1-skeleton of \mathcal{D}_F is finite and the path we are constructing is p -increasing, a vertex on the boundary of F will eventually be reached.

This concludes the proof of Claim 4.22. \square

PROOF OF CLAIM 4.23. Let c be a vertex of \mathcal{D}_F with $c \notin B_F$. Assume that $c \in \text{conv } B_F$, i.e., c can be written as a convex combination $c = \sum_{j=1}^k t_j b_j$ with $\varphi(b_j)^\top$ a row of B for all $j = 1, \dots, k$. Clearly, c cannot be a vertex of F , so $\varphi^{-1}(c) \bullet x \geq 1$ defines a non-NR facet of P by Remark 4.4. We compute

$$c - \sum_{v \in V} \lambda_v(c) d_v = \sum_j t_j \left(b_j - \sum_{v \in V} \lambda_v(b_j) d_v \right) - \sum_{v \in V} \left(\lambda_v(c) - \sum_j t_j \lambda_v(b_j) \right) d_v.$$

Letting $\sigma := 1 - \sum_v \lambda_v(c)$, $\tau_j := 1 - \sum_v \lambda_v(b_j)$, and $\mu_v := \lambda_v(c) - \sum_j t_j \lambda_v(b_j)$, we see that

$$\begin{aligned} \sigma \varphi(c) &= \sum_j t_j \tau_j \varphi(b_j) - \sum_v \mu_v d_v \\ \sigma &= \sum_j t_j \tau_j - \sum_{v \in V} \mu_v \end{aligned}$$

This means that the inequality $\sigma \varphi(c) \bullet x \geq \sigma$ can be written as a non-negative linear combination of the inequalities $\varphi(b_j) \bullet x \geq 1$, $j = 1, \dots, k$ plus a linear combination of degree vertices as in (24). Since the former inequality defines a facet of P by Theorems 4.2 and 4.3, and the inequalities forming the non-negative linear combination are taken from the system $Bx \geq \mathbf{1}$, Lemma 4.21 yields a contradiction. \square

4.5. Outlook

We conjecture that the necessary condition for parsimonious property in Theorem 4.5 is also sufficient.

Conjecture 4.24. *If every connected component of \mathcal{G}_B contains vertices corresponding to NR-facets of P , then the relaxation \mathcal{R}_B of has the parsimonious property.*

The conjecture holds for the known relaxations of S consisting of NR-inequalities described in [ORT07] which fail the parsimonious property.

Acknowledgments

The author would like to thank the *Deutsche Forschungsgemeinschaft*, DFG, for funding this research, and the *Communauté française de Belgique – Actions de Recherche Concertées* for supporting the author during the time the paper was written down.

Moreover, thanks are extended to Jean-Paul Doignon and Samuel Fiorini, U.L.B., for helpful discussions on the topic of this paper, and to Marcus Oswald for inspiring discussions on the topic of Section 4.4.

On a class of metrics related to graph layout problems

Jointly with

Adam N. Letchford (Lancaster University, GB),
 Gerhard Reinelt (Uni Heidelberg),
 and Hanna Seitz (Uni Heidelberg)

Abstract. We examine the metrics that arise when a finite set of points is embedded in the real line, in such a way that the distance between each pair of points is at least 1. These metrics are closely related to some other known metrics in the literature, and also to a class of combinatorial optimization problems known as graph layout problems. We prove several results about the structure of these metrics. In particular, it is shown that their convex hull is not closed in general. We then show that certain linear inequalities define facets of the closure of the convex hull. Finally, we characterise the unbounded edges of the convex hull and of its closure.

5.1. Introduction

For a given positive integer n , let $[n]$ denote $\{1, \dots, n\}$. A *metric* on $[n]$ is a mapping $d: [n] \times [n] \rightarrow \mathbb{R}_+$ which satisfies the following three conditions:

- $d(i, j) = d(j, i)$ for all $\{i, j\} \subset [n]$,
- $d(i, k) + d(j, k) \geq d(i, j)$ for all ordered triples $(i, j, k) \subset [n]$,
- $d(i, j) = 0$ if and only if $i = j$.

Metrics are a special case of *semimetrics*, which are obtained by dropping ‘and only if’ from the third condition. There is a huge literature on metrics and semimetrics; see for example [DL97]. The inequalities in the second condition are the well-known *triangle inequalities*.

In this paper we study the metrics d on $[n]$ that arise when n points are embedded in the real line, in such a way that the distance between each pair of points is at least 1. More formally, we require that d satisfies the following two properties:

- there exist real numbers r_1, \dots, r_n such that $d(i, j) = |r_i - r_j|$ for all $\{i, j\} \subset [n]$;
- $d(i, j) \geq 1$ for all $\{i, j\} \subset [n]$.

We remark that one could easily replace the value 1 with some arbitrary constant $\epsilon > 0$; the results in this paper would remain essentially unchanged.

We call the metrics in question ‘ \mathbb{R} -embeddable 1-separated’ metrics. We believe that these metrics are a natural object of study, and of interest in their own right. We have, however, two specific motives for studying them. First, they are closely related to certain well-known metrics that have appeared in the literature. Second, they are also closely related to an important class of combinatorial optimization problems, known as *graph layout problems*.

As well as studying the metrics themselves, we also study their convex hull. It turns out that the convex hull is not always closed, which leads us to study also the closure of the convex hull. Among other things, we characterise some of the $(n - 1)$ -dimensional faces (i. e., facets) of the closure, and some of the 1-dimensional faces (i. e., edges) of both the convex hull and its closure.

The structure of the paper is as follows. In Section 5.2, we review some of the relevant literature on metrics and graph layout problems. In Section 5.3, we present various results

CUT_n	ℓ_1 -embeddable semimetrics (cut cone)
HYP_n	hypermetrics, see (25)
NEG_n	negative-type cone, see (26)
M_n^{L2}	ℓ_2 -embeddable semimetrics
M_n^R	\mathbb{R} -embeddable semimetrics
M_n^{R1}	\mathbb{R} -embeddable 1-separated metrics
Q_n	convex hull of M_n^{R1}
$\overline{Q_n}$	closure of Q_n
P_n	permutation metrics polytope, see (29)

TABLE 1. Sets of matrices

concerned with the structure of the metrics and their convex hull. Next, in Section 5.4, we present some inequalities that define facets of the closure of the convex hull. In Section 5.5, we give a combinatorial characterisation of the unbounded edges of the convex hull and of its closure. Finally, some concluding remarks are given in Section 5.6.

We close this section with a word on notation. To study convex geometric properties, we view metrics as points in a vector space \mathbb{S}_n^0 . In our notation, \mathbb{S}_n^0 will be either the vector space of all symmetric functions $[n] \times [n] \rightarrow \mathbb{R}$ or the vector space of all real symmetric $(n \times n)$ -matrices whose diagonal entries are zero, and we will switch freely between them. For the latter, the inner product is defined as usual by

$$A \bullet B := \text{tr}(A^\top B) = \sum_{k=1}^n \sum_{l=1}^n A_{k,l} B_{k,l}.$$

We understand a metric both as a function and a matrix, and we will switch between the two concepts without further mentioning.

By $S(n)$ we denote the set of all permutations of $[n]$. We occasionally view $S(n)$ as a subset of \mathbb{R}^d by identifying the permutation π with the point $(\pi(1), \dots, \pi(n))^\top$. Furthermore we let $\iota_n := (1, \dots, n)$ the identity permutation in $S(n)$. We omit the index n when no confusion can arise. $\mathbf{1}$ is a column vector of appropriate length consisting of ones. Similarly $\mathbf{0}$ is a vector whose entries are all zero. If appropriate, we will use a subscript $\mathbf{1}_k$, $\mathbf{0}_k$ to identify the length of the vectors. The symbol \mathcal{K} denotes an all-zeros matrix not necessarily square, and we also use it to say “this part of the matrix consists of zeros only.” By \mathcal{K}_n we denote the square matrix of order n whose (k, l) -entry is 1 if $k \neq l$ and 0 otherwise. As above we will omit the index n when appropriate. We denote by $\complement U$ the complement of the set U .

5.2. Literature Review

In this section, we review some of the relevant literature. We cover related semimetrics in Subsection 5.2.1 and graph layout problems in Subsection 5.2.2. To facilitate reading we have summarized all matrix sets discussed in Table 1.

5.2.1. Some related semimetrics. The following four classes of semimetrics on $[n]$, which are closely related to the \mathbb{R} -embeddable 1-separated metrics, have been extensively studied in the literature (see [DL97] for a detailed survey):

- The ℓ_1 -embeddable semimetrics, i. e., those for which there exist a positive integer m and points $x_1, \dots, x_n \in \mathbb{R}^m$ such that $d(i, j) = |x_i - x_j|_1 := \sum_{k=1}^m |x_{ik} - x_{jk}|$ for all $\{i, j\} \subset [n]$.
- The ℓ_2 -embeddable semimetrics, which are defined as in the ℓ_1 case, except that $d(i, j) = |x_i - x_j|_2 := \sqrt{\sum_{k=1}^m (x_{ik} - x_{jk})^2}$.
- The \mathbb{R} -embeddable semimetrics, which are the special case of ℓ_1 - (or ℓ_2 -) embeddable semimetrics obtained when $m = 1$.
- The *hypermetrics*, which are semimetrics that satisfy the following *hypermetric* inequalities [Dez61]:

$$(25) \quad \sum_{\{i,j\} \subset [n]} b_i b_j d(i, j) \leq 0 \quad (\forall b \in \mathbb{Z}^n : \sum_{i=1}^n b_i = 1).$$

It is known [Ass80] that the set of ℓ_1 -embeddable semimetrics on $[n]$ is a polyhedral cone in $\mathbb{R}^{\binom{n}{2}}$. In fact, it is nothing but the well-known *cut cone*, denoted by CUT_n . The set of all hypermetrics on $[n]$, called the *hypermetric cone* and denoted by HYP_n , is also polyhedral [DGL93].

We will let M_n^{L2} and M_n^R denote the set of ℓ_2 - and \mathbb{R} -embeddable semimetrics, respectively. It is known that M_n^{L2} and M_n^R are not convex (unless n is small), and that the convex hull of M_n^{L2} and M_n^R is CUT_n . It is also known [Sch35] that a symmetric function d lies in M_n^{L2} if and only if d^2 (i. e., the symmetric function obtained by squaring each value) lies in the so-called *negative-type cone*. The negative-type cone, denoted by NEG_n , is the (non-polyhedral) cone defined by the following *negative-type* inequalities:

$$(26) \quad \sum_{\{i,j\} \subset [n]} b_i b_j d(i, j) \leq 0 \quad (\forall b \in \mathbb{R}^n : \sum_{i=1}^n b_i = 0).$$

The structure of M_n^R and related sets is studied in [BD92].

In recent years, there has been a stream of papers on so-called *negative-type* semimetrics (also known as ℓ_2^2 -semimetrics) [ALN07, ALN08, CGR08, KV05, KR06, Lee05]. These are simply semimetrics that lie in NEG_n . They have been used to derive approximation algorithms for various combinatorial optimisation problems, including the graph layout problems that we mention in the next subsection.

The following inclusions are known: $M_n^R \subset M_n^{L2} \subset \text{CUT}_n \subset \text{HYP}_n \subset \text{NEG}_n$. Denoting the set of all \mathbb{R} -embeddable 1-separated metrics by M_n^{R1} , we obtain from their definition $M_n^{R1} \subset M_n^R$. We will explore the relationship between M_n^{R1} , M_n^R and CUT_n further in Subsection 5.3.1.

5.2.2. Graph layout problems. Given a graph $G = (V, E)$, with $V = [n]$, a *layout* is simply a permutation of $[n]$. If we view a layout $\pi \in S(n)$ as a placing of the vertices on points $1, \dots, n$ along the real line, the quantity $|\pi(i) - \pi(j)|$ corresponds to the Euclidean distance between vertices i and j . Several important combinatorial optimization problems, collectively known as *graph layout problems*, call for a layout minimising a function of these distances (see the survey [DPS02]). For example, in the *Minimum Linear Arrangement Problem* (MinLA), the objective is to minimize $\sum_{\{i,j\} \in E} |\pi(i) - \pi(j)|$. In the *Bandwidth Problem*, the objective is to minimise $\max_{\{i,j\} \in E} |\pi(i) - \pi(j)|$.

Now, let $d(i, j)$ for $\{i, j\} \subset [n]$ be a decision variable, representing the quantity $|\pi(i) - \pi(j)|$. It has been observed by several authors that interesting relaxations of graph layout problems can be formed by deriving valid linear inequalities that are satisfied by all feasible symmetric functions d . To our knowledge, the first paper of this kind was [LV95], which presented

the following *star* inequalities:

$$(27) \quad \sum_{j \in S} d(i, j) \geq \lfloor (|S| + 1)^2 / 4 \rfloor.$$

Here, $i \in [n]$ and $S \subset [n] \setminus \{i\}$ is such that every node in S is adjacent to i .

Apparently independently, Even *et al.* [ENRS00] defined the so-called *spreading metrics*. These are metrics that satisfy the following *spreading* inequalities:

$$(28) \quad \sum_{j \in S} d(i, j) \geq |S|(|S| + 2)/4 \quad (\forall i \in [n], \forall S \subseteq [n] \setminus \{i\}).$$

Note that the spreading inequalities are more general than the star inequalities, but have a slightly weaker right-hand side when n is odd. Spreading metrics were used in [ENRS00, RR05] to derive approximation algorithms for various graph layout problems.

In [CHKR08, FL07], it was noted that one can get a tighter relaxation of graph layout problems by requiring the spreading metrics to lie in the negative-type cone NEG_n . The authors called the resulting metrics ℓ_2^2 -*spreading* metrics.

A natural way to derive further valid linear inequalities for graph layout problems is to study the following *permutation metrics polytope*:

$$(29) \quad P_n = \text{conv} \left\{ d \mid \exists \pi \in S(n) : d(i, j) = |\pi(i) - \pi(j)| \forall \{i, j\} \subset [n] \right\}.$$

Surprisingly, this was not done until very recently [AL09]. In [AL09], it is shown that P_n is of dimension $\binom{n}{2} - 1$ and that its affine hull is defined by the equation $\sum_{\{i, j\} \subset [n]} d(i, j) = \binom{n+1}{3}$. It is also shown that the following four classes of inequalities define facets of P_n under mild conditions:

- *pure hypermetric* inequalities, which are simply the hypermetric inequalities (25) for which $b \in \{0, \pm 1\}^n$;
- *strengthened pure negative-type* inequalities, which are like the negative-type inequalities (26) for which $b \in \{0, \pm 1\}^n$, except that the right-hand side is increased from 0 to $\frac{1}{2} \sum_{i \in [n]} |b_i|$;
- *clique* inequalities, which take the form

$$(30) \quad \sum_{\{i, j\} \subset S} d(i, j) \geq \binom{|S| + 1}{3},$$

where $S \subset [n]$ satisfies $2 \leq |S| < n$;

- *strengthened star* inequalities, which take the form

$$(31) \quad (|S| - 1) \sum_{i \in S} d(r, i) - \sum_{\{i, j\} \subset S} d(i, j) \geq \lfloor (|S| + 1)^2 (|S| - 1) / 12 \rfloor,$$

where $r \in V$ and $S \subseteq V \setminus \{r\}$ with $|S| \geq 2$.

It is pointed out in the same paper that each star inequality (27) with $|S| \geq 2$ is dominated by a clique inequality (30) and a strengthened star inequality (31). Therefore, very few of the star inequalities define facets of P_n .

Finally, we mention that some more valid inequalities were presented recently by Caprara *et al.* [CLSG09]. Some of them were proved to define facets of the *dominant* of P_n , though not of P_n itself.

We will establish an interesting connection between M_n^{R1} , CUT_n and P_n in Subsection 5.3.2.

5.3. On M_n^{R1} and its Convex Hull

5.3.1. On M_n^{R1} and related sets. We now study M_n^{R1} and its relationship with M_n^R , P_n and CUT_n . We will find it helpful to recall the definition of a *cut metric*:

Definition 5.1. For a set $U \subset [n]$, we let d_U be the metric which assigns to two points on different sides of the bipartition $U, \complement U$ of $[n]$ a value of 1 and to points on the same side a value of 0.

We will say that the set U *induces* the associated cut metric. In other words, if we let $D_{k,l}(x) := |x_k - x_l|$ for every vector $x \in \mathbb{R}^n$ (and identify, as promised, functions and matrices), then $d_U = D(\chi^U)$. With this notation, CUT_n is the convex cone with apex 0 in \mathbb{S}_n^R generated by the points d_U , i. e.,

$$\text{CUT}_n := \text{cone} \left\{ d_U \mid d_U \text{ is the cut metric for } U \subset [n] \right\}.$$

It is known [BM86] that each cut metric defines an extreme ray of CUT_n .

We will also need the following notation. For a given permutation $\pi \in S(n)$, let N_π be the set of $x \in \mathbb{R}^n$ which satisfy $x_{\pi(i)} \leq x_{\pi(i+1)}$ for $i = 1, \dots, n-1$. Now let $M(\pi)$ denote the set of metrics d for which there exists an $x \in N_\pi$ with $d = D(x)$. Also, for a given π and for $k = 1, \dots, n-1$, we emphasize that $D(\chi^{\pi^{-1}([k])})$ is the cut metric induced by the set $U = \{\pi^{-1}(1), \dots, \pi^{-1}(k)\}$. (So, for example, if $n = 4$ and $\pi = \{2, 3, 1, 4\}$, then $D(\chi^{\pi^{-1}([2])})$ is the cut metric induced by the set $\{2, 3\}$.)

We have the following lemma:

Lemma 5.2. $M(\pi)$ is a polyhedral cone of dimension $n-1$ defined by the $n-1$ cut metrics $D(\chi^{\pi^{-1}([1])}), \dots, D(\chi^{\pi^{-1}([n-1])})$.

PROOF. Let $d^* \in M(\pi)$ and let x_1, \dots, x_n be the corresponding points in \mathbb{R} . One can check that:

$$d^* = \sum_{k=1}^{n-1} (x_{k+1} - x_k) D(\chi^{\pi^{-1}([k])}).$$

From the definition of $M(\pi)$, we have $x_{k+1} - x_k \geq 0$ for $k = 1, \dots, n-1$. Thus, d^* is a conical combination of the $n-1$ cut metrics mentioned. This shows that $M(\pi)$ is contained in the cone mentioned. The reverse direction is similar. \square

This enables us to describe the structure of M_n^R :

Proposition 5.3. M_n^R is the union of $n!/2$ polyhedral cones, each of dimension $n-1$.

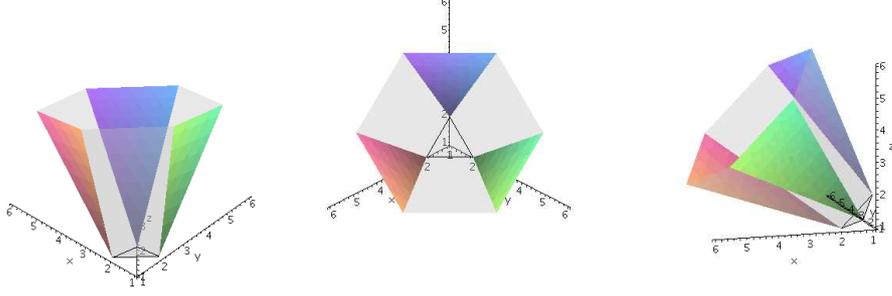
We define the *antipodal* permutation of $\pi \in S(n)$ by

$$\pi^- := (n+1) \cdot \mathbf{1} - \pi.$$

This is the permutation obtained by reversing π . A swift computation shows that $D(\pi) = D(\pi^-)$.

PROOF. From the definitions, we have $M_n^R = \bigcup_{\pi \in S(n)} M(\pi)$. From the above lemma, the set $M(\pi)$ is a polyhedral cone of dimension $n-1$. Now, note that, for any $\pi \in S(n)$, we have $M(\pi) = M(\pi^-)$. Thus, the union can be taken over $n!/2$ permutations, instead of over all permutations. \square

We note in passing that every cut metric belongs to $M(\pi)$ for some $\pi \in S(n)$. This explains the well-known fact, mentioned in Subsection 5.2.1, that the convex hull of M_n^R is equal to CUT_n .

FIGURE 5.1. The convex set Q_3

Now, we adapt these results to the case of M_n^{R1} . We define $M^1(\pi)$ similar to $M(\pi)$: we denote by $M^1(\pi)$ the set of all metrics d which are of the form $D(x)$ for an $x \in \mathbb{R}^n$ which satisfies $x_{\pi(i)} + 1 \leq x_{\pi(i+1)}$ for $i = 1, \dots, n-1$.

Note that the $D(\pi)$ are nothing but the metrics associated with feasible layouts, which by a result in [AL09] are the extreme points of P_n . Note also that the sets $M^1(\pi)$ are disjoint.

We have the following lemma:

Lemma 5.4. $M^1(\pi)$ is the Minkowski sum of the point $D(\pi)$ and the cone $M(\pi)$:

$$M^1(\pi) = D(\pi) + D(N_\pi).$$

PROOF. This can be proven in the same way as Lemma 5.2. The only difference is that we decompose $d^* \in M^1(\pi)$ as:

$$d^* = D(\pi) + \sum_{k=1}^{n-1} (r_{k+1} - r_k - 1) D(\chi^{\pi^{-1}([k])}),$$

and note that $r_{k+1} - r_k - 1 \geq 0$ for $k = 1, \dots, n-1$. \square

We can now derive an analog of Proposition 5.3:

Proposition 5.5. M_n^{R1} is the union of $n!/2$ disjoint translated polyhedral cones, each of dimension $n-1$.

PROOF. From the definitions, we have $M_n^{R1} = \bigcup_{\pi \in S(n)} M^1(\pi)$. From Lemmas 5.2 and 5.4, each set $M^1(\pi)$ is a translated polyhedral cone of dimension $n-1$. As in the proof of Proposition 5.3, the union can be taken over only $n!/2$ permutations. \square

5.3.2. On the convex hull of M_n^{R1} and related sets. We now turn our attention to the convex hull of M_n^{R1} , which we denote by Q_n . To give some intuition, we present in Fig. 5.1 drawings of M_n^{R1} and Q_3 from three different angles. (Of course, the drawing is truncated, since Q_3 is unbounded.) The three co-ordinates represent $d(1,2)$, $d(1,3)$ and $d(2,3)$. The three coloured regions represent the three disjoint subsets of M_3^{R1} mentioned in Proposition 5.5.

One can see that Q_3 is a three-dimensional polyhedron, with one bounded facet, six unbounded facets, three bounded edges and six unbounded edges.

For $n \leq 3$, Q_n is closed (and therefore a polyhedron). We will show in Section 5.5, however, that Q_n is not closed for $n \geq 4$. Therefore, we are led to look at the closure of Q_n , which we denote by $\overline{Q_n}$.

Our next result shows that there is a close connection between the polyhedron $\overline{Q_n}$, the polytope P_n , and the cone CUT_n :

Proposition 5.6. $\overline{Q_n}$ is the Minkowski sum of P_n and CUT_n .

PROOF. We use the same notation as in the previous subsection. By definition, every point in M_n^{R1} belongs to $M^1(\pi)$ for some $\pi \in S(n)$. From Lemma 5.4, every point in $M^1(\pi)$ is the sum of the point $D(\pi)$ and a point in the cut cone CUT_n . Moreover, the point $D(\pi)$ is an extreme point of P_n . Thus, every point in M_n^{R1} is the sum of an extreme point of P_n and a point in CUT_n . Since $\overline{Q_n}$ is the closure of the convex hull of M_n^{R1} , it must be contained in the Minkowski sum of P_n and CUT_n . The reverse direction is proved similarly, noting that every cut metric is of the form $D(\chi^{\pi^{-1}([k])})$ for some $\pi \in S(n)$ and some $k \in [n-1]$. \square

This immediately implies the following result:

Corollary 5.7. $\overline{Q_n}$ is full-dimensional (i. e., of dimension $\binom{n}{2}$).

We also have the following result:

Proposition 5.8. P_n is the unique bounded facet of $\overline{Q_n}$.

PROOF. As mentioned in the previous section, all points in P_n satisfy the equation

$$\sum_{\{i,j\} \subset [n]} d(i,j) = \binom{n+1}{3}.$$

Moreover, every point in CUT_n satisfies $\sum_{\{i,j\} \subset [n]} d(i,j) > 0$. Since $\overline{Q_n}$ is the Minkowski sum of P_n and CUT_n , it follows that the inequality $\sum_{\{i,j\} \subset [n]} d(i,j) \geq \binom{n+1}{3}$ is valid for $\overline{Q_n}$ and that P_n is the face of $\overline{Q_n}$ exposed by this inequality. Since $\overline{Q_n}$ and P_n are of dimension $\binom{n}{2}$ and $\binom{n}{2} - 1$, respectively, P_n is a facet of $\overline{Q_n}$. It must be the unique bounded facet, since all extreme points of $\overline{Q_n}$ are in P_n . \square

In the next section, we will explore the connection between $\overline{Q_n}$, P_n and CUT_n in more detail. To close this section, we make an observation about how the individual ‘pieces’ of M_n^{R1} , called the $M^1(\pi)$ in the previous subsection, are positioned within $\overline{Q_n}$:

Proposition 5.9. For any $\pi \in S(n)$, the set $M^1(\pi)$ is an $(n-1)$ -dimensional face of $\overline{Q_n}$.

PROOF. By definition, $\overline{Q_n}$ satisfies all triangle inequalities. Now, without loss of generality, suppose that π is the identity permutation. Every point in $M^1(\pi)$ satisfies all of the following triangle inequalities at equality:

$$d(i,j) + d(j,k) \geq d(i,k) \quad (\forall 1 \leq i < j < k \leq n).$$

Moreover, no other point in M_n^{R1} does so. Thus, $M^1(\pi)$ is a face of $\overline{Q_n}$. It was shown to be $(n-1)$ -dimensional in the previous subsection. \square

5.4. Inequalities Defining Facets of $\overline{Q_n}$

In this section, we study linear inequalities that define *facets* of $\overline{Q_n}$, i. e., faces of dimension $\binom{n}{2} - 1$. Subsection 5.4.1 presents some general results about such inequalities, whereas Subsection 5.4.2 lists some specific inequalities.

5.4.1. General results on facet-defining inequalities. In this subsection, we prove a structural result about inequalities that define facets of $\overline{Q_n}$, and show how this can be used to construct facets of $\overline{Q_n}$ in a mechanical way from facets of either P_n or CUT_n .

We will need the following definition, taken from [AL09]:

Definition 5.10 (Amaral & Letchford, 2009). Let $\alpha^T d \geq \beta$ be a linear inequality, where $\alpha, d \in \mathbb{R}^{\binom{n}{2}}$. The inequality is said to be ‘canonical’ if:

$$(32) \quad \min_{\emptyset \neq S \subset [n]} \sum_{i \in S} \sum_{[n] \setminus S} \alpha_{ij} = 0.$$

By definition, an inequality $\alpha^T d \geq 0$ defines a proper face of CUT_n if and only if it is canonical. In [AL09], it is shown that every facet of P_n is defined by a canonical inequality. The following lemma is the analogous result for \overline{Q}_n :

Lemma 5.11. *Every unbounded facet of \overline{Q}_n is defined by a canonical inequality.*

PROOF. Suppose that the inequality $\alpha^T d \geq \beta$ defines an unbounded facet of \overline{Q}_n . Since \overline{Q}_n is the Minkowski sum of P_n and CUT_n , the inequality must be valid for CUT_n . Therefore, the left-hand side of (32) must be non-negative. Moreover, since the inequality defines an unbounded facet, there must be at least one extreme ray of CUT_n satisfying $\alpha^T d = 0$. Therefore the left-hand side of (32) cannot be positive. \square

We remind the reader that only one facet of \overline{Q}_n is bounded (Proposition 5.8).

Now, we show how to derive facets of \overline{Q}_n from facets of P_n :

Proposition 5.12. *Let F be any facet of P_n , and let $\alpha^T d \geq \beta$ be the canonical inequality that defines it. This inequality defines a facet of \overline{Q}_n as well.*

PROOF. The fact that the inequality is valid for \overline{Q}_n follows from the fact that \overline{Q}_n is the Minkowski sum of P_n and CUT_n . Now, since F is a facet of P_n , there exist $\binom{n}{2} - 1$ affinely-independent vertices of P_n that satisfy the inequality at equality. Moreover, since the inequality is canonical, there exists at least one extreme ray of CUT_n that satisfies $\alpha^T d = 0$. Since \overline{Q}_n is the Minkowski sum of P_n and CUT_n , there exist $\binom{n}{2}$ affinely-independent points in \overline{Q}_n that satisfy the inequality at equality. Thus, the inequality defines a facet of \overline{Q}_n . \square

Now, we show how to derive facets of \overline{Q}_n from facets of CUT_n :

Proposition 5.13. *Let $\alpha^T d \geq 0$ define a facet of CUT_n , and let β be the minimum of $\alpha^T d$ over all $d \in P_n$. Then the inequality $\alpha^T d \geq \beta$ define a facet of \overline{Q}_n .*

PROOF. As before, the fact that the inequality $\alpha^T d \geq \beta$ is valid for \overline{Q}_n follows from the fact that \overline{Q}_n is the Minkowski sum of P_n and CUT_n . Now, since the inequality $\alpha^T d \geq 0$ defines a facet of CUT_n , there exist $\binom{n}{2} - 1$ linearly-independent extreme rays of CUT_n that satisfy $\alpha^T d = 0$. Moreover, from the definition of β , there exists at least one extreme point of P_n that satisfies $\alpha^T d = \beta$. Since \overline{Q}_n is the Minkowski sum of P_n and CUT_n , there exist $\binom{n}{2}$ affinely-independent points in \overline{Q}_n that satisfy $\alpha^T d = \beta$. Thus, the inequality $\alpha^T d \geq \beta$ defines a facet of \overline{Q}_n . \square

5.4.2. Some specific facet-defining inequalities. The results in the previous subsection enable one to derive a wide variety of facets of \overline{Q}_n . In this subsection, we briefly examine some specific valid inequalities; namely, the inequalities mentioned in [AL09].

First, we deal with the clique and pure hypermetric inequalities:

Proposition 5.14. *The clique inequalities (30) define facets of \overline{Q}_n for all $S \subseteq [n]$ with $|S| \geq 2$.*

PROOF. It was shown in [AL09] that the clique inequalities define facets of P_n when S is a proper subset of $[n]$. In this case, the inequalities are canonical and so, by Proposition 5.12, they define facets of \overline{Q}_n as well. The case $S = [n]$ is covered in the proof of Proposition 5.8. \square

Proposition 5.15. *All pure hypermetric inequalities define facets of \overline{Q}_n .*

PROOF. It was shown in [BM86] that all pure hypermetric inequalities define facets of CUT_n . It was also shown in [AL09] that every pure hypermetric inequality is satisfied at equality by at least one extreme point of P_n . The result then follows from Proposition 5.13. \square

As for the strengthened pure negative-type and strengthened star inequalities, it was shown in [AL09] that they define facets of P_n under certain conditions. Since they are canonical, they define facets of \overline{Q}_n under the same conditions. In fact, using the same proof technique used in [AL09], one can show the following two results:

Proposition 5.16. *All strengthened pure negative-type inequalities define facets of $\overline{Q_n}$.*

Proposition 5.17. *Strengthened star inequalities define facets of $\overline{Q_n}$ if and only if $|S| \neq 4$.*

We omit the proofs, for the sake of brevity.

5.5. Unbounded Edges of Q_n and $\overline{Q_n}$

5.5.1. Unbounded edges of Q_n . We now investigate how the polyhedral cones $M^1(\pi) = D(\pi) + D(N_\pi)$ as subsets of Q_n . In Fig. 5.1, it can be seen that in the case $n = 3$, the three cones are faces of Q_3 (recall that Q_3 is a polyhedron, which means that we can safely speak of faces). In the following proposition, we show that this is the case for all n , and we also characterize the extremal half-lines of Q_n . This will be useful in comparing Q_n with its closure: We will characterize the unbounded edges issuing from each vertex for the polyhedron $\overline{Q_n} = P_n + \text{CUT}_n$ in the following subsection.

We are dealing with an unbounded convex set of which we do not know whether it is closed or not. (In fact, we will show that Q_n is almost never closed). For this purpose, we supply the following fact for easy reference.

Fact 5.18. For $k = 1, \dots, m$ let K_k be a (closed) polyhedral cone with apex x_k . Suppose that the K_k are pairwise disjoint and define $S := \bigsqcup_{k=1}^m K_k$. Let x, y be vectors such that $x + \mathbb{R}_+ y$ is an extremal subset of $\text{conv}(S)$. It then follows that there exists a $\lambda_0 \in \mathbb{R}_+$ and a k such that $x + \lambda y \in K_k$ for all $\lambda \geq \lambda_0$. Since $x + \mathbb{R}_+ y$ is extremal, this implies that there exists a $\lambda_1 \in \mathbb{R}_+$ such that $x_k = x + \lambda_1 y$ and $x_k + \mathbb{R}_+ y = \{x + \lambda y \mid \lambda \geq \lambda_1\}$ is an extreme ray of the polyhedral cone K_k .

Definition 5.19. We say that a permutation π and a non-empty set $U \subsetneq [n]$ are *incident*, if $U = \{\pi^{-1}(1), \dots, \pi^{-1}(k)\}$, where $k := |U|$.

Proposition 5.20.

- (i) *For every $\pi \in S(n)$, each edge of the cone $D(\pi) + D(N_\pi)$ is an exposed subset of Q_n .*
- (ii) *The unbounded one dimensional extremal sets of Q_n are exactly the defining half-lines. In other words, every half-line $X + \mathbb{R}_+ Y$ which is an extremal subset of Q_n is of the form $D(\pi) + \mathbb{R}_+ D(\chi^U)$ for a $\pi \in S(n)$ and a set U incident to π . In particular, for every vertex $D(\pi)$ of Q_n , the unbounded one-dimensional extremal subsets of Q_n containing $D(\pi)$ are in bijection with the non-empty proper subsets of $[n]$ incident to π . Thus there are precisely $n - 1$ of them.*

PROOF. *i.* By symmetry it is sufficient to treat the case $\pi = \iota := (1, \dots, n)^\top$, the identity permutation. Consider the matrix

$$C := \begin{pmatrix} 0 & 1 & & & -1 \\ 1 & 0 & 1 & & \not\prec \\ & 1 & & & \\ & & & \ddots & \\ & & & & 1 \\ -1 & \not\prec & & & 1 & 0 & 1 \\ & & & & 1 & 0 & 0 \end{pmatrix} \in \mathbb{S}_n^0.$$

It is easy to see that the minimum over all $C \bullet D(\pi)$, $\pi \in S(n)$, is attained only in $\pi = \iota, \iota^-$ with the value 0. Moreover, for any non-empty proper subset U of $[n]$, we have $C \bullet D(\chi^U) = 0$ if U is incident to ι and $C \bullet D(\chi^U) > 0$ otherwise. Hence, we have that $D(\iota) + D(N_\iota)$ is equal to the set of all points in Q_n which satisfy the valid inequality $C \bullet X \geq 0$ with equality. Out of this matrix C we will now construct a matrix C' and a right hand side such that only some of the subsets incident to ι fulfill the inequality with equality. To do so let U_0 be a subsets of $[n]$ incident to ι . If, for each $U \subset [n]$ incident to ι but different from U_0 , we increase the matrix entries $C_{\max U, \max U+1}$ and $C_{\max U+1, \max U}$ by one, we obtain an inequality $C' \bullet X \geq 0$

which is valid for Q_n and such that the set of all points of Q_n which are satisfied with equality is precisely the edge of $D(\iota) + D(N_\iota)$ generated by the half-lines $D(\iota) + \mathbb{R}_+D(\chi^{U_0})$.

ii. That the defining half lines are extremal has just been proved in *i*. The converse statement follows from Fact 5.18 and the fact that the extreme points of Q_n are precisely the vertices of P_n , which are of the form $D(\pi)$, for $\pi \in S(n)$. \square

5.5.2. Unbounded edges in $\overline{Q_n}$. We have just identified some unbounded edges of $\overline{Q_n} = P_n + \text{CUT}_n$ starting at a particular vertex $D(\pi)$ of this polyhedron. We now set off to characterize *all* unbounded edges of $\overline{Q_n}$. Clearly, the unbounded edges are of the form $D(\pi) + \mathbb{R}_+D(\chi^U)$, but not all these half-lines are edges. For a permutation π and a non-empty subset $U \subsetneq [n]$, we say that $D(\pi) + \mathbb{R}_+D(\chi^U)$ is the half-line *defined by the pair* $\pi \nearrow U$. In this section, we characterize the pairs $\pi \nearrow U$ which have the property that the half-lines they define are edges. For this, we make the following definition.

Definition 5.21. Let π be a permutation, and let U be a subset of $[n]$. We say that U is *almost incident* to π , if there exists a $k \in [n-1]$ such that $U = \pi^{-1}([k-1] \cup \{k+1\})$.

We can now state our theorem.

Theorem 5.22. *For all $n \geq 3$, the unbounded edges of $\overline{Q_n}$ are precisely the half-lines defined by those pairs $\pi \nearrow U$, for which neither U nor $\complement U$ is almost incident to π .*

From Theorem 5.22, we have the following consequences.

Corollary 5.23. *For $n \geq 4$, the number of unbounded edges issuing from a vertex of $\overline{Q_n} = P_n + C_n$ is $2^{n-1} - n$.*

Corollary 5.24. *For $n \geq 4$, the extremal half-lines containing an extreme point of Q_n are a proper subset of the unbounded edges issuing from the same vertex of $\overline{Q_n}$.*

PROOF. We have $n-1 < 2^{n-1} - n$ if $n \geq 4$. \square

Corollary 5.25. *The convex set Q_n is closed if and only if $n \leq 3$.*

Major parts of the proof of the above stated theorem work in an inductive fashion by reducing to the case when $n \in \{3, 4, 5, 6\}$. We will present the cases $n = 3$ and $n = 4$ as examples, which also helps motivating the definitions we require for the proof.

We will switch to a more “visual” notation of the subsets of $[n]$ by identifying a set U with a “word” of length n over $\{0, 1\}$ having a 1 in the j th position iff $j \in U$ — it is just the row-vector $(\chi^U)^\top$.

Example 5.26 (Unbounded edges of $\overline{Q_3}$). We deal with the case $n = 3$ “visually” by regarding Fig. 5.1. There are two edges starting at each vertex. In fact, with some computation, it can be seen that the unbounded edges containing $D(\iota)$ are

$$\begin{aligned} M\begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} + \mathbb{R}_+M\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} &= \begin{pmatrix} 0 & 1 & 2 \\ 1 & 0 & 1 \\ 2 & 1 & 0 \end{pmatrix} + \mathbb{R}_+\begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \quad \text{and} \\ M\begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} + \mathbb{R}_+M\begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix} &= \begin{pmatrix} 0 & 1 & 2 \\ 1 & 0 & 1 \\ 2 & 1 & 0 \end{pmatrix} + \mathbb{R}_+\begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}; \quad \text{while} \\ M\begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} + \mathbb{R}_+M\begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} &= \begin{pmatrix} 0 & 1 & 2 \\ 1 & 0 & 1 \\ 2 & 1 & 0 \end{pmatrix} + \mathbb{R}_+\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} \end{aligned}$$

is not an edge. This agrees with Proposition 5.20, because the sets 100 and 110 are incident to ι , while 101 and 010 are not. Moreover, the set 101 is almost incident to ι and 010 is its complement. Thus, Theorem 5.22 is true for the special case when $\pi = \iota$. For the other permutations, the easiest thing to do is to use symmetry. We describe this in the next remark.

Remark 5.27. For every $\sigma, \pi \in S(n)$ and $U \subset [n]$ we have the following.

- (i) Due to symmetry the pair $\pi \nearrow U$ defines an edge of $\overline{Q_n}$ if and only if the pair $\pi \circ \sigma \nearrow \sigma^{-1}(U)$ defines an edge of $\overline{Q_n}$.
- (ii) U is incident to π if and only if $\sigma^{-1}(U)$ is incident to $\pi \circ \sigma$.
- (iii) U is almost incident to a permutation π if and only if $\sigma^{-1}(U)$ is almost incident $\pi \circ \sigma$.
- (iv) $\mathbb{C}U$ is almost incident to a permutation π if and only if U is almost incident to π^- .

PROOF. Can be checked using the definitions of $\pi \nearrow U$ and U being incident respectively almost incident of π . \square

We now give the first general result as a step towards the proof of Theorem 5.22.

Lemma 5.28. *If $\pi \in S(n)$ and $U \subset [n]$ is almost incident π , then the half-line $D(\pi) + \mathbb{R}_+D(\chi^U)$ defined by the pair $\pi \nearrow U$ is not an edge of $\overline{Q_n}$.*

PROOF. By the above remarks on symmetry, it is sufficient to prove the claim for the identical permutation $\iota \in S(n)$. Consider a $k \in [n-1]$, and let $\pi' := \langle k, k+1 \rangle$ be the transposition exchanging k and $k+1$, and let $U := [k-1] \cup \{k+1\}$. Then a little computation shows that $D(\chi^U)$ can be written as a conic combination of vectors defining rays issuing from $D(\iota)$ as follows:

$$D(\chi^U) = D(\chi^{[k]}) + (D(\pi') - D(\iota)).$$

Hence $D(\iota) + \mathbb{R}_+D(\chi^U)$ is not an edge. \square

Note that by applying Remark 5.27, the Lemma 5.28 implies that if $\mathbb{C}U$ is almost incident π , then the pair $\pi \nearrow \mathbb{C}U$ does not define an edge of $\overline{Q_n}$.

Before we proceed, we note the following easy consequence of Farkas' Lemma.

Lemma 5.29. *The following are equivalent:*

- (i) *The half-line $D(\iota) + \mathbb{R}_+D(\chi^U)$ defined by the pair $\iota \nearrow U$ is an edge of $\overline{Q_n}$.*
- (ii) *There exists a matrix D satisfying the following constraints:*

$$(33a) \quad D \bullet D(\pi) > D \bullet D(\iota) \quad \forall \pi \neq \iota, \iota^-,$$

$$(33b) \quad D \bullet D(\chi^{U'}) > D \bullet D(\chi^U) = 0 \quad \forall U' \neq U, \mathbb{C}U.$$

- (iii) *There exists a matrix C satisfying*

$$(34a) \quad C \bullet D(\pi) \geq C \bullet D(\iota) \quad \forall \pi \neq \iota, \iota^-,$$

$$(34b) \quad C \bullet D(\chi^{U'}) \geq 0 \quad \forall U' \neq U, \mathbb{C}U,$$

$$(34c) \quad C \bullet D(\chi^U) < 0.$$

Condition (33) is easier to check for individual matrices, but condition (34) will be needed in a proof below.

We move on to the next example which both provides some cases needed for the proof of Theorem 5.22 and motivates the following definitions.

Let U be a subset of $[n]$ and consider its representation as a word of length n . We say that a maximal sequence of consecutive 0s in this word is a *valley* of U . In other words, a valley is an inclusion wise maximal subset $[l, l+j] \subset \mathbb{C}U$. Accordingly, a maximal sequence of consecutive 1s is called a *hill*. A valley and a hill meet at a *slope*. Thus the number of slopes is the number of occurrences of the patterns 01 and 10 in the word, or in other words, the number of $k \in [n-1]$ with $k \in U$ and $k+1 \notin U$ or vice versa. If all valleys and hills of a subset U of $[n]$ consist of only one element (as for example in 10101) or, equivalently, if U has the maximal possible number $n-1$ of slopes, or, equivalently, if U consists of all odd or all even numbers in $[n]$, we speak of an *alternating set*.

Lemma 5.30. *For every set $\{W_1, \dots, W_r\}$ of non-empty proper subsets of $[n]$ incident on π , there is a matrix C such that the minimum $C \bullet D(\sigma)$ over all $\sigma \in S(n)$ is attained solely in π and π^- , and that $C \bullet D(\chi^{U'}) \geq 0$ for every non-empty proper subset U' of $[n]$ where equality holds precisely for the sets W_i and their complements. This implies that $D(\pi) + \text{cone}\{D(\chi^{W_1}), \dots, D(\chi^{W_r})\}$ is a face of the polyhedron $\overline{Q_n} = P_n + \text{CUT}_n$.*

PROOF. Follows from Proposition 5.9. \square

Example 5.31 (Unbounded edges of $\overline{Q_4}$). We consider the edges of $\overline{Q_4}$ containing $D(i) = D(i^-)$ (this is justified by Remark 5.27). We distinguish the sets U by their number of slopes. Clearly, a set U with a single slope is incident either to i or to i^- , and we have already dealt with that case in Lemma 5.30. The following sets have two slopes: 0100, 0110, 0010, 1011, 1001, and 1101. We only have to consider 1011, 1001, and 1101, because the others are their complements. The first one, 1011, is almost incident i^- , and the last one, 1101, is almost incident to i , so we know that the pairs $i \nearrow 1011$ and $i \nearrow 1101$ do not define edges of $\overline{Q_4}$ by Lemma 5.28. For the remaining set with two slopes, 1001, the following matrix satisfies property (34) with C replaced by C^{1001} and U by 1001:

$$C^{1001} := \begin{pmatrix} 0 & 1 & -2 & 1 \\ 1 & 0 & 3 & -2 \\ -2 & 3 & 0 & 1 \\ 1 & -2 & 1 & 0 \end{pmatrix}.$$

The two alternating sets (i. e., sets with tree slopes) are 1010 and 0101, which are almost incident to i and i^- respectively. This concludes the discussion of $\overline{Q_4}$.

Having settled some of the cases for small values of n , we give the result by which the reduction to smaller n is performed, which is an important ingredient for settling Theorem 5.22. The following lemma shows that unbounded edges of $\overline{Q_n}$ can be “lifted” to a larger polyhedron $\overline{Q_{n+k}}$.

Lemma 5.32. *Let U_0 be a non-empty proper subset of $[n]$ whose word has the form $a1b$ for two (possibly empty) words a, b . For any $k \geq 0$ define the subset U_k of $[n+k]$ by its word*

$$U_k := a \underbrace{1 \dots 1}_{k+1} b.$$

If the pair $i_n \nearrow U_0$ defines an edge of $\overline{Q_n}$, then the pair $i_{n+k} \nearrow U_k$ defines an edge of $\overline{Q_{n+k}}$.

Note that the lemma also applies to consecutive zeroes, by exchanging the respective set by its complement.

PROOF. Let $C \in \mathbb{S}_n$ be a matrix satisfying conditions (34) for $U := U_0$. Fix $k \geq 1$ and let $n' := n+k$. We will construct a matrix $C' \in \mathbb{S}_{n'}$ satisfying (34) for $U := U_k$. For a “big” real number $\omega \geq 1$ define a matrix $B_\omega \in \mathbb{S}_{k+1}$ whose entries are zero except for those connecting j and $j+1$, for $j \in [k]$:

$$B_\omega := \begin{pmatrix} 0 & \omega & & & \\ \omega & 0 & \omega & & \nearrow \\ & \omega & & & \\ & & \ddots & & \\ \nearrow & & & \omega & 0 & \omega \\ & & & \omega & 0 & \omega \\ & & & & \omega & 0 \end{pmatrix}.$$

We use this matrix to put a heavy weight on the “path” which we “contract.” For our second ingredient, let l_a denote the length of the word a and l_b the length of the word b (note that $l_a = 0$

and $l_b = 0$ are possible). Then we define

$$B_- := \begin{pmatrix} +1 & \dots & +1 \\ \mathbf{0}_{k-1} & \dots & \mathbf{0}_{k-1} \\ -1 & \dots & -1 \end{pmatrix} \in \mathbb{M}((k+1) \times l_a) \quad \text{and}$$

$$B_+ := \begin{pmatrix} -1 & \dots & -1 \\ \mathbf{0}_{k-1} & \dots & \mathbf{0}_{k-1} \\ +1 & \dots & +1 \end{pmatrix} \in \mathbb{M}((k+1) \times l_b),$$

where $\mathbf{0}_{k-1}$ stands for a column of $k-1$ zeros. Putting these matrices together we obtain an $n' \times n'$ -matrix B :

$$B := \begin{pmatrix} \not\! / & B_-^\top & \not\! / \\ B_- & B_\omega & B_+ \\ \not\! / & B_+^\top & \not\! / \end{pmatrix}.$$

Now it is easy to check that for any $\pi' \in \pi[n']$ we have $B \bullet D(\pi') \geq B \bullet D(\iota)$. Moreover let $\pi' \in \pi[n']$ satisfy $B \bullet D(\pi') < B \bullet D(\iota) + 1$. By exchanging π' with π'^- , we can assume that $\pi'(1) < \pi'(n')$. It is easy to see that such a π' then has the following ‘‘coarse structure’’

$$(35) \quad \begin{aligned} \pi'([l_a]) &\subset [l_a] \\ \pi'([n'] \setminus [n' - l_b]) &\subset [n'] \setminus [n' - l_b] \\ \pi'(j) &= j \quad \forall j \in \{l_a + 1, \dots, l_a + k + 1\}. \end{aligned}$$

Thus the matrix B enforces that the ‘‘coarse structure’’ of a $\pi' \in \pi[n']$ minimizing $B \bullet D(\pi')$ coincides with ι . We now modify the matrix C to take care of the ‘‘fine structure’’. For this, we split C into matrices $C_{11} \in \mathbb{S}_{l_a}^0$, $C_{22} \in \mathbb{S}_{l_b}^0$, $C_{12} \in \mathbb{M}(l_a \times l_b)$, $C_{21} = C_{12}^\top \in \mathbb{M}(l_b \times l_a)$, and vectors $c \in \mathbb{R}^{l_a}$, $d \in \mathbb{R}^{l_b}$ as follows:

$$C = \begin{pmatrix} C_{11} & c & C_{12} \\ c^\top & 0 & d^\top \\ C_{21} & d & C_{22} \end{pmatrix}.$$

Then we define the ‘‘stretched’’ matrix $\check{C} \in \mathbb{S}_{n'}^0$ by

$$\check{C} := \begin{pmatrix} C_{11} & c & \not\! / & \mathbf{0} & C_{12} \\ c^\top & 0 & & 0 & \mathbf{0}^\top \\ \not\! / & & \not\! / & & \not\! / \\ \mathbf{0}^\top & 0 & & 0 & d^\top \\ C_{21} & \mathbf{0} & \not\! / & d & C_{22} \end{pmatrix}$$

where the middle $\not\! /$ has dimensions $(k-1) \times (k-1)$. Finally we let $C' := B + \varepsilon \check{C}$, where $\varepsilon > 0$ is small. We show that C' satisfies (34).

We first consider $C' \bullet D(\chi^{U'})$ for non-empty subsets $U' \subsetneq [n']$. Note that, if U' contains $\{l_a + 1, \dots, l_a + k + 1\}$, then for $U := U' \setminus \{l_a + 1, \dots, l_a + k + 1\}$, we have $C' \bullet D(\chi^{U'}) = C \bullet D(\chi^U)$. Thus we have $C' \bullet D(\chi^{U_k}) = C \bullet D(\chi^{U_0}) < 0$ proving (34c) for C' and U_k . For every other U' with $C' \bullet D(\chi^{U'}) < 0$, if ω is big enough, then either U' or $\mathcal{C}U'$ contains $\{l_a + 1, \dots, l_a + k + 1\}$, and w.l.o.g. we assume that U' does. By (34b) applied to C and U , we know that this implies $U = U_0$ or $U = \mathcal{C}U_0$ and hence $U' = U_k$ or $\mathcal{C}U' = U_k$. Thus, (34b) holds for C' and U_k .

Second, we address the permutations. To show (34a), let $\pi' \in S(n)$ be given which minimizes $C' \bullet D(\pi')$. Again, by replacing π' by π'^- if necessary, we assume $\pi'(1) < \pi'(n')$ w.l.o.g. If ε is small enough, we know that π' has the coarse structure displayed in (35). This

implies that we can define a permutation $\pi \in S(n)$ by letting

$$\pi(j) := \begin{cases} \pi'(j) & \text{if } j \in [l_a], \\ \pi'(j) = j & \text{if } j = l_a + 1, \\ \pi'(j - k) + k & \text{if } j \in [n] \setminus [l_a + 1]. \end{cases}$$

An easy but lengthy computation (see [Sei09] for the details) shows that

$$\begin{aligned} C' \bullet D(\pi') - C' \bullet D(\iota_{n'}) &\geq \varepsilon \left[C \bullet D(\pi) + k \cdot C \bullet \begin{pmatrix} \times_{l_a \times l_a} & \times \\ \times & \times_{l_b \times l_b} \end{pmatrix} \right. \\ &\quad \left. - \left(C \bullet D(\iota_n) + k \cdot C \bullet \begin{pmatrix} \times_{l_a \times l_a} & \times \\ \times & \times_{l_b \times l_b} \end{pmatrix} \right) \right] \\ &= \varepsilon [C \bullet D(\pi) - C \bullet D(\iota_n)] \geq 0. \end{aligned}$$

Thus (34a) holds. \square

Example 5.33. We give an example for the application of Lemma 5.32. For $n = 5$, consider the half-line defined by the pair $\iota \nearrow 11001$. The set 11001 can be reduced to 1001 by contracting the hill $1 - 2$. To do so we set

$$C^{11001} := \varepsilon \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -2 & 1 \\ 0 & 1 & 0 & 3 & -2 \\ 0 & -2 & 3 & 0 & 1 \\ 0 & 1 & -2 & 1 & 0 \end{pmatrix} + \begin{pmatrix} 0 & \omega & -1 & -1 & -1 \\ \omega & 0 & 1 & 1 & 1 \\ -1 & 1 & 0 & 0 & 0 \\ -1 & 1 & 0 & 0 & 0 \\ -1 & 1 & 0 & 0 & 0 \end{pmatrix}$$

for a small $\varepsilon > 0$ and a big $\omega \geq 1$.

After these preparations we can tackle the proof of the theorem.

PROOF OF THEOREM 5.22. By Remark 5.27, we only need to consider $\pi = \iota$. We distinguish the sets U by their numbers of slopes.

One slope. This is equivalent to U or $\complement U$ being incident to ι . We have treated this case in Lemma 5.30.

Two slopes. The complete list of all possibilities, up to complements, and how they are dealt with is summarized in Table 2. In this table, 0 stands for a valley consisting of a single zero while $0 \dots 0$ stands for a valley consisting of at least two zeros (the same with hills). The matrices for the reduced words satisfying (34) can be found in the appendix on page 61. The condition (34) can be verified by some case distinctions.

TABLE 2. List of all sets with two slopes (up to complement)

Word			Edge?	Why?
Hill 1	Valley	Hill 2		
1	0	1	no	almost incident to ι
1	0	$1 \dots 1$	no	almost incident to ι^-
1	$0 \dots 0$	1	yes	reduce to $n = 4$, 1001, by Lemma 5.32
1	$0 \dots 0$	$1 \dots 1$	yes	reduce to $n = 4$, 1001, by Lemma 5.32
$1 \dots 1$	0	1	no	almost incident to ι
$1 \dots 1$	0	$1 \dots 1$	yes	reduce to $n = 5$, 11011, by Lemma 5.32
$1 \dots 1$	$0 \dots 0$	1	yes	reduce to $n = 4$, 1001, by Lemma 5.32
$1 \dots 1$	$0 \dots 0$	$1 \dots 1$	yes	reduce to $n = 5$, 11011, by Lemma 5.32

Three slopes. This case can be tackled using the same methods we applied in the case above. Table 3 gives the results.

TABLE 3. List of all sets with three slopes (up to complement)

Word				Edge?	Why?
Hill 1	Valley 1	Hill 2	Valley 2		
1	0	1	0	no	almost incident to ι
1	0	1	0...0	no	almost incident to ι
1	0	1...1	0	yes	reduce to $n = 5$, 10110, by Lemma 5.32
1	0	1...1	0...0	yes	reduce to $n = 5$, 10110, by Lemma 5.32
1	0...0	1	0	yes	reduce to $n = 5$, 10010, by Lemma 5.32
1	0...0	1	0...0	yes	reduce to $n = 5$, 10010, by Lemma 5.32
1	0...0	1...1	0	yes	reduce to $n = 5$, 10010, by Lemma 5.32
1	0...0	1...1	0...0	yes	reduce to $n = 5$, 10110, by Lemma 5.32
1...1	0	1	0	no	almost incident to ι
1...1	0	1	0...0	no	almost incident to ι
1...1	0	1...1	0	yes	reduce to $n = 5$, 10110, by Lemma 5.32
1...1	0	1...1	0...0	yes	reduce to $n = 5$, 10110, by Lemma 5.32
1...1	0...0	1	0	yes	reduce to $n = 5$, 10010, by Lemma 5.32
1...1	0...0	1	0...0	yes	reduce to $n = 5$, 10010, by Lemma 5.32
1...1	0...0	1...1	0	yes	reduce to $n = 5$, 10010, by Lemma 5.32
1...1	0...0	1...1	0...0	yes	reduce to $n = 5$, 10010, by Lemma 5.32

$s \geq 4$ slopes. Using Lemma 5.32, we reduce such a set to an alternating set with s slopes showing that for all these sets U the pair $\iota \nearrow U$ defines an edge of $\overline{Q_n}$. This is in accordance with the statement of the theorem because sets which are almost incident to ι can have at most three slopes. The statement for alternating sets is proven by induction on n in Lemma 5.34 below. Note that the starts of the inductions in the proof of that lemma are $n = 5$ and $n = 6$ for even or odd s respectively.

This concludes the proof of the theorem. \square

We now present the inductive construction which we need for the case of an even number of $s \geq 4$ slopes.

Lemma 5.34. *For an integer $n \geq 5$ let U be an alternating subset of $[n]$. The pair $\iota \nearrow U$ defines an edge of $\overline{Q_n}$.*

PROOF. We first prove the case when n is odd.

The proof is by induction over n . For the start of the induction we consider $n = 5$ and offer the matrix $C^{10101} \in \mathfrak{S}_5$ in Table 4 of the appendix satisfying (33). We will need this matrix in the inductive construction.

Now set $E^5 := C^{10101}$ and assume that the pair $\iota \nearrow U^-$ defines an edge of $\overline{Q_n}$ where U^- is an alternating subset of $[n]$. W.l.o.g., we assume that $U^- = 10\dots 01$. There exists a

matrix $E^- \in \mathbb{S}_n^0$ for which (33) holds. We will construct a matrix $E \in \mathbb{S}_{n+2}^0$ satisfying (33) for $U := 010 \dots 010$.

We extend E^- to a $(n+2) \times (n+2)$ -Matrix

$$\widehat{E} := \begin{pmatrix} E^- & \mathbf{0} & \mathbf{0} \\ \mathbf{0}^\top & \mathbf{0} & \mathbf{0} \\ \mathbf{0}^\top & \mathbf{0} & \mathbf{0} \end{pmatrix}.$$

We do the same with E^5 , except on the other side:

$$\widehat{E}^5 := \begin{pmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0}^\top \\ \mathbf{0} & \mathbf{0} & \mathbf{0}^\top \\ \mathbf{0} & \mathbf{0} & E^5 \end{pmatrix}.$$

Now we let $E := \widehat{E} + \widehat{E}^5$ and check the conditions (33) on E . These are now easily verified.

For the even case we guarantee the start of induction investigating $n = 6$. We give a matrix C^{101010} satisfying (33) in Table 4 in the appendix. (Note that 101010 is the only set which is not incident to ι , is not almost incident to ι or ι^- , cannot be reduced by Lemma 5.32 and is no complement of sets of any of these three types.) The induction is proved in the same way by using the matrix $E^6 := C^{101010}$. \square

5.6. Concluding Remarks

The \mathbb{R} -embeddable 1-separated metrics are a natural and fascinating class of metrics, which are also of some practical importance due to their connection with graph layout problems. We have established some fundamental properties of such metrics, and also initiated a study of their convex hull and its closure.

There are several possible avenues for future research. First, one could search for new valid or facet-defining inequalities. Second, one could study the complexity of the separation problems associated with various families of inequalities, which would be essential if one wished to use the inequalities within a cutting-plane algorithm. Third, it would be interesting to know whether the *bounded* edges of the convex hull, or its closure, have a simple combinatorial interpretation.

5.7. Appendix: Table of cases

TABLE 4. Matrices certifying unbounded edges of Q_n

n	Slopes	Matrix
4	2	$C^{1001} := \begin{pmatrix} 0 & 1 & -2 & 1 \\ 1 & 0 & 3 & -2 \\ -2 & 3 & 0 & 1 \\ 1 & -2 & 1 & 0 \end{pmatrix}$
5	2	$C^{11011} := \begin{pmatrix} 0 & 8 & -6 & -1 & -1 \\ 8 & 0 & 2 & 9 & -3 \\ -6 & 2 & 0 & 5 & -7 \\ -1 & 9 & 5 & 0 & 11 \\ -1 & -3 & -7 & 11 & 0 \end{pmatrix}$
5	3	$C^{10110} := \begin{pmatrix} 0 & 2 & 2 & 1 & -3 \\ 2 & 0 & 0 & -2 & 2 \\ -2 & 0 & 0 & 2 & 0 \\ 1 & -2 & 2 & 0 & 1 \\ -3 & 2 & 0 & 1 & 0 \end{pmatrix}$
5	3	$C^{10010} := \begin{pmatrix} 0 & 2 & -2 & 2 & -2 \\ 2 & 0 & 4 & -3 & 1 \\ -2 & 4 & 0 & 1 & 1 \\ 2 & -3 & 1 & 0 & 1 \\ -2 & 1 & 1 & 1 & 0 \end{pmatrix}$
5	4	$C^{10101} := \begin{pmatrix} 0 & 0 & 3 & -2 & -1 \\ 0 & 0 & 1 & 1 & -2 \\ 3 & 1 & 0 & 1 & 3 \\ -2 & 1 & 1 & 0 & 0 \\ -1 & -2 & 3 & 0 & 0 \end{pmatrix}$
6	5	$C^{101010} := \begin{pmatrix} 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & 1 & -2 & 0 \\ 1 & 1 & 0 & 1 & 3 & -2 \\ -1 & 1 & 1 & 0 & 0 & 1 \\ 0 & -2 & 3 & 0 & 0 & 1 \\ 0 & 0 & -2 & 1 & 1 & 0 \end{pmatrix}$

The VPN problem with concave costs

Jointly with

Samuel Fiorini (U.L.B., Brussels),

Gianpaolo Oriolo (Università di Roma Tor Vergata, Rome),

and Laura Sanità (Università di Roma Tor Vergata, Rome)

Abstract. Only recently Goyal, Olver and Shepherd (*Proc. STOC*, 2008) proved that the symmetric Virtual Private Network Design (sVPN) problem has the tree routing property, namely, that there always exists an optimal solution to the problem whose support is a tree. Combining this with previous results by Fingerhut, Suri and Turner (*J. Alg.*, 1997) and Gupta, Kleinberg, Kumar, Rastogi and Yener (*Proc. STOC*, 2001), sVPN can be solved in polynomial time.

In this paper we investigate an APX-hard generalization of sVPN, where the contribution of each edge to the total cost is proportional to some non-negative, concave and non-decreasing function of the capacity reservation. We show that the tree routing property extends to the new problem, and give a constant-factor approximation algorithm for it. We also show that the undirected uncapacitated single-source minimum concave-cost flow problem has the tree routing property when the cost function has some property of symmetry.

6.1. Introduction

All the problems considered in this paper involve a (finite) simple, undirected, connected graph $G = (V, E)$ that represents a communication network. The graph comes with a vector $c \in \mathbb{Q}_+^E$ describing edge costs, and a vector $b \in \mathbb{Z}_+^V$ pertaining to the traffic departing from or arriving at each vertex; the exact interpretation depends on the problem. A vertex v with $b_v > 0$ is referred to as a *terminal*. We denote the set of terminals by W . Also, we let B be the sum of all components of b . Thus, $W = \{v \in V \mid b_v > 0\}$ and $B = \sum_{v \in V} b_v$.

In the *symmetric Virtual Private Network design* (sVPN) problem, the vertices want to communicate with each other. However, the exact amount of traffic between pairs of vertices is not known in advance. Instead, for each vertex v the cumulative amount of traffic that it can send or receive is bounded from above by b_v . The aim is to install minimum cost capacities on the edges of the graph supporting any possible communication scenario, where the cost for installing one unit of capacity on edge e equals c_e .

A *set of traffic demands* $D = \{d_{uv} \mid \{u, v\} \subseteq W\}$ specifies for each unordered pair of terminals $\{u, v\} \subseteq W$ the amount $d_{uv} \in \mathbb{Q}_+$ of traffic between u and v . A set D is *valid* if it respects the upper bounds on the traffic of the terminals. That is,

$$\sum_{u \in W} d_{uv} \leq b_v \quad \text{for all terminals } v \in W.$$

A solution to the instance of sVPN defined by the triple (G, b, c) consists of a collection of paths \mathcal{P} containing exactly one u - v path P_{uv} in G for each unordered pair u, v of terminals, and a vector $\gamma \in \mathbb{Q}_+^E$ describing the capacity to be installed on each edge. Such a set of paths \mathcal{P} , together with capacity reservations γ , is called a *virtual private network*. A virtual private network is *feasible* if all valid sets of traffic demands can be routed without exceeding the

reserved capacities, in case all traffic between terminals u and v is routed along path P_{uv} , that is,

$$\gamma_e \geq \sum_{\{u,v\} \subseteq W: e \in P_{uv}} d_{uv} \quad \text{for all edges } e \in E.$$

Given a collection of paths \mathcal{P} as above, one may compute in polynomial time the capacity reservations γ_e for $e \in E$ in order to obtain a feasible virtual private network [GKK⁺01, ILO06].

The *concave symmetric Virtual Private Network Design* (cSVPN) problem is defined similarly as sVPN. The total cost of virtual private network (\mathcal{P}, γ) is now

$$(36) \quad \sum_{e \in E} c_e f(\gamma_e),$$

where $f: [0, B] \rightarrow \mathbb{R}_+$ is concave, non-decreasing and such that $f(0) = 0$. (We assume we are given oracle access to f , see Section 6.1.4 below.) An instance of cSVPN is described by a quadruple (G, b, c, f) .

In the *Concave Routing* (CR) problem, one of the terminals is marked as *root*. We denote the root by r . For each vertex v , the number b_v describes the *demand* at the vertex. We remark that, by the choice of r , there is a demand $b_r > 0$ at the root. This is a dummy demand that does not play any role in the problem¹.

A solution to CR consists of a collection \mathcal{P} of simple r - v paths P_v , one path for each terminal v distinct from the root. We call such a collection a *routing*. We denote by $x_e(\mathcal{P})$ the amount of flow routed on the edge e by \mathcal{P} . Thus, $x_e(\mathcal{P}) = \sum_{v \in W \setminus \{r\}: e \in P_v} b_v$. The cost of a routing is then:

$$(37) \quad \sum_{e \in E} c_e g(x_e(\mathcal{P})),$$

where $g: [0, B] \rightarrow \mathbb{R}_+$ is a concave function such that $g(0) = 0$. (Once again, we assume that we are given oracle access to g .) An instance of CR is then defined by a quintuple (G, r, b, c, g) . We remark that CR may be viewed as an undirected uncapacitated single-source minimum concave-cost flow problem [GP90].

We are interested in the following restrictions of CR. The instances of the *non-decreasing Concave Routing* (ndCR) problem are those for which g is non-decreasing. In this case, we use the letter f instead of g whenever possible. The instances of the *axis-symmetric Concave Routing* (sCR) problem are those for which g is (*axis*-)symmetric, that is, $g(B - x) = g(x)$ for all $x \in [0, B]$. In this case, we use the letter h instead of g whenever possible. Finally, the instances of the *Pyramidal Routing* (PR) problem [GKOS08] are those for which $g(x) = \min\{x, B - x\}$ for all $x \in [0, B]$. In this case, we use the letter p instead of g .

The various problems considered here and their relationships are illustrated in Fig. 6.1. Notice that cSVPN, sCR and ndCR are all APX-hard because they admit the minimum Steiner tree problem as a special case.

A feasible solution to one of the problems described above is a *tree solution* if the support of the capacity vector γ or the union of the paths in \mathcal{P} induces a tree in G . To make the terminology concise, we say that an instance of either cSVPN or sCR has the tree routing property provided one of its optimal solution is a tree solution.

¹We use this convention in order to be consistent with previous published work [GKOS08].

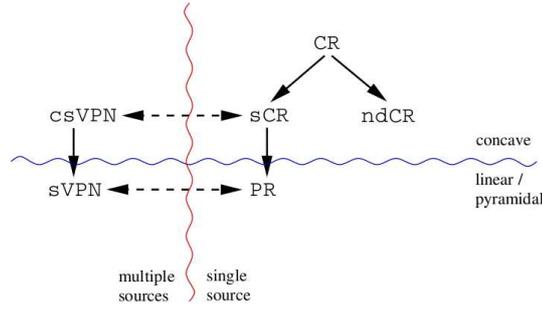


FIGURE 6.1. The problems considered in this article. Bold arrows indicate specialization, dashed arrows indicate equivalence.

6.1.1. Previous work. It was shown by Fingerhut *et al.* [FST97] and later, independently, by Gupta *et al.* [GKK⁺01] that $sVPN$ can be solved in polynomial time if it has the tree routing property, that is, each instance has an optimal solution that is a tree solution². Subsequently, it has been discussed [GK02] and then conjectured [ER04, ILO06] that $sVPN$ has the tree routing property. This has become known as the *VPN tree routing conjecture*. The conjecture has first been proved for the case of cycles [HKS07, GKOS08], and then in general graphs [GOS08b].

Goyal *et al.* [GOS08b] prove the VPN tree routing conjecture by proving that PR has the tree routing property. This result was initially proposed as a conjecture by Grandoni, Kaibel, Oriolo and Skutella [GKOS08], together with a proof that it implies the VPN tree routing conjecture. Remarkably, Goyal *et al.* [GOS08b] also show that two results are equivalent, that is, $sVPN$ has the tree routing property if and only if PR has the tree routing property.

6.1.2. Our contribution. First, we show that $csVPN$ has the tree routing property. Our proof goes as follows. On the one hand, we build upon the result by Goyal *et al.* [GOS08b] to show that sCR has the tree routing property. On the other hand, we show that there is an equivalence between $csVPN$ and sCR , so that $csVPN$ has the tree routing property too.

Second, we study approximation algorithms for $csVPN$. For general f , using known results on the so-called Single Source Buy at Bulk (SSBB) problem [GKPR07, GI06], we give a 24.92-approximation algorithm. For a restricted class of functions f , by reducing to the so-called Single Source Rent or Buy (SSRB) problem [EGRS08], we show that a 2.92-approximation algorithm exists.

Third, although sCR and $ndCR$ both have the tree routing property, we show that this is not the case for the general CR problem.

6.1.3. Outline. In Section 6.2 we prove our main statements: $csVPN$ and sCR have the tree routing property. The proof uses as a basis an equivalence, stated in Section 6.2.1, between $csVPN$ and sCR . We show that, when b is a 0-1 vector, solving an $csVPN$ instance (G, b, c, f) amounts to solving an sCR instance of the form (G, r, b, c, h) where r is one of the terminals and h is obtained by symmetrizing f . Moreover, the $csVPN$ instance has an optimal solution that is a tree solution if and only if the corresponding sCR instance has an optimal solution that is a tree solution. This allows us to focus only on sCR . By combining one decisive polyhedral observation with the fact that PR has the tree routing property [GOS08b], we show that sCR has the tree routing property, which then implies that $csVPN$ also has the tree routing property.

In Section 6.3 we give a constant factor approximation algorithm for $csVPN$. Our approximation algorithm works by reduction to the Single Source Buy at Bulk (SSBB) problem. The reduction is in two steps. First, we observe in Section 6.3.1 that the approximation algorithm for SSBB due to Grandoni and Italiano [GI06], that is a variation of the algorithm of Gupta,

²Such a solution can be obtained in polynomial time by solving a single all-pair shortest paths problem.

Kumar, Pål and Roughgarden [GKPR07], gives an approximation algorithm for ndCR with the same approximation factor. Then, we show in Section 6.3.2 how to turn any approximation algorithm for ndCR into an approximation algorithm for cSVPN with the same approximation factor. Combining both steps, we obtain a ρ -approximation algorithm for cSVPN from the ρ -approximation algorithm for SSBB [GI06], where $\rho = 24.92$. Using a subset of the tools developed, we give a 2.92-approximation algorithm for cSVPN when the function f is to be of the type $f(x) = \min\{\mu x, M\}$ for positive constants μ and M . Here, we resort to the Single Source Rent or Buy (SSRB) problem, for which the best known approximation factor currently is 2.92 [GKPR07, EGRS08].

In Section 6.4, we give an instance of CR such that no tree solution is optimal, thereby showing that CR does not have the tree property.

6.1.4. Fractional problems and value-giving oracles. Before starting Section 6.2, we conclude this section by providing necessary extra details.

We define the *fractional* version of CR (denoted by frac-CR) where we allow, for each terminal $v \neq r$, to fractionally split the b_v units of flow from r to v along several r - v paths. Formally, a fractional routing \mathcal{P} specifies, for each terminal $v \neq r$, a set \mathcal{P}_v of simple r - v paths and, for each path $P \in \mathcal{P}_v$, an amount of flow $\beta_v(P) \in \mathbb{R}_+$ such that $b_v = \sum_{P \in \mathcal{P}_v} \beta_v(P)$. The cost of a routing is as in Eq. (37) above, with $x_e(\mathcal{P}) := \sum\{\beta_v(P) \mid v \in W \setminus \{r\}, P \in \mathcal{P}_v, e \in P\}$.

It results from the concavity of g (see, e.g., Goyal *et al.* [GOS08b, Lemma 2.2]) that there always exists an optimal solution to CR that is unsplittable, i.e., that routes all flows from the source to a terminal on a unique path, even when we allow fractional flows. Therefore, the frac-CR problem and CR problem are essentially equivalent.

The problem frac-ndCR is defined similarly. This last problem is closely related to a known variant of the Single Source Buy at Bulk problem, see Section 6.3.1 for details.

Finally, in the cSVPN (resp. CR) problem, we assume that we are given oracle access to the function f (resp. g). That is, we are given access to a subroutine that, given a rational $x \in [0, B]$, returns a non-negative rational $f(x)$ (resp. $g(x)$) whose size is polynomial in the size of x . The computation is assumed to take constant time.

6.2. The tree routing property

We show here that both cSVPN and sCR have the tree routing property. We start by proving, in Section 6.2.1, an equivalence between the two problems when b is a 0-1 vector. Then, in Section 6.2.2, we prove the tree routing property for sCR, and thus also for cSVPN.

6.2.1. Equivalence of cSVPN and sCR instances in the binary case. Here we restrict ourselves to instances where b is a 0-1 vector. In this case, the number of terminals is B and, for any routing \mathcal{P} , there are precisely $x_e(\mathcal{P})$ paths in \mathcal{P} using the edge e . For $f: [0, B] \rightarrow \mathbb{R}_+$ concave and non-decreasing with $f(0) = 0$, we define

$$(38) \quad h: [0, B] \rightarrow \mathbb{R}_+ : x \mapsto \begin{cases} f(x) & \text{if } x \leq B/2, \\ f(B-x) & \text{if } x > B/2. \end{cases}$$

Then h is concave and axis-symmetric and has $h(0) = 0$. The proof of the next lemma builds upon previous results of Gupta *et al.* [GKK⁺01], Grandoni *et al.* [GKOS08] and Goyal *et al.* [GOS08b].

Lemma 6.1. *Let (G, b, c, f) be a cSVPN instance with $b \in \{0, 1\}^V$, and h as in (38). There exists a choice of a root $r \in W$ such that the sCR instance (G, r, b, c, h) has the same optimum value as the cSVPN instance. Moreover, for any such choice of root r , the corresponding sCR instance has the tree routing property if and only if the cSVPN instance has the tree routing property.*

PROOF. Let (\mathcal{P}, γ) be a feasible virtual private network for (G, b, c, f) , with $\mathcal{P} = \{P_{uv} \mid \{u, v\} \subseteq W\}$. For each possible root $r \in W$, let \mathcal{P}_r denote the routing consisting of all paths of \mathcal{P} one of whose ends is r . So $\mathcal{P}_r := \{P_{rv} : v \in W \setminus \{r\}\}$. It is known [GKK⁺01, Theorem 3.2] [GKOS08, Lemma 3] that the following holds:

$$\gamma_e \geq \frac{1}{B} \sum_{r \in W} \min\{x_e(\mathcal{P}_r), B - x_e(\mathcal{P}_r)\}.$$

Since f is concave and non-decreasing we have:

$$\begin{aligned} \sum_{e \in E} c_e f(\gamma_e) &\geq \sum_{e \in E} c_e f\left(\frac{1}{B} \sum_{r \in W} \min\{x_e(\mathcal{P}_r), B - x_e(\mathcal{P}_r)\}\right) \\ &\geq \frac{1}{B} \sum_{e \in E} c_e \sum_{r \in W} f(\min\{x_e(\mathcal{P}_r), B - x_e(\mathcal{P}_r)\}) = \frac{1}{B} \sum_{r \in W} \sum_{e \in E} c_e h(x_e(\mathcal{P}_r)). \end{aligned}$$

Hence, the optimum value for the cSVPN instance (G, b, c, f) is at least the optimum value of the sCR instance (G, r, b, c, h) for some choice of root $r \in W$. Note that, if (\mathcal{P}, γ) is a tree solution, then \mathcal{P}_r is also a tree solution for any $r \in W$. It is not difficult to see that, in this case, the cost of the routing \mathcal{P}_r is not dependent on the root r . It follows that, given a tree solution to the cSVPN instance (G, b, c, f) , we can construct a tree solution to the sCR instance (G, r, b, c, h) that is not more costly, for any choice of root r .

Conversely, take any $r \in W$ and suppose that we are given a routing \mathcal{P}_r for some sCR instance (G, r, b, c, h) , where this time $\mathcal{P}_r := \{P_v \mid v \in W \setminus \{r\}\}$. Following [GOS08b], we define a collection of paths $\mathcal{Q} = \{Q_{uv} \mid \{u, v\} \subseteq W\}$, where Q_{uv} is any u - v path in the component of the symmetric difference $P_u \Delta P_v$ containing u and v . Let δ_e be the minimum amount of capacity that we must install on each edge e so that (\mathcal{Q}, δ) is a feasible virtual private network for (G, b, c, f) . Goyal *et al.* [GOS08b] show that the following holds:

$$\delta_e \leq \min\{x_e(\mathcal{P}_r), B - x_e(\mathcal{P}_r)\}.$$

Since f is non-decreasing, we have

$$\sum_{e \in E} c_e f(\delta_e) \leq \sum_{e \in E} c_e f(\min\{x_e(\mathcal{P}_r), B - x_e(\mathcal{P}_r)\}) = \sum_{e \in E} c_e h(x_e(\mathcal{P}_r)).$$

Hence, the optimum value of the cSVPN instance (G, b, c, f) is at most the optimum value of any sCR instance of the form (G, r, b, c, h) for $r \in W$. Again, note that if \mathcal{P}_r is a tree solution to (G, r, b, c, h) , then (\mathcal{Q}, δ) is a tree solution to the cSVPN instance (G, b, c, f) . Therefore, given a tree solution to the sCR instance (G, r, b, c, h) , we can construct a tree solution to the cSVPN instance (G, b, c, f) , that is not more costly. The statement easily follows. \square

6.2.2. Proof of the tree routing property for sCR. In this section, we will show how the tree routing property for sCR follows from the tree routing property for PR.

Theorem 6.2. *The tree routing property holds for sCR.*

Our approach is simple and geometric: We associate polyhedra with instances of sCR in such a way that the tree routing property for an instance can be expressed as a property of the extreme points of the associated polyhedron. We then show how the transition from the pyramidal function to an arbitrary concave axis-symmetric function h amounts to a transformation of the corresponding polyhedra, which preserves the property of the extreme points.

Recall that, for a set $Z \subseteq \mathbb{R}_+^E$, the dominant $\text{dom } Z$ of Z is defined as follows:

$$\text{dom } Z := \{z' \in \mathbb{R}^E \mid \text{there exists some } z \in Z \text{ with } z \leq z'\}.$$

Here, and below, comparisons between vectors are component-wise. Given G, r, b , and h as above in the definition of sCR, a routing \mathcal{P} defines a point $y(h, \mathcal{P}) \in \mathbb{R}_+^E$ by $y_e(h, \mathcal{P}) :=$

$h(x_e(\mathcal{P}))$ for all $e \in E$. We define the sCR-polyhedron $P_{(G,r,b,h)}$ as the dominant of the convex hull of the points $y(h, \mathcal{P})$, where \mathcal{P} ranges over all routings. Now, finding a routing that is minimum w.r.t. some non-negative cost vector c is equivalent to minimizing the linear function $y \mapsto c^T y$ over the sCR-polyhedron . We note an easy consequence of this fact.

Lemma 6.3. *Given G, r, b and h , as above the following are equivalent:*

- (i) *For every extreme point y of $P_{(G,r,b,h)}$, there exists a tree solution \mathcal{T} such that $y = y(h, \mathcal{T})$.*
- (ii) *For every $c \geq 0$, the sCR instance (G, r, b, c, h) has the tree routing property. \square*

We say that a mapping $\Phi: \mathbb{R}_+^E \rightarrow \mathbb{R}_+^E$ is *concave* if $\Phi(tx + (1-t)y) \geq t\Phi(x) + (1-t)\Phi(y)$ holds for every $t \in [0, 1]$ and $x, y \in \mathbb{R}_+^E$. Similarly, we say that such a mapping is *non-decreasing* if $x \leq y$ implies $\Phi(x) \leq \Phi(y)$. The key observation to realizing that the tree routing property for sCR is a consequence of the tree routing property for PR is the following.

Lemma 6.4. *Let p denote the pyramidal function, and h be as above. There exists a non-decreasing concave function $\Phi: \mathbb{R}_+^E \rightarrow \mathbb{R}_+^E$ such that $\Phi(y(p, \mathcal{P})) = y(h, \mathcal{P})$ for all routings \mathcal{P} .*

PROOF. For every e , we define $\Phi_e(y) := h(y_e)$ whenever $y_e \leq B/2$ and $\Phi_e(y) := h(B/2)$ if $y_e \geq B/2$. The properties are readily verified, since any axis-symmetric concave function $h: [0, B] \rightarrow \mathbb{R}_+$ is non-decreasing in the interval $[0, B/2]$, and $y_e(p, \mathcal{P})$ is always at most $B/2$. \square

The final ingredient is the following elementary geometric fact.

Lemma 6.5. *If $\Phi: \mathbb{R}_+^E \rightarrow \mathbb{R}_+^E$ is non-decreasing and concave, and Y is a finite set of points in \mathbb{R}_+^E , then every extreme point of $\text{dom conv } \Phi(Y)$ is the image under Φ of an extreme point of $\text{dom conv } Y$. In other words, Φ maps a subset of the extreme points of $\text{dom conv } Y$ onto the extreme points of $\text{dom conv } \Phi(Y)$.*

PROOF. Consider an extreme point z of $\text{dom conv } \Phi(Y)$. If some point in $\Phi^{-1}(z)$ is an extreme point of $\text{dom conv } Y$, then we are done. Otherwise, pick any point y in $Y \cap \Phi^{-1}(z)$. By assumption, there exist extreme points $y_1, \dots, y_n \in Y \setminus \Phi^{-1}(z)$ and coefficients $\lambda_1, \dots, \lambda_n \geq 0$ with $\sum \lambda_j = 1$ such that $y \geq \sum_{j=1}^n \lambda_j y_j$. Hence, the assumptions on Φ imply $z = \Phi(y) \geq \Phi(\sum_{j=1}^n \lambda_j y_j) \geq \sum_{j=1}^n \lambda_j \Phi(y_j)$. Because $\Phi(y_j) \neq z$ for all j , the point z is not an extreme point of $\text{dom conv } \Phi(Y)$, a contradiction. \square

Combining the previous two lemmas and this fact we obtain our theorem.

PROOF OF THEOREM 6.2. We give the proof for 0-1 demands first. For this situation, Goyal *et al.* [GOS08b] have proven the tree routing property for all instances of PR . Lemma 6.3 implies that for every extreme point of $P_{(G,r,b,p)}$ there exists a tree solution defining it. By Lemmas 6.4 and 6.5, we know that this is also true for the extreme points of $P_{(G,r,b,h)}$. Another application of Lemma 6.3 yields the result for 0-1 demands.

Now consider an sCR instance (G, r, b, c, h) such that b is not a 0-1 vector. We define a new instance $(\tilde{G}, \tilde{r}, \tilde{b}, \tilde{c}, h)$, as follows. For each terminal v with $b_v \geq 2$, we add $k := b_v$ pendant edges vu_1, \dots, vu_k with cost zero to the graph. Then, we let $\tilde{b}_v := 0$ and $\tilde{b}_{u_i} := 1$ for $i = 1, \dots, k$. Finally, we let \tilde{r} be one of the new vertices pending from r except if $b_r = 1$ in which case we let $\tilde{r} = r$. Since the new instance has an optimal solution that is a tree solution, it follows that also the original instance has an optimal solution that is a tree solution. \square

Corollary 6.6. *The tree routing property holds for csVPN .*

PROOF. First, consider an csVPN instance (G, b, c, f) with $b_v \in \{0, 1\}$ for each $v \in V$. Here the statement follows from Lemma 6.1 and Theorem 6.2. The case where some terminals have demand greater than 1 can be reduced to the previous one by the same arguments as in the proof of Theorem 6.2. \square

Remark. As pointed out by an anonymous referee, the results of this section still hold in case a concave function f_e (resp. h_e) is associated to each edge e of the graph, and allowing different edges to have different functions associated to them.

6.3. Approximation algorithms

6.3.1. An approximation algorithm for ndCR . Our approximation algorithm for csVPN is based on an approximation algorithm for ndCR . The approximation algorithm for ndCR is, in its turn, related to an approximation algorithm for the *Single Source Buy at Bulk* (SSBB) problem.

The latter problem is defined as follows: we are given a (finite, simple, undirected, connected) graph $G = (V, E)$ with edge costs $c \in \mathbb{Q}_+^E$, where each vertex $v \in V$ wants to exchange an amount of flow $b_v \in \mathbb{Z}_+$ with a common source vertex r . In order to support the traffic, we can install cables on edges. Specifically we can choose among k different cables: each cable $i \in \{1, \dots, k\}$ provides $\mu(i) \in \mathbb{Q}_+ \setminus \{0\}$ units of capacity at price $p(i) \in \mathbb{Q}_+ \setminus \{0\}$. For each $i \in \{1, \dots, k-1\}$, it is assumed that $\mu(i) < \mu(i+1)$ and $\frac{p(i)}{\mu(i)} \geq \frac{p(i+1)}{\mu(i+1)}$. The latter inequality is referred to as the *economy of scale principle*. An instance of SSBB is therefore defined by a quintuple (G, r, b, c, K) , where $K = \{(\mu(i), p(i)) \mid i = 1, \dots, k\}$ describes the different cable types.

A solution to SSBB consists of a multiset κ_e of cables to install on each edge $e \in E$. Repetitions are allowed, that is, several cables of the same type can be installed on some edge.

We point out that there is some confusion in the literature in the definition of SSBB, because in some papers SSBB is defined as above, and in some other papers the SSBB problem is defined as the problem we call frac-ndCR . In this paper, when referring to SSBB we always mean the version with cables. It is a known fact (see, e.g., Gupta *et al.* [GKPR07]) that from an approximation viewpoint, the two formulations are equivalent up to a factor of 2. However, we here show how to adapt the 24.92-approximation algorithm for SSBB described in [GI06], in order to obtain an algorithm with *the same* approximation ratio for ndCR .

Theorem 6.7. *There exists a 24.92-approximation algorithm for ndCR .*

PROOF. We start with a description of a simple approximation preserving reduction from ndCR to SSBB. Let $I = (G, r, b, c, f)$ be an instance of ndCR . Consider the instance $J = (G, r, b, c, K)$ of SSBB obtained by setting $K := \{(1, f(1)), (2, f(2)) \dots, (B, f(B))\}$. The capacity of the cables are non-decreasing because f is non-decreasing. Since $f(0) = 0$ and f is concave, $x \mapsto f(x)/x$ is non-increasing, and thus the economy of scale principle holds. It is easy to see that (i) given a solution to I there exists a solution to J of the same cost; (ii) from a solution κ to J one can build, in time polynomial in the sizes of I and κ , a solution to I that does not cost more. In other words, we could run the 24.92-approximation algorithm for SSBB on J and obtain a 24.92-approximate solution to I .

However, we point out that the size of J is not always bounded by a polynomial in the size of I , because B could be exponentially large. To address this issue, we rely on a key fact used in the analysis of Grandoni and Italiano [GI06], which we now describe. Given any instance $(\tilde{G}, \tilde{r}, \tilde{b}, \tilde{c}, \tilde{K})$ of SSBB, they select a subset $\{i_1, \dots, i_{k'}\} \subseteq \{1, \dots, k\}$ of cables with the following properties: $i_1 = 1$, $i_{k'} = k$ and, for all $t \in \{1, \dots, k' - 2\}$, cable i_{t+1} is the

smallest such that

$$(39a) \quad p(i_{t+1} + 1) \geq \alpha p(i_t)$$

$$(39b) \quad \frac{p(i_{t+1})}{\mu(i_{t+1})} \leq \frac{1}{\beta} \frac{p(i_t)}{\mu(i_t)}.$$

with $\alpha := 3.1207$ and $\beta := 2.4764$. Then, they find a 24.92-approximate solution to the SSBB instance using only cables in the following subset:

$$(40) \quad \tilde{K}' := \{(\mu(i_1), p(i_1)), \dots, (\mu(i_{k'}), p(i_{k'}))\}.$$

with a running time polynomial in the size of $(\tilde{G}, \tilde{r}, \tilde{b}, \tilde{c}, \tilde{K}')$.

For our purpose, the point is therefore to find a list of cables K' as in (40) satisfying (39a) and (39b), with respect to the instance J , in time polynomial in $\log B$. To construct the list of cables K' , we let $i_1 := 1$. If i_t has been found, we search for the $(t+1)$ th cable i_{t+1} as follows.

Firstly, since f is increasing, given $p(i_t)$, a binary search in $\{i_t + 1, \dots, B\}$ finds the smallest value i' satisfying (39a) with i_{t+1} replaced by i' . If no such i' satisfies (39a), we let $i_{t+1} := k$ and $k' := t + 1$. If i' does exist, since $x \mapsto f(x)/x$ is non-increasing, the smallest possible value for i_{t+1} satisfying (39b) in the range $\{i', \dots, B\}$ can be found by binary search. Again, if no i_{t+1} satisfies (39b), we let $i_{t+1} := k$ and $k' := t + 1$.

Recalling that $\mu(i_t) = i_t$, from (39a) and (39b) it follows: $i_{t+1} \geq \beta \cdot i_t \cdot \frac{f(i_{t+1})}{f(i_t)} \geq \beta \cdot i_t \cdot \frac{i_{t+1}}{i_{t+1}+1} \frac{f(i_{t+1}+1)}{f(i_t)} \geq \frac{1}{2} \alpha \cdot \beta \cdot i_t$. Therefore the number of selected cables is $O(\log_{\frac{\alpha\beta}{2}} B) = O(\log B)$ and each cable can be found in time $O(\log B)$. The result follows. \square

6.3.2. An approximation algorithm for csVPN . In order to state our approximation algorithm for csVPN we need two further results from the literature.

First, let (G, b, c, f) be an instance of the csVPN problem. Consider a tree T spanning all the terminals in W . For each pair of terminals $\{u, v\} \subseteq W$ there is a unique u - v path in T . These paths form a collection of paths that we denote \mathcal{P}^T . It is straightforward to compute the minimum amount of capacity γ_e^T we have to reserve on each edge e of T in order to obtain a feasible virtual private network from \mathcal{P}^T . We denote $z(\mathcal{P}^T, \gamma^T)$ the cost of this virtual private network.

For any choice of root $r \in V(T)$, one can similarly derive from T a tree solution to the ndCR instance (G, r, b^r, c, f) , where we let $b_v^r := b_v$ for all vertices $v \neq r$, and $b_r^r := \max\{b_r, 1\}$ ³. We denote the resulting routing by \mathcal{P}_r^T and its cost by $z(\mathcal{P}_r^T)$. The next lemma is known [GKK⁺01, Lemma 2.1], [ILO06, Lemma 2.4]. For the sake of completeness, we give a sketch of its proof.

Lemma 6.8. *Let T , \mathcal{P}^T , γ^T and \mathcal{P}_r^T (for $r \in V(T)$) be as above. Then, there exists a vertex r of T such that $\gamma_e^T = x_e(\mathcal{P}_r^T)$ for all edges e of T . For that choice of r , we have $z(\mathcal{P}^T, \gamma^T) = z(\mathcal{P}_r^T)$.*

PROOF SKETCH. Consider an edge e of T . The removal of e from T determines a partition of the set of terminals W into two of its subsets, say $W_1(e)$ and $W_2(e)$. For definiteness, we assume that $W_1(e)$ and $W_2(e)$ are chosen in such a way that $\sum_{v \in W_1(e)} b_v \leq \sum_{v \in W_2(e)} b_v$. Then, the minimum capacity reservation γ_e^T for edge e is simply $\sum_{v \in W_1(e)} b_v$. By breaking ties consistently and orienting each edge $e \in E(T)$ towards $W_1(e)$, we can turn T into an arborescence. Letting r denote the root of this arborescence, we have $\gamma_e^T = x_e(\mathcal{P}_r^T)$ for all edges e of T . \square

Second, suppose that we are given a solution \mathcal{P}_r to an instance (G, r, b^r, c, f) of ndCR . As observed by Goyal *et al.* [GOS08b] and used in Lemma 6.1 above, we can build a feasible

³Recall that in the definition of CR , we assume to have a positive (dummy) demand at the root.

solution (\mathcal{Q}, δ) to the instance (G, b, c, f) of csVPN as follows: for each pair of terminals u, v , choose the path Q_{uv} to be any path in $P_u \Delta P_v$ from u to v , where P_u and P_v respectively denote the unique r - u and r - v paths in \mathcal{P}_r . Define \mathcal{Q} as the collection formed by all the paths Q_{uv} . As mentioned in the introduction, we may efficiently deduce from \mathcal{Q} the minimum capacity reservation δ such that (\mathcal{Q}, δ) is a feasible virtual private network. Let $z(\mathcal{Q}, \delta)$ denote the cost of this virtual private network. We will need the next lemma. We omit its proof because it is not difficult (see Goyal *et al.* [GOS08b] for a stronger result):

Lemma 6.9. *Let \mathcal{P}_r , \mathcal{Q} and δ be as above. Then, we have $\delta_e \leq x_e(\mathcal{P}_r)$ for all edges e of G . Thus $z(\mathcal{Q}, \delta) \leq z(\mathcal{P}_r)$.*

We are now ready to complete the description and analysis of our approximation algorithm for csVPN . The input to the algorithm is a csVPN instance (G, b, c, f) . In the proof below, we use $\text{OPT}(\cdot)$ to denote the cost of an optimal solution to the corresponding csVPN or ndCR instance.

Algorithm 1 Approximation algorithm for csVPN

- (1) For each $r \in V$, compute a ρ -approximate sol. \mathcal{P}_r to the ndCR instance (G, r, b^r, c, f) .
 - (2) Let r^* be such that $z(\mathcal{P}_{r^*}) = \min_{r \in V} z(\mathcal{P}_r)$.
 - (3) From \mathcal{P}_{r^*} , build a solution (\mathcal{Q}, δ) to the csVPN instance (G, b, c, f) as in Lemma 6.9.
 - (4) Output (\mathcal{Q}, δ) .
-

Theorem 6.10. *Algorithm 1 is a ρ -approximation algorithm for csVPN .*

PROOF. From Corollary 6.6, we know that there exists a tree T such that $z(\mathcal{P}^T, \gamma^T) = \text{OPT}(G, b, c, f)$. By Lemma 6.8, $\min_{r \in V(T)} z(\mathcal{P}_r^T) \leq z(\mathcal{P}^T, \gamma^T)$. Since \mathcal{P}_r^T is a solution to the ndCR instance (G, r, b^r, c, f) , it follows $\min_{r \in V(T)} z(\mathcal{P}_r^T) \geq \min_{r \in V} \text{OPT}(G, r, b^r, c, f)$. Let $\tilde{r} \in V$ be such that $\min_{r \in V} \text{OPT}(G, r, b^r, c, f) = \text{OPT}(G, \tilde{r}, b^{\tilde{r}}, c, f)$. By choice of r^* , $z(\mathcal{P}_{r^*}) \leq z(\mathcal{P}_{\tilde{r}}) \leq \rho \text{OPT}(G, \tilde{r}, b^{\tilde{r}}, c, f)$. From Lemma 6.9, $z(\mathcal{Q}, \delta) \leq z(\mathcal{P}_{r^*})$. Putting everything together, we conclude $z(\mathcal{Q}, \delta) \leq \rho \text{OPT}(G, b, c, f)$, as desired. \square

By Theorem 6.7, there exists a ρ -approximation algorithm for csVPN with $\rho = 24.92$.

Notice that Algorithm 1 preserves the function f when the approximation algorithm for ndCR is invoked. In particular, if f belongs to a restricted class of functions where ndCR has a small approximation factor, our algorithm will have same factor on the corresponding instances. In particular, if f is defined as $f(x) := \min\{\mu x, M\}$, for two positive numbers μ, M , then the ndCR instance constructed in Algorithm 1 from a csVPN instance is, except for decomposing into paths, just an instance of the so-called Single Source Rent or Buy (SSRB) problem [GKPR07, EGRS08]. Hence, our results imply an approximation-preserving reduction from csVPN —restricted to instances such that $f(x) := \min\{\mu x, M\}$ for some positive numbers μ and M —to SSRB. The best known approximation algorithm for SSRB known to us is the one by Gupta *et al.* [GKPR07], which has an approximation factor of 2.92, as was shown by Eisenbrand, Grandoni, Rothvoß, and Schäfer [EGRS08].

6.4. A remark on general concave funtions

It is known (see, e.g., [KM00]) that the tree routing property is satisfied by every CR instance such that g is non-decreasing, and it follows from our results that this also holds when g is axis-symmetric. A natural question arises: is the tree routing property satisfied by all CR instances?

The example below shows that this is not the case, even if $g(x) \leq g(B - x)$, for each $x \in [0, B/2]$, and G is a cycle.

Example 6.11. Consider an instance (G, r, b, c, g) of the CR problem, where $G = (V, E)$ is a cycle with vertex set $V := \{0, 1, 2, 3, 4\}$ and edge set $E := \{\{i, i + 1\} \mid i \in V\}$ (the sum is modulo 5). Let $r := 0$; let $b_i := 1$ for $i \in V$; let $c_e := M$ for $e = \{3, 4\}$, $c_e := M + \epsilon$ for $e = \{0, 1\}$, $c_e := 0$ otherwise. Finally, let g be defined as the linear interpolation of the following points: $g(0) = 0$, $g(2) = 2$, $g(3) = 2 + 2\epsilon$, $g(5) = 0$. It is easy to check that g is concave, non-negative, non-axis-symmetric and $g(x) \leq g(B - x)$, for each $x \in [0, B/2]$.

Consider the routing \mathcal{P} where the paths from 0 to i go counterclockwise (that is, have the edge $\{0, 4\}$ as their first edge) for $i = 1, 2, 3$, while the path from 0 to 4 goes clockwise. The cost of this solution is $(2 + \epsilon)M + \epsilon$, and it is easy to check that taking ϵ and M respectively small and big enough, every tree solution costs more.

Acknowledgments

We thank the three anonymous referees for providing remarks that guided us when revising this manuscript.

Note added in preparation: Following the results in this manuscript, an alternative proof of the fact that the tree routing property holds for cSVPN has been given [GOS08a]. This proof, however, does not show that it also holds also for sCR .

CHAPTER 7

Small minors

Jointly with
Samuel Fiorini (U.L.B., Brussels),
Gwenaël Joret (U.L.B., Brussels),
and David Wood (University of Melbourne)

Abstract. A fundamental result in structural graph theory states that every graph with large average degree contains a large complete graph as a minor. We prove this result with the extra property that the minor is small with respect to the order of the whole graph. More precisely, we describe functions f and h such that every graph with n vertices and average degree at least $f(t)$ contains a K_t -model with at most $h(t) \cdot \log n$ vertices. The logarithmic dependence on n is best possible. In general, we prove that $f(t) \leq 2^{t-1} + \varepsilon$. For $t \leq 4$, we determine the least value of $f(t)$; in particular $f(3) = 2 + \varepsilon$ and $f(4) = 4 + \varepsilon$. For $t \leq 4$, we establish similar results for graphs embedded on surfaces, where the size of the K_t -model is bounded.

7.1. Introduction

A fundamental result in structural graph theory states that every sufficiently dense graph contains a large complete graph as a minor¹. More precisely, there is a minimum function $f(t)$ such that every graph with average degree at least $f(t)$ contains a K_t -minor. Mader [Mad67] first proved that $f(t) \leq 2^{t-2}$, and later proved that $f(t) \in O(t \log t)$ [Mad68]. Kostochka [Kos82, Kos84] and Thomason [Tho84, Tho01] proved that $f(t) \in \Theta(t\sqrt{\log t})$; see [Tho06] for a survey of related results.

Here we prove similar results with the extra property that the K_t -minor is ‘small’. This idea is evident when $t = 3$. A graph contains a K_3 -minor if and only if it contains a cycle. Every graph with average degree at least 2 contains a cycle, whereas every graph G with average degree at least 3 contains a cycle of length $O(\log |G|)$. That is, high average degree forces a short cycle, which can be thought of as a small K_3 -minor.

In general, we measure the size of a K_t -minor via the following definition. A K_t -model in a graph G consists of t connected subgraphs B_1, \dots, B_t of G , such that $V(B_i) \cap V(B_j) = \emptyset$ and some vertex in B_i is adjacent to some vertex in B_j for all distinct $i, j \in \{1, \dots, t\}$. The B_i are called *branch sets*. Clearly a graph contains a K_t -minor if and only if it contains a K_t -model. We measure the size of a K_t -model by the total number of vertices, $\sum_{i=1}^t |B_i|$. Our main result states that every sufficiently dense graph contains a small model of a complete graph.

Theorem 7.1. *There are functions f and h such that every graph G with average degree at least $f(t)$ contains a K_t -model with at most $h(t) \cdot \log |G|$ vertices.*

For fixed t , the logarithmic upper bound in Theorem 7.1 is within a constant factor of optimal, since every K_t -model contains a cycle, and for all $d \geq 3$ and $n > 3d$ such that nd

¹We consider simple, finite, undirected graphs G with vertex set $V(G)$ and edge set $E(G)$. Let $|G| := |V(G)|$ and $\|G\| := |E(G)|$. A graph H is a *minor* of a graph G if H is isomorphic to a graph obtained from a subgraph of G by contracting edges.

is even, Chandran [Cha03] constructed a graph with n vertices, average degree d , and girth at least $(\log_d n) - 1$. (The *girth* of a graph is the length of a shortest cycle.)

In this paper we focus on minimising the function f in Theorem 7.1 and do not calculate h explicitly. In particular, Theorem 7.8 proves Theorem 7.1 with $f(t) \leq 2^{t-1} + \varepsilon$ for any $\varepsilon > 0$ (where the function h also depends on ε). Note that for Theorem 7.1 and all our results, the proofs can be easily adapted to give polynomial algorithms that compute the small K_t -model.

For $t \leq 4$, we determine the least possible value of $f(t)$ in Theorem 7.1. The $t = 2$ case is trivial—one edge is a small K_2 -minor. To force a small K_3 -model, average degree 2 is not enough, since every K_3 -model in a large cycle uses every vertex. On the other hand, we prove that average degree $2 + \varepsilon$ forces a cycle of length $O_\varepsilon(\log |G|)$; see Lemma 7.4. For $t = 4$ we prove that average degree $4 + \varepsilon$ forces a K_4 -model with $O_\varepsilon(\log |G|)$ vertices; see Theorem 7.5. This result is also best possible. Consider the square of an even cycle C_{2n}^2 , which is a 4-regular graph illustrated in Figure 7.1. If the base cycle is (v_1, \dots, v_{2n}) then $C_{2n}^2 - \{v_i, v_{i+1}\}$ is outerplanar for each i . Since outerplanar graphs contain no K_4 -minor, every K_4 -model in C_{2n}^2 contains v_i or v_{i+1} , and thus contains at least n vertices.

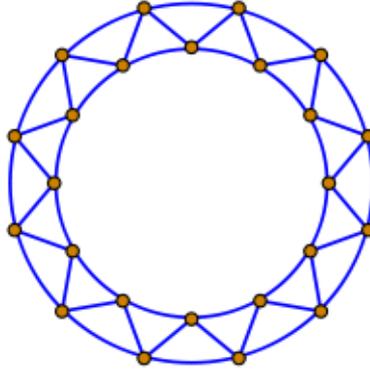


FIGURE 7.1. C_{24}^2

Motivated by Theorem 7.1, we then consider graphs that contain K_3 and K_4 -models of bounded size (not just small with respect to $|G|$). First, we prove that planar graphs satisfy this property. In particular, every planar graph with average degree at least $2 + \varepsilon$ contains a K_3 -model with $O(\frac{1}{\varepsilon})$ vertices (Theorem 7.11). This bound on the average degree is best possible since a cycle is planar and has average degree 2. Similarly, every planar graph with average degree at least $4 + \varepsilon$ contains a K_4 -model with $O(\frac{1}{\varepsilon})$ vertices (Theorem 7.19). Again, this bound on the average degree is best possible since C_{2n}^2 is planar and has average degree 4. These results generalise for graphs embedded on other surfaces (Theorems 7.21 and 7.24).

Finally, we mention three other results in the literature that force a model of a complete graph of bounded size.

- Kostochka and Pyber [KP88] proved that for every integer t and $\varepsilon > 0$, every n -vertex graph with at least $4^{t^2} n^{1+\varepsilon}$ edges contains a subdivision of K_t with at most $\frac{1}{\varepsilon} 7t^2 \log t$ vertices; see [Jia11] for recent related results.

Note that Theorem 7.1 can be proved by adapting the proof in [KP88]. As far as we can tell, this method does not give a bound better than $f(t) \leq 16^t + \varepsilon$ (ignoring lower order terms). This bound is inferior to our Theorem 7.8, which proves $f(t) \leq 2^{t-1} + \varepsilon$. Also note that the method in [KP88] can be adapted to prove the following.

Theorem 7.2. *There is a function h such that for every integer $t \geq 2$ and real $\varepsilon > 0$, every graph G with average degree at least $4^{t^2} + \varepsilon$ contains a subdivision of K_t with at most $h(t, \varepsilon) \cdot \log |G|$ division vertices per edge.*

- Kühn and Osthus [KO06] proved that every graph with minimum degree at least t and girth at least 27 contains a K_{t+1} -subdivision. Every graph with average degree at least $2t$ contains a subgraph with minimum degree at least t . Thus every graph with average degree at least $2t$ contains a K_{t+1} -subdivision or a K_3 -model with at most 26 vertices.
- Krivelevich and Sudakov [KS09] proved that for all integers $s' \geq s \geq 2$, there is a constant $c > 0$, such that every $K_{s,s'}$ -free graph with average degree r contains a minor with average degree at least $cr^{1+1/(2s-2)}$. Applying the result of Kostochka [Kos82, Kos84] and Thomason [Tho84] mentioned above, for every integer $s \geq 2$ there is a constant c such that every graph with average degree at least $c(t\sqrt{\log t})^{1-1/(2s-1)}$ contains a K_t -minor or a $K_{s,s}$ -subgraph, in which case there is a K_{s+1} -model with $2s$ vertices.

7.2. Definitions and Notations

See [Die05] for undefined graph-theoretic terminology and notation. For $S \subseteq V(G)$, let $G[S]$ be the subgraph of G induced by S . Let $e(S) := \|G[S]\|$. For disjoint sets $S, T \subseteq V(G)$, let $e(S, T)$ be the number of edges between S and T in G .

A *separation* in a graph G is a pair of subgraphs $\{G_1, G_2\}$, such that $G = G_1 \cup G_2$ and $V(G_1) \setminus V(G_2) \neq \emptyset$ and $V(G_2) \setminus V(G_1) \neq \emptyset$. The *order* of the separation is $|V(G_1) \cap V(G_2)|$. A separation of order 1 corresponds to a cut-vertex v , where $V(G_1) \cap V(G_2) = \{v\}$. A separation of order 2 corresponds to a cut-pair v, w , where $V(G_1) \cap V(G_2) = \{v, w\}$.

See [MT01] for background on graphs embedded in surfaces. Let \mathbb{S}_h be the orientable surface obtained from the sphere by adding h handles. The *Euler genus* of \mathbb{S}_h is $2h$. Let \mathbb{N}_c be the non-orientable surface obtained from the sphere by adding c cross-caps. The *Euler genus* of \mathbb{N}_c is c .

An *embedded graph* means a connected graph that is 2-cell embedded in \mathbb{S}_h or \mathbb{N}_c . A *plane graph* is a planar graph embedded in the plane. Let $F(G)$ denote the set of faces in an embedded graph G . For a face $f \in F(G)$, let $|f|$ be the length of the facial walk around f . For a vertex v of G , let $F(G, v)$ be the multiset of faces incident to v , where the multiplicity of a face f in $F(G, v)$ equals the multiplicity of v in the facial walk around f . Thus $|F(G, v)| = \deg(v)$.

Euler's formula states that $|G| - \|G\| + |F(G)| = 2 - g$ for a connected graph G embedded in a surface with Euler genus g . Note that $g \leq \|G\| - |G| + 1$ since $|F(G)| \geq 1$. The *Euler genus* of a graph G is the minimum Euler genus of a surface in which G embeds.

We now review some well-known results that will be used implicitly (see [Die05, Section 7.3]). If a graph G contains no K_4 -minor then $\|G\| \leq 2|G| - 3$, and if $|G| \geq 2$ then G contains at least two vertices with degree at most 2. Hence, if $\|G\| > 2|G| - 3$ then G contains a K_4 -minor. Similarly, if $|G| \geq 2$ and at most one vertex in G has degree at most 2, then G contains a K_4 -minor.

Throughout this paper, logarithms are binary unless stated otherwise.

7.3. Small K_3 - and K_4 -Models

In this section we prove tight bounds on the average degree that forces a small K_3 - or K_4 -model. The following lemma is at the heart of many of our results. It is analogous to Lemma 1.1 in [KP88]

Lemma 7.3. *There is a function f such that for every two reals $d > d' \geq 2$, every graph G with average degree at least d contains a subgraph with average degree at least d' and diameter at most $f(d, d') \cdot \log |G|$.*

PROOF. We may assume that every subgraph of G has average degree strictly less than d (otherwise, simply consider a minimal subgraph with that property). Let

$$\beta := \frac{d}{d'} > 1 \quad \text{and} \quad f(d, d') := \frac{2}{\log \beta} + 2 .$$

Let v be an arbitrary vertex of G . Let $B_k(v)$ be the subgraph of G induced by the set of vertices at distance at most k from v . Let $k \geq 1$ be the minimum integer such that $|B_k(v)| < \beta \cdot |B_{k-1}(v)|$. (There exists such a k , since $\beta > 1$ and G is finite.) It follows that $\beta^{k-1} \leq |B_{k-1}(v)| \leq |G|$. Thus $B_k(v)$ has diameter at most $2k \leq 2(\log_\beta |G| + 1) \leq f(d, d') \cdot \log |G|$.

We now show that $B_k(v)$ also has average degree at least d' . Let

$$\begin{aligned} A &:= V(B_{k-1}(v)), \\ B &:= V(B_k(v)) \setminus V(B_{k-1}(v)), \\ C &:= V(G) \setminus (A \cup B) . \end{aligned}$$

If $C = \emptyset$, then $B_k(v) = G[A \cup B] = G$, and hence $B_k(v)$ has average degree at least $d \geq d'$. Thus, we may assume that $C \neq \emptyset$. Let t be the average degree of $B_k(v)$. Thus,

$$(41) \quad 2(e(A) + e(B) + e(A, B)) = t \cdot (|A| + |B|) .$$

Since C is non-empty, $G - A$ is a proper non-empty subgraph of G . By our hypothesis on G , this subgraph has average degree strictly less than d ; that is,

$$(42) \quad 2(e(B) + e(C) + e(B, C)) < d \cdot (|B| + |C|) .$$

By (41) and (42) and since $e(A, C) = 0$,

$$\begin{aligned} 2\|G\| &= 2(e(A) + e(B) + e(C) + e(A, B) + e(B, C)) \\ &= t(|A| + |B|) + 2e(C) + 2e(B, C) \\ &< t(|A| + |B|) + d(|B| + |C|) - 2e(B) \\ &\leq d|G| - d|A| + t(|A| + |B|) . \end{aligned}$$

Thus $t(|A| + |B|) > d|A|$ (since $2\|G\| \geq d|G|$). On the other hand, by the choice of k ,

$$\frac{|A|}{|A| + |B|} > \frac{1}{\beta} .$$

Hence

$$t > d \frac{|A|}{|A| + |B|} > \frac{d}{\beta} = d' ,$$

as desired. \square

Lemma 7.4. *There is a function g such that for every real $\varepsilon > 0$, every graph G with average degree at least $2 + \varepsilon$ has girth at most $g(\varepsilon) \cdot \log |G|$,*

PROOF. By Lemma 7.3, G contains a subgraph G' with average degree at least 2 and diameter at most $f(2 + \varepsilon, 2) \cdot \log |G|$. Let T be a breadth-first search tree in G' . Thus T has diameter at most $2f(2 + \varepsilon, 2) \cdot \log |G|$. Since G' has average degree at least 2, G' is not a tree, and there is an edge $e \in E(G') \setminus E(T)$. Thus T plus e contains a cycle of length at most $2f(2 + \varepsilon, 2) \cdot \log |G| + 1$. \square

Theorem 7.5. *There is a function h such that for every real $\varepsilon > 0$, every graph G with average degree at least $4 + \varepsilon$ contains a K_4 -model with at most $h(\varepsilon) \cdot \log |G|$ vertices.*

PROOF. By Lemma 7.3, G contains a subgraph G' with average degree at least $4 + \frac{\varepsilon}{2}$ and diameter at most $f(4 + \varepsilon, 4 + \frac{\varepsilon}{2}) \cdot \log |G|$. Let v be an arbitrary vertex of G' . Let T be a breadth-first search tree from v in G' . Let k be the depth of T . Thus $k \leq f(4 + \varepsilon, 4 + \frac{\varepsilon}{2}) \cdot \log |G|$.

Let $H := G' - E(T)$. Since $\|T\| = |G| - 1$, the graph H has average degree at least $2 + \frac{\varepsilon}{2}$. By Lemma 7.4, H contains a cycle C of length at most $g(\frac{\varepsilon}{2}) \cdot \log |G|$. We will prove the theorem with $h(\varepsilon) := g(\frac{\varepsilon}{2}) + 3f(4 + \varepsilon, 4 + \frac{\varepsilon}{2})$.

Observe that $v \notin V(C)$, since v is isolated in H . A vertex w of C is said to be *maximal* if, in the tree T rooted at v , no other vertex of C is an ancestor of w .

First, suppose that C contains three maximal vertices x, y, z . For $w \in \{x, y, z\}$, let P_w be the unique v - w path in T . Then $C \cup P_x \cup P_y \cup P_z$ contains a K_4 -model with at most $|C| + |P_x - x| + |P_y - y| + |P_z - z| \leq |C| + 3k \leq h(\varepsilon) \cdot \log |G|$ vertices. Now assume that at most two vertices of C are maximal.

Next, suppose that there is a unique maximal vertex x in C . Let i be the distance between v and x in T . Let y be a neighbour of x in C . The vertex y is not maximal, implying there is an ancestor of y in C . Since T is a breadth-first search tree, y is at distance at most $i + 1$ from v in T . However, $xy \notin E(T)$, which implies that x is not an ancestor of y in T , a contradiction.

Finally, suppose there are exactly two maximal vertices x and y in C . If one is closer to v than the other in T , say x is closer than y , then considering a neighbour x' of x in C that is distinct from y again yields a contradiction: x' is not maximal, thus x' has an ancestor in C , and this ancestor must be x . However, this cannot be since $xx' \notin E(T)$. Hence, x and y are at the same distance from v in T .

Let P be an x - y path in C that is not the edge xy . Let x' be the neighbour of x in P , and let y' be the neighbour of y in P . The ancestor of x' in T must be y , since otherwise there would be a path in T between x' and v that avoids both x and y . For the same reason, x is the ancestor of y' in T . Thus, $x'y$ and $y'x$ are both edges of T , and hence $x' \neq y'$. Now, the cycle C plus these two edges gives a K_4 -model with $|C| \leq g(\frac{\varepsilon}{2}) \cdot \log |G| \leq h(\varepsilon) \cdot \log |G|$ vertices. \square

7.4. Small K_t -Models

The following theorem establishes our main result (Theorem 7.1).

Theorem 7.6. *There is a function h such that for every integer $t \geq 2$ and real $\varepsilon > 0$, every graph G with average degree at least $2^t + \varepsilon$ contains a K_t -model with at most $h(t, \varepsilon) \cdot \log |G|$ vertices.*

PROOF. We prove the following slightly stronger statement: Every graph G with average degree at least $2^t + \varepsilon$ contains a K_t -model with at most $h(t, \varepsilon) \cdot \log |G|$ vertices such that each branch set of the model contains at least two vertices.

The proof is by induction on t . For $t = 2$, let $h(t, \varepsilon) := 2$. Here we need only assume average degree at least $2 + \varepsilon$. Some component of G is neither a tree nor a cycle, as otherwise G would have average degree at most 2. It is easily seen that this component contains a path on 4 vertices, yielding a K_2 -model in which each branch set contains two vertices. This model has $4 \leq h(t, \varepsilon) \cdot \log |G|$ vertices, as desired. (Observe that $|G| \geq 4$, since G contains a vertex with degree at least 3.)

Now assume $t \geq 3$ and the claim holds for smaller values of t . Using Lemma 7.3, let G' be a subgraph of G with average degree at least $2^t + \frac{\varepsilon}{2}$ and diameter at most $f(2^t + \varepsilon, 2^t + \frac{\varepsilon}{2}) \cdot \log |G|$. Let $h(t, \varepsilon) := 2 + (t - 1)f(2^t + \varepsilon, 2^t + \frac{\varepsilon}{2}) + h(t - 1, \frac{\varepsilon}{4})$.

Choose an arbitrary edge uv of G' . Define the *depth* of a vertex $w \in V(G')$ to be the minimum distance in G' between w and a vertex in $\{u, v\}$. The *depth* of an edge $xy \in E(G')$ is the minimum of the depth of x and the depth of y .

Considering edges of G' with even depth on one hand, and with odd depth on the other, we obtain two edge-disjoint spanning subgraphs of G' . Since G' has average degree at least $2^t + \frac{\varepsilon}{2}$, one of these two subgraphs has average degree at least $2^{t-1} + \frac{\varepsilon}{4}$. Let H be a component of this subgraph with average degree at least $2^{t-1} + \frac{\varepsilon}{4}$. Observe that every edge of H has the same depth k in G .

If $k = 0$, then $E(H)$ is precisely the set of edges incident to u or v (or both). Thus, every vertex in $V(H) \setminus \{u, v\}$ has degree at most 2 in H . Hence H has average degree less than $4 < 2^{t-1} + \frac{\varepsilon}{4}$, a contradiction. Therefore $k \geq 1$.

Now, by induction, H contains a K_{t-1} -model with at most $h(t-1, \frac{\varepsilon}{4}) \cdot \log |G'|$ vertices such that each of the $t-1$ branch sets B_1, \dots, B_{t-1} has at least two vertices. Thus, each B_i contains an edge of H . Hence, there is a vertex v_i in B_i having depth k in G' . Therefore, there is a path P_i of length k in G' between v_i and some vertex in $\{u, v\}$. Let P_{uv} be the trivial path consisting of the edge uv . Let

$$B_t := P_{uv} \cup \bigcup_{1 \leq i \leq t-1} (P_i - v_i) .$$

The subgraph B_t is connected, contains at least two vertices (namely, u and v), and is vertex disjoint from B_i for all $i \in \{1, \dots, t-1\}$. Moreover, there is an edge between B_t and each B_i , and

$$\begin{aligned} \sum_{1 \leq i \leq t} |B_i| &\leq |B_t| + h(t-1, \frac{\varepsilon}{4}) \cdot \log |G'| \\ &\leq 2 + \sum_{1 \leq i \leq t-1} |P_i - v_i| + h(t-1, \frac{\varepsilon}{4}) \cdot \log |G| \\ &\leq 2 + (t-1)k + h(t-1, \frac{\varepsilon}{4}) \cdot \log |G| \\ &\leq 2 + (t-1)f(2^t + \varepsilon, 2^t + \frac{\varepsilon}{2}) \cdot \log |G| + h(t-1, \frac{\varepsilon}{4}) \cdot \log |G| \\ &\leq h(t, \varepsilon) \cdot \log |G| . \end{aligned}$$

Hence, adding B_t to our K_{t-1} -model gives the desired K_t -model of G . \square

Observe that the obstacle to reducing the lower bound on the average degree in Theorem 7.6 is the case $t = 3$, which we address in the following result.

Lemma 7.7. *There is a function h such that for every real $\varepsilon > 0$, every graph G with average degree at least $4 + \varepsilon$ contains a K_3 -model with at most $h(\varepsilon) \cdot \log |G|$ vertices, such that each branch set contains at least two vertices.*

PROOF. The proof is by induction on $|G| + \|G\|$. We may assume that no proper subgraph of G has average degree at least $4 + \varepsilon$, since otherwise we are done by induction. This implies that G is connected. Note that $|G| \geq 6$ since G has average degree > 4 .

First, suppose that G contains a K_4 subgraph with vertex set X .

Case 1. All edges between X and $V(G) \setminus X$ in G are incident to a common vertex $v \in X$: Let $Y := X \setminus \{v\}$. Then

$$2\|G - Y\| = 2\|G\| - 12 \geq (4 + \varepsilon)|G| - 12 \geq (4 + \varepsilon)|G - Y| ,$$

implying that $G - Y$ also has average degree at least $4 + \varepsilon$, a contradiction.

Case 2. There are two independent edges uu' and vv' between X and $V(G) \setminus X$ in G , where $u, v \in X$: Then $\{u, u'\}, \{v, v'\}, X \setminus \{u, v\}$ is the desired K_3 -model.

Case 3. Some vertex $w \in V(G) \setminus X$ is adjacent to two vertices $u, v \in X$: No vertex in X has a neighbour in $V(G) \setminus (X \cup \{w\})$, as otherwise Case 2 would apply. Since G is connected and $|G| \geq 6$, it follows that w has a neighbour w' outside X . Let x, y be the two vertices in $X \setminus \{u, v\}$. Then $\{w, w'\}, \{u, x\}, \{v, y\}$ is the desired K_3 -model.

This concludes the case in which G contains a K_4 subgraph. Now, assume that G is K_4 -free. By Theorem 7.5, G contains a K_4 -model B_1, \dots, B_4 with at most $h(\varepsilon) \cdot \log |G|$ vertices. Without loss of generality, $|B_1| \geq |B_2| \geq |B_3| \geq |B_4|$ and $|B_1| \geq 2$.

Case 1. $|B_2| \geq 2$: Then $B_1, B_2, B_3 \cup B_4$ is the desired K_3 -model. Now assume that $B_i = \{x_i\}$ for all $i \in \{2, 3, 4\}$.

Case 2. Some x_i is adjacent to some vertex w not in $B_1 \cup B_2 \cup B_3 \cup B_4$: If $i = 2$ then $\{x_2, w\}, B_1, B_3 \cup B_4$ is the desired K_3 -model. Similarly for $i \in \{3, 4\}$.

Case 3. $|B_1| \geq 3$. Then there are two independent edges in G between B_1 and $\{x_2, x_3, x_4\}$, say ux_2 and vx_3 with $u, v \in B_1$ (otherwise, there would be a K_4 subgraph). There is a vertex $w \in B_1 \setminus \{u, v\}$ adjacent to at least one of u, v , say u . Let C be the vertex set of the component of $G[B_1] - \{u, w\}$ containing v . Then $\{u, w\}, C \cup \{x_3\}, \{x_2, x_4\}$ is the desired K_3 -model.

Case 4. $B_1 = \{u, v\}$. As in the previous cases, there are two independent edges in G between $\{u, v\}$ and $\{x_2, x_3, x_4\}$, say ux_2 and vx_3 . At least one of u, v , say u , is adjacent to some vertex w outside $\{u, v, x_2, x_3, x_4\}$, because G is connected with at least 6 vertices, and none of x_2, x_3, x_4 has a neighbour outside $\{u, v, x_2, x_3, x_4\}$. Then $\{u, w\}, \{v, x_3\}, \{x_2, x_4\}$ is the desired K_3 -model. \square

Note that average degree greater than 4 is required in Lemma 7.7 because of the disjoint union of K_5 's. Lemma 7.7 enables the following improvement to Theorem 7.6.

Theorem 7.8. *There is a function h such that for every integer $t \geq 2$ and real $\varepsilon > 0$, every graph G with average degree at least $2^{t-1} + \varepsilon$ contains a K_t -model with at most $h(t, \varepsilon) \cdot \log |G|$ vertices.*

PROOF. As before, we prove the following stronger statement: Every graph G with average degree at least $2^{t-1} + \varepsilon$ contains a K_t -model with at most $h(t, \varepsilon) \cdot \log |G|$ vertices such that each branch set of the model contains at least two vertices.

The proof is by induction on t . The $t = 2$ case is handled in the proof of Theorem 7.6. Lemma 7.7 implies the $t = 3$ case. Now assume $t \geq 4$ and the claim holds for smaller values of t . The proof proceeds as in the proof of Theorem 7.6. We obtain a subgraph G' of G with average degree at least $2^{t-1} + \frac{\varepsilon}{2}$ and diameter at most $f(2^{t-1} + \varepsilon, 2^{t-1} + \frac{\varepsilon}{2}) \cdot \log |G|$. Choose an edge uv of G' and define the depth of edges with respect to uv . We obtain a connected subgraph H with average degree at least $2^{t-2} + \frac{\varepsilon}{4}$, such that every edge of H has the same depth k . If $k = 0$, then $E(H)$ is precisely the set of edges incident to u or v (or both), implying H has average degree less than $4 < 2^{t-2} + \frac{\varepsilon}{4}$. Now assume $k \geq 1$. The remainder of the proof is the same as that of Theorem 7.6. \square

Thomassen [Tho83] first observed that high girth (and minimum degree 3) forces a large complete graph as a minor; see [KO03] for the best known bounds. We now show that high girth (and minimum degree 3) forces a *small* model of a large complete graph.

Theorem 7.9. *Let k be a positive integer. Let G be a graph with girth at least $8k + 3$ and minimum degree $r \geq 3$. Let t be an integer such that $r(r-1)^k \geq 2^{t-1} + 1$. Then G contains a K_t -model with at most $h(k, r) \cdot \log |G|$ vertices, for some function h .*

PROOF. Mader [Mad98] proved that G contains a minor H of minimum degree at least $r(r-1)^k$, such that each branch set has radius at most $2k$; see [Die05, Lemma 7.2.3]. Let $V(H) = \{b_1, \dots, b_{|H|}\}$, and let $B_1, \dots, B_{|H|}$ be the corresponding branch sets in G . Let r_i be a centre of B_i . For each vertex v in B_i , let $P_{i,v}$ be a path between r_i and v in B_i of length at most $2k$.

By Theorem 7.8, H contains a K_t -model with at most $h(t) \cdot \log |H|$ vertices. Let C_1, \dots, C_t be the corresponding branch sets. Say C_i has n_i vertices. Thus $\sum_{i=1}^t n_i \leq h(t) \cdot \log |H|$. We now construct a K_t -model X_1, \dots, X_t in G .

For $i \in \{1, \dots, t\}$, let T_i be a spanning tree of C_i . Each edge $b_j b_\ell$ of T_i corresponds to an edge vw of G , for some v in B_j and w in B_ℓ . Add to X_i the $r_i r_j$ -path $P_{j,v} \cup \{vw\} \cup P_{\ell,w}$. This path has at most $4k + 2$ vertices. Thus X_i is a connected subgraph of G with at most $(4k + 2)(n_i - 1)$ vertices (since T_i has $n_i - 1$ edges).

For distinct $i, i' \in \{1, \dots, t\}$ there is an edge between C_i and $C_{i'}$ in H . This edge corresponds to an edge vw of G , where v is in some branch set B_j in C_i , and w is in some branch

set $B_{j'}$ in $C_{i'}$. Add the path $P_{j,v}$ to X_i , and add the path $P_{j',w}$ to $X_{i'}$. Thus v in X_i is adjacent to w in $X_{i'}$.

Hence X_1, \dots, X_t is a K_t -model in G with at most $\sum_{i=1}^t (4k+2)(n_i-1) \leq (4k+2) \cdot h(t) \cdot \log |H|$ vertices from the first step of the construction, and at most $\binom{t}{2} (4k+2)$ vertices from the second step. Since t is bounded by a function of r and k , there are at most $h'(k, r) \cdot \log |G|$ vertices in total, for some function h' . \square

Corollary 7.10. *Let k be a positive integer. Let G be a graph with girth at least $8k+3$ and minimum degree at least 3. Then G contains a K_k -model with at most $h(k) \cdot \log |G|$ vertices, for some function h .* \square

7.5. Planar Graphs

In this section we prove that sufficiently dense planar graphs have K_3 - and K_4 -models of bounded size. We start with the K_3 case.

Theorem 7.11. *Let $\varepsilon \in (0, 4)$. Every planar graph G with average degree at least $2 + \varepsilon$ has girth at most $1 + \lceil \frac{4}{\varepsilon} \rceil$.*

PROOF. Let H be a connected component of G with average degree at least $2 + \varepsilon$. Thus H is not a tree. Say H has n vertices and m edges. Fix an embedding of H in the plane with r faces. Let ℓ be the minimum length of a facial walk. Thus $\ell \geq 3$ and $2m \geq r\ell = (2 + m - n)\ell$, implying

$$n - 2 \geq m(1 - \frac{2}{\ell}) \geq \frac{1}{2}(2 + \varepsilon)n(1 - \frac{2}{\ell}) > \frac{1}{2}(2 + \varepsilon)(n - 2)(1 - \frac{2}{\ell}).$$

It follows that $\ell < 2 + \frac{4}{\varepsilon}$. Since ℓ is an integer, $\ell \leq 1 + \lceil \frac{4}{\varepsilon} \rceil$. Since H is not a tree, every facial walk contains a cycle. Thus H and G have girth at most $1 + \lceil \frac{4}{\varepsilon} \rceil$. \square

To prove our results for K_4 -models in embedded graphs, the notion of visibility will be useful (and of independent interest). Distinct vertices v and w in an embedded graph are *visible* if v and w appear on a common face; we say v sees w .

Lemma 7.12. *Let v be a vertex of a plane graph G , such that $\deg(v) \geq 3$, v is not a cut-vertex, and v is in no cut-pair. Then v and the vertices seen by v induce a subgraph containing a K_4 -minor.*

PROOF. We may assume that G is connected. Since v is not a cut-vertex, $G - v$ is connected. Let f be the face of $G - v$ that contains v in its interior. Let F be the facial walk around f . Suppose that F is not a simple cycle. Then F has a repeated vertex w . Say $(a, w, b, \dots, c, w, d)$ is a subwalk of F . Then there is a Jordan curve C from v to w , arriving at w between the edges wa and wb , then leaving w from between the edges wc and wd , and back to v . Thus C contains b in its interior and a in its exterior. Hence v, w is a cut-pair. This contradiction proves that F is a simple cycle. Hence v and the vertices seen by v induce a subdivided wheel with $\deg(v)$ spokes. Since $\deg(v) \geq 3$ this subgraph contains a subdivision of K_4 . \square

Recall that $F(G, v)$ is the multiset of faces incident to a vertex v in an embedded graph G , where the multiplicity of a face f in $F(G, v)$ equals the multiplicity of v in the facial walk around f .

Lemma 7.13. *Each vertex v in an embedded graph G sees at most $\sum_{f \in F(G, v)} (|f| - 2)$ other vertices.*

PROOF. The vertex v only sees the vertices in the faces in $F(G, v)$. Each $f \in F(G, v)$ contributes at most $|f| - 1$ vertices distinct from v . Moreover, each neighbour of v is counted at least twice. Thus v sees at most $\sum_{f \in F(G, v)} (|f| - 1) - \deg(v)$ other vertices, which equals $\sum_{f \in F(G, v)} (|f| - 2)$. \square

The 4-regular planar graph C_{2n}^2 has an embedding in the plane, in which each vertex sees $n + 1$ other vertices; see Figure 7.1. On the other hand, we now show that every plane graph with minimum degree 5 has a vertex that sees a bounded number of vertices.

Lemma 7.14. *Every plane graph G with minimum degree 5 has a vertex that sees at most 7 other vertices.*

PROOF. For each vertex v of G , associate a charge of

$$2 - \deg(v) + \sum_{f \in F(G, v)} \frac{2}{|f|} .$$

By Euler's formula, the total charge is $2|G| - 2\|G\| + 2|F(G)| = 4$. Thus some vertex v has positive charge. That is,

$$2 \sum_{f \in F(G, v)} \frac{1}{|f|} > \deg(v) - 2 .$$

Now $\frac{1}{|f|} \leq \frac{1}{3}$. Thus $\frac{2}{3} \deg(v) > \deg(v) - 2$, implying $\deg(v) < 6$ and $\deg(v) = 5$. If some facial walk containing v has length at least 6, then

$$3 = 2 \left(\frac{4}{3} + \frac{1}{6} \right) \geq 2 \sum_{f \in F(G, v)} \frac{1}{|f|} > 3 ,$$

which is a contradiction. Hence each facial walk containing v has length at most 5. If two facial walks containing v have length at least 4, then

$$3 = 2 \left(\frac{3}{3} + \frac{2}{4} \right) \geq 2 \sum_{f \in F(G, v)} \frac{1}{|f|} > 3 ,$$

which is a contradiction. Thus no two facial walks containing v each have length at least 4. Hence all the facial walks containing v are triangles, except for one, which has length at most 5. Thus v sees at most 7 vertices. \square

The bound in Lemma 7.14 is tight since there is a 5-regular planar graph with triangular and pentagonal faces, where each vertex is incident to exactly one pentagonal face (implying that each vertex sees exactly 7 vertices). The corresponding polyhedron is called the *snub dodecahedron*; see Figure 7.2 and [Wik10].

Lemmas 7.12 and 7.14 imply:

Theorem 7.15. *Every 3-connected planar graph with minimum degree 5 contains a K_4 -model with at most 8 vertices.*

Theorem 7.15 is best possible since it is easily seen that every K_4 -model in the snub dodecahedron contains at least 8 vertices. Also note that no result like Theorem 7.15 holds for planar graphs with minimum degree 4 since every K_4 -model in the 4-regular planar graph C_{2n}^2 has at least n vertices.

We now generalise Lemma 7.14 for graphs with average degree greater than 4.

Lemma 7.16. *Let $\varepsilon \in (0, 2)$. Every plane graph G with minimum degree at least 3 and average degree at least $4 + \varepsilon$ has a vertex v that sees at most $1 + \lceil \frac{8}{\varepsilon} \rceil$ other vertices.*

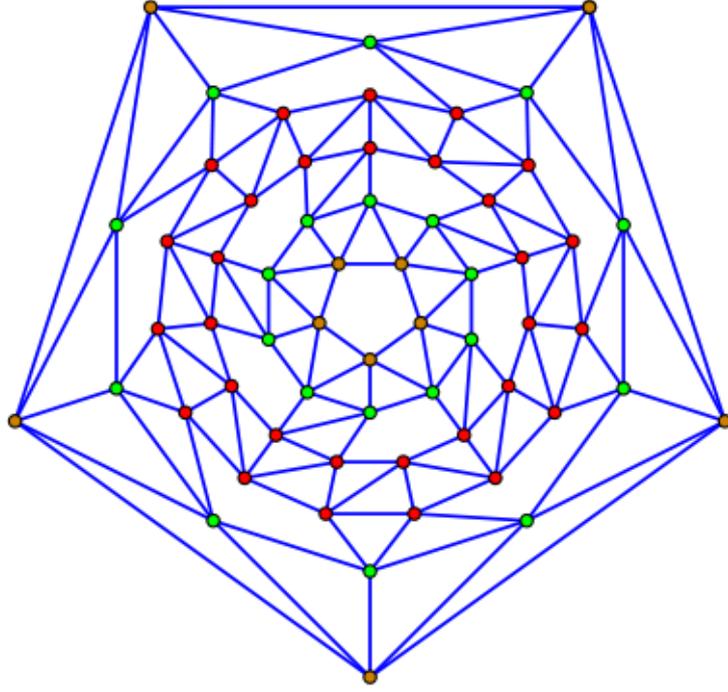


FIGURE 7.2. The snub dodecahedron.

PROOF. For each vertex v of G , associate a charge of

$$(8 + 2\varepsilon) - (8 + 3\varepsilon) \deg(v) + (24 + 6\varepsilon) \sum_{f \in F(G,v)} \frac{1}{|f|} .$$

By Euler's formula, the total charge is

$$\begin{aligned} & (8 + 2\varepsilon)|G| - (16 + 6\varepsilon) \|G\| + (24 + 6\varepsilon) |F(G)| \\ &= (8 + 2\varepsilon)|G| - (16 + 6\varepsilon) \|G\| + (24 + 6\varepsilon) (\|G\| - |G| + 2) \\ &= 4(2\|G\| - (4 + \varepsilon)|G|) + 2(24 + 6\varepsilon) \\ &\geq 2(24 + 6\varepsilon) . \end{aligned}$$

Thus some vertex v has positive charge. That is,

$$(24 + 6\varepsilon) \sum_{f \in F(G,v)} \frac{1}{|f|} > (8 + 3\varepsilon) \deg(v) - (8 + 2\varepsilon) .$$

That is,

$$\sum_{f \in F(G,v)} \frac{1}{|f|} > \left(\frac{1}{3} + \frac{1}{\alpha} \right) \deg(v) - \frac{1}{3} ,$$

where $\alpha := 6 + \frac{24}{\varepsilon}$. We have proved that $\deg(v)$ and the lengths of the facial walks incident to v satisfy Lemma 7.17 below. Thus

$$\sum_{f \in F(G,v)} (|f| - 2) \leq \left\lceil \frac{\alpha}{3} \right\rceil - 1 = 1 + \left\lceil \frac{8}{\varepsilon} \right\rceil .$$

The result follows from Lemma 7.13. □

Lemma 7.17. *Let $\alpha > 0$. Let d, f_1, \dots, f_d be integers, each at least 3, such that*

$$\sum_{i=1}^d \frac{1}{f_i} > \left(\frac{1}{3} + \frac{1}{\alpha} \right) d - \frac{1}{3} .$$

Then

$$\sum_{i=1}^d (f_i - 2) \leq \left\lceil \frac{\alpha}{3} \right\rceil - 1 .$$

PROOF. We may assume that f_1, \dots, f_d firstly maximise $\sum_i (f_i - 2)$, and secondly maximise $\sum_i \frac{1}{f_i}$. We claim that $f_i = 3$ for all $i \in \{1, \dots, d\}$ except perhaps one. Suppose on the contrary that $f_j \geq f_k \geq 4$ for distinct $j, k \in \{1, \dots, d\}$. Let $f'_i := f_i$ for $i \in \{1, \dots, d\} \setminus \{j, k\}$, $f'_j := f_j + 1$, and $f'_k := f_k - 1$. Then

$$\sum_{i=1}^d f'_i = \sum_{i=1}^d f_i \quad \text{but} \quad \sum_{i=1}^d \frac{1}{f'_i} > \sum_{i=1}^d \frac{1}{f_i} ,$$

implying f_1, \dots, f_d do not maximise $\sum_j \frac{1}{f_j}$. Thus the claim holds and we may assume $f_i = 3$ for $i \in \{1, \dots, d-1\}$. Hence

$$\frac{d-1}{3} + \frac{1}{f_d} > \left(\frac{1}{3} + \frac{1}{\alpha} \right) d - \frac{1}{3} .$$

Thus $\frac{1}{f_d} > \frac{d}{\alpha}$, implying $f_d \leq \left\lceil \frac{\alpha}{d} \right\rceil - 1$. Since $\frac{\alpha}{d} > f_d \geq 3$ and since $d \geq 3$,

$$\frac{\alpha}{3} = \frac{\alpha}{d} \left(\frac{d}{3} - 1 \right) + \frac{\alpha}{d} \geq 3 \left(\frac{d}{3} - 1 \right) + \frac{\alpha}{d} = d - 3 + \frac{\alpha}{d} .$$

Hence

$$\left\lceil \frac{\alpha}{3} \right\rceil \geq \left\lceil d - 3 + \frac{\alpha}{d} \right\rceil = d - 3 + \left\lceil \frac{\alpha}{d} \right\rceil .$$

Therefore

$$\sum_{i=1}^d (f_i - 2) \leq (d-1)(3-2) + \left\lceil \frac{\alpha}{d} \right\rceil - 3 = d - 3 + \left\lceil \frac{\alpha}{d} \right\rceil - 1 \leq \left\lceil \frac{\alpha}{3} \right\rceil - 1 .$$

This completes the proof. \square

Lemmas 7.16 and 7.12 imply:

Theorem 7.18. *Let $\varepsilon \in (0, 2)$. Every 3-connected planar graph G with average degree at least $4 + \varepsilon$ contains a K_4 -model with at most $2 + \left\lceil \frac{8}{\varepsilon} \right\rceil$ vertices.*

We now prove that the 3-connectivity assumption in Theorem 7.18 can be dropped, at the expense of a slightly weaker bound on the size of the K_4 -model.

Theorem 7.19. *Let $\varepsilon \in (0, 2)$. Every planar graph G with average degree at least $4 + \varepsilon$ contains a K_4 -model with at most $\left\lceil \frac{8}{\varepsilon} \right\rceil + \left\lceil \frac{2}{\varepsilon} \right\rceil$ vertices. Moreover, this bound is within a constant factor of optimal.*

PROOF. If G has at most $2 + \left\lceil \frac{2}{\varepsilon} \right\rceil$ vertices, then we are done since $m > 2n$ implies G contains a K_4 -model, which necessarily has at most $2 + \left\lceil \frac{2}{\varepsilon} \right\rceil < \left\lceil \frac{8}{\varepsilon} \right\rceil + \left\lceil \frac{2}{\varepsilon} \right\rceil$ vertices.

We now proceed by induction on n with the following hypothesis: Let G be a planar graph with $n \geq 2 + \left\lceil \frac{2}{\varepsilon} \right\rceil$ vertices and m edges, such that

$$(43) \quad 2m > (4 + \varepsilon)(n - 2) .$$

Then G contains a K_4 -model with at most $\left\lceil \frac{8}{\varepsilon} \right\rceil + \left\lceil \frac{2}{\varepsilon} \right\rceil$ vertices.

This will imply the theorem since $2m \geq (4 + \varepsilon)n > (4 + \varepsilon)(n - 2)$.

Suppose that $n \leq \lceil \frac{8}{\varepsilon} \rceil + \lceil \frac{2}{\varepsilon} \rceil$. Since $n \geq 2 + \frac{2}{\varepsilon}$,

$$2m > (4 + \varepsilon)(n - 2) = 4n - 8 + \varepsilon(n - 2) \geq 4n - 6 .$$

Thus $m > 2n - 3$, implying G contains a K_4 -model, which has at most $n \leq \lceil \frac{8}{\varepsilon} \rceil + \lceil \frac{2}{\varepsilon} \rceil$ vertices.

Now assume that $n \geq \lceil \frac{8}{\varepsilon} \rceil + \lceil \frac{2}{\varepsilon} \rceil + 1$.

Suppose that $\deg(v) \leq 2$ for some vertex v . Thus $G - v$ satisfies (43) since

$$2\|G - v\| = 2(m - \deg(v)) > (4 + \varepsilon)(n - 2) - 4 > (4 + \varepsilon)(n - 3) .$$

Now $n - 1 \geq \lceil \frac{8}{\varepsilon} \rceil + \lceil \frac{2}{\varepsilon} \rceil > 2 + \lceil \frac{2}{\varepsilon} \rceil$. Thus, by induction, $G - v$ and hence G contains the desired K_4 -minor. Now assume that $\deg(v) \geq 3$ for every vertex v .

Suppose that G contains a separation $\{G_1, G_2\}$ of order at most 2. Let $S := V(G_1 \cap G_2)$. Say each G_i has n_i vertices and m_i edges. Thus $n_1 + n_2 \leq n + 2$ and $m_1 + m_2 \leq m$. Equation (43) is satisfied for G_1 or G_2 , as otherwise

$$(4 + \varepsilon)(n - 2) < 2m \leq 2m_1 + 2m_2 \leq (4 + \varepsilon)(n_1 + n_2 - 4) \leq (4 + \varepsilon)(n - 2) .$$

Without loss of generality, G_1 satisfies (43). Thus we are done by induction if $n_1 \geq 2 + \lceil \frac{2}{\varepsilon} \rceil$. Now assume that $n_1 \leq 1 + \lceil \frac{2}{\varepsilon} \rceil$. Also assume that $m_1 \leq 2n_1 - 3$, as otherwise G_1 contains a K_4 -model, which has at most $n_1 \leq 1 + \lceil \frac{2}{\varepsilon} \rceil$ vertices.

Suppose that $S = \{v\}$ for some cut-vertex v . Since every vertex in G has degree at least 3, every vertex in G_1 , except v , has degree at least 3 in G_1 . Since $n_1 \geq 2$, G_1 contains a K_4 -model, which has at most $n_1 \leq 1 + \lceil \frac{2}{\varepsilon} \rceil$ vertices. Now assume that G is 2-connected.

Suppose that $S = \{v, w\}$ for some adjacent cut-pair v, w . Thus $n_1 + n_2 = n + 2$ and $m = m_1 + m_2 - 1$ and

$$\begin{aligned} 2m_2 = 2m + 2 - 2m_1 &> (4 + \varepsilon)(n - 2) + 2 - 2(2n_1 - 3) = (4 + \varepsilon)(n_1 + n_2 - 4) - 4n_1 + 8 \\ &= (4 + \varepsilon)(n_2 - 4) + \varepsilon n_1 + 8 \\ &\geq (4 + \varepsilon)(n_2 - 4) + 2(4 + \varepsilon) \\ &= (4 + \varepsilon)(n_2 - 2) . \end{aligned}$$

That is, G_2 satisfies (43). Also,

$$n_2 = n - n_1 + 2 \geq \left(\left\lceil \frac{8}{\varepsilon} \right\rceil + \left\lceil \frac{2}{\varepsilon} \right\rceil \right) + 1 - \left(1 + \left\lceil \frac{2}{\varepsilon} \right\rceil \right) + 2 = 2 + \left\lceil \frac{8}{\varepsilon} \right\rceil > 2 + \left\lceil \frac{2}{\varepsilon} \right\rceil .$$

Hence, by induction G_2 and thus G contains the desired K_4 -model. Now assume that every cut-pair of vertices are not adjacent.

Suppose that $S = \{v, w\}$ for some non-adjacent cut-pair v, w and $m_1 \leq 2n_1 - 4$: Thus $n_1 + n_2 = n + 2$ and $m_1 + m_2 = m$ and

$$\begin{aligned} 2m_2 = 2m - 2m_1 &> (4 + \varepsilon)(n - 2) - 2(2n_1 - 4) = (4 + \varepsilon)(n_1 + n_2 - 4) - 4n_1 + 8 \\ &= (4 + \varepsilon)(n_2 - 4) + \varepsilon n_1 + 8 \\ &\geq (4 + \varepsilon)(n_2 - 4) + 2\varepsilon + 8 \\ &= (4 + \varepsilon)(n_2 - 2) . \end{aligned}$$

That is, G_2 satisfies (43). As proved above, $n_2 > 2 + \lceil \frac{2}{\varepsilon} \rceil$. Hence, by induction G_2 and thus G contains the desired K_4 -model. Now assume that for every cut-pair v, w we have $vw \notin E(G)$, and if $\{G_1, G_2\}$ is the corresponding separation with G_1 satisfying (43), then $m_1 = 2n_1 - 3$ and $n_1 \leq 1 + \lceil \frac{2}{\varepsilon} \rceil$.

Fix an embedding of G . By Lemma 7.16, there is a vertex v in G that sees at most $1 + \lceil \frac{8}{\varepsilon} \rceil$ other vertices. If v is in no cut-pair then by Lemma 7.12 and since G is 2-connected, v plus the vertices seen by v induce a subgraph that contains a K_4 -model, which has at most $2 + \lceil \frac{8}{\varepsilon} \rceil \leq \lceil \frac{8}{\varepsilon} \rceil + \lceil \frac{2}{\varepsilon} \rceil$ vertices. Now assume that v, w is a cut-pair. Thus $vw \notin E(G)$, and if $\{G_1, G_2\}$ is the corresponding separation, then $m_1 = 2n_1 - 3$ and $n_1 \leq 1 + \lceil \frac{2}{\varepsilon} \rceil$. Since v, w is a cut-pair,

there is a vw -path P contained in G_2 , such that P is contained in a single face of G . Every vertex in P is seen by v , and v sees at least 2 vertices in $G_1 - w$. Thus P has at most $\lceil \frac{8}{\varepsilon} \rceil - 2$ internal vertices. Let H be the minor of G obtained by contracting P into the edge vw , and deleting all the other vertices in G_2 . Thus H has n_1 vertices and $2n_1 - 2$ edges. Hence H contains a K_4 -minor. The corresponding K_4 -model in G is contained in $G_1 \cup P$, and thus has at most $(1 + \lceil \frac{2}{\varepsilon} \rceil) + (\lceil \frac{8}{\varepsilon} \rceil - 2) < \lceil \frac{2}{\varepsilon} \rceil + \lceil \frac{8}{\varepsilon} \rceil$ vertices.

We now prove the lower bound. Assume that $\varepsilon \in (0, 1]$ and $k := \frac{1}{\varepsilon} - 1$ is a non-negative integer. Let H be a cubic plane graph in which the length of every facial walk is at least 5 (for example, the dual of a minimum degree 5 plane triangulation). Say H has p vertices. Let G be the plane graph obtained by replacing each vertex of H by a triangle, and replacing each edge of H by $2k$ vertices, as shown in Figure 7.3. Thus G has $3p$ vertices with degree 5 and $3kp$ vertices with degree 4. Thus $|G| = 3p + 3kp = \frac{3p}{\varepsilon}$ and $2\|G\| = 3p \cdot 5 + 3kp \cdot 4 = 4|G| + 3p = (4 + \varepsilon)|G|$. Thus G has average degree $4 + \varepsilon$. Every K_4 -model in G includes a cycle that surrounds a ‘big’ face with more than $5k$ vertices. Thus every K_4 -model has more than $5k = \frac{5}{\varepsilon} - 5$ vertices. Similar constructions are possible for $\varepsilon > 1$ starting with a 4- or 5-regular planar graph. \square

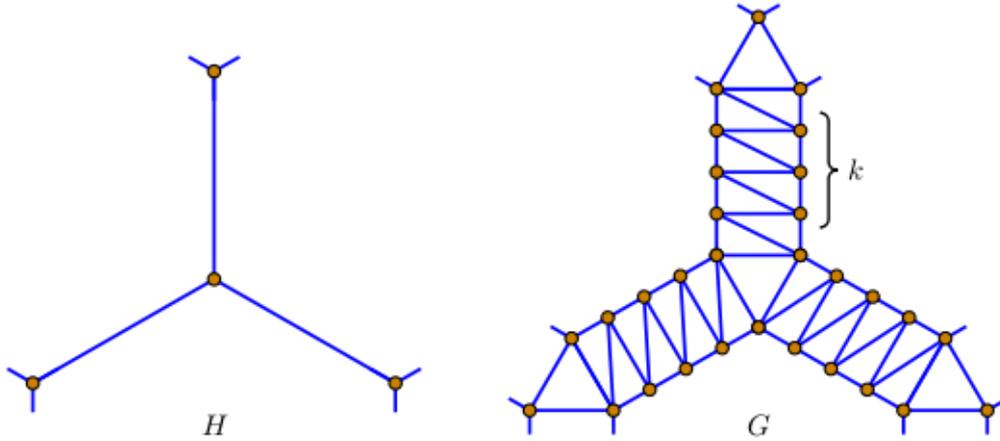


FIGURE 7.3. Construction of G .

7.6. Higher Genus Surfaces

We now extend our results from Section 7.5 for graphs embedded on other surfaces.

Lemma 7.20. *Let $\varepsilon > 0$. Let G be a graph with average degree at least $2 + \varepsilon$. Suppose that G is embedded in a surface with Euler genus at most g . Then some facial walk has length at most $(\frac{4}{\varepsilon} + 2)(g + 1)$. Moreover, this bound is tight up to lower order terms.*

PROOF. Say G has n vertices, m edges, and r faces. Let ℓ be the minimum length of a facial walk. Thus $2m \geq r\ell$. By Euler’s formula, $n - m + r = 2 - g$. Hence

$$\begin{aligned} (2 + \varepsilon)n &\leq 2m \\ (2 + \varepsilon)(2 - g) &= (2 + \varepsilon)(n - m + r) \\ \frac{\varepsilon}{2}(r\ell) &\leq \frac{\varepsilon}{2}(2m) . \end{aligned}$$

Summing gives $\frac{\varepsilon}{2}(r\ell) \leq (2 + \varepsilon)(g + r - 2)$. Since $r \geq 1$,

$$\ell \leq \frac{2}{\varepsilon r} (2 + \varepsilon)(g + r - 2) = \left(\frac{4}{\varepsilon} + 2\right) \left(\frac{g}{r} + \frac{r - 2}{r}\right) < \left(\frac{4}{\varepsilon} + 2\right) (g + 1) .$$

Hence some facial walk has length at most $(\frac{4}{\varepsilon} + 2)(g + 1)$.

Now we prove the lower bound. Assume that $g = 2h \geq 2$ is a positive even integer, and that $0 < \varepsilon \leq 1 - \frac{3}{2g+1}$. Let $k := \lfloor \frac{2}{\varepsilon} - \frac{2}{\varepsilon g} - \frac{1}{g} \rfloor$. Thus $k \geq 2$. Let G be the graph consisting of g cycles of length $k + 1$ with exactly one vertex in common. Thus

$$\begin{aligned} 2\|G\| &= 2g(k + 1) = 2gk + 2 + \varepsilon + \varepsilon g \left(\frac{2}{\varepsilon} - \frac{2}{\varepsilon g} - \frac{1}{g} \right) \geq 2gk + 2 + \varepsilon + \varepsilon gk \\ &= (2 + \varepsilon)(gk + 1) \\ &= (2 + \varepsilon)|G|. \end{aligned}$$

Hence G has average degree at least $2 + \varepsilon$. As illustrated in Figure 7.4(a), G has an embedding in \mathbb{S}_h (which has Euler genus $2h = g$) with exactly one face. Thus every facial walk in G has length $2\|G\| = 2g(k + 1) > 2g(\frac{2}{\varepsilon} - \frac{2}{\varepsilon g} - \frac{1}{g}) \geq \frac{4(g-1)}{\varepsilon} - 2$. \square

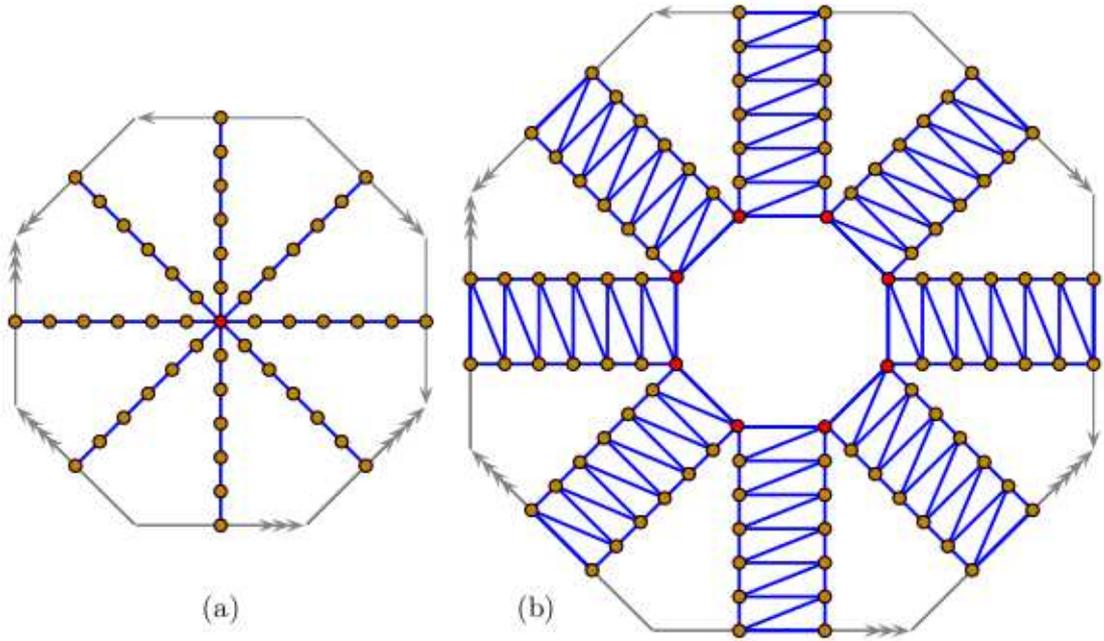


FIGURE 7.4. Graphs embedded in \mathbb{S}_2 : (a) average degree $2 + \varepsilon$ and one face, and (b) average degree $4 + \varepsilon$ and every vertex on one face.

Theorem 7.21. *There is a function h , such that for every real $\varepsilon > 0$, every graph G with average degree at least $2 + \varepsilon$ and Euler genus g has girth at most $h(\varepsilon) \cdot \log(g + 2)$. Moreover, for fixed ε , this bound is within a constant factor of optimal.*

PROOF. Say G has n vertices and m edges. We may assume that every proper subgraph of G has average degree strictly less than $2 + \varepsilon$. This implies that G has minimum degree at least 2. Fix an embedding of G with Euler genus g . Let ℓ be the minimum length of a facial walk. By Euler's formula, there are $m - n + 2 - g$ faces. Thus $2m \geq (m - n + 2 - g)\ell$, implying $\ell(n + g - 2) \geq m(\ell - 2) \geq \frac{1}{2}(2 + \varepsilon)(\ell - 2)n$. Thus $\ell(n + g - 2) \geq \frac{1}{2}(2 + \varepsilon)(\ell - 2)n$, implying $\ell(g - 2) \geq (\frac{\varepsilon}{2}(\ell - 2) - 2)n$. First suppose that $\ell < 6 + \frac{12}{\varepsilon}$. Since G has no degree-1 vertices, every facial walk contains a cycle. Thus G has girth at most $6 + \frac{12}{\varepsilon}$, which is at most $h(\varepsilon) \cdot \log(g + 2)$ for some function h . Now assume that $\ell \geq 6 + \frac{12}{\varepsilon}$, which implies that

$\ell(g-2) \geq (\frac{\varepsilon}{2}(\ell-2) - 2)n \geq \frac{\varepsilon}{3}\ell n$. Thus $n \leq \frac{3}{\varepsilon}(g-2)$. By Lemma 7.4, the girth of G is at most $g(\varepsilon) \cdot \log n \leq g(\varepsilon) \cdot \log(\frac{3}{\varepsilon}(g-2))$, which is at most $h(\varepsilon) \cdot \log(g+2)$ for some function h .

Now we prove the lower bound. Let d be the integer such that $d-3 < \varepsilon \leq d-2$. Thus $d \geq 3$. For all $n > 3d$ such that nd is even, Chandran [Cha03] constructed a graph G with n vertices, average degree $d \geq 2 + \varepsilon$, and girth at least $(\log_d n) - 1$. Now G has Euler genus $g \leq \frac{dn}{2} - n + 1 \leq dn - 2$. Thus G has girth at least $(\log_d \frac{g+2}{d}) - 1$. Since $d < 3 + \varepsilon$, the girth of G is at least $h(\varepsilon) \cdot \log(g+2)$ for some function h . \square

We now extend Lemma 7.16 for sufficiently large embedded graphs.

Lemma 7.22. *Let $\varepsilon \in (0, 2)$. Let G be a graph with minimum degree 3 and average degree at least $4 + \varepsilon$. Assume that G is embedded in a surface with Euler genus g , such that $|G| \geq (\frac{24}{\varepsilon} + 6)g$. Then G has a vertex v that sees at most $2 + \lceil \frac{12}{\varepsilon} \rceil$ other vertices.*

PROOF. For each vertex v of G , associate a charge of

$$(8 + 2\varepsilon) - (8 + 3\varepsilon) \deg(v) + (24 + 6\varepsilon) \frac{g}{|G|} + (24 + 6\varepsilon) \sum_{f \in F(G,v)} \frac{1}{|f|} .$$

Thus the total charge is

$$\begin{aligned} & (8 + 2\varepsilon)|G| - (16 + 6\varepsilon) \|G\| + (24 + 6\varepsilon)g + (24 + 6\varepsilon) |F(G)| \\ &= (8 + 2\varepsilon)|G| - (16 + 6\varepsilon) \|G\| + (24 + 6\varepsilon)g + (24 + 6\varepsilon) (\|G\| - |G| - g + 2) \\ &= 4(2\|G\| - (4 + \varepsilon)|G|) + 2(24 + 6\varepsilon) \\ &\geq 2(24 + 6\varepsilon) . \end{aligned}$$

Thus some vertex v has positive charge. That is,

$$(8 + 2\varepsilon) - (8 + 3\varepsilon) \deg(v) + (24 + 6\varepsilon) \frac{g}{|G|} + (24 + 6\varepsilon) \sum_{f \in F(G,v)} \frac{1}{|f|} > 0 .$$

Since $\frac{(24+6\varepsilon)g}{|G|} \leq \varepsilon$,

$$(24 + 6\varepsilon) \sum_{f \in F(G,v)} \frac{1}{|f|} > (8 + 3\varepsilon)(\deg(v) - 1) .$$

That is,

$$\sum_{f \in F(G,v)} \frac{1}{|f|} > \left(\frac{1}{3} + \frac{1}{\alpha} \right) (\deg(v) - 1) ,$$

where $\alpha := 6 + \frac{24}{\varepsilon}$. We have proved that $\deg(v)$ and the lengths of the facial walks incident to v satisfy Lemma 7.23 below. Thus

$$\sum_{f \in F(G,v)} (|f| - 2) \leq \left\lceil \frac{\alpha}{2} \right\rceil - 1 = 2 + \left\lceil \frac{12}{\varepsilon} \right\rceil .$$

The result follows from Lemma 7.13. \square

Lemma 7.23. *Let $\alpha > 0$. Let d, f_1, \dots, f_d be integers, each at least 3, such that*

$$\sum_{i=1}^d \frac{1}{f_i} > \left(\frac{1}{3} + \frac{1}{\alpha} \right) (d-1) .$$

Then

$$\sum_{i=1}^d (f_i - 2) \leq \left\lceil \frac{\alpha}{2} \right\rceil - 1 .$$

PROOF. As in the proof of Lemma 7.17, we may assume that $f_j = 3$ for all $j \in \{3, \dots, d-1\}$. Hence

$$\frac{d-1}{3} + \frac{1}{f_d} > \left(\frac{1}{3} + \frac{1}{\alpha}\right)(d-1).$$

Thus $\frac{1}{f_d} > \frac{d-1}{\alpha}$, implying $f_d \leq \lceil \frac{\alpha}{d-1} \rceil - 1$. Since $\frac{\alpha}{d-1} > f_d \geq 3$ and since $d \geq 3$,

$$\frac{\alpha}{2} \geq \frac{\alpha d}{3(d-1)} = \left(\frac{\alpha}{d-1}\right) \left(\frac{d}{3} - 1\right) + \frac{\alpha}{d-1} \geq 3 \left(\frac{d}{3} - 1\right) + \frac{\alpha}{d-1} = d - 3 + \frac{\alpha}{d-1}.$$

Hence

$$\left\lceil \frac{\alpha}{2} \right\rceil \geq \left\lceil d - 3 + \frac{\alpha}{d-1} \right\rceil = d - 3 + \left\lceil \frac{\alpha}{d-1} \right\rceil.$$

Therefore

$$\sum_{i=1}^d (f_i - 2) \leq (d-1)(3-2) + \left\lceil \frac{\alpha}{d-1} \right\rceil - 3 = d - 3 + \left\lceil \frac{\alpha}{d-1} \right\rceil - 1 \leq \left\lceil \frac{\alpha}{2} \right\rceil - 1.$$

This completes the proof. \square

We now prove that the assumption that $n \in \Omega(\frac{g}{\varepsilon})$ in Lemma 7.22 is needed. Assume we are given $\varepsilon \in (0, 1]$ such that $k := \frac{1}{\varepsilon} - 1$ is an integer. Hence $k \geq 0$. Consider the graph G shown in Figure 7.4(b) with $2g$ vertices of degree 5 and $2gk$ vertices of degree 4. Thus $|G| = 2g(k+1)$ and $2\|G\| = 10g + 8gk = 2g(5 + 4k) = \frac{|G|}{k+1}(4k+5) = (4 + \frac{1}{k+1})|G| = (4 + \varepsilon)|G|$. Thus G has average degree $4 + \varepsilon$. Observe that every vertex lies on a single face. Thus each vertex sees $|G| - 1 = \frac{2g}{\varepsilon} - 1$ other vertices.

A k -noose in an embedded graph G is a noncontractible simple closed curve in the surface that intersects G in exactly k vertices. The *facewidth* of G is the minimum integer k such that G contains a k -noose.

Theorem 7.24. *Let $\varepsilon > 0$. Let G be a 3-connected graph with average degree at least $4 + \varepsilon$, such that G has an embedding in a surface with Euler genus g and with facewidth at least 3. Then G contains a K_4 -model with at most $f(\varepsilon) \cdot \log(g+2)$ vertices, for some function f . Moreover, for fixed ε , this bound is within a constant factor of optimal.*

PROOF. If $|G| \leq (\frac{24}{\varepsilon} + 6)g$ then the result follows from Theorem 7.5. Otherwise, by Lemma 7.22 some vertex v sees at most $2 + \lceil \frac{12}{\varepsilon} \rceil$ other vertices. The graph $G - v$ is 2-connected and has facewidth at least 2. Thus every face of $G - v$ is a simple cycle [MT01, Proposition 5.5.11]. In particular, the face of $G - v$ that contains v in its interior is bounded by a simple cycle C . The vertices in C are precisely the vertices that v sees in G . Thus $G[C \cup \{v\}]$ is a subdivided wheel with $\deg(v) \geq 3$ spokes. Hence G contains a K_4 -model with at most $2 + \lceil \frac{12}{\varepsilon} \rceil$ vertices, which is at most $f(\varepsilon) \cdot \log(g+2)$ for an appropriate function f .

Now we prove the lower bound. Let d be the integer such that $d - 5 < \varepsilon \leq d - 4$. Thus $d \geq 5$. For every integer $n > 3d$ such that nd is even, Chandran [Cha03] constructed a graph G with n vertices, average degree $d \geq 4 + \varepsilon$, and girth greater than $(\log_d n) - 1$. Thus G has Euler genus $g \leq \frac{dn}{2} \leq dn - 2$. Since every K_4 -model contains a cycle, every K_4 -model in G has at least $(\log_d n) - 1$ vertices. Since $n \geq \frac{g+2}{d}$ and $d < 5 + \varepsilon$, every K_4 -model in G has at least $f(\varepsilon) \cdot \log(g+2)$ vertices, for some function f . \square

For a class of graphs, an edge is ‘light’ if both its endpoints have bounded degree. For example, Wernicke [Wer04] proved that every planar graph with minimum degree 5 has an edge vw such that $\deg(v) + \deg(w) \leq 11$; see [Bor89, Kot55, JM96, JV05] for extensions. For a class of embedded graphs, we say an edge is ‘blind’ if both its endpoints see a bounded number of vertices. In a triangulation, a vertex only sees its neighbours, in which case the notions of ‘light’ and ‘blind’ are equivalent. But for non-triangulations, a ‘blind edge’ theorem

is qualitatively stronger than a ‘light edge’ theorem. Hence the following result is a qualitative generalisation of the above theorem of Wernicke [Wer04] (and of Lemma 7.14), and is thus of independent interest. No such result is possible for minimum degree 4 since every edge in C_{2n}^2 sees at least n vertices.

Proposition 7.25. *Let G be a graph with minimum degree 5 embedded in a surface with Euler genus g , such that $|G| \geq 240g$. Then G has an edge vw such that v and w each see at most 12 vertices. Moreover, for plane graphs (that is, $g = 0$), v and w each see at most 11 vertices.*

PROOF. Consider each vertex x . Let ℓ_x be the maximum length of a facial walk containing x . Let t_x be the number of triangular faces incident to x , unless every face incident to x is triangular, in which case let $t_x := \deg(x) - 1$. Say x is *good* if x sees at most 12 vertices, otherwise x is *bad*. Let

$$c_x := 240 - 120 \deg(x) + 240 \frac{g}{|G|} + 240 \sum_{f \in F(G,x)} \frac{1}{|f|}$$

be the charge at x . By Euler’s formula, the total charge is

$$240(|G| - \|G\| + g + |F(G)|) = 480 .$$

Observe that (since $\ell_x \geq 3$ and $t_x \leq \deg(x) - 1$ and $\deg(x) \geq 5$)

$$\begin{aligned} c_x &\leq 240 - 120 \deg(x) + 240 \frac{g}{|G|} + 240 \left(\frac{1}{\ell_x} + \frac{t_x}{3} + \frac{\deg(x) - t_x - 1}{4} \right) \\ (44) \quad &\leq 181 - 60 \deg(x) + \frac{240}{\ell_x} + 20t_x \\ (45) \quad &\leq 241 - 40 \deg(x) \leq 41 . \end{aligned}$$

For each good vertex x , equally distribute the charge on x to its neighbours. (Bad vertices keep their charge.) Let c'_x be the new charge on each vertex x . Since the total charge is positive, $c'_v > 0$ for some vertex v . If v is good, then all the charge at v was received from its neighbours during the charge distribution phase, implying some neighbour w of v is good, and we are done. Now assume that v is bad. Let D_v be the set of good neighbours of v . By (44) and (45), and since $\deg(w) \geq 5$,

$$(46) \quad 0 < c'_v = c_v + \sum_{w \in D_v} \frac{c_w}{\deg(w)} \leq 181 - 60 \deg(v) + \frac{240}{\ell_v} + 20t_v + \frac{41}{5} |D_v| .$$

We may assume that no two good neighbours of v are on a common triangular face.

Claim 7.26. $|D_v| \leq \deg(v) - \frac{t_v}{2}$. Moreover, if $|D_v| = \deg(v) - \frac{t_v}{2}$ then some face incident to v is non-triangular, and for every bad neighbour w of v , the edge vw is incident to two triangular faces.

PROOF. First assume that every face incident to v is triangular. Thus no two consecutive neighbours of v are good. Hence $|D_v| \leq \frac{\deg(v)}{2} < \frac{\deg(v)+1}{2} = \deg(v) - \frac{t_v}{2}$, as claimed. This also proves that if $|D_v| = \deg(v) - \frac{t_v}{2}$ then some face incident to v is non-triangular.

We prove the case in which some face incident to v is non-triangular by a simple charging scheme. If w is a good neighbour of v , then charge vw by 1. Charge each triangular face incident to v by $\frac{1}{2}$. Thus the total charge is $|D_v| + \frac{t_v}{2}$. If uvw is a triangular face incident to v , then at least one of u and w , say w , is bad; send the charge of $\frac{1}{2}$ at uvw to vw . Each good edge incident to v gets a charge of 1, and each bad edge incident to v gets a charge of at most $\frac{1}{2}$ from each of its two incident faces. Thus each edge incident to v gets a charge of at most 1. Thus the total charge, $|D_v| + \frac{t_v}{2}$, is at most $\deg(v)$, as claimed.

Finally, assume that $|D_v| = \deg(v) - \frac{t_v}{2}$. Then for every bad neighbour w of v , the edge vw gets a charge of exactly 1, implying vw is incident to two triangular faces. \square

Claim 7.26 and (46) imply

$$\begin{aligned} 0 &< 181 - 60 \deg(v) + \frac{240}{\ell_v} + 20t_v + \frac{41}{5} \deg(v) - \frac{41t_v}{10} \\ &= 181 - \frac{259}{5} \deg(v) + \frac{240}{\ell_v} + \frac{159}{10} t_v . \end{aligned}$$

Since $t_v \leq \deg(v) - 1$ and $\deg(v) \geq 5$,

$$0 < \frac{1651}{10} - \frac{359}{10} \deg(v) + \frac{240}{\ell_v} \leq -\frac{144}{10} + \frac{240}{\ell_v} .$$

implying $\ell_v \in \{3, 4, \dots, 16\}$. Since $\ell_v \geq 3$,

$$0 < \frac{2451}{10} - \frac{359}{10} \deg(v) ,$$

implying $\deg(v) \in \{5, 6\}$ and $t_v \in \{0, 1, \dots, \deg(v) - 1\}$.

We have proved that finitely many values satisfy (46). We now strengthen this inequality in the case that $|D_v| = \deg(v) - \frac{t_v}{2}$.

Let f be a face of length ℓ_v incident to v . Let x and y be two distinct neighbours of v on f . Suppose on the contrary that x is bad. By Claim 7.26, vx is incident to two triangular faces, one of which is vxy . Thus $\ell_v = 3$, and every face incident to v is a triangle, which contradicts the Claim. Hence x is good. Similarly y is good.

Thus $\ell_x \geq \ell_v$. By (44),

$$c_x \leq 181 - 60 \deg(x) + \frac{240}{\ell_v} + 20t_x \leq 161 - 40 \deg(x) + \frac{240}{\ell_v} \leq \frac{240}{\ell_v} - 39 .$$

Similarly, $c_y \leq \frac{240}{\ell_v} - 39$. Hence (assuming $|D_v| = \deg(v) - \frac{t_v}{2}$),

$$\begin{aligned} (47) \quad 0 < c'_v &\leq 181 - 60 \deg(v) + \frac{240}{\ell_v} + 20t_v + \frac{c_x}{\deg(x)} + \frac{c_y}{\deg(y)} + \sum_{w \in D_v \setminus \{x, y\}} \frac{c_w}{\deg(w)} \\ &\leq 181 - 60 \deg(v) + \frac{240}{\ell_v} + 20t_v + \frac{\frac{240}{\ell_v} - 39}{\deg(x)} + \frac{\frac{240}{\ell_v} - 39}{\deg(y)} + \sum_{w \in D_v \setminus \{x, y\}} \frac{41}{\deg(w)} \\ &\leq 181 - 60 \deg(v) + \frac{240}{\ell_v} + 20t_v + 2 \left(\frac{48}{\ell_v} - \frac{39}{5} \right) + \frac{41}{5} (|D_v| - 2) . \end{aligned}$$

Checking all values of $\deg(v)$, t_v and ℓ_v that satisfy (46) and (47) proves that

$$t_v + (\deg(v) - t_v)(\ell_v - 2) \leq 12$$

(which is tight for $\deg(v) = 5$ and $t_v = 4$ and $\ell_v = 10$ and $|D_v| = 2$). Thus

$$\sum_{f \in F(G, v)} (|f| - 2) \leq t_v(3 - 2) + (\deg(v) - t_v)(\ell_v - 2) \leq 12 .$$

By Lemma 7.13, v sees at most 12 vertices. Therefore v is good, which is a contradiction.

In the case of planar graphs, we define a vertex to be *good* if it sees at most 11 other vertices. Since $g = 0$, (44) and (45) can be improved to

$$(48) \quad c_x \leq 180 - 60 \deg(x) + \frac{240}{\ell_x} + 20t_x \leq 240 - 40 \deg(x) \leq 40 .$$

Subsequently, (46) is improved to

$$(49) \quad 0 < c'_v = 180 - 60 \deg(v) + \frac{240}{\ell_v} + 20t_v + 8|D_v| ,$$

and (47) is improved to

$$(50) \quad 0 < c'_v \leq 180 - 60 \deg(v) + \frac{240}{\ell_v} + 20t_v + 2 \left(\frac{48}{\ell_v} - 8 \right) + 8(|D_v| - 2) .$$

Checking all values of $\deg(v)$, t_v and ℓ_v that satisfy (49) and (50) proves that $t_v + (\deg(v) - t_v)(\ell_v - 2) \leq 11$. As in the main proof, it follows that v is good. \square

We now prove that the assumption that $|G| \in \Omega(g)$ in Proposition 7.25 is necessary. Let G be the graph obtained from C_{2n}^2 by adding a perfect matching, as shown embedded in \mathbb{S}_n in Figure 7.5 (where there is one handle for each pair of crossing edges). This graph is 5-regular, but each vertex is on a facial walk of length n . Thus no vertex sees a bounded number of vertices.

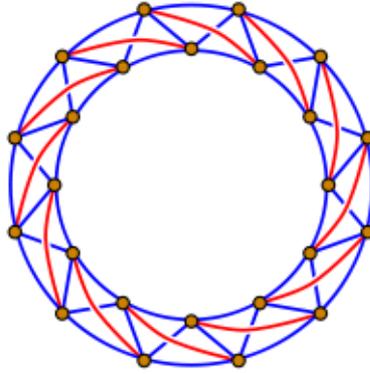


FIGURE 7.5. C_{24}^2 plus a perfect matching, embedded on \mathbb{S}_{12} .

7.7. Open Problems

The first open problem that arises from this work is to determine the best possible function f in Theorem 7.1. In particular, does average degree at least some polynomial in t force a small K_t -model? Even stronger, is there a function h , such that every graph G with average degree at least $f(t) + \varepsilon$ contains a K_t -model with $h(t, \varepsilon) \cdot \log |G|$ vertices, where $f(t)$ is the minimum number such that every graph with average degree at least $f(t)$ contains a K_t -minor? We have answered this question in the affirmative for $t \leq 4$. The case $t = 5$ is open. It follows from Wagner's characterisation of graphs with no K_5 -minor that average degree at least 6 forces a K_5 -minor [Wag37]. Theorem 7.8 proves that average degree at least $16 + \varepsilon$ forces a K_5 -model with at most $h(\varepsilon) \cdot \log n$ vertices. We conjecture the following improvement:

Conjecture 7.27. *There is a function h such that for all $\varepsilon > 0$, every graph G with average degree at least $6 + \varepsilon$ contains a K_5 -model with at most $h(\varepsilon) \cdot \log |G|$ vertices.*

This degree bound would be best possible: Let G_n be the 6-regular $n \times 3$ triangulated toroidal grid, as illustrated in Figure 7.6. Every K_5 -model in G_n intersects every column (otherwise K_5 is planar). Thus every K_5 -model in G_n has at least n vertices.

Note that while in this paper we have only studied small K_t -models, the same questions apply for small H -models, for arbitrary graphs H . This question was studied for $H = K_4 - e$ in [FJP10]. See [Tho06, MT05, Mye03, KP08, KO05] for results about forcing H -minors.

Acknowledgments. Thanks to Michele Conforti for suggesting to study the relationship between average degree and small models. Thanks to Paul Seymour for suggesting the example following Lemma 7.7. Thanks to Alexandr Kostochka for pointing out reference [KP88].

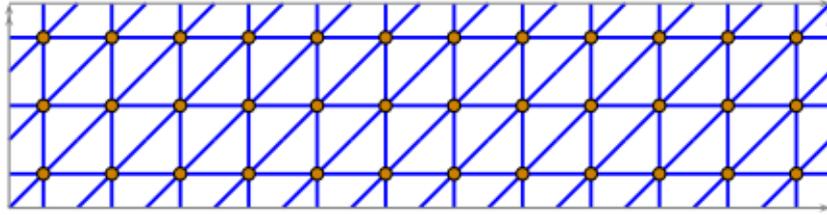


FIGURE 7.6. 6-regular 12×3 triangulated toroidal grid

Good edge labelings and graphs with girth at least five

Jointly with

Michel Bode (stud. math. at OvGU, Magdeburg)
and Babak Farzad (Brock University, St. Catharines, ON)

Abstract. A good edge-labeling of a graph [Araújo, Cohen, Giroire, Havet, Discrete Appl. Math., forthcoming] is an assignment of numbers to the edges such that for no pair of vertices, there exist two non-decreasing paths. In this paper, we study edge-labeling on graphs with girth at least 5. In particular we verify, under this additional hypothesis, a conjecture by Araújo et al. This conjecture states that if the average degree of G is less than 3 and G is minimal without an edge-labeling, then $G \in \{C_3, K_{2,3}\}$. (For the case when the girth is 4, we give a counterexample.)

8.1. Introduction

All graphs are finite and simple. We refer to Diestel [Die06] for most of our graph theory terminology.

A *good edge-labeling* [JCBP09] of a graph G is a labeling of its edges $\phi: E(G) \rightarrow \mathbb{R}$ such that, for any ordered pair of vertices u and v , there is at most one nondecreasing path from u to v . We will mostly use the following characterization of a good edge-labeling, which involves cycles instead of pairs of paths:

An edge-labeling is good, if, and only if, every cycle has at least two local minima.

Here, by a local minimum we mean an edge e whose label is strictly less than the labels of the two edges incident to e on the cycle (this differs from the definition in the next section because at this point, unlike later in the paper, for convenience, we assume that all labels are distinct).

Good edge-labelings have first been studied by Bermond, Cosnard, and Pérennes [JCBP09] in the context of so-called Wavelength Division Multiplexing problems [BCCP06]. There, given a network, the so-called Routing and Wavelength Assignment Problem asks for finding routes and associated wavelengths, such that a set of traffic requests is satisfied, while minimizing the number of used wavelengths.

Araujo, Cohen, Giroire, and Havet [ACGH09, ACGH12] have studied good edge-labelings in more depth. They call a graph with no good edge-labeling *bad*, and say that a *critical* graph is a minimal bad graph, that is, every proper subgraph has a good edge-labeling. It is easy to see that C_3 and $K_{2,3}$ are critical. Araujo et al.'s [ACGH12] paper comprises an infinite family of critical graphs; results that graphs in some classes always have a good edge-labeling (planar graphs with girth at least 6, $(C_3, K_{2,3})$ -free outerplanar graphs, $(C_3, K_{2,3})$ -free sub-cubic graphs); the algorithmic complexity of recognizing bad graphs; and a connection to matching-cuts. (A *matching-cut*, aka “simple cut” [Gra70], is a set of independent edges which is an edge-cut.)

In fact, all their arguments for proving non-criticality rely on the existence of matching-cuts. One of the central contributions of our paper is that we move beyond using matching-cuts.

Araujo et al. also pose a number of problems and conjectures. In particular, they have the following conjecture, which is one of the two motivations behind our paper.

Conjecture 8.1 (Araujo et al. [ACGH12]). *There is no critical graph with average degree less than 3, with the exception of C_3 and $K_{2,3}$.*

Araujo et al. [ACGH12] prove a weaker version of this conjecture. They establish the existence of a matching-cut, relying in part on a theorem by Farley & Proskurowski [FP84, BFP11] stating that a graph with sufficiently few edges always has a matching cut. They also use a characterization of extremal graphs with no matching-cut by Bonsma [Bon05, BFP11]. From the proofs in Araujo et al. [ACGH12], it appears that the depths of the arguments increases rapidly as the upper bound 3 is approached.

In this paper, we show that there is no critical graph with average degree less than three and girth at least five. Put differently, we prove Conjecture 8.1 in the case when the graph has girth at least five.

Theorem 8.2. *There is no critical graph with average degree less than three and girth at least five.*

Moreover, we falsify Conjecture 8.1 for the case of girth 4: Fig. 8.1 shows a graph with girth 4 and average degree $\frac{26}{9} < 3$ (9 vertices, 13 edges), which does not contain either C_3 or $K_{2,3}$ as a subgraph. We leave to the reader as an exercise to argue that the graph has no good edge labeling. It can easily be verified that every proper subgraph has a good edge labeling, so the shown graph is critical. In other words, Fig. 8.1 shows a counterexample to Conjecture 8.1 for the case of girth 4.

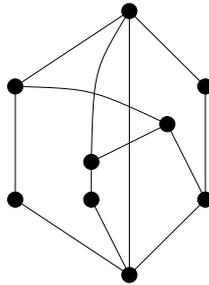


FIGURE 8.1. Critical graph with girth 4 and average degree < 3

Another motivation behind our paper is to demonstrate how large girth makes labeling arguments easier.¹ In Theorem 8.42, roughly speaking, we prove that a critical graph with girth at least five cannot contain a “windmill”. A windmill essentially consists of a number of shortest paths meeting in an “axis”, with the paths originating from vertices of degree two and having in their interior only vertices of degree three. Theorem 8.2 is a corollary of Theorem 8.42: using an approach inspired by the discharging method from topological graph theory, we argue that a hypothetical critical graph with girth at least five and average degree less than three always contains a windmill.

For our proof of Theorem 8.42, we define a class of graphs which we call “decent”, which have the property that they cannot be contained in a critical graph. More importantly, we give a *gluing* operation which preserves “decency”. Starting from a small family of basic “decent” graphs, by gluing inductively, this approach allows us to show that certain more complicated configurations cannot be contained in critical graphs, which leads to the proof of Theorem 8.42.

¹Indeed, until very recently, no bad graph with girth larger than four was known. In particular, the bad graphs in Araujo et al.’s construction contain many 4-cycles. This fact had led us to conjecture, that there exists a finite number g such that every graph with girth at least g has a good edge-labeling; as mentioned above, Araujo et al. [ACGH12] have shown that with the additional restriction that the graphs be planar the conjecture holds true for $g := 6$. The conjecture was refuted in [Meh12].

This paper is organized as follows. In the next section, we will discuss some notation as well as basic facts on good edge-labelings. In Section 8.3, we define windmills, and commence upon the proof of their non-existence. Section 8.4 contains the definition of “decent” graphs and the gluing mechanism. Theorem 8.42 is stated and proved in Section 8.5, and Theorem 8.2 is derived in Section 8.6.

8.2. Basic facts about good edge-labelings

We will heavily rely on the above-mentioned characterization of a good edge-labeling using cycles instead of paths. For this, we use the following definitions. Let H be a path or a cycle, and $\phi: E(H) \rightarrow \mathbb{R}$ an edge-labeling of H . Let Q be a proper sub-path of H (i.e., a path contained in H which is not equal to H) with at least one edge. For a real number μ , we say that Q is a *local minimum with value μ in H* , if $\phi(e) = \mu$ for all $e \in E(Q)$, and for every edge $e' \in E(H) \setminus E(Q)$ sharing a vertex with Q we have $\mu < \phi(e')$.

Distinct minima must necessarily be vertex disjoint. Good edge-labelings can be characterized in terms of local minima of cycles. We leave the verification of the following easy lemma to the reader (or see [Bod11]).

Lemma 8.3. *An edge-labeling ϕ of a graph G is good, if, and only if, every cycle C in H has two local minima.* \square

Obviously, the property of an edge-labeling being good depends only on the order relation between the labels of the edges. In particular, scaling (multiplying each label by a strictly positive constant), and translation (adding a constant to each label) do not change whether a labeling is good or not.

We say that a k -vertex is a vertex of degree k ; a k^- -vertex is a vertex of degree at most k ; and a k^+ -vertex is a vertex of degree at least k .

Araujo et al. [ACGH12] proved the following property of critical graphs.

Lemma 8.4 ([ACGH12]). *A critical graph does not contain a matching-cut.*

In particular, the minimum degree of a critical graph is at least two, and, unless it is a triangle C_3 , it contains no two adjacent 2-vertices.

For the rest of the section, let G be a critical graph other than C_3 and $K_{2,3}$. We prove some basic properties of G .

Lemma 8.5. *Let C be a cycle in G whose every vertex has degree at most three. Then there are two vertices of C with a common neighbour in $G - C$.*

PROOF. We proceed by contradiction: let C' be a shortest cycle whose every vertex has degree at most three. If $G - C' \neq \emptyset$, then it can be easily seen that the set of edges with exactly one endpoint in C' forms a matching-cut, contradicting Lemma 8.4. If $G - C' = \emptyset$, then G is a cycle. Since $G \neq C_3$, there is a good edge-labeling for this cycle, contradicting the criticality of G . \square

Lemma 8.6. *Let P be a path of length at least one in G whose end vertices have degree two and internal vertices have degree at most three. Then two vertices of P have a common neighbour in $G - P$.*

PROOF. We proceed by contradiction: let P' be a shortest path between two vertices of degree two with inner vertices of degree three. If $G - P' \neq \emptyset$, then the set of edges with exactly one endpoint in P' forms a matching cut; contradicting Lemma 8.4. If $G - P' = \emptyset$, then P' cannot be a shortest such path. (We note that the proof goes through if the length of P is 1.) \square

8.3. Windmills

To motivate the definition of windmills, let us take a look at how they will be used in the proof of Theorem 8.2. The proof uses a discharging type argument. We assign “charges” to the vertices: vertex v receives charge $6 - 2d(v)$. Note that only 2-vertices have positive charge. Since the average degree of G is less than 3, the total charge of the graph is positive. Now, we “discharge” 2-vertices. Applying Lemma 8.6, charges are sent from 2-vertices to 4^+ -vertices via shortest paths consisting of only 3-vertices. Later on, we will show that these paths are internally disjoint. Since no charge is lost during the discharging phase, if after discharging all vertices have non-positive charge, then we have a contradiction. However, there may be some vertices with positive charge. These vertices are the centers of the structures which we refer to as “windmills.”

In the remainder of this section, let G be a critical graph of girth at least five.

For a tree H and vertices x, y of H , we denote by xHy the unique path between x and y in H . An internally shortest 3-path is a path $P = x_0 \dots x_\ell$ with $\ell \geq 1$ and $d(x_j) = 3$ for $j \in \{1, \dots, \ell - 1\}$, such that, for $e := x_0x_1$, the path x_1Px_ℓ is a shortest path in $G - e$. In particular, the path x_1Px_ℓ is induced in G . We say that P starts in x_0 and ends in x_ℓ .

Remark 8.7. By Lemma 8.6, the graph G has no internally shortest 3-path that starts and ends in 2-vertices.

So for an edge x_0x_1 , a *sail* with *tip* x_0x_1 is defined to be an internally shortest 3-path $P = x_0x_1 \dots x_\ell$ that starts in a 2-vertex and ends in a 4^+ -vertex which has minimum length ℓ among all such internally shortest 3-paths.

Remark 8.8. Among the vertices v of degree $\deg(v) \neq 3$, the ending vertex x_ℓ of a sail is among those which have minimum distance from x_1 in $G - e$. Note that, in $G - e$, the vertex x_0 has larger distance from x_1 than x_ℓ : otherwise we would have a contradiction (either from having a 4^+ vertex closer to x_1 or by Lemma 8.5).

Definition 8.9. Let $k \geq 3$ be an integer, and y a vertex of degree $\max(4, k)$ or $k + 1$. A k -windmill with axis y in G is an induced subgraph H of G , spanned by the union of k sails beginning in k distinct tips, and each of them ending in y . A windmill H is called *complete* in G , if it is not a proper subgraph of another windmill.

Note that it is possible that two sails of the same windmill start in the same 2-vertex (but have different tips).

Lemma 8.10. Let $P = x_0x_1 \dots x_\ell$ and $P' = x'_0x'_1 \dots x'_{\ell'}$ be any two sails in G . Then

- (a) if a vertex is adjacent to two vertices of P , then one of the two is the starting vertex x_0 ;
- (b) if P and P' have distinct tips but identical ending, i.e., $x_\ell = x'_{\ell'}$, then no vertex is adjacent to two (or more) of the vertices of the path $x_1 \dots x_{\ell-1}x_\ell x'_{\ell'-1} \dots x'_1$.
- (c) P and P' either share the same tip or they are internally disjoint;

PROOF. The facts that G has girth at least 5 and that P is a shortest path from its tip x_0x_1 to a 4^+ -vertex easily imply (a). To prove (b), assume otherwise, that is, there are vertices adjacent to two vertices of the path $Q_1 = x_1 \dots x_{\ell-1}x_\ell x'_{\ell'-1} \dots x'_1$. Of all such vertices, let w_1 be a vertex whose neighborhood's intersection with P , say x_i , is closest to x_1 on P . If there are ties, then take w_1 to be the vertex whose neighborhood's intersection with P' , say x'_j , is closest to x'_1 on P' . Since G has girth at least 5, vertex w_1 is well defined. Now recalling that P and P' are shortest paths from their tips to a 4^+ -vertex, it can be easily seen that w_1 must be a 3-vertex (cf. Remark 8.7). By Lemma 8.6 and the choice of w_1 , two vertices of the path $Q_2 = x_1 \dots x_iw_1x'_j \dots x'_1$ have a common neighbor (one of which must be w_1). Notice that $|Q_2| < |Q_1|$, since the girth of G is at least 5. Choose w_2 for Q_2 the same way that we chose w_1

for Q_1 and we continue the process. Since the lengths of Q_i 's are decreasing, the process cannot be repeated forever, so at some point we get a contradiction to Lemma 8.6.

Item (c) follows along the same lines as Item (b). \square

Since a windmill contains a unique 4^+ -vertex, it can only be a subgraph of another windmill if their axes coincide, and the larger one has at least one more sail. The sails of a windmill are pairwise internally disjoint by Lemma 8.10 (and the definition of a windmill, by which all tips have to be distinct). Moreover, there are no edges between vertices of the same sail, except if one of them is the starting vertex and the other is the axis. Note that every edge between the axis and the starting vertex of a sail is itself a sail.

8.3.1. Flags. For a fixed complete windmill, we now study vertices that are themselves outside the windmill, but that have two or more neighbors inside the windmill. We call such a vertex a *flag*. We need to classify the flags. For this, we make the following notational convention. Let H be a complete windmill and w an H -flag (i.e., a vertex not in H that has at least two neighbors in H). We say that w has *signature* $(d_0 \mid d_1, d_2, \dots)$, if $d(w) = d_0$, and the neighbors of w in H have degrees d_i , $i = 1, 2, \dots$, listed with multiplicities. We will conveniently replace sub-lists with an asterisk $*$: For example, w has signature $(d_0 \mid d_1, *)$ if it has degree d_0 and at least one of the neighbors of w in H has degree d_1 . Or, similarly, w has signature $(d_0 \mid d_1, d_2, d_3, *)$, if w has degree d_0 , at least three neighbors in H , and these three are of degrees d_i , $i = 1, 2, 3$. We will also replace the degree of w with a joker: w has signature $(* \mid d_1, d_2, \dots)$, if the neighbors of w in H have degrees d_i , $i = 1, 2, \dots$.

The concept of signature is only needed to reduce the possible occurrences of flags to a very small number of cases. It will not be used beyond this section.

Lemma 8.11. *Let H be a k -windmill. The graph G has no flag with either of the following signatures:*

- (a) $(2 \mid 2, *)$,
- (b) $(3 \mid 2, 2, *)$
- (c) $(* \mid 3, 3, *)$
- (d) $(* \mid 3, 4^+, *)$
- (e) $(3 \mid 2, 4^+, *)$

PROOF. Lemma 8.4 implies that there is no flag with signature $(2 \mid 2, *)$. Lemma 8.6 gives (b). Lemma 8.10(b) implies that there are no flags with signatures $(* \mid 3, 3, *)$ and $(* \mid 3, 4^+, *)$. For $(3 \mid 2, 4^+, *)$, let w be the flag and let y the axis of the (complete) windmill, and x the starting vertex of a sail such that $x \sim w \sim y$. (We use the symbol " \sim " for the adjacency relation in G). Since w is a 3-vertex, xwy is a sail. Adding the vertex w to H thus gives a larger windmill, contradicting the maximality of H . \square

Lemma 8.12. *Let H be a complete windmill with axis y and let w be an H -flag of signature $(3 \mid 2, 3, *)$, so that there are $x, v \in V(H)$ with $d(x) = 2$, $d(v) = 3$, and $x \sim w \sim v$. Then x and v are not in the same sail of H .*

PROOF. Suppose x and v are on the same sail P . Then $xPv + vw + wx$ is a cycle consisting of 2- and 3-vertices only. By Lemma 8.5, there must be a vertex z which is a common neighbor of two vertices on the cycle. By Lemma 8.10(b), one of these two is w . Denote the other by u and note that u is on xPv , but $u \neq x, v$.

Since v has degree 3, z must also have degree 3, because P is a sail starting in x , and z having degree different from 3 would contradict the minimality of the distance from the tip to the end-vertex (by Remark 8.8). Now consider any sail with tip xw . Since w is of degree 3, it must contain either v or z . But v cannot be contained in such a sail, because otherwise it would have a non-empty interior intersection with P , contradicting Lemma 8.10(c). It follows that

there is a sail Q with tip xw containing z such that

$$(*) \quad |wz + zQ| < |wv + vPy|.$$

However, the length of P is at most the length of $xPu + uz + zQ$, and the inequality must be strict, because otherwise the sail $xPu + uz + zQ$ would have a non-empty interior intersection with the sail Q . Thus, it follows that

$$(**) \quad |uPy| < |uz + zQ|.$$

Now $(*)$ and $(**)$ together imply that $|uPy| < 1 + |zQ| < 1 + |vPy|$, from which we conclude $|uPy| < |vPy|$, a contradiction to the fact that u is between v and x . \square

The remaining cases are more complex. We start with the following fact.

Lemma 8.13. *Let w be an H -flag with signature $(* | 2, 3, *)$, so that there are $x, v \in V(H)$ with $d(x) = 2$, $d(v) = 3$, and $x \sim w \sim v$. If x and v are not in the same sail of H , then v is adjacent to the axis y of H .*

PROOF. By Lemma 8.11(a), w has degree at least 3. Suppose that v is on the sail P' of H with starting vertex x' , and $v \not\sim y$. (Note that $x' \neq x$.)

On one hand, if the degree of w is 3, then we have an internally shortest path $x'P'v + vw + wx$ ending in a 2-vertex, contradicting Remark 8.7. On the other hand, if the degree of w is 4^+ , then $x'P'v + vw$ is a path shorter than P' , but ends in a vertex not of degree 3, contradicting Remark 8.8. \square

Lemma 8.14. *No flag can have signature $(3 | *)$.*

PROOF. The cases $(3 | 2, 2)$, $(3 | 4^+, *)$ and $(3 | 3, 3)$ are dealt with in Lemma 8.11.

Let us consider $(3 | 2, 3)$. By Lemmas 8.12 and 8.13, denoting the flag by w , the 2-vertex by x , the 3-vertex by v , and by x' the starting vertex of the sail P' containing v , the start and end vertices of the path $Q := x'P'v + vw + wx$ have degree 2. By Lemma 8.6, there is a vertex z which is a common neighbor of two vertices on Q . By Lemma 8.11(c), one of the two must be w . Denote the other one by u .

If z has degree 4 or more, then the length of $x'Qu + uz$ is shorter than that of P' , contradicting Remark 8.8.

If z has degree 3, then the length of $x'Qu + uz + zw + wx$ is at most the length of P' (because G has no C_3 or C_4), contradicting Remark 8.8. \square

Lemma 8.15. *Let w be an H -flag with signature $(4^+ | 2, 3, *)$, so that there are $x, v \in V(H)$ with $d(w) = 4^+$, $d(x) = 2$, $d(v) = 3$, and $x \sim w \sim v$. If x and v are on the same sail of H , then v is adjacent to the axis y of H .*

PROOF. Denote the sail containing both x and v by P . If v is not adjacent to the axis y , then $xPv + vw$ is a path shorter than P that ends in a vertex not of degree 3, contradicting Remark 8.8. \square

We conclude the subsection with the following important consequence of our investigation of flags. Per se, windmills are subgraphs of G induced by sails, but the next lemma shows that in a complete windmill, every edge already belongs to some sail.

Lemma 8.16. *All edges of a complete windmill are on sails.*

PROOF. To argue by contradiction, consider an edge $e = uv$ whose vertices are on a windmill, but which is not on a sail. By Lemma 8.4, vertices u and v cannot be both 2-vertices. By Lemma 8.10, vertices u and v cannot be both 3-vertices on the same sail. The same lemma, shows a 3-vertex and the axis cannot be the endpoints of e . Now let u be a 2-vertex and v be a 3-vertex. Vertices u and v can either be in the same sail or not. The former contradicts Lemma 8.5 and Lemma 8.10. For the latter case, let P be a sail which contains a 3-vertex v

which is adjacent to a 2-vertex u on another sail. Let x be the 2-vertex of the sail P . Then $xPv + vu$ is an internally shortest path, contradicting Lemma 8.10.

The only remaining case is when u_1 and u_2 are two 3-vertices on distinct sails P_1, P_2 , of the windmill. Say $u_1 \in P_1$ and $u_2 \in P_2$, and that the starting vertex of P_i is x_i (note that $x_1 = x_2$ is possible). Denote by $Q := x_1P_1u_1 + u_1u_2 + u_2P_2x_2$ the path (or cycle) starting with the tip of one of the two sails, taking the edge e , and ending in the tip of the other – the two starting vertices x_i of the sails may coincide, in which case Q is a cycle.

By Lemma 8.6 (or Lemma 8.5 if Q is a cycle), there must be a vertex w which is adjacent to two vertices y_1, y_2 on Q . The vertex w cannot be the axis as it contradicts Remark 8.8 or entails that there is a triangle. Hence, w is a flag.

By Lemma 8.11(c), one of the y_i must be a 2-vertex, the other may be a 2- or a 3-vertex. If, say y_2 is a 3-vertex, then either by Lemma 8.13 or Lemma 8.15, y_2 must be adjacent to the axis — a contradiction, since the only vertex on Q which might be adjacent to the axis has degree at most three, and two neighbors on Q , the third is the axis (and w is not the axis).

Hence, we conclude that y_1 and y_2 are both 2-vertices (in particular, $x_1 \neq x_2$). Since, by what we have just said, w is not adjacent to a third vertex of Q , by Lemma 8.5, there is another vertex w' which is adjacent to two 3-vertices on Q . But such a vertex would be a flag with signature $(* | 3, 3, *)$, which is impossible by Lemma 8.11(c). \square

8.3.2. The flag graph. We have narrowed down the possible configurations involving flags of windmills. To summarize the results above, a flag can be adjacent to

- several 2-vertices on the tips of sails,
- and at most one of the following:
 - the axis, or
 - one 3-vertex which is adjacent to the axis on a sail.

Moreover,

- only one flag can be adjacent to the axis,
- every 3-vertex as above can be adjacent to at most one flag (obviously),
- there are no edges except those in the sails of the windmill or incident to the flags.

This structure can be nicely dealt with in an inductive manner (rather than delving into a humongous list of case distinctions). In the remainder of this section, we show how the structure of flags on windmills can be modelled by a directed graph which we call *flag graph*, which has a tree-like structure. The possibility of a flag which is adjacent to the axis is a complication. Such a flag, if existent, is omitted from the construction of the flag graph.

Let W be a complete windmill contained in G . A flag which is adjacent to the axis of W is called *irregular* (recall that there can be at most one); the other flags are called *regular*. It is important to realize that a sail whose tip is adjacent to an irregular flag has length at least 3, because the girth of G is at least 5.

The flag graph $F = F(W, G)$ of a windmill W is a directed bipartite graph. One side of the bipartition of the vertex set of F comprises the *flag-vertices*, which are in one-to-one correspondence with the regular flags of F . The other side of the bipartition consists of the *sail-vertices*, which are in one-to-one correspondence with the 2-vertices at the tips of the sails of W . We say that a sail-vertex which corresponds to a 2-vertex contained in two sails is *degenerate*; a sail-vertex corresponding to a 2-vertex contained in only one sail is called *non-degenerate*. There are two types of arcs:

- 2-arc* Whenever a regular flag w is adjacent to a 2-vertex on the tip of a sail P of W , we have an arc from the sail-vertex corresponding to P to the flag-vertex corresponding to w . Note that, in this case, the sail-vertex is non-degenerate.

3-arc Whenever a regular flag w is adjacent to a 3-vertex x in W , then there is an arc from the flag-vertex corresponding to w to the sail-vertex which represents the sail containing x .

The flag graph may contain anti-parallel arcs: A regular flag might be adjacent to a 2- and a 3-vertex of the same sail.

We note the following observations which follow directly from the construction and the earlier results of this section (see the summary above).

Lemma 8.17. *Let G contain a windmill W , and let $F = F(W, G)$ be its flag graph.*

- (a) *Degenerate sail-vertices have out-degree 0; non-degenerate sail-vertices have out-degree at most 1.*
- (b) *Flag-vertices have out-degree at most 1.*
- (c) *Degenerate sail-vertices have in-degree at most 2; non-degenerate sail-vertices have in-degree at most 1.*
- (d) *Flag-vertices have in-degree at least 1.*

Moreover, the undirected connected components of F are in one-to-one correspondence with the blocks in a block-decomposition of H . In particular, only non-degenerate sail-vertices can be contained in directed cycles. \square

We now show how the flag graph can be constructed inductively from basic elements and construction rules.

- **Basic element S** A single non-degenerate sail-vertex.
- **Basic element S₋** A single degenerate sail-vertex.
- **Basic element S⁺** A flag-vertex and a sail-vertex with an arc from the sail-vertex to the flag-vertex.
- **Basic element C2** A cycle of length 2.
- **Basic element C4** A cycle of even length at least 4.

If F is a flag graph, it can be extended with the following construction rules:

- **Construction rule U** Start a new connected component (in the undirected sense) by adding one of the basic elements without connection to F .
- **Construction rule A** Add a sail-vertex with an outgoing 2-arc linking it to an arbitrary flag-vertex of F .
- **Construction rule B** Add a flag-vertex and a sail-vertex, together with a 2- and a 3-arc: the 3-arc goes from the new flag-vertex to an arbitrary sail-vertex in F , and the 2-arc goes from the new sail-vertex to the new flag-vertex.

Lemma 8.18. *Let G contain a windmill W , and let $F = F(W, G)$ be its flag graph. Then F can be constructed using the above basic elements and construction rules.*

PROOF. We show that each connected component of the flag graph can be constructed in an inductive manner as follows. Let C be a connected component of the flag graph. Assume that C cannot be obtained in one step by Construction rule U. If C has a sail vertex s with out-degree 1 and in-degree 0 whose flag neighbour has either out-degree or in-degree at least 2, then construct $C \setminus s$ first and then obtain C by applying Construction rule A (notice that by the definition, the flag neighbour of s must be in $C \setminus s$ as C cannot be constructed by Construction rule U). Otherwise, C has a sail vertex s with out-degree 1 and in-degree 0 whose flag neighbour f has both out-degree and in-degree 1. In this case, construct $C - \{s, f\}$ first and then obtain C by applying Construction rule B. \square

In Section 8.5, this construction will be used to inductively “glue together” the graph induced by a windmill and its flags. Before we can do that, in the following section, we explain the gluing operation.

8.4. Typed graphs and gluing

Let P be a path and Q a local minimum with value μ in P . We say Q is an *imin* if $\mu < 0$ or Q contains no endvertices of P .

Definition 8.19. A *typed graph with types* τ is a graph together with a mapping $\tau: V(G) \rightarrow \{0, 1, 2\}$. In other words, every vertex has one of three possible types: it can be either a type-0, a type-1, or a type-2 vertex. The figures in this section show graphs with the types of the vertices in square brackets.

A *decent labeling* of a typed graph is a good edge-labeling with the following properties:

- (a) If P is a path between two type-2 vertices, then the length of P is at least three, and at least one of the following two conditions hold:
 - (a.1) there is an imin on P such that between each endpoint of P and this imin, there is an edge with strictly positive label;
 - (a.2) there are (at least) two imins.
- (b) If P is a path between a type-1 vertex v and a type-2 vertex w , then the length of P is at least two, and
 - (b.1) there is an imin on P which does not contain v .

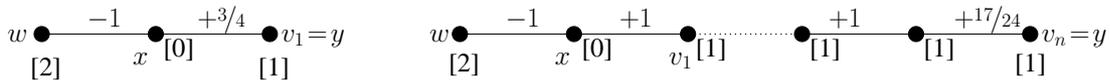
A typed graph is *decent* if it has a decent edge-labeling.

Note that if a path P satisfies either of the conditions (a.1) or (a.2), then P also satisfies (b.1). Moreover, if a typed graph G has no type-2 vertex, then any good edge-labeling of G is also decent.

For $t \in \{1, 2\}$, a path in a typed graph is called *t-simple*, if the type of every interior vertex is strictly less than t . We leave it to the reader to convince himself that, in order to verify that a good edge-labeling is decent, it suffices to check 2-simple paths in (a), and 1-simple paths in (b), respectively.

Before we continue discussing typed graphs in general, we discuss several examples which we will need in the remainder of the paper: We describe typed graphs and define concrete decent edge-labelings on them (where they exist). The graphs are indeed *rooted graphs* with the root denoted by y — this is owed to the fact that we will later apply them to windmills, with the root corresponding to the axis of the windmill. Hence, we will discuss multiple versions of some graphs, with the difference lying only in the location of the root vertex y .

Example 8.20. *Decent labeling of typed paths ending in a type-2 vertex.* Let P be a path of length at least two with root vertex y as shown in Fig. 8.2, and ending in a type-2 vertex w . Denote the vertex adjacent to w by x . The root y is type-0 or type-1, x is type-0, and all remaining vertices of P are type-1 vertices. Then P is a decent typed graph. Fig. 8.2 shows a decent edge-labeling ϕ . If P has length two, then $\phi(wx) = -1$ and $\phi(xy) = 3/4$. If P has length at least three, then let $\phi(wx) = -1$, $\phi(v_{n-1}v_n) = 17/24$ and the rest of the edges have label $+1$.

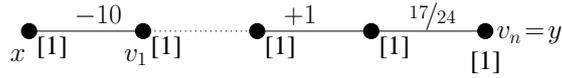


typed 2-path w/ decent labeling

typed 3⁺-path w/ decent labeling

FIGURE 8.2. Typed paths with decent edge-labelings

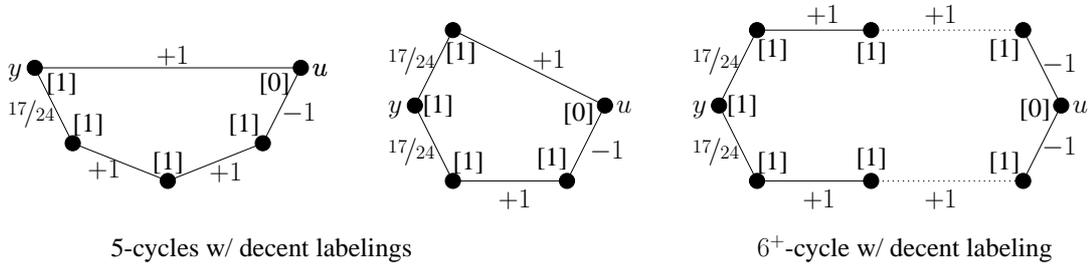
Example 8.21. *Decent labeling of typed paths ending in a type-1 vertex.* Let P be a path of length at least three with root y as shown in Fig. 8.3. The type of the root vertex is 0 or 1, and all other vertices have type 1. Then P is a decent typed graph. Fig. 8.2 shows a decent edge-labeling. The additional edges (dots) have label $+1$.



typed path w/ decent labeling

FIGURE 8.3. Typed paths with decent edge-labelings

Example 8.22. *Decent labelings of cycles without type-2 vertices.* Let C be a cycle on type-0 and type-1 vertices as shown in Fig. 8.4. Let u be a type-0 vertex and all remaining vertices of C be type-0 or type-1 vertices. Then C is a decent typed graph. Fig. 8.4 shows decent edge-labelings. If C has length five, then based on the position of vertex y , we present two different decent labeling for later applications, either of which is a decent edge-labeling of the typed 5-cycle C independent of the position of y .



5-cycles w/ decent labelings

6⁺-cycle w/ decent labeling

FIGURE 8.4. Typed cycles with decent edge-labelings

Example 8.23. *Decent labelings of cycles with one type-2 vertex, part I.* Consider a typed graph consisting of a cycle and one extra edge with one end on the cycle as shown in the top part of Fig. 8.5. The root y is off the cycle and has type 0 or 1. There is a type 2 vertex on the cycle, and it is adjacent to two type 0 vertices. The rest of the vertices have type 1. The figure shows a decent labeling.

Example 8.24. *Decent labelings of cycles with one with type-2 vertex, part II.* Consider a typed graph consisting of a cycle and one extra edge with one end on the cycle as shown in the middle part of Fig. 8.5. The root y is on the cycle, and has type 0 or 1. There is a type-2 vertex on the cycle, and it is adjacent to two type-0 vertices. The rest of the vertices have type 1. For given α, β satisfying $2/3 < \alpha < 3/4 < \beta$, a decent labeling can be constructed, as shown in the figure.

NEW TODO: CHECK!

Example 8.25. *Decent labelings of cycles with one with type-2 vertex, part III.* Consider a typed graph consisting of a cycle as shown in the bottom part of Fig. 8.5. The root y is on the cycle, and has type 0 or 1. There is a type-2 vertex on the cycle, and it is adjacent to two type-0 vertices. The rest of the vertices have type 1. For given α, γ satisfying $2/3 < \alpha < 3/4$ and $\gamma < 0$, a decent labeling can be constructed, as shown in the figure.

A k -wheel is a typed graph H which is the union of a cycle and a center vertex connected to $k \geq 2$ of the vertices on the cycle, called *anchors*. The distance on the cycle of any two anchors must be at least three. Fix an orientation of the cycle. A successor of an anchor is called a *bogey*; the successor of a bogey is called a *spectator*; all other vertices on the cycle are called *boobies*. A path contained in the cycle connecting successive anchors is called a *segment*. The vertices are to have the following types:

- anchors and spectators have type 0;

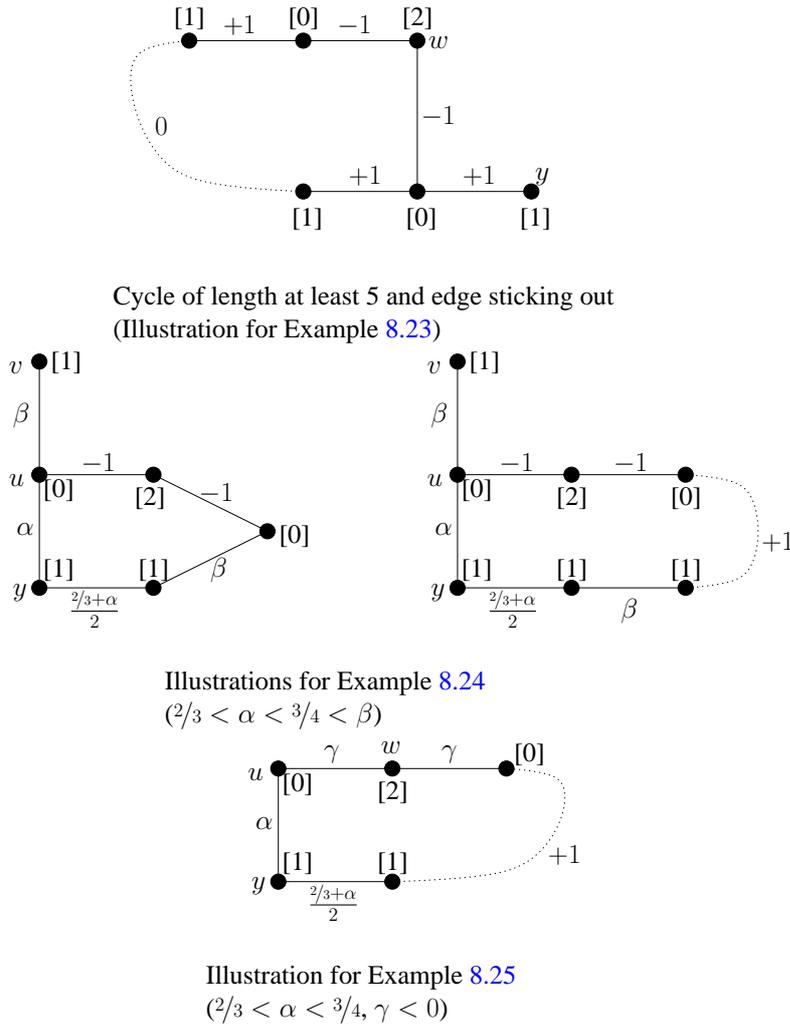


FIGURE 8.5. Cycles having a type-2 vertex with decent edge-labelings.

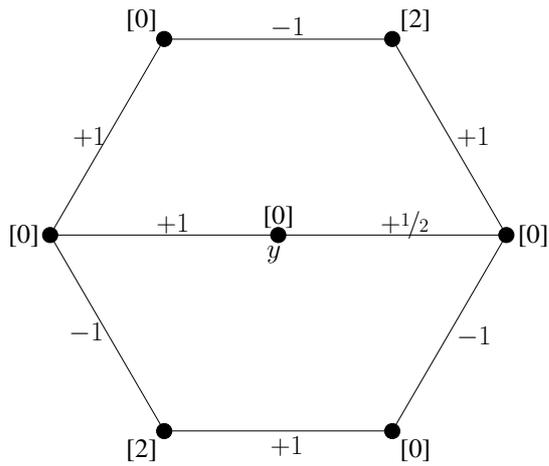
- the bogies have type 2;
- the boobies have type 1; and
- the center vertex either type 0 or type 1.

We divide the class of wheels into 3 subclasses: Benign wheels, almost evil wheels, and evil wheels. The first two kinds are decent, while the third is not.

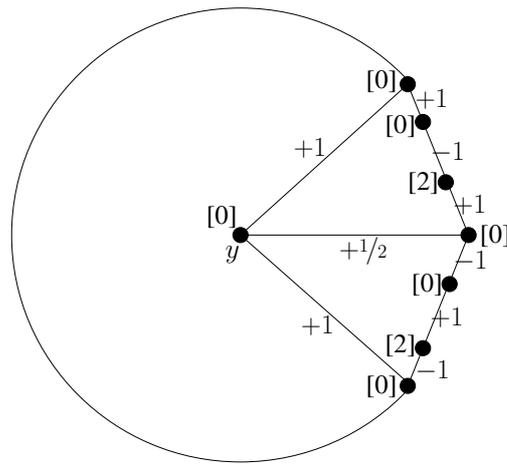
Example 8.26. *Decent labelings of benign wheels.* Consider a wheel in which the center is a type-1 vertex but contains at least one pair of consecutive anchors whose distance is at least four (this is the “benign” segment of the wheel). These typed graphs are called *benign wheels*. Fig. 8.6 shows decent edge-labelings of benign wheels.

Example 8.27. *Decent labeling of almost evil wheels.* Consider a wheel in which the distance of every pair of consecutive anchors is exactly three and the center is a type-0 vertex. These typed graphs that are only different from evil wheels in the type of the center, are called *almost evil wheels*. Fig. 8.7 shows decent edge-labelings of almost evil wheels.

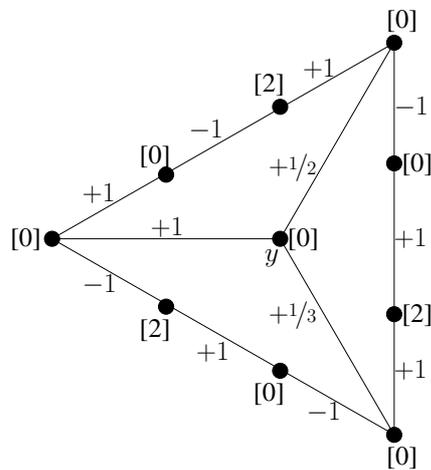
Example 8.28. *Evil wheels are not decent.* If the distance of every pair of consecutive anchors in a wheel is exactly three and the center is a type-1 vertex, as shown in Fig. 8.8, then wheel is called an *evil wheel*. It can be easily seen that evil wheels have a good edge-labeling. However, they are not decent.



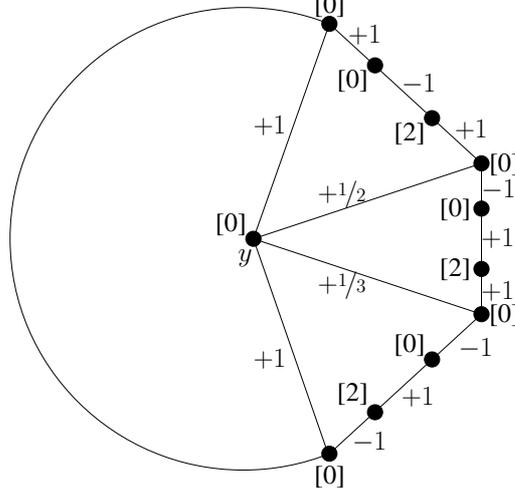
Almost evil 2-wheel w/ decent labeling



Almost evil even wheel w/ decent labeling



Almost evil 3-wheel w/ decent labeling



Almost evil odd wheel w/ decent labeling

FIGURE 8.7. Almost evil wheels with decent labelings

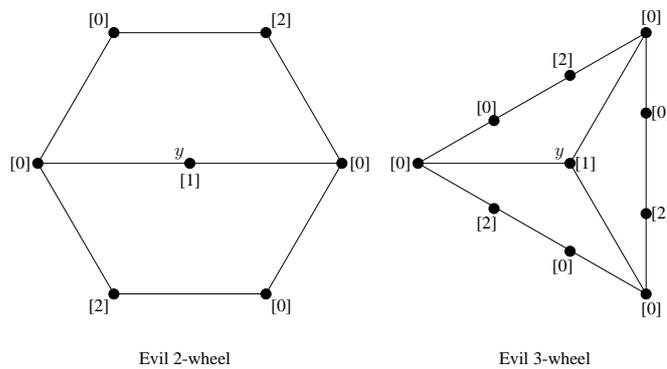


FIGURE 8.8. Evil wheels are *not* decent.

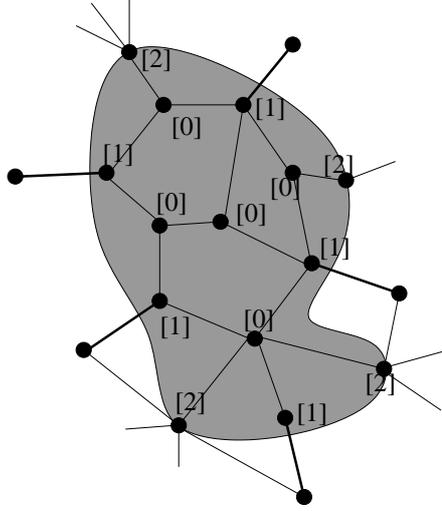


FIGURE 8.9. A swell subgraph

Notice that with the above mentioned relabeling, an imin on P is in fact a local minimum on C .

First, assume that both x and y are of type 2. If P has at least two imins, then those two are in fact two local minima in C . Otherwise P has an imin such that between each endpoint of P and this imin, there is an edge with strictly positive label. Moreover, by the scaling of labels in H , these two labels have value at least 2. Considering the scale of the labels in G' , there is a local minimum of C that belongs to G' . This local minimum in addition to the imin on P are two local minima of C .

If x has type 2 and y has type 1, then the edge e of $C \setminus P$ adjacent to y has label $-\infty$. Moreover, by the definition of a decent labeling, there is an imin on P which is not incident to e . By the same argument as above, this imin in P is a local minimum in C . Hence, we have two local minima on C .

Finally, if both x and y have type 1, let e and f be the edges of $C \setminus P$ incident to x and y , respectively. By the definition of a swell subgraph, e and f cannot be adjacent. So e and f are local minima of C as their labels are $-\infty$, \square

8.4.2. Gluing. In order to use decent typed graphs in inductive arguments, we have the following construction which allows to “glue” two decent typed graphs and obtain a new one.

Definition 8.32. Let G_1 and G_2 be typed graph with types τ_i , $i = 1, 2$, let H be an induced subgraph of both G_1 and G_2 , and $V(H) = V(G_1) \cap V(G_2)$. We say that the typed graph G with types τ is the result of *gluing* G_1 and G_2 along H , if $V(G) = V(G_1) \cup V(G_2)$, $E(G) = E(G_1) \cup E(G_2)$, and

$$\tau(v) = \begin{cases} \tau_1(v), & \text{if } v \in V(G_1) \setminus V(G_2), \\ \tau_2(v), & \text{if } v \in V(G_2) \setminus V(G_1), \\ \min(\tau_1(v), \tau_2(v)), & \text{if } v \in V(G_1) \cap V(G_2). \end{cases}$$

We wish to have conditions which ensure that if G_1 and G_2 are decent, then G is, too. As a motivating example, the reader might want to verify the following fact (which we do not need in this paper):

Lemma 8.33. *If for all $v \in V(H)$ we have $\tau_1(v) = \tau_2(v) = 2$, and if G_1 and G_2 are decent, then G is decent.*

Our aim is to decompose windmills into elementary parts—indeed, all parts we need have been discussed in Examples 8.20–8.24 and 8.27–8.26. For this, we need a considerably more powerful gluing mechanism than that of Lemma 8.33. We define the class of “gluable” typed graphs, which can be glued to each other by 1- and 2-sum operations.

We need to first classify certain special type-2 vertices.

Definition 8.34. For a given quadruple (G, τ, ϕ, y) consisting of a typed graph G with types τ , a decent edge labeling ϕ of G , and a root vertex y of G we say that a type-2 vertex w is *locked* if the distance $d_G(w, y)$ between w and y is two, the (unique) path P between w and y of length two has an imin, and the edge incident to y on P has label in $]2/3, 3/4[$.

We call P the *locking path* of w . If w is not locked, we call it *connectable*.

Now we are ready to give the complete definition of the gluing operation.

Definition 8.35. We say that a quadruple (G, τ, ϕ, y) consisting of a typed graph G with types τ , a decent edge-labeling ϕ of G , and a root vertex y of G is *gluable*, if the following conditions hold.

- (a) Every path y, v_1, v_2 of length 2 originating from y and containing a type-1 vertex v_1 and a vertex v_2 of type 0 or 1 is *admissible*: With $\alpha := \phi(yv_1)$ and $\beta := \phi(v_1v_2)$, we have

$$2/3 < \alpha \leq 3/4 < \beta.$$

- (b) Every 1-simple path² of length at least one between a type-1 vertex and y contains an edge with value at least $2/3$.
(c) Not type-2 vertex is adjacent to the root y .
(d) Let w be a type-2 vertex in G . If the distance $d_G(w, y)$ between w and y is two, then every 2-simple path P between w and y except for the locking path of w , if it exists, satisfies one of the following:

- (d2.i) P has an imin, and the edge incident to y on P has label at least $3/4$; or
(d2.ii) The edge incident to y on P is a local minimum with value in $]0, 1/2]$.

If the distance $d_G(w, y)$ between w and y is at least three, then every 2-simple path P between w and y satisfies

- (d3) The edge incident to y on P is a local minimum with value in $]2/3, 3/4[$; and there is a second imin of P between this edge and w .

Before we prove that gluable graphs can be glued to each other, we review the examples from the beginning of this section.

Example 8.36. The typed graphs with the decent labelings and root-vertices y described in the examples in the previous subsection are all gluable. Checking this amounts to mechanically going through all the t -simple paths of the graphs. We omit it here.

Let us now prove that gluing really works.

Lemma 8.37. *Let (G_1, τ_1, ϕ_1, y) and (G_2, τ_2, ϕ_2, y) be gluable, and let G result from gluing G_1 and G_2 along $\{y\}$. Moreover, for all $e \in E$, let $\phi(e) := \phi_1(e)$, if $e \in E(G_1)$ and $\phi(e) := \phi_2(e)$, otherwise. Then (G, τ, ϕ, y) is gluable.*

The proof is purely mechanical and can be found in the appendix.

Lemma 8.38. *Let $(G_1, \tau_1, \phi_1, y_1)$ and $(G_2, \tau_2, \phi_2, y_2)$ be gluable, w_1 a connectable type-2 vertex of G_1 and w_2 a connectable type-2 vertex of G_2 . Let G be the typed graph resulting from identifying y_1 with y_2 to y and w_1 with w_2 . If G_1 and G_2 are gluable, and $d_{G_1}(y_1, w_1) + d_{G_1}(y_2, w_2) \geq 5$, then (G, τ, ϕ, y) is gluable.*

²Recall the definition of t -simple from page 101.

The condition on the distances, which means that identifying $y_1 = y_2$ and $w_1 = w_2$ cannot create a C_4 in G , is needed because if the labels on two paths satisfy the condition (d2.ii), then gluing them does not give a good edge-labeling. The proof of Lemma 8.38 can be found in the appendix.

The operation which adds a graph of the kind described in Example 8.24 differs from the above two.

Let $(G_1, \tau_1, \phi_1, y_1)$ be a gluable graph, and (y_1, u_1, v_1) a path in G as in Definition 8.35(a). Let H be a typed graph with types τ , as described in Example 8.24. To specify the edge labeling of H , we let $\alpha := \phi_1(y_1 u_1)$ and $\beta := \phi_1(u_1 v_1)$. By Definition 8.35(a), these values satisfy the conditions in Example 8.24 to define the decent edge-labeling ϕ_2 of H . The proof of the following lemma is in the appendix.

Lemma 8.39. *The typed graph G' resulting from gluing G and H along $\{y = y_1, u = u_1, v = v_1\}$ is gluable.*

This is the only lemma that can create flags that are locked type-2 vertices. The following lemma will give us the option to add sails that connect to these flags, the proof is analogous to that of Lemma 8.39.

Let (G, τ_1, ϕ_1, y_1) be a gluable graph, and (y_1, u_1, w_1) a locking path in G . Let H be a typed graph with types τ , as described in Example 8.25. To specify the edge labeling of H , we let $\alpha := \phi_1(y_1 u_1)$ and $\gamma := \phi_1(u_1 w_1)$. By Definition 8.34, these values satisfy the conditions in Example 8.25 to define the decent edge labeling ϕ_2 of H .

Lemma 8.40. *The typed graph G' resulting from gluing G and H along $\{y = y_1, u = u_1, w = w_1\}$ is gluable.*

8.5. Non-existence of windmills

In this section, we prove the following theorem mentioned in the introduction.

Again, in this section, G is a critical graph of girth at least five. Let W be a windmill in G with axis y and k sails, and denote by \overline{W} be the subgraph of G induced by W and all of its flags, regular or not. We say that \overline{W} is the *closure* of W . Define types $\bar{\tau}$ for \overline{W} as follows:

$$(*) \quad \bar{\tau}(v) = \begin{cases} 2, & \text{if } v \text{ is a flag,} \\ \deg_G(v) - \deg_{\overline{W}}(v) & \text{otherwise.} \end{cases}$$

We will prove the following.

Lemma 8.41. *The typed graph \overline{W} with types $\bar{\tau}$ is decent, unless*

- *it contains an evil wheel, and there is no irregular flag;*
- *it contains an almost evil wheel, and there is an irregular flag.*

We will prove this lemma below. Disregarding the types, from this lemma, we can immediately derive the following main result.

Theorem 8.42. *For every windmill W in G , the closure \overline{W} of W contains an induced subgraph as depicted in Fig. 8.7, i.e., an (almost or not) evil wheel.*

PROOF. This follows from Lemma 8.41 by noting that \overline{W} is a swell subgraph of G , and invoking Lemma 8.31. \square

The proof of Lemma 8.41 is performed in two steps. We first prove that the ‘‘regular’’ part of H is gluable, and then add the irregular flag, if existent. For this, let H be the subgraph of G induced by W and its regular flags, and define types τ for H as in (*). We now prove the following.

Lemma 8.43. *There exists an edge-labeling ϕ such that (H, τ, ϕ, y) is gluable, unless it contains an evil wheel.*

PROOF. Suppose that H does not contain an evil wheel with axis y . Recall that this implies that H does not contain an evil wheel as a subgraph.

We prove that H is gluable. To do this, we associate to each of the basic elements (as laid down in Lemma 8.18) a gluable graph (one of the examples of the previous section); and to each of the construction rules, we associate one of the operations of Lemmas 8.37–8.39. By induction, this implies that H is gluable.

- **Basic element S** This corresponds to a typed path as in Example 8.20.
- **Basic element S₋** This corresponds to a cycle as in Example 8.22.
- **Basic element S⁺** This corresponds to a typed path as in Example 8.21.
- **Basic element C2** This corresponds to a cycle in H as in Example 8.23.
- **Basic element C4** This either corresponds to an almost evil wheel in H , as in Example 8.27, or to a benign wheel, as in Example 8.26, because evil wheels are excluded.

Suppose that the graph H' represented by a partial flag graph F' is gluable. We perform one of the construction rules to obtain an extended new flag graph F , and explain how we use the gluing lemmas to extend H' to a gluable graph H .

- **Construction rule U** This corresponds to taking a 1-sum as in Lemma 8.37. The identification takes place at the axes of the components.
- **Construction rule A** This corresponds to adding a path as in Example 8.20 or Example 8.25 via the 2-sum operation of Lemma 8.38 or Lemma 8.40, respectively. In the first case, the new sail-vertex from which the arc initiates corresponds to the path; the old flag-vertex which is the target of the arc identifies a flag w_1 of H' . This flag w_1 is identified with the vertex w of the path. The axis y_1 is identified with the root vertex y_2 of the path. In the second case, the flag is the type-2 vertex w in the bottom part of Fig. 8.5, with the bottom path connecting y and w corresponding to the new sail-vertex.
- **Construction rule B** This corresponds to adding a cycle as in Example 8.24 via Lemma 8.39. The sail-vertex of F to which the new vertices are attached, identifies a sail (degenerate or not) in H . The two edges in this sail (or, on one path of the sail in the case when it is degenerate) which are closest to the axis y correspond to the two vertically drawn edges in the middle part of Fig. 8.5, yu, uv . The new flag-vertex is the type-2 vertex in that picture, and the new sail-vertex corresponds to the path between the root y and the type-0 vertex to the right of the type-2 vertex (the path which does not use the vertex u).

□

We point out the following property of the edge-labeling ϕ constructed in this proof.

Remark 8.44. If W has an irregular flag w_0 , then on every sail P whose tip is adjacent to w_0 , the edge-labeling ϕ for H has the labels shown in Fig. 8.3.

In the second step, if necessary, we will need to add the irregular flag to H . This step will complete the proof of Lemma 8.41.

PROOF OF LEMMA 8.41. If no irregular flag exists, this lemma is just a weaker form of Lemma 8.43. Suppose that an irregular flag in \overline{W} exists; denote it by w_0 . We take the labeling ϕ from Lemma 8.43, and extend it to a decent labeling $\bar{\phi}$ of \overline{W} . To do this, let y, x_1, \dots, x_r be the neighbors of w_0 in \overline{W} . We let $\bar{\phi}(e) := \phi(e)$ for all $e \in E(H)$; $\bar{\phi}(yw_0) := -10$; and $\bar{\phi}(w_0x_j) := +1, j = 1, \dots, r$.

We now verify that the resulting labeling is decent, using the above Remark 8.44. Since there exists an irregular flag, by (*), we have $\tau(y) = 0$. Since, in Definition 8.19, we only need to check 2-simple paths, the only paths we need to check are those starting or ending in w_0 . Consider first paths starting in a type-2 vertex of H and ending in w_0 . Such a path enters

the sails whose tips are incident to w_0 either through the axis or through a vertex adjacent to the axis. In each of the two cases, the path touches a type-1 vertex before it reaches w_0 . By condition (b.1) of Definition 8.19, and using the fact that the edge with label -10 on the path is an imin, we find that such a path has at least two imins, i.e., it satisfies condition (a.2).

Secondly, consider a path P starting in w_0 and ending in a type-1 vertex. Since neither the axis nor the vertices adjacent to w_0 are type-1 vertices, the edge with label -10 is an imin not incident to the type-1 vertex in which P ends, and thus (b.1) is satisfied.

We leave it to the reader to verify that $\bar{\phi}$ is in fact a good edge-labeling of \bar{W} . \square

8.6. Proof of Theorem 8.2

Let G be a minimum counter example to Theorem 8.2, i.e., G is a critical graph of girth at least 5 and with average degree less than 3. Let $\deg(v)$ denote the degree of vertex v . To every vertex v assign a charge of $6 - 2 \deg(v)$. The total charge of the graph is $\sum_v (6 - 2 \deg(v)) > 0$, because the average degree of G is less than 3. Note that after the assignment of initial charges, only 2-vertices have positive charge.

Now we discharge the graph according to the following rule:

- For every 2-vertex u and every neighbour v of u , if there are k sails with tip uv , then u sends $\frac{1}{k}$ charge (via these sails) to each 4^+ -vertex at the end of these k sails.

It can be seen that charges are sent from 2-vertices to 4^+ -vertices via paths consisting of only 3-vertices. Now we show that after the discharging phase, every vertex of the graph has nonpositive charge, a contradiction. Indeed, let v be a vertex. We consider the following cases.

- (i) v is a 2-vertex. Then it has an initial charge of 2. In the discharging, v sends a total of 1 unit of charge out via each of the two tips, and v does not receive any charge in the discharging phase. So v has 0 charge after the graph is discharged.
- (ii) v is a 3-vertex. Then v has an initial charge of 0. Moreover, v does not gain or lose any charge in the discharging phase. So v has 0 charge after the graph is discharged.
- (iii) v is a 4-vertex. Then v has an initial charge of -2 . To become positive, it must receive charges via at least three incident edges, implying that v is the axis of a windmill. (We note that this holds true even if two sails share a common tip and both end in v .) By Lemma 8.41, such a windmill must contain an evil or almost evil wheel as shown in Fig. 8.8 and Fig. 8.7. It can be seen that in both cases, the axis of the windmill is at the same distance from the 2-vertices of the windmill as one of the flags. Hence, vertex v receives at most $\frac{1}{2}$ via each sail of the wheel. Thus, the charge of v after discharging is either at most $-2 + 3 \cdot \frac{1}{2} = -1/2$ or at most $-2 + 2 \cdot 1/2 + 1 = 0$.
- (iv) v is a 5-vertex. Then v has an initial charge of -4 . To become positive, it must receive charges via every incident edge, implying that v is the axis of a 5-windmill in G . Again, similar to the above case, Lemma 8.41 implies that such a windmill contains an evil or almost evil wheel in both of which cases, the axis of the windmill is of the same distance from the 2-vertices of the windmill as one of the flags. Hence, vertex v receives at most $\frac{1}{2}$ via each sail of the wheel. Thus, the charge of v after discharging is either at most $-4 + 4 \cdot 1 = 0$ or at most $-4 + k \cdot \frac{1}{2} + (5 - k) \times 1 \geq 0$ where $k \geq 2$.
- (v) v is a 6^+ -vertex. Then v has an initial charge of $6 - 2 \deg(v)$. Since v receives at most 1 unit of charge via each incident edge, v has at most $(6 - 2 \deg(v)) + \deg(v) \leq 0$ charge after the graph is discharged.

8.7. Conclusion

We have seen that imposing a lower bound on the girth facilitates the construction of good edge-labelings, or even decent edge-labelings. In this paper, we have used this approach together with a degree-bound. It seems probable that high girth benefits other open problems

about good edge-labeling. For example, Araújo et al. [ACGH12] conjecture that for every $c < 4$, the number of (pairwise non-isomorphic) critical graphs with average degree at most c is finite. We propose the following weakening of their conjecture.

Conjecture 8.45. *For every $c < 4$, the number of (pairwise non-isomorphic) critical graphs with girth at least five and average degree at most c is finite.*

This paper settles the case $c = 3$. For $c = 3$, but without restriction to the girth, a modification of Conjecture 8.1 proposes itself naturally:

Conjecture 8.46 (Araújo-Cohen-Giroire-Havet/modified). *There is no critical graph with average degree less than 3, with the exception of C_3 , $K_{2,3}$, and the graph displayed in Fig. 8.1.*

Acknowledgments. We would like to thank the anonymous referees for their thorough work and insightful comments.

Appendix: Deferred proofs

PROOF OF LEMMA 8.37. The conditions of Definition 8.35 are satisfied, since they require to check paths originating from the root vertex y only. Moreover, ϕ is a good edge-labeling. It remains to show that ϕ is decent.

To verify property (a) of Definition 8.30, let Q be a 2-simple path between two type-2 vertices $w_1 \in V(G_1) \setminus V(G_2)$ and $w_2 \in V(G_2) \setminus V(G_1)$. Let $P_1 := w_1Qy$ and $P_2 := yQw_2$. If both Q_1 and Q_2 satisfy (d2.i) then the edges incident to y in neither Q_1 nor Q_2 are part of the respective imins, and (a.2) of Definition 8.30 holds. If Q_1 satisfies (d2.i) and Q_2 satisfies (d2.ii) then (a.2) of Definition 8.30 holds: one of the imins is the imin of Q_1 , the other is the edge of Q_2 incident on y . If both Q_1 and Q_2 satisfy (d2.ii) then (a.1) holds, the imin there being the path of length two consisting of the two edges of Q incident on y . If Q_1 satisfies (d2.i) and Q_2 satisfies (d3), then (a.2) holds for P . If Q_1 satisfies (d2.ii) and Q_2 satisfies (d3), then (a.2) holds for P . If both Q_1 and Q_2 satisfy (d3), then (a.2) holds for P .

To verify property (b), let Q be a 1-simple path between a type-1 vertex $w_1 \in V(G_1) \setminus V(G_2)$ and a type-2 vertex $w_2 \in V(G_2) \setminus V(G_1)$. Note that the property in (b) of Definition 8.35 holds for Q_1 . If (d2.i) holds for Q_2 , then the imin of Q_2 is an imin of P not incident on w_1 . If (d2.ii) holds for Q_2 , then the edge incident on y in Q_2 is an imin of P not incident on w_1 . If (d3) holds for Q_2 , then the imin of Q_2 closer to w_2 is an imin of P not incident on w_1 . \square

PROOF OF LEMMA 8.38. Denote the vertex of G resulting from identifying y_1 and y_2 by y , and the one resulting from identifying w_1 and w_2 by w .

Let us first check Definition 8.35(a–d). Property Definition 8.35(a) is satisfied because no new path of this kind is added. The conditions of Definition 8.35(b–d) are satisfied, since they require to check 1- and 2-simple paths originating from the root vertex y only: these paths cannot contain w , and are thus contained entirely in either G_1 or G_2 .

We have to make sure that ϕ is good, and that it satisfies the conditions (a) and (b) of Definition 8.30. We may assume w.l.o.g. that $d_{G_2}(y_2, w_2) \geq 3$.

We first prove that ϕ is good. For this, let Q_1 be a path in G_1 between y and w_1 , and let Q_2 be a path in G_2 between y and w_2 . We have to verify that the cycle $C := Q_1 + Q_2$ has two local minima.

If (d2.i) holds for Q_1 and (d3) for Q_2 , then C has two local minima. If (d2.ii) holds for Q_1 and (d3) for Q_2 , then C has two local minima. If both Q_1 and Q_2 satisfy (d3), then C has two local minima.

Secondly, we prove properties (a) and (b) of Definition 8.30 hold. Note that for both these properties, we do not need to consider paths containing w as an interior vertex, because those are not 2-simple (in the case of (a)) or even 1-simple (for (b)). But this leaves us with the same situation which we have checked in the previous lemma. \square

PROOF OF LEMMA 8.39. We start by proving that ϕ is a good edge-labeling. For this, let C be a cycle in G containing edges of both $E(G_1) \setminus E(H)$ and $E(H) \setminus E(G_1)$. Such a cycle can be turned into a cycle C' in G_1 by replacing the path of C in $E(H) \setminus E(G_1)$ by the single edge u_1y_1 . We show that the fact that there are two local minima Q_1, Q_2 , on C' implies that there are two local minima on C .

Obviously, if any of the local minima of C' contains neither u_1 nor y_1 , then it is a local minimum of C . On the other extreme, if one of the two, say Q_2 , contains the edge u_1y_1 , then Q_1 and the path P_{-1} formed by the two edges with label -1 in $C \setminus E(G_1)$ are two distinct local minima, because $-1 < \alpha$. Thus, we have to make sure that if any of the local minima of C' contains exactly one of the vertices u_1 or y_1 , then it can be modified to be a local minimum of C . Firstly, suppose Q_1 has value μ_1 and contains u_1 but not y_1 . If $\mu_1 < -1$, then Q_1 is a local minimum of C ; if $\mu_1 > -1$ then P_{-1} is a local minimum of C ; if the two are equal, then $Q_1 + P_{-1}$ is a local minimum of C . Secondly, suppose Q_2 has value μ_2 and contains y_1 but not u_1 . If $\mu_2 < \frac{2/3+\alpha}{2}$, then Q_2 is a local minimum of C ; if $\mu_2 > \frac{2/3+\alpha}{2}$, then the path \dot{P} formed by the edge of $C \setminus E(G_1)$ with label $\frac{2/3+\alpha}{2}$ is a local minimum of C ; if the two are equal, then $Q_2 + \dot{P}$ is a local minimum of C .

Next, we have to show that the edge-labeling ϕ satisfies the properties (a) and (b) of Definition 8.19. For property (a), let w be the type-2 vertex of H , let w_1 be any type-2 vertex of G_1 , and let P be a w - w_1 -path in G . On the one hand, if $u = u_1$ is on P , then by Definition 8.19(b.1) applied to $P(w_1, u_1)$, Q has one imin not incident on u , and the edge of Q incident to w is a second, distinct, imin. On the other hand, if $y = y_1$ is on P , we use (d2.i), (d2.ii), or (d3) of Definition 8.35 for the path $Q' := Q(y_1, w_1)$. Indeed, if Q' satisfies (d2.i) the length of $Q(w, y)$ is two (i.e., it contains the edge uy), then Q has two imins; if Q' satisfies (d2.i) the length of $Q(w, y)$ is at least three, then Q has two imins; if Q' satisfies (d2.ii) the length of $Q(w, y)$ is two, then Q has two imins; if Q' satisfies (d2.ii) the length of $Q(w, y)$ is three, then Q has two imins; if Q' satisfies (d3) the length of $Q(w, y)$ is two, then Q has two imins; if Q' satisfies (d3) the length of $Q(w, y)$ is at least three, then Q has two imins.

Finally, to check the conditions of Definition 8.35, the only kind of paths which are added beyond those which were present in G_1 and H are those which result from taking a path Q_1 in G_1 from y_1 to $u_1 = u$, and adding the edge uw of H . Invoking the condition (b.1) of Definition 8.30, $Q_1 + uw$ contains two imins: one on Q_1 and the other being the edge uw . \square

Random lifts

Jointly with
Babak Farzad (Brock University, St. Catharines)

Abstract. Amit, Linial, and Matoušek (Random lifts of graphs III: independence and chromatic number, *Random Struct. Algorithms*, 2001) have raised the following question: Is the chromatic number of random h -lifts of K_5 asymptotically (for $h \rightarrow \infty$) almost surely (a.a.s.) equal to a single number? In this paper, we offer the following partial result: The chromatic number of a random lift of $K_5 \setminus e$ is a.a.s. three.

9.1. Introduction

Let G be a graph, and h a positive integer. An h -lift of G is a graph \tilde{G} which is an h -fold covering of G in the topological sense. Equivalently, there is a graph homomorphism $\phi: \tilde{G} \rightarrow G$ which maps the neighbourhood of any vertex v in \tilde{G} one-to-one onto the neighbourhood of the vertex $\phi(v)$ of G . The graph G is called the *base graph* of the lift.

More concretely, we may say that an h -lift of G has vertex set $V(G) \times [h]$ (where we let $[h] := \{1, \dots, h\}$ as usual). The set $\{v\} \times [h]$ is called the *fibre over v* . Fixing an orientation of the edges of G , the edge set of an h -lift is of the following form: There exist permutations σ_e of $[h]$, $e \in E(G)$, such that for every two adjacent vertices u and v of G , if the edge uv is oriented $u \rightarrow v$, the edges between the fibres $\{v\} \times [h]$ and $\{u\} \times [h]$ are $(u, j)(v, \sigma_{uv}(j))$, $j \in [h]$. Changing the orientation of the edges in the graph does not change the lift, provided that permutations on edges on which the orientation is changed are replaced by their respective inverses. In this spirit, for an edge uv in G , regardless of its orientation, we denote by σ_{uv} the permutation for which the edges between the fibres are $\{(u, j)(v, \sigma_{uv}(j)) \mid j \in [h]\}$.

By a *random h -lift* we mean a graph chosen uniformly at random from the graphs just described, which amounts to choosing a permutation, uniformly at random, independently for every edge of G .

Random lifts of graphs have been proposed in a seminal paper by Amit, Linial, Matoušek, and Rozenman [[ALMR01](#)]. Their paper sketched results on connectivity, independence number, chromatic number, perfect matchings, and expansion of random lifts, and was followed by a series of papers containing broader and more detailed results by the same and other authors [[AL02](#), [AL06](#), [ALM02](#), [LR05](#)], and e.g. [[BL06](#), [DL06](#)], [[BCCF06](#)].

In [[ALM02](#)] Amit, Linial, and Matoušek focused on independence and chromatic numbers of random lifts of graphs. They asked the following question.

Is there a zero-one law for the chromatic number of random lifts? In particular, is the chromatic number of a random lift of K_5 a.a.s. (for $h \rightarrow \infty$) equal to a single number (which may be either 3 or 4)?

A random h -lift \tilde{G} of K_5 a.a.s. has an odd cycle, whence a.a.s. we have $\chi(\tilde{G}) \geq 3$. Moreover, \tilde{G} a.a.s. does not contain a 5-clique. Brooks' theorem implies that a.a.s. $\chi(\tilde{G}) \leq 4$. So, a.a.s. $\chi(\tilde{G}) \in \{3, 4\}$.

In their paper, Amit, Linial, and Matoušek [[ALM02](#)] conjectured that the chromatic number of random lifts of any fixed base graph obeys a zero-one law, i.e., it is asymptotically almost

surely equal to a fixed number (depending only on the base graph). In the case when the base graph is K_n , they prove that $\chi(\tilde{G}) = \Theta(n/\log n)$ a.a.s. (the constant in the Θ notation may depend neither on h nor on n). Five is the smallest value for n , for which this is not trivial.

In this paper, we contribute the following to this problem.

Theorem 9.1. *A random lift of $K_5 \setminus e$ is a.a.s. 3-colorable.*

9.2. Notation and Terminology

Let $G := K_5 \setminus e$. Clearly, G is obtained by joining a cycle $C := [x_1, x_2, x_3]$ to a stable set $S := \{y_1, y_2\}$. Here, by *join* we mean that every vertex of C is made adjacent to every vertex of S . From now on, \tilde{G} will be a random h -lift of G . Let \tilde{G}_C and \tilde{G}_S denote the subgraphs of \tilde{G} induced by the fibres over the vertices of C and those over vertices of S , respectively. Moreover, for $x \in V(G)$, we denote by $V_x = \{x\} \times [h]$ the set of vertices of \tilde{G} over x . Similarly, for any set U of vertices of \tilde{G} and $x \in V(G)$, we let $U_x := U \cap V_x$.

As an *hors d'œuvre* intended to familiarise the reader with the most basic random lift arguments, we serve the following easy lemma.

Lemma 9.2. *The graph \tilde{G}_C is a union of cycles, each of which is divisible by three. A.a.s., the number of cycles in \tilde{G}_C is at most $\log^2 h$.*

PROOF. The cycles with length 3ℓ of \tilde{G}_C correspond to the cycles with length ℓ of the permutation $\sigma_{x_1x_2} \circ \sigma_{x_2x_3} \circ \sigma_{x_3x_1}$. The latter is a uniformly distributed random permutation of $[h]$. It is a folklore fact (e.g., [Lov07]) that the average number of cycles of a random permutation of $[h]$ is $\log h + o(1)$. The statement of the lemma now follows from Markov's inequality. \square

Lemma 9.2 allows us to assume that \tilde{G}_C has at most $\log^2 h$ cycles. As a matter of fact, this is the only statement about \tilde{G}_C which we need.

9.3. The 3-colouring algorithm

Our colouring algorithm is detailed in the box Algorithm 2. We use the colours *red*, *black*, and *white*, where the colour red will have a special significance. We point the reader to the fact that, once Algorithm 2 has coloured a vertex, the vertex never changes its colour or becomes uncoloured again. A vertex of \tilde{G}_S which is adjacent to precisely one red vertex is called *pale* (this is not a colour).

The algorithm works in three phases. In phase I, Steps (1–2), we destroy the uncoloured cycles of \tilde{G}_C by colouring one vertex per cycle red. By Lemma 9.2, a.a.s., we colour at most $\log^2 h$ vertices red in Phase I, i.e., Phase I fails with probability $o(1)$.

In Phase II, more accurately in the loop (4), the algorithm successively chooses uncoloured vertices of \tilde{G}_C and colours them red. This is done by maintaining the set $P(\cdot)$ of pale vertices (i.e., those vertices of \tilde{G}_S which are adjacent to precisely one red vertex).

In Phase III, Steps (5–7), the remaining vertices are coloured in a straight forward way.

The rationale behind the algorithm is as follows.

At any fixed time between Steps (3) and (5), consider the connected components of \tilde{G}_C after deleting all red vertices. These are uncoloured paths of different lengths in \tilde{G}_C , separated by red vertices. We call them *chunks*. These chunks can be thought of as the vertices of a multi-graph, which we call the *chunk-graph*, whose edges are the pale vertices in \tilde{G}_C : Every pale vertex has precisely two uncoloured neighbours in \tilde{G}_C , thus connecting the corresponding chunks. We refer to such a connection between chunks via a pale vertex as a *chunk-edge*. A chunk-edge may be a loop, which happens when a pale vertex have both uncoloured neighbours in the same chunk. Furthermore, there may be parallel chunk-edges in the chunk-graph, which

Algorithm 2 Three-Colour \tilde{G} **Phase I:**

- (1) The algorithm starts with all edges in \tilde{G}_C exposed, but no edge in between \tilde{G}_C and \tilde{G}_S exposed. If \tilde{G}_C has more than $\log^2 h$ cycles, **fail**.
- (2) Choose exactly one red vertex in each cycle of \tilde{G}_C .

Phase II:

- (3) Expose all edges incident to red vertices. If there exists a vertex in \tilde{G}_S which has two or more red neighbours, **fail**. Otherwise, let $P(0)$ be the set of pale vertices before the first iteration.
- (4) For $t = 1, \dots, \lfloor h^{1/3} \rfloor$:
 - (4.1) Let v be chosen arbitrarily from the set $P(t-1)$.
 - (4.2) From the two non-exposed edges incident to v , expose one arbitrarily (the other edge remains unexposed). Let u be the end-vertex in \tilde{G}_C of the exposed edge.
 - (4.3) Expose the other edge incident to u , and let v' be the corresponding neighbour of u in \tilde{G}_S . If $v' \in \bigcup_{s=0}^{t-1} P(s)$, **fail**. Otherwise $P(t) = P(t-1) \cup \{v'\} \setminus \{v\}$ (this is now the new set of pale vertices).
 - (4.4) Colour u red.

Phase III:

- (5) Expose all remaining edges.
- (6) Colour every vertex red which is in \tilde{G}_S and does not have a red neighbour.
- (7) If the graph induced by the non-red vertices is acyclic, colour it black and white, otherwise **fail**.

happens when two pale vertices connect the same pair of chunks. The reason why, in Step 3 of the algorithm, we abort if a vertex has two or more red neighbors, is only because such vertices would not correspond to edges of the chunk-graph. Indeed, at the end of Phase II, there are only two kinds of uncolored vertices left: Those making up the chunk graph, and those being colored red in Step 6.

The chunk-graph is a random multi-graph. At Step (3), it has as many vertices as there are cycles in \tilde{G}_C (at most $\log^2 h$ by Lemma 9.2), and as many edges as there are pale vertices. If the algorithm does not fail in Step (3), then to every red vertex there are two pale vertices, and they are all distinct. Hence, at this time, there are twice as many chunk-edges as there are chunks.

When the algorithm proceeds through loop (4), the number of chunks is increased as we colour more vertices of \tilde{G}_C red. However, the number of pale vertices stays constant, and hence so does the number of chunk-edges.

The reasoning at this point is a heuristic analogy with the random (simple) graph model $G(n, m)$, where a set of m edges is drawn uniformly at random from the set of all possible m -sets of edges between n vertices. For us, n is the number of chunks and m is the number of chunk-edges. At Step (3), where $m = 2n$, we expect the chunk-graph to contain lots of cycles (including loops and parallel edges), which makes it unlikely that it can be coloured with just the two remaining colours. However, when n grows and m stays constant, a random graph $G(n, m)$ will be acyclic as soon as $m \ll n$, and we expect the same to be true for the chunk-graph.

There are complications in making this heuristic analogy work rigorously, the foremost being that the distribution of the edges in the chunk-graph is not uniform but instead depends on the sizes of the chunks. We will address these issues in the next section.

9.4. Proof of correctness of the 3-colouring algorithm

We prove that a.a.s. Algorithm 2 properly 3-colours \tilde{G} .

Lemma 9.3. *A.a.s., Algorithm 2 does not fail in Steps (1), (3), or (4.3).*

PROOF. Lemma 9.2 implies that, a.a.s., the algorithm does not fail in Step (1).

For Step (3), note that, at this point in the algorithm, the probability that a fixed vertex in \tilde{G}_S has two or more red neighbours is $O((\log^4 h)/h^2)$. Hence, the probability that there exists such a vertex having two or more red neighbours is $O((\log^4 h)/h) = o(1)$.

For Step (4.3), we see that for each fixed t , the probability that $v' \in \bigcup_{s=0}^{t-1} P(s)$ is $O(h^{-2/3})$. Thus, the probability that the algorithm fails after at most t iterations is $O(th^{-2/3})$. Consequently, the probability that the algorithm fails at Step (4.3) before completing $t := \lfloor h^{1/3} \rfloor$ iterations is $o(1)$. \square

Denote by T the last iteration (value of t) of the loop (4) which is completed (without failing). We let $R(t)$, $t = 0, 1, \dots, T$ be the set of vertices which are red after t iterations of the loop (4). In particular, $R(0)$ is the set of vertices coloured red in Step (2). Let $R^+(t) := R(t) \setminus R(0)$. Recall that adding an index to a letter denoting a set refers to taking its intersection with the corresponding fibre, for example $R_x(t)$ refers to $V_x \cap R(t)$. Moreover, we use the following notation to refer to the cardinalities of each of these sets: If a set is denoted by an upper-case letter (possibly with sub- or superscript or followed by parentheses), the corresponding lower-case letter (with the same sub- or superscripts or parentheses) denotes its cardinality. For example $r_x(t) = |R_x(t)|$. We have the following.

Lemma 9.4. *For each $x \in C$ and $t = 1, \dots, T$, set $R_x^+(t)$ is uniformly distributed in the set of all $(r_x^+(t))$ -element subsets of $V_x \setminus R_x(0)$.*

PROOF. Fix an $x \in C$. In every iteration of the loop (4) in which the fibre over $x \in C$ is selected in Step (4.3), when exposing the edge in Step (4.3), the vertex u is selected uniformly at random from the set of all previously uncoloured vertices in V_x . In other words, for every fixed value of $R_x^+(t-1)$, the distribution of u is uniform. By induction, $R_x^+(t)$ is uniformly distributed. \square

Lemma 9.5. *In the loop (4) of Algorithm 2, a.a.s. no two adjacent vertices are coloured red.*

PROOF. Let $x_1, x_2 \in C$, and consider the situation after T iterations, i.e., when the algorithm leaves the loop (4). By Lemma 9.4, at this time, the expected number of edges between V_{x_1} and V_{x_2} both of whose end vertices are red is at most

$$h \cdot \frac{T}{h - r_{x_1}(0)} \cdot \frac{T}{h - r_{x_2}(0)} = O\left(\frac{h^{5/3}}{(h - \log^2 h)^2}\right) = o(1).$$

\square

Now, it only remains to show that when Step (7) of Algorithm 2 is reached, the graph consisting of the yet uncoloured vertices is a.a.s. acyclic.

Now, suppose that the algorithm has completed Phase II without failing, i.e., we find ourselves just before Step (5). Let H denote the chunk graph as we defined in Section 3. Thus H is a random multi-graph with $n \leq r(T) = T + r(0) = \Theta(h^{1/3})$ vertices and $m := p(T) = 2r(0) = O(\log^2 h)$ edges. In fact, if no two red vertices are adjacent, the first inequality becomes an equation, cf. Lemma 9.5. The distribution of H can be described in terms of random permutations taking into account the edges which have already been exposed, and the sizes of the chunks. It appears sensible to guess that H has no cycles. That is in fact correct.

Sizes of the chunks. The first thing we require to turn this analogy into a rigorous proof is an upper bound on the sizes of the chunks. We find it convenient to reduce the question to the distribution of the gaps between n points drawn uniformly at random from the interval $[0, 1]$. There, the probability that two consecutive points enclose a gap of size a is $(1 - a)^n$, which yields an upper bound of, say, $(2h \log n)/n$ for the largest gap, a.a.s. In the following lemmas, we put this plan into action.

Let n numbers Y_1, \dots, Y_n be drawn independently uniformly at random from $[N]$, where N is a function of n . Let S_k be the k -th order statistics (i.e., $0 \leq S_1 \leq \dots \leq S_n \leq 1$, and $\{S_1, \dots, S_n\} = \{Y_1, \dots, Y_n\}$) and set $S_0 := 0$ and $S_{n+1} := N$.

We determine the distribution of $S_{k+1} - S_k$. This can be done directly, but it can also easily be derived from the Bapat-Beg theorem, of which the following is a special case (see the appendix for a proof).

Lemma 9.6. *Let X_1, \dots, X_n be points drawn independently uniformly at random in $[0, 1]$ and denote by S'_k the k -th order statistics. With $S'_0 := 0$ and $S'_{n+1} := 1$, for each $k = 0, \dots, n$, the distribution of $S'_{k+1} - S'_k$ is as follows: $\mathbf{P}[S'_{k+1} - S'_k > a] = (1 - a)^n$. \square*

For the discrete version we obtain the following.

Lemma 9.7. *For every $a > 0$, we have*

$$\mathbf{P}[S_{k+1} - S_k > \frac{aN}{n}] \leq e^{-a+O(n/N)},$$

(with an absolute constant in the $O(\cdot)$).

PROOF. Let X_1, \dots, X_n be drawn independently uniformly at random from $[0, 1]$. We can assume that the Y s are the X s multiplied by N and then rounded up: $Y_j = \lceil NX_j \rceil$. We also assume that the permutation taking the X s to the S' s is equal to the permutation taking the Y s to the S s (this condition makes sense when two Y s coincide). By Lemma 9.6, we conclude that

$$\begin{aligned} \mathbf{P}[S_{k+1} - S_k > \frac{aN}{n}] &\leq \mathbf{P}[S'_{k+1} - S'_k > (\frac{aN}{n} - 2)/N] \\ &= (1 - (a/n - 2/N))^n \leq e^{-a+2n/N}. \end{aligned} \quad \square$$

From this, we conclude the following.

Lemma 9.8. *Let an n -subset R be drawn uniformly at random from all the n -subsets of $[N]$, and $a > 0$. The probability that there are $\lceil aN/n \rceil$ consecutive numbers not in R is at most $(n+1)e^{-a+O(n/N)}$.*

PROOF. Let $b := \lceil aN/n \rceil$, and let Y_1, \dots, Y_n be drawn independently uniformly at random from $[N]$. Let A be the event that the Y_j 's are all distinct, \bar{A} its complement, and let B be the event that there are b consecutive numbers not containing any of the Y_j 's. Since $\mathbf{P}(B)$ is a convex combination of $\mathbf{P}(B|A)$ and $\mathbf{P}(B|\bar{A})$, and $\mathbf{P}(B) \leq (n+1)e^{-a+O(n/N)}$ by Lemma 9.7, this upper bound must also be true for the smaller of the two conditional probabilities. But, clearly $\mathbf{P}(B|A) \leq \mathbf{P}(B|\bar{A})$. \square

We can now prove the upper bound on the sizes of the chunks.

Lemma 9.9. *Let $\omega \xrightarrow{h} \infty$ arbitrarily slowly. If n is the number of red vertices in \tilde{G}_C at the completion of Phase II of the algorithm, a.a.s. as $h \rightarrow \infty$, there is no chunk with size larger than $6(\omega + \log n)h/n$.*

PROOF. Choose an arbitrary $x \in C$. By Lemma 9.4, the conditions of Lemma 9.8 are satisfied if we let $n := r_x^+(T)$ and $N := |V_x \setminus R_x(0)|$. The vertices in $V_x \setminus R_x(0)$ are numbered in the following way.

For each cycle of \tilde{G}_C , choose an orientation. The numbers associated to the vertices in the intersection of $V_x \setminus R_x(0)$ and this cycle are then taken consecutively: starting with the vertex in $V_x \setminus R_x(0)$ which, in positive orientation, is next to the $R(0)$ -vertex of the cycle, and continuing to number in positive orientation.

If there is a path in \tilde{G}_C of length greater than $6(\omega + \log n)h/n$ not containing a red vertex, then there is a gap in $[N]$ larger than $(\omega + \log n)N/n$. (Notice that every third vertex of the path belongs to V_x . The factor 2 comes from the left and right end strips, i.e., the vertices which are close to the $R(0)$ -vertex on a cycle but which do not have consecutive numbers.) By Lemma 9.8, the probability of this happening is at most

$$(n+1)e^{-\omega - \log n + O(n/N)} = \frac{n+1}{n}e^{-\omega + O(1)} = o(1).$$

□

Bounding the expected number of cycles in H . We now come to the classical first-moment argument which shows that, a.a.s., our random multi-graph H has no cycles. For the remainder of this section, we condition on the event that the algorithm does not fail before Step (5), and that no two adjacent vertices have been coloured red (cf. Lemmas 9.3 and 9.5 respectively).

Lemma 9.10. *The probability that the edge set of H contains a fixed set F of edges with $|F| = \ell$ is at most*

$$O\left(\ell! \binom{m}{\ell} \frac{\log^{2\ell} n}{n^{2\ell}}\right).$$

PROOF. Recall that n denotes the number of vertices of H , which is equal to the number of chunks in \tilde{G}_C . This is equal to the number of red vertices at the end of Phase II, which is $\Theta(h^{1/3})$. The number m of edges of H is equal to the number $p(T)$ of pale vertices after termination of Phase II, which is $O(\log^2 h)$. The edges come in six different types, depending on which fibre V_y , $y \in S$, contains the corresponding pale vertex, and also which fibres contain the end-vertices of the two non-exposed edges adjacent to the pale vertex.

For each edge of H , one by one, we draw the two end-vertices one by one. An edge corresponding to a pale vertex v of \tilde{G} connects two fixed vertices of H if the two yet unexposed edges incident to v end turn out to be contained in the chunks corresponding to the fixed vertices of H . Since the sizes of the chunks are a.a.s. $O(h \log^2 n/n)$ by Lemma 9.9, and the number of possible neighbors of v is between h and $h - n - m + O(1) = \Theta(h)$, the probability that the edge of H connects the two fixed vertices is $O(\frac{\log^2 n}{n^2})$.

From this, the statement of the lemma follows. □

Now we adapt the classical first-moment calculation to prove that there are no cycles in H , and therefore, no cycles in the graph induced on uncoloured vertices in Step (7).

Lemma 9.11. *A.a.s. H contains no cycles.*

PROOF. By Lemma 9.10, the expected number of cycles of length $\ell \geq 1$ is

$$\sum_{\substack{C \text{ cycle} \\ |C|=\ell}} \mathbf{P}[C \subseteq H] = O\left(\binom{n}{\ell} \ell! \binom{m}{\ell} \frac{\log^{2\ell} n}{n^{2\ell}}\right).$$

Summing over all possible values of ℓ , we obtain an upper bound for the expected number of cycles in H : With $t := (C \log^2 n)/n$ for a suitable constant C , we have

$$\begin{aligned} \sum_{\ell=1}^m \binom{n}{\ell} \ell! \binom{m}{\ell} \frac{\log^{2\ell} n}{n^{2\ell}} \\ \leq \sum_{\ell=1}^m \binom{m}{\ell} t^\ell = -1 + (1+t)^m \leq -1 + e^{mt} = \\ = -1 + e^{(C \log^4 n)/n} = o(1). \end{aligned}$$

□

9.5. Conclusions

The argument for 3-colourability of random lifts of $K_5 \setminus e$ in this manuscript can be extended to a more general class of base graphs. Let $G := G_{k,s}$ be a graph obtained by joining a stable set S of size s to a cycle C of size k , where $k \geq 3$ and $s \geq 1$. For $k = 3$ and $s = 2$ we recover $K_5 \setminus e$. The proof of Theorem 9.1 extends with hardly any changes to the following.

Theorem 9.12. *The chromatic number of a random lift of $G_{k,s}$ is a.a.s. three.*

It is known that the chromatic number of random 4-regular graphs (with uniform distribution) is three [SW07]. Even though random lifts of K_{d+1} have some similarity to random d -regular graphs, adapting the methods of the latter to obtain results for random lifts of K_{d+1} appears to be a challenging task.

Appendix: Distribution of the gaps between n points drawn in $[0, 1]$

As mentioned above, Lemma 9.6 is a special case of the Bapat-Beg theorem. For the sake of completeness, we give an elementary proof.

PROOF OF LEMMA 9.6. Clearly, $\min(X_1, \dots, X_n)$ has cumulative distribution function $t \mapsto 1 - (1-t)^n$. This settles the easy cases when $k = 0$ or $k = n$.

Partitioning $\otimes_{j=1}^n [0, 1]$ into $n!$ sets we need to compute

$$(51) \quad \mathbf{P}[S'_{k+1} - S'_k \leq a] = n! \int_{\mathbb{R}^n} 1_{\{0 \leq \text{pr}_1 \leq \dots \leq \text{pr}_n \leq 1\}} 1_{\{\text{pr}_k \leq \text{pr}_{k+1} \leq \text{pr}_k + a\}} d\lambda^n.$$

Denoting

$$v(\ell, t) := \int_{\mathbb{R}^\ell} 1_{\{0 \leq \text{pr}_1 \leq \dots \leq \text{pr}_\ell \leq t\}} d\lambda^\ell = \frac{t^\ell}{\ell!}$$

we have that (51) is equal to

$$\begin{aligned} (52) \quad \int_0^1 \int_0^1 v(s, k-1) v(1-t, n-k-1) 1_{s \leq t \leq s+a} dt ds = \\ = \int_0^1 v(s, k-1) \int_s^{\min(1, s+a)} v(1-t, n-k-1) dt ds = \\ = \frac{1}{(k-1)!(n-k-1)!} \int_0^1 s^{k-1} \int_s^{\min(1, s+a)} (1-t)^{n-k-1} dt ds \end{aligned}$$

We evaluate the inner integral

$$\begin{aligned} \int_s^{\min(1,s+a)} (1-t)^{n-k-1} dt &= \\ &= \int_s^{\min(1,s+a)} (1-t)^{n-k-1} dt = \\ &= \begin{cases} \frac{1}{n-k} (1-s)^{n-k} & \text{if } s \leq 1-a \\ \frac{1}{n-k} (1-s)^{n-k} - \frac{1}{n-k} (1-a-s)^{n-k} & \text{if } s \geq 1-a. \end{cases} \end{aligned}$$

Then the integral in (52) (without the factorial factor) becomes

$$\begin{aligned} \frac{1}{n-k} \int_0^1 s^{k-1} (1-s)^{n-k} ds - \frac{1}{n-k} \int_0^{1-a} s^{k-1} (1-a-s)^{n-k} ds &= \\ = -\frac{(k-1)!(n-k-1)!}{n!} (0-1) + \frac{(k-1)!(n-k-1)!}{n!} (0 - (1-a)^n) &= \\ = \frac{(k-1)!(n-k-1)!}{n!} (1 - (1-a)^n). \end{aligned}$$

Hence, (51) is equal to

$$n! \frac{1}{(k-1)!(n-k-1)!} \frac{(k-1)!(n-k-1)!}{n!} (1 - (1-a)^n) = 1 - (1-a)^n.$$

□

Random 3-SAT with interval constraints

Jointly with
Kathrin Ballerstein (ETHZ)

Abstract. In signed k -SAT problems, one fixes a set M and a set \mathcal{S} of subsets of M , and is given a formula consisting of a conjunction of m clauses, each of which is a disjunction of k literals. Each literal is of the form “ $\mathbf{x} \in S$ ”, where $S \in \mathcal{S}$, and \mathbf{x} is one of n variables.

For Interval-SAT (iSAT), M is an ordered set and \mathcal{S} the set of intervals in M .

We propose an algorithm for 3-iSAT, and analyze it on uniformly random formulas. The algorithm follows the Unit Clause paradigm, enhanced by a (very limited) backtracking option. Using Wormald’s ODE method, we prove that, if $m/n \leq 2.3$, with high probability, our algorithm succeeds in finding an assignment of values to the variables satisfying the formula.

10.1. Introduction

Let M be a (usually finite) set, \mathcal{S} a set of subsets of M , and X a set of variables. A (*signed*) *literal* is the pair $(\mathbf{x}, S) \in X \times \mathcal{S}$, which we will denote as $\mathbf{x} \in S$, and for a positive integer k , a *k -clause* (or simply *clause*) is the disjunction (\vee) of at most k literals. The conjunction (\wedge) of finitely many k -clauses is called the *signed k conjunctive normal form (k -CNF)*. In this setting the central question is the *signed k -satisfiability problem*, or *signed k -SAT*, which asks for a satisfying *interpretation*, that is, an assignment of values to the variables such that in each clause there is at least one literal (\mathbf{x}, S) for which \mathbf{x} takes a value in S .

This setting includes as a special case the classical satisfiability (SAT) problem. There, one chooses for M the 2-element set $\{\text{TRUE}, \text{FALSE}\}$ and $\mathcal{S} = \{\{\text{TRUE}\}, \{\text{FALSE}\}\}$. In case M is an ordered set (a chain) and the set \mathcal{S} is the set of all intervals in M , we speak of *Interval SAT*, or *iSAT*. In our contribution, we set $M := [0, 1]$, because this includes all iSAT settings with finite M . In particular, we consider formulas of the type

$$\bigwedge_{i=1}^t \bigvee_{j \in \mathcal{J}_i} \mathbf{x}_j \in \mathbb{I}_j^i,$$

where, for all $i = 1, \dots, t$, \mathcal{J}_i , with $|\mathcal{J}_i| \leq 3$, is an index set of variables in X , and $\mathbb{I}_j^i \subseteq [0, 1]$ are intervals for all i and j . Then, an interpretation of a clause i is satisfying if there is a variable \mathbf{x}_j taking a value in the interval \mathbb{I}_j^i . Identifying a satisfying interpretation of the complete 3-CNF is related to the study of random interval graphs [Sch88, JSW90]. Our notation and terminology on signed SAT follows [CCHS10].

Signed SAT problems originated in the area of so-called multi-valued logic [Lu20], where variables can take a (usually finite) number of so-called *truth values*, not just TRUE or FALSE. Work on signed CNF formulas started in earnest with the work of Hähnle and Manyà and their coauthors. We refer the reader to the survey paper [BHM00b], and the references therein.

The motivation for studying signed formulas was to extend algorithmic techniques developed for deductive systems in multi-valued logic to better cover practical applications [H91]. Indeed, on the one hand, a number of papers show how combinatorial problems can be solved

using signed SAT algorithms [BM99b, BCF⁺01, FP01, BMC⁺07]; on the other hand, a large number of heuristic and exact algorithms have been studied (see [AM03, Bri04] and the references therein), and a number of polynomially solvable subclasses of signed SAT have been identified [EIM94, BHM00b, Man00, BHM00a, ABCM04, AM03, CCHS10]. While in the works of Manyà and his collaborators, order-theoretic properties of the ground set M are exploited to make conclusions on the complexity of signed SAT, Chepoi et al. [CCHS10] completely settle the complexity question in the general case by reverting to combinatorial properties of the set system \mathcal{S} . In particular, they prove that: signed k -SAT, $k \geq 3$, is polynomial, if $\bigcap_{S \in \mathcal{S}} S \neq \emptyset$ and NP-complete otherwise; signed 2-SAT is polynomial if, and only if, \mathcal{S} has the Helly property (if no two sets in a subfamily are disjoint, then the subfamily has non-empty intersection), and NP-complete otherwise.

For the case when \mathcal{S} has the Helly property, Chepoi et al. give a non-satisfiability certificate for signed 2-SAT in the spirit of Aspvall, Plass, and Tarjan’s famous result for classical 2-SAT [APT79].

Most applications and a great deal of the earlier complexity results [BHM00b] focus on *regular* signed SAT, where M is a poset, and the formulas may only involve sets of the form $S = \{j \mid j \geq i\}$ or $S = \{j \mid j \leq i\}$. Regular iSAT (or just regular SAT) is regular signed SAT for posets M which are chains.

In particular, for regular iSAT, random formulas have been investigated from a heuristic point of view. Manyà et al. [MBE198] study uniformly generated random regular 3-iSAT instances, and observe a phase transition similar to that observed in classical SAT (see [AP04] and the references therein): (i) the most computationally difficult instances tend to be found near the threshold, (ii) there is a sharp transition from satisfiable to unsatisfiable instances at the threshold and (iii) the value of the threshold increases as the number of truth values considered increases. Their results are confirmed and extended by further papers exploring uniformly random regular 3-iSAT instances [BM99a, BHM00b, BMC⁺07].

Further, in [BM99a, BMC⁺07] a bound on the ratio m/n is given, beyond which a random formula is with high probability (whp) unsatisfiable. To our knowledge, however, ours is the first rigorous analysis of an algorithm for random signed SAT.

Our interest in the particular version of signed SAT arises from applications in computational systems biology, where iSAT yields a generalization of modeling with Boolean networks [Kau69], where biological systems are represented by logical formulas with variables corresponding to biological components like proteins. Reactions are modeled as logical conditions which have to hold simultaneously, and then transferred into CNF. The model is widely used by practitioners (see e.g. [Dow01, KSRL⁺06, HNTW09] and the references therein). Often, though, this binary approach is not sufficient to model real life behavior or even accommodate all known data. Due to new measurement techniques, a typical situation is that an experiment yields several “activation levels” of a component. Thus, one wants to make statements of the form: If the quantity of component A reaches a certain threshold but does not exceed another, and component B occurs in sufficient quantity, then another component C is in a certain frame of activation levels. The collection of such rules accurately models the global behavior of the system. We refer to [Bal12] for details of models and applications.

In this paper we present and analyze an algorithm which solves uniformly random 3-iSAT instances with high probability, provided that the ratio between the number m of clauses and the number n of variables is at most 2.3. Our algorithm is an adaption of the well-known Unit Clause algorithm from classical SAT [CF86, Ach01], where, in an inner loop, 1-clauses are treated if any exist, and in an outer loop, a variable is chosen freely and assigned some value. This Unit Clause approach is enhanced with a “repair” subroutine (a very simple backtracking mechanism). The algorithm in [FS96] is currently the best known algorithm that succeeds

with high probability, although other algorithms (e.g., [KKL06, HS03]) can be outfitted with a backtracking routine to provide better results. See also [CO10] for general $k \rightarrow \infty$.

Unlike the algorithms in [Ach00, AS00, KKL06], we prove that our algorithm succeeds with high probability. To obtain a whp result, the “repair” subroutine is essential, cf. e.g., [FS96], where the range in which the algorithm succeeds increases dramatically, once such a routine part is added. In the case of iSAT, the repair mechanism needs to be considerably more subtle than the one in [FS96] for classical 3-SAT.

In the analysis of the algorithm, we use Wormald’s differential equations method [Wor95]. ODE methods have been used for the analysis of algorithms for classical SAT with great success [CF86, CF90, FS96, Ach00, AS00]. In our analysis, we combine the idea of Achlioptas and Sorkin [AS00] to consider as a time step an iteration of the outer loop, but we use Wormald’s theorem [Wor99] where they use a Markov-chain based approach. The analysis of the inner loop requires to study the first busy period of a certain stable server system [Ach00, Ach01], or, in our case, more accurately, the total population size in a type of branching process. The value 2.3 arises from the numerical solution to an initial value problem (IVP). Extending the results for k -iSAT for $k \geq 4$ is conceptually easy; we briefly discuss it in the conclusions.

The outline of the paper is as follows: In the next section, we present our algorithm for random 3-iSAT in detail. In Section 10.3, we prove some facts about uniformly at random chosen sub-intervals of $[0, 1]$. In Section 10.4 we take a brief excursion to random 2-iSAT as our algorithm for 3-iSAT ultimately relies on solving a 2-iSAT instance. In Section 10.5, we compile the required facts about total population sizes of a kind of branching system, which are then applied in Section 10.6 to the study of the inner loop of our algorithm. Finally, in Section 10.7, we prove the whp result for our algorithm. We raise some issues for future research in the final section. Several technical arguments have been moved into the appendix.

Throughout the paper, we hide absolute constants in the big- O -notation. If the constant depends on other parameters, we make this clear by adding an index, e.g., $O_\varepsilon(\cdot)$. As customary, we use the abbreviation iid for “independent and identically distributed” and uar for “uniformly at random”. Whp and wpp are to be understood for $n \rightarrow \infty$, with $m = m(n)$ depending on n .

10.2. An algorithm for random 3-iSAT

In this section, we describe an algorithm which finds a satisfying interpretation if the number of clauses is $m = cn$ with $c \leq 2.3$.

10.2.1. The random model; exposure. For our random model, we assume that each 3-clause consists of three distinct variables. We choose a formula uar from the set of all possible classical 3-CNF formulas on n variables with m 3-clauses, each containing three distinct variables. Then, we choose an interval for each literal uar from the subintervals of $[0, 1]$: We select uar two points x and y from $[0, 1]$ and determine the interval as $[a, b]$ with $a = \min\{x, y\}$ and $b = \max\{x, y\}$. In this context, note that due to Scheinerman [Sch88] the endpoints x and y can be arbitrary reals. In fact, he proves that this strategy is equivalent to choosing $2l$ endpoints for l intervals uar from the finite set $\{1, \dots, 2l\}$ without repetition as the probability that all chosen endpoints from $[0, 1]$ are distinct is 1. For the distribution of a random interval $[a, b]$ chosen as $a = \min\{x, y\}$ and $b = \max\{x, y\}$ for $x, y \in [0, 1]$ uar, we find with $u, v \in [0, 1]$

$$\mathbf{P}([a, b] \subseteq [u, v]) = 2 \cdot \mathbf{P}(a \geq u, b \leq v) = 2 \cdot (1 - u) \cdot v.$$

As is customary in the context of random SAT, we use the language of “exposing” literals. Intuitively, the idea is that the information about each literal is written on a card which lies face down, until the information is exposed. Clearly, the unexposed part of the formula is uar conditioned on which literals have been exposed and which have not. We refer to the elegant description in Achlioptas’ paper [Ach01].

10.2.2. Brief description of the algorithm. The basic framework of our algorithm is the same as for most algorithms for classical k -SAT. A formerly unused variable is selected, and a value is assigned to it. Then, clauses containing the variable are updated: if the literal of the clause involving the variable is satisfied, the clause is deleted; otherwise the literal is deleted from the clause, leaving a shorter clause. The variable is removed from the set of *unused variables*, and declared a *used variable*. The algorithm fails if, and only if, it creates an empty clause.

However, to a certain extent, our algorithm is able to repair bad choices it has made. Thus, it occasionally only assigns *tentative* values to variables. As long as it is not certain that a variable keeps its tentative value, no deletions of clauses or literals from clauses are performed. Instead, we assign colors to the clauses, which code the number of satisfied, unsatisfied, and unexposed literals they contain. The meaning of the colors will be explained in Table 1 but at this point it suffices to know that red clauses correspond to unexposed 1-clauses, i.e., clauses with one unexposed literal and the variables in any other literal of the clause have tentative values which render the literals false.

As said before, the basic approach is that of the Unit-Clause algorithm. The *outer loop* of the algorithm will maintain the property that there is no 1-clause. In each iteration of the outer loop, a variable is selected from the set of unused variables. Such a variable selected in the outer loop is referred to as a *free variable*. The *inner loop* is initialized by assigning a tentative value to this free variable, and then repeats as long as there are red clauses. In each iteration of the inner loop, a red clause is selected and *serviced*: the variable contained in the clause (the *current variable* of the iteration) is tentatively set to some value in such a manner that the serviced red clause becomes true. We refer to the variables selected in the inner loop as *constrained variables*.

If, during a run of the inner loop, a situation is reached in which it is probable that an empty clause will be created, it backtracks. This happens when the following *fatality* is suffered: The current variable occurs in another red clause, other than the one serviced. If that happens, there is a $1/3$ probability that the two intervals occurring in the two red clauses are disjoint [Sch88], so that creating an empty clause is inevitable.

For this situation, the inner loop maintains a rooted tree G of decisions it has taken so far. The nodes of the tree correspond to variables to which tentative values have been assigned and those which occur in the unexposed part of red or blue 2-clauses. The root of the tree is the free variable with which the run of the inner loop was initialized. The edges correspond to 2-clauses. For every 2-clause in which the current variable of an iteration occurs, the unexposed variable is added as a node and an edge is added connecting the current variable with this new node. Doing so in every iteration constructs a tree. If the current variable of an iteration occurs in two red clauses, then this implies that a cycle is closed in G , because there must exist two paths from x_0 to the current variable. The tree G is in detail defined in the algorithm. If a fatality occurs, the values of the variables along the paths from the root to the serviced literal are changed so that all 2-clauses along the path are fulfilled and only one red clause remains which is satisfied. Then, all other tentative values are made permanent, and the inner loop is restarted with the new formula, but this time without a free variable in the initialization. We call *Phase I* the run of the inner loop before a repair occurs (or if no repair occurs), and as *Phase II* to the run of the inner loop after a repair has been performed. In Phase II, no further repair is attempted. Instead, if fatalities occur, the inner loop just moves on (without repair). In Phase I, if a fatality occurs, there's the possibility that a repair is not possible. In this case, too, the inner loop just moves on without repair. In order to be able to refer to these situations in the proofs, we indicate these positions in the code by the pseudo-command "raise a flag".

After all red clauses have been dealt with in either Phase I or Phase II, the tentative values are made permanent, and control is returned to the outer loop, which selects another free variable, and so on.

The outer loop terminates, if the number of 2-clauses plus the number of 3-clauses drops below a certain factor c' of the number of unused variables. Then, it deletes an arbitrary literal from every 3-clause and invokes the exact polynomial algorithm by Chepoi et al. [CCHS10] to decide whether the resulting 2-iSAT formula has a satisfying interpretation. We will prove in Section 10.4 that this is always the case if the ratio of the number of resulting 2-clauses over the number of unused variables is below $\frac{3}{2}$.

The complete algorithm is shown below as Algorithm 3 (the outer loop), Algorithm 4 (the inner loop), and Algorithm 5 (the repair procedure). Throughout the course of the algorithm, for $i = 0, 1, 2, 3$, we denote by $Y_i(t)$ the number of i -clauses, and by $X(t)$ the number of unused variables, respectively, at the beginning of iteration t of the outer loop. Moreover, for an interval I , we denote by

$$(53) \quad \bar{x}(I) := \operatorname{argmin}_{x \in I} |x - 1/2|$$

the point in I which is closest to $1/2$. We refer to the variable x_j which is selected in iteration j of the inner loop as the *current variable* of that iteration.

Below, we will prove the following fact.

Lemma 10.1. *A single run of Algorithm 4 (including a possible repair and consequent Phase II) produces an empty clause, only if it “raises a flag”.*

The performance of the algorithm on random 3-iSAT instances is analyzed in Sections 10.6 and 10.7. There, we will prove the following theorem.

Theorem 10.2. *Let $c := 2.3$, and suppose Algorithm 3 is applied to a uniformly random iSAT formula on n variables with m 3-clauses. If $m \leq cn$, then, whp, Algorithm 3 creates no empty clause, i.e., it finds a satisfying interpretation.*

The value 2.3 is determined through the numerical solution of an initial value problem. It corresponds to the point in which the increase in red clauses in each iteration of the inner loop would become so large that the inner loop will not terminate.

Algorithm 3 *UC w/ backtracking (outer loop)*

(o-1) Given: 3-CNF-formula; positive constant c' .

(o-2) $t := 0$

(o-3) While $Y_2(t) + Y_3(t) > c'X(t)$:

(o-3.1) Choose a variable x uar.

(o-3.2) Invoke *Inner loop* (Phase I).

(o-3.3) $t := t + 1$

(o-4) In every 3-clause, remove one literal at random.

(o-5) Invoke Chepoi et al.’s algorithm (cf. Section 10.4) for the remaining 2-iSAT formula.

10.2.3. Comparison to algorithms for classical SAT. For classical SAT, if a variable x is set to a value, the probability that a random literal containing x evaluates to true is $1/2$ — independent of the value. As will become apparent in the next section, this is far from true for random interval literals. There, the value $1/2$ is the single, most likely value to be contained in a random interval (the probability is $1/2$) and all other values are less likely. Hence, we will assign $1/2$ to the variables as long as possible which is for all free variables.

The rationale behind assigning the value $1/2$ to free variables is two-fold. Firstly, it makes the analysis a lot more easy than if one tries to find a maximum cardinality subset of literals containing x all of whose intervals have pairwise non-empty intersection. Secondly, for large numbers of literals containing x , the maximum cardinality of a subset with pairwise intersecting intervals is asymptotically attained by taking all literals with intervals containing $1/2$ (this is

Color	Meaning
Uncolored	All literals in the clause are unexposed.
Black	All literals are exposed.
Red	The clause has precisely one unexposed literal. The tentative values of any other variables in the clause make the corresponding literals false. In particular, unexposed 1-clauses are red.
Blue	The clause contains precisely one unexposed literal and at least one exposed literal which evaluates to true for the tentative value of its variable.
Pink	The clause is a 3-clause, precisely one of its literals is exposed, and this literal evaluates to false for the tentative value of its variable.
Turquoise	The clause is a 3-clause, precisely one of its literals is exposed, and this literal evaluates to true for the tentative value of its variable.

TABLE 1. Semantics of the colors of the clauses.

Theorem 4.7 of Scheinerman’s paper [Sch88]). This, in particular, implies that assigning an interval of values to a variable does asymptotically not lead to a satisfying interpretation of the formula which is not satisfying if assigning the single value $1/2$.

The situation for constrained variables is similar, but a bit more complicated. For constrained variables, we are free only to choose the value for the variable within the interval \mathbb{I} for the literal $L = x \in \mathbb{I}$ which we wish to satisfy. Unlike to classical SAT, where this does not change the probability that other random literals containing x are satisfied, depending on \mathbb{I} , this probability may change considerably. Moreover, for two literals containing x , the two events of both being satisfied simultaneously with L are not independent.

However, an adaption of Scheinerman’s argument mentioned above shows that, asymptotically, the best choice is to take the point \mathbb{I} which is closest to $1/2$ as we do in our algorithm.

Concerning the backtracking part of the algorithm, we would like to point out the difference to the approach in [FS96]. If the (essentially identical) fatality is suffered, a very elegant remedy is to simply flip the values of all variables with tentative values: if the tentative value of a variable is TRUE, make it FALSE, and vice versa. Needless to say, for variable values in a larger set, there is no obvious choice for the new value of a variable. Thus, in our approach, we have to choose the variable values in a smart manner, with the single aim to undo the fatality. Namely, those variables that led to the fatality are assigned $\bar{x}(I)$ as described in *Repair Path* (Algorithm 5).

10.2.4. Proof of the “raise a flag”-lemma.

PROOF OF LEMMA 10.1. Assume that Algorithm 4 does not “raise a flag”.

The only place where a 0-clause can be generated without having “raised a flag” is in the final step 4 of the repair, Algorithm 5. Clearly, none of the clauses on the path will become empty.

Moreover, setting the final variable, x_k , cannot create an empty clause, because of the conditions in steps (i-9.1) and (i-9.2).

For a 3-clause to become empty, it is necessary that when the repair is invoked in Algorithm 4, all three of its literals have been exposed (possibly in the same iteration). In other words, it must have been red, blue, or black in step (i-9.1), a contradiction.

For a 2-clause to become empty, both literals must have been exposed, one of them possibly in the iteration where the repair occurs. Moreover, if it was blue, the value of the variable satisfying one of its literals must change during the repair. In other words, the following three scenarios are possible:

Algorithm 4 *Inner loop*

-
- (i-1) Given:
- In **Phase I**: formula consisting of 2- and 3-clauses only; a (free) variable x_0 .
 - In **Phase II**: formula consisting of 1-, 2- and 3-clauses.
- (i-2) $j := 0$
- (i-3) Initialize: Expose the occurrences of x_0 in all clauses.
- In **Phase I** only:
 - (i-3.1) Tentatively set x_0 to $1/2$.
 - (i-3.2) Initialize the graph $G := (\{x_0\}, \emptyset)$.
 - In **Phase II** only:
 - (i-3.1) Color all 1-clauses red.
- (i-4) Expose the intervals associated with x_0 . Color clauses containing x_0 according to Tab. 1.
- (i-5) $j := j + 1$
- (i-6) If there is no red clause, exit inner loop: Set all variables to their tentative values; remove satisfied clauses and remove violated literals from their clauses; return to outer loop.
- (i-7) Select a red clause C_j at random; let L_j be the unexposed literal in C_j ; expose current variable x_j of L_j
- (i-8) Expose all occurrences of x_j in colored clauses.
- (i-9) If x_j is contained in a red clause other than C_j :
- In **Phase I** only:
 - (i-9.1) If there is a red, blue, or black 3-clause: “raise a flag”!
 - (i-9.2) If the graph G contains a cycle, or x_j is in a blue clause: “raise a flag”!
 - (i-9.3) If x_j occurs in three or more red clauses (including C_j): “raise a flag”!
 - (i-9.4) Otherwise: **Phase I** is completed. Let C' be the unique red clause different from C_j containing x_j in a literal $L' = x_j \in J'$. Repair the unique path between x_0 and C_j ; then initiate **Phase II**.
 - In **Phase II** only: “raise a flag”!
- (i-10) Expose all occurrences of x_j in all uncolored clauses.
- (i-11) For every uncolored 2-clause $x_j \in I \vee y \in J$ containing x_j , add to G the vertex y and the edge $x_j \in I \vee y \in J$ between x_j and y .
- (i-12) Tentatively set x_j to $\bar{x}(I_j)$.
- (i-13) Update the colors of all clauses containing x_j .
- (i-14) Goto step (i-5).
-

Algorithm 5 *Repair path*

-
- (r-1) Given: Set of colored 1-, 2- and 3-clauses; a literal $L' = x_k \in J'$; a path of the form
- $$x_0, \quad x_0 \in J_0 \vee x_1 \in I_1, \quad x_1 \in J_1 \vee x_2 \in I_2, \quad \dots, \quad x_{k-1} \in J_{k-1} \vee x_k \in I_k;$$
- (r-2) For $j = 0, \dots, k - 1$:
- (r-2.1) Set x_j (permanently) to $\bar{x}(J_j)$
- (r-3) Set x_k (permanently) to $\bar{x}(J')$
- (r-4) Set all variables from Phase I, except those which have just been set in (r-2) and (r-3), to their tentative values; remove satisfied clauses and remove violated literals from their clauses.
-

- (i) it was black before the repair was invoked
- (ii) it was red before the repair was invoked, but it contains x_j
- (iii) it was blue before the repair was invoked, it is of the form $x_i \in I_i \vee x_j \in I_j$ for some $i < j$, and x_i is one of the variables set in step (r-2) of Algorithm 5.

In case (i), if the black 2-clause becomes an empty clause, either it was red when its final literal was exposed, a contradiction, or it was blue, which means that at least one of its variables lies on the path which is repaired. If the whole clause lies on the path, we have already noted that it cannot become empty. If only one of its variables is on the path, then it must be an edge in the tree having one end vertex on the path and the other lying further away from the root than the path. The fact that it is black means that the variable which is not on the path was the current variable of some earlier iteration $i < j$. But then the corresponding literal was either the selected literal L_i , in which case it was satisfied by the tentative value of x_i , or the if-condition in step (i-9) for iteration i held, which is a contradiction (either a repair occurred, or the algorithm has “raised a flag”).

In case (ii), if the 2-clause is on the path, it does not become empty. If it is the unique other red clause C' , then it will be satisfied in the initialization of Phase II.

Case (iii), is not possible because of the condition in step (i-9.2) □

10.2.5. Random formulas. The following easy facts (see the discussion at the beginning of this section) underlies the analysis of the algorithm on random formulas.

Lemma 10.3. *If Algorithm 3 is invoked with a uar random 3-iSAT formula, then*

- (a) *at the beginning of each iteration of the outer loop, the current formula is distributed uar conditioned on the number of unused variables, 2-clauses, and 3-clauses;*
- (b) *at the beginning of each iteration of the inner loop, the current formula is distributed uar conditioned on the number of unused variables, 1-clauses, 2-clauses, 3-clauses, and the colors of the clauses.*
- (c) *at the beginning of Phase II in the inner loop, the current formula is distributed uar not only conditioned on the number of unused variables, 1-clauses, 2-clauses, 3-clauses, the colors of the clauses, and the list L of clauses which are known not to contain x_0 and the list of clauses in which an occurrence of x_0 has been exposed.*

By Lemma 10.3, the history of the random process defined by the outer loop, that is, for each t , the state of the formula and all other information relevant to how the algorithm will proceed, available at the beginning of iteration t , is completely determined by

$$(54) \quad \mathcal{H}(t) := (X(t), Y_2(t), Y_3(t));$$

in particular it is Markov.

10.3. Computations for random intervals

In this section, we make some computations regarding intervals chosen uar from the subintervals of $[0, 1]$ as described before. We refer to [Sch88, JSW90] for further background.

We aim to study the event $\bar{x}(I) \in J$, with two random intervals I and J (\bar{x} is defined in (53)). We start with the following observation.

Lemma 10.4 ([Sch88]). *For $x \in [0, 1]$ and for a random interval I , we have*

$$\mathbf{P}[x \in I] = 2x(1 - x).$$

In particular, the probability that a random interval contains the point $1/2$ is $1/2$.

The cumulative distribution function of $\bar{x}(I)$ can be written down.

Lemma 10.5. *For a random interval I , the random variable $\bar{x}(I)$ has cumulative distribution function*

$$(55) \quad F(t) := \begin{cases} 0 & \text{if } t \leq 0 \\ t^2, & \text{if } t < 1/2 \\ 1 - (1 - t)^2, & \text{if } t \geq 1/2 \\ 1 & \text{if } t \geq 1. \end{cases}$$

PROOF. Direct computation. \square

Let X be a random variable with cumulative distribution function F as in (55), and define

$$(56) \quad P := 1 - 2X(1 - X).$$

Thus, by the previous two lemmas, the probability that, for two random intervals I and J we have $\bar{x}(I) \in J$, is

$$\mathbf{E}(\mathbf{P}[\bar{x}(I) \in J \mid P]) = \mathbf{E}(1 - P) = 1 - \mathbf{E} P.$$

The following computations are straightforward, see 10.8.1.

Lemma 10.6.

$$(a) \quad \mathbf{E} P = 13/24$$

$$(b) \quad \mathbf{E} P^2 = 3/10 \quad \square$$

Lemma 10.7. *For two random intervals I, J , the following is true.*

$$\mathbf{P}[\bar{x}(I) \in J] = \frac{11}{24}.$$

PROOF. Immediate from Lemmas 10.4, 10.5, and 10.6(a). \square

Remark 10.8. It could be interesting to choose the intervals in a different way rather than uniformly at random, for instance, to reflect certain realistic structures. However, the strategy of choosing intervals does not change the main analysis of the algorithm. The only adaptations to be made are the previous computations of the probabilities, and thus the new constants need to be used in the analysis, which can lead to different results.

10.4. 2-iSAT

In this section, we take a brief glance at the situation for random 2-iSAT. The reason is that, ultimately, our 3-iSAT algorithm reduces the 3-iSAT formula to one with exactly two literals per clause, and then invokes the polynomial time algorithm by Chepoi et al. [CCHS10] to find a solution. We need to make sure that the resulting random 2-iSAT instance is satisfiable.

For this, we proceed along the same lines as [CR92], using Chepoi et al.'s Aspvall-Plass-Tarjan-type [APT79] certificate for the non-satisfiability of signed 2-SAT formulas for set systems satisfying the Helly-property. We describe the certificate now.

For a 2-iSAT formula F , define a digraph G_F which contains two vertices labeled xIt and xIf , respectively, for every literal $x \in I$ occurring in F . For every clause $x \in I \vee x' \in I'$ of F , the digraph G_F contains two arcs $xIf \rightarrow x'I't$ and $x'I'f \rightarrow xIt$. We refer to these arcs as *clause arcs*. Moreover, for every two literals $x \in I$ and $x \in J$ occurring in F , if $I \cap J = \emptyset$, the digraph G_F contains the two arcs $xIt \rightarrow xJf$ and $xJt \rightarrow xIf$. These arcs we call *disjointness arcs*.

For a literal $x \in I$ occurring in F , we refer to the vertex xIt as a *positive* vertex, and to xIf as a *negative* vertex. Moreover, we say that these two vertices are *complements* of each other; in other words, the complement of the (positive) vertex xIt is the (negative) vertex xIf and vice versa. Note that arcs originating from negative vertices are clause arcs, while arcs originating from positive vertices are disjointness arcs.

Chepoi et al. relate the satisfiability of F to the strongly connected components (SCCs) of G_F .

Proposition 10.9 (Aspvall-Plass-Tarjan-type certificate, [CCHS10]). *The formula F is satisfiable if, and only if, no SCC of G_F contains a pair of vertices which are complements of each other.*

Remark 10.10. A path in G_F of length ℓ contains $\lfloor \ell/2 \rfloor$ or $\lceil \ell/2 \rceil$ disjointness arcs, and no two of them are incident.

Chepoi et al. also give an algorithm which determines, in polynomial time, whether a formula F is satisfiable, and if it is, produces a satisfying interpretation. We refer to their paper for details.

From Proposition 10.9, we obtain the following corollary.

Corollary 10.11. *If F is not satisfiable, then G_F contains a bicycle, i.e., a directed walk*

$$u_0 \rightarrow \cdots \rightarrow u_{\ell+1},$$

with at least one clause arc, and the following properties:

- (a) *the literals in the vertices u_1, \dots, u_ℓ are all distinct;*
- (b) *the literals in the vertices u_0 and $u_{\ell+1}$ occur among the literals in the other vertices;*
- (c) *the clauses in the arcs are all distinct.*

PROOF. For a vertex v , we denote its complement by \bar{v} . By what we said about the different types of arcs, on every path from v to \bar{v} , there is at least one clause arc.

Choose an SCC and take a pair of complementing vertices v and \bar{v} in the SCC such that the distance from v to \bar{v} in G_F is minimal. Then, on the shortest path P from v to \bar{v} , no literal appears twice. Denote by \mathbb{L} the literal defining v and \bar{v} .

Now take a shortest path Q in G_F from \bar{v} to v . If there is no literal other than \mathbb{L} which appears twice on $P \cup Q$, then $P \cup Q$ is a bicycle starting and ending in v . On the other hand, if there is a literal \mathbb{L}' other than \mathbb{L} which appears twice on $P \cup Q$, then the desired bicycle is constructed by taking the path P from v to \bar{v} , and then the path Q until the first vertex whose literal already occurred earlier. \square

Suppose a 2-iSAT formula with n variables and $m = cn$ clauses is drawn uniformly at random from the set of all such formulas (with the intervals all in $[0, 1]$). We estimate the asymptotic probability that such a formula is satisfiable.

Proposition 10.12. *Let $c' < 3/2$. If $m \leq c'n$ then, whp as $n \rightarrow \infty$, a randomly drawn 2-iSAT instance is satisfiable.*

The proof mimics that of Chvátal & Reed [CR92] for the classical 2-SAT very closely; we include it here just to point out where the number $3/2$ comes in.

PROOF. Given a fixed bicycle $u_0 \rightarrow \cdots \rightarrow u_{\ell+1}$ with r clause arcs, the probability that it occurs in G_F is at most

$$\left(\frac{m}{\binom{n}{2}} \right)^r p^{r-1},$$

where $p := 1/3$ is the probability that two independently chosen intervals are disjoint [Sch88]. Hence, the expected number of bicycles with r clause arcs occurring in G_F is at most

$$n^{r-1} (r-1)^2 \left(\frac{m}{\binom{n}{2}} \right)^r p^{r-1} = \frac{2m}{n-1} (r-1)^2 \left(\frac{2pm}{n-1} \right)^{r-1}.$$

Thus, the expected total number of bicycles is at most

$$\frac{2m}{n-1} \sum_{r=1}^{\infty} r^2 \left(\frac{2pm}{n-1} \right)^{r-1}.$$

With $m \leq c'n$ the sum is finite if, and only if, $c' < 3/2$ for $n \rightarrow \infty$. Thus, in this case, the probability that a bicycle exists is $O_{c'}(1)$. \square

Thus, for every $c' < 3/2$, whp, a satisfying interpretation can be found by Chepoi et al.'s algorithm [CCHS10]. We make no attempt at optimizing this bound as we indeed conjecture that this is the threshold for 2-iSAT.

10.5. Total population size of our branching system

As is done in classical SAT, the sub-routine eliminating the unit clauses can be viewed as a “discrete time” queue in which customers (i.e., unit clauses) arrive per time unit, the number depending on the customer currently serviced, and the single server, corresponding to one run of the inner loop of the algorithm, can process at least one customer per time unit. The number of iterations of the sub-routine then roughly corresponds to the length of the (first) busy period of the server.

Here, since, we are only interested in the length of the first busy period, the “queue” is really a branching system, for which we need to know the total number of individuals which are born before extinction. Compared to classical SAT, the interval-version poses several small challenges which we address in this section.

Let a be a non-negative integer, and $B(j)$, $j = 0, 1, 2, \dots$, random variables taking values in the non-negative integers. We say the following sequence of random variables $Q(j)$ a *discrete queue*:

$$\begin{aligned} Q(0) &= 0 \\ Q(1) &= a \\ Q(j+1) &= \begin{cases} a, & \text{if } Q(j) = 0 \\ Q(j) - 1 + B(j+1) & \text{if } Q(j) > 0 \end{cases} \end{aligned}$$

The number $Q(j+1)$ is the number of individuals of the branching system after the j th individual has reproduced and died.

Denote by Z the length of the first busy period of the server, that is, the total population size of the branching process:

$$Z := \sup\{j \geq 0 \mid Q(i) > 0 \quad \forall i = 1, \dots, j\} = \inf\{j > 0 \mid Q(j) = 0\} - 1.$$

A straightforward adaption of the branching-process based textbook arguments for continuous-time M/G/1-queues gives the following (see 10.8.2).

Lemma 10.13. *Suppose the $B(j)$, $j = 1, 2, \dots$, are iid with mean λ_B and common probability generating function g_B . The probability generating function h of Z satisfies*

$$(57a) \quad h\left(\frac{y}{g_B(y)}\right) = y^a$$

for every y for which the power series $g_B(y)$ converges and does not vanish. In particular, if $\lambda_B < 1$, we obtain

$$(57b) \quad \mathbf{E} Z = \frac{a}{1 - \lambda_B}.$$

Moreover, we have

$$(57c) \quad \mathbf{P}[Z \geq \alpha] \leq \frac{g_B(y)^\alpha}{y^{\alpha-a}}$$

for all $\alpha > 0$ and $y > 0$ with $y \geq g_B(y)$. □

Remark 10.14. Since we are only interested in the first busy period, we make the following modification to the definition of Q : If $Q(j) = 0$ but $j > 0$, then we let $Q(j+1) = 0$ (and not $Q(j+1) = a$ as above). This makes some inequalities less cumbersome to write down.

10.5.1. Bounding the tail probability for iid binomial B . Let P be a random variable with values in $[0, 1]$. We say that a random variable B has binomial distribution with random parameter P , or $\text{Bin}(m, P)$, if

$$\mathbf{P}[B = k \mid P = p] = \binom{m}{k} p^k (1-p)^{m-k}.$$

In our setting n is a (large) integer, and $m = m(n)$ is an integer depending on n . Define $\lambda = \lambda(n) := \frac{m}{n}$. Let P be as in (56), and suppose that B is $\text{Bin}(m, 2P/n)$.

Lemma 10.15. *If $\lambda(y-1) \leq 1/2$ we have*

$$g_B(y) \leq \exp\left(\frac{13}{12}\lambda(y-1) + \frac{6}{5}\lambda^2(y-1)^2\right)$$

PROOF. We have $e^t \leq 1 + t + t^2$ for all $t \leq 1$. For ease of notation, let $\tau := \mathbf{E} P = 13/24$ and $\tau_2 := \mathbf{E}(P^2) = 3/10$, by Lemma 10.6. Since $(y-1)\lambda 2P \leq 1$ with probability one, the following estimate holds:

$$\begin{aligned} g_B(y) &= \sum_{k=0}^m \mathbf{E}\left(\binom{m}{k} p^k \left(1 - \frac{2P}{n}\right)^{m-k}\right) = \mathbf{E}\left(\sum_{k=0}^m \binom{m}{k} p^k \left(1 - \frac{2P}{n}\right)^{m-k}\right) \\ &= \mathbf{E}\left(\left(1 + (y-1)\frac{2P}{n}\right)^m\right) \leq \mathbf{E}\left(e^{2(y-1)\lambda P}\right) \leq \mathbf{E}\left(1 + 2(y-1)\lambda P + 4(y-1)^2\lambda^2 P^2\right) \\ &= 1 + 2\tau(y-1)\lambda + 4\tau_2(y-1)^2\lambda^2 \leq e^{2\tau(y-1)\lambda + 4\tau_2(y-1)^2\lambda^2} = \exp\left(\frac{13}{12}\lambda(y-1) + \frac{6}{5}\lambda^2(y-1)^2\right), \end{aligned}$$

as claimed. \square

Now suppose that $P(j)$, $j = 1, 2, \dots$, are iid random variables distributed as P defined in (56), and that $B(j)$, $j = 1, 2, \dots$, are iid random variables distributed as $\text{Bin}(m, 2P(j)/m)$.

Lemma 10.16. *For every $\varepsilon > 0$ there exist $\delta > 0$ and $C \geq 1$ such that, if $1/2 \leq \frac{13}{12}\lambda \leq 1 - \varepsilon$, the following is true.*

For all $\alpha \geq Ca$, there exists a y with $1 < g_B(y) < y \leq 2$ such that

$$(58) \quad \frac{g_B(y)^\alpha}{y^{\alpha-a}} \leq e^{-\delta\alpha}.$$

PROOF. For ease of notation, let $u := y - 1$ and $r := \frac{13}{12}\lambda$, so that $1/2 \leq r \leq 1 - \varepsilon$. If $0 < u < \frac{1-r}{r} \leq 1$, by Lemma 10.15, we may estimate

$$g_B(y) \leq \exp\left(\frac{13}{12}\lambda u + \frac{6}{5}\lambda^2 u^2\right),$$

and thus obtain

$$\mathbf{P}[Z \geq \alpha] \leq \exp\left(\alpha\left(\frac{13}{12}\lambda u + \frac{6}{5}\lambda^2 u^2\right) - (\alpha - a) \log(u + 1)\right).$$

Using Lemma 10.6, we write the exponent as

$$(*) \quad \alpha r u + \frac{6 \cdot 12^2}{5 \cdot 13^2} \alpha r^2 u^2 - (\alpha - a) \log(u + 1).$$

In order to find a u minimizing (*), we take the derivative and solve the resulting quadratic equation

$$(**) \quad \frac{12^3}{5 \cdot 13^2} r^2 u^2 + \left(r + \frac{12^3}{5 \cdot 13^2} r^2\right) u - (1 - r) + a/\alpha = 0$$

The value of u which works is the larger one of the two roots:

$$(***) \quad u_r := \frac{-\left(1 + \frac{12^3}{5 \cdot 13^2} r\right) + \sqrt{\left(1 - \frac{12^3}{5 \cdot 13^2} r\right)^2 + \frac{4 \cdot 12^3}{5 \cdot 13^2}}}{\frac{2 \cdot 12^3}{5 \cdot 13^2} r} - O(a/\alpha),$$

with an absolute constant in the $O(\cdot)$ (see 10.8.3 for the computation). The numerator is greater than zero if, and only if, $4 \cdot \frac{12^3}{5 \cdot 13^2} r < \frac{4 \cdot 12^3}{5 \cdot 13^2}$, which is equivalent to $r < 1$. Thus, there exists a C depending only on r , such that $u_r > 0$ whenever $\alpha \geq Ca$. Moreover, by letting $u = \frac{1-r}{r}$ in (**), we see that $u_r < \frac{1-r}{r} \leq 1$, as required. Letting $u = u_r$ in (*), we obtain, for $\alpha \geq Ca$,

$$(***) \quad (\delta_r(u_r) + O(1/C))\alpha,$$

with an absolute constant in the $O(\cdot)$, where

$$\delta_r(u) = ru + \frac{6 \cdot 12^2}{5 \cdot 13^2} r^2 u^2 - \log(u + 1)$$

(see 10.8.3 for the computation). We have $\delta_r(u_r) < 0$, because $\delta_r(0) = 0$ and since, by the choice of u_r , the derivative of δ_r in the open interval $[0, u_r[$ is negative. This also implies that $y > g_B(y)$. Let

$$\delta_* := \max\{\delta_r(u_r) \mid 1/2 \leq r \leq 1 - \varepsilon\} < 0,$$

Finally, increase C , if necessary, to take care of the dependence on $O(1/c)$ in (***) and (****), and define $\delta := -\delta_*/2$. This completes the proof of the lemma. \square

Lemma 10.17. *If $\lambda \leq (1 - \varepsilon)\frac{12}{13}$, then*

$$(59a) \quad \mathbf{E} Z = \frac{a}{1 - \frac{13}{12}\lambda}$$

and there exist $\delta > 0$ and $C \geq 1$ depending only on ε , such that for all $\alpha \geq Ca$ we have the upper tail inequality

$$(59b) \quad \mathbf{P}[Z \geq \alpha] \leq e^{-\delta\alpha}.$$

PROOF. Equation (59a) is directly from Lemma 10.13.

Lemmas 10.13 and 10.16 together imply the tail inequality in the case when $\frac{13}{12}\lambda \geq 1/2$. For smaller values of λ , we just note that increasing λ increases the length of the first busy period, so that the probability for $\lambda := 6/13$ gives an upper bound for the probability for smaller values of λ . \square

10.5.2. Not-independent binomial. The arrivals at the queue in the context of our algorithm are not completely independent. Here we deal with the small amount of dependence.

We now describe what kind of $B(j)$ we allow. The setting is that n is a (large) integer, and that $m = m(n) = \Theta(n)$. Let $r > 1$ and

$$(60) \quad z = z_r = z_r(n) := \frac{r}{\delta} \log n,$$

where δ is as in Lemma 10.17. Suppose that $M(j), N(j)$ are random variables satisfying

$$(61a) \quad n - j \leq N(j) \leq n + j \quad \text{for all } j,$$

$$(61b) \quad 0 \leq M(j) \leq m \quad \text{for all } j,$$

with probability one, and

$$(61c) \quad m^- \leq M(j) \leq m^+ \quad \text{for all } j = 1, \dots, z$$

with probability at least $1 - O(n^{-r})$. Let the $B(j)$ be distributed as $\text{Bin}(M(j), \frac{P(j)}{N(j)})$ for all j . More accurately, we assume that there is an iid family of $P(j), j = 1, 2, 3, \dots$, distributed as P above, and an independent family of random variables $U(j, i), j = 1, 2, 3, \dots, i = 1, 2, 3, \dots$ each having uniform distribution on $[0, 1]$, and that the joint distribution of the $B(j)$ is the same as for the family of sums

$$(62) \quad \sum_{i=1}^{M(j)} \mathbf{I}\left[U(j, i) \leq \frac{P(j)}{N(j)}\right].$$

The $P(j)$ and $U(j, i)$ are assumed to be jointly independent, but we make no assumptions about independence regarding the $M(j)$ and $N(j)$ among themselves or from the $U(j, i)$ and $P(j)$. However, we do assume that a , the $M(j)$, and the $N(j)$ are such that

$$(63) \quad a + \sum_{j=1}^{\infty} B(j) = O(n)$$

holds with probability one.

Lemma 10.18. *Let $\lambda^\pm = \lambda^\pm(n) := \frac{m^\pm}{n - (\pm z)}$, and suppose $z \geq Ca$. If*

$$(64) \quad \lambda^+ \leq (1 - \varepsilon) \frac{12}{13},$$

then with the δ and C from Lemma 10.17, the following holds for large enough n :

$$(65a) \quad \frac{a}{1 - \frac{13}{12}\lambda^-} - O(n^{1-r}) \leq \mathbf{E} Z \leq \frac{a}{1 - \frac{13}{12}\lambda^+} + O(n^{1-r});$$

and for all $\alpha \geq Ca$

$$(65b) \quad \mathbf{P}[Z \geq \alpha] \leq e^{-\delta\alpha} + O(n^{-r}).$$

The proof can be found in the appendix: 10.8.4.

Remark 10.19. There is no danger in assuming $\delta \leq 1$ and $C \geq 1$, and we will do that from this point on.

10.6. The inner loop

Here we analyze Algorithm 4. Conditioning on $X(t)$, $Y_2(t)$, and $Y_3(t)$, we analyze the changes of the parameters X , Y_2 , and Y_3 during the $t + 1$ st run of the inner loop, and bound the probability that an empty clause is generated.

From now on, n and m denote the number of variables and clauses, respectively, in the initial random CNF formula, with $m = cn$ for some constant c . We assume $c \leq 10$, to get rid of some of the letter c in the expressions below. For any $\varepsilon > 0$, we say that $(x, y_2, y_3) \in \mathbb{R}^3$ is ε -good, if

$$(66) \quad \varepsilon n < x \quad \text{and} \quad \frac{y_2}{x} < (1 - \varepsilon) \frac{12}{13},$$

and that $\mathcal{H}(t)$ is ε -good if $(X(t), Y_2(t), Y_3(t))$ is ε -good.

10.6.1. Setup of the queues for Phases I and II. We now define the queues corresponding to the Phases I and II. We will suppress the dependency of the random processes on $\mathcal{H}(t)$ in the notation.

We define the queues Q_I and Q_{II} for the Phases I and II, respectively, by modifying Algorithm 4 a little bit. We will then analyze (the original) Algorithm 4 with the help of the queues Q_I and Q_{II} defined via this modification. The changes we make are the following: replace step (i-7) by

(i-7') If there are unused variables left, choose one var;

and step (i-8) by

(i-8') Expose all occurrences of the current variable x_j in clauses colored with a color different from red;

moreover, in the modification, we do not initiate a repair (since that would kill the queueing process).

Since, with these modifications, red clauses can contain used variables, it is possible to run out of variables before running out of clauses. It can be easily verified that this can only happen when all clauses are red. Hence, in this situation, the modified algorithm will just eat up the red clauses one per iteration.

In the Phase-I queue Q_I , the number of customers arriving in the first time interval, A_I , is the number of red clauses generated by setting the free variable x_0 (tentatively) to $1/2$. Thus, A_I is distributed as $\text{Bin}(Y_2(t), \frac{1}{X(t)})$. For the iterations $j = 1, 2, 3, \dots$, we find that $B_I(j + 1)$ is the number of uncolored 2-clauses which become red, plus the number of pink 3-clauses which become red, when setting the current variable x_j (tentatively) to $\bar{x}(\mathbb{I}_j)$. Thus, if we

denote by $Y_2'(j)$ the number of uncolored 2-clauses plus the number of pink 3-clauses at the beginning of iteration j , then conditioned on $Y_2'(j)$, the distribution of $B_I(j+1)$ is that of $\text{Bin}(Y_2'(j), \frac{P(j+1)}{X(t)-j})$, where as in the previous section, the $P(j+1)$ are iid random variables distributed as P defined in (56). If we agree on the convention that a $\text{Bin}(0, p/0)$ -variable is deterministically 0, this also holds when the queue runs out of variables.

In the Phase-II queue, the number of customers arriving in the first time interval, A_{II} , is the number of unit-clauses generated at the end of Phase I by setting the variables to their tentative values. The $B_{II}(j)$ are defined analogous to the $B_I(j)$.

At this point, note that the condition (63), which is needed for Lemma 10.18, is satisfied for both queues.

10.6.2. Bounds for the probabilities of some essential events. Below, we repeatedly use the following simple Chernoff-type inequality (e.g. equation (2.11) in [JLR00]): if U is a binomially distributed random variable with mean μ , then

$$(67) \quad \mathbf{P}[U \geq \alpha] \leq e^{-\alpha} \quad \text{for } \alpha \geq 7\mu.$$

Lemma 10.20. *Let $r > 1$, $1 \leq z = z(n) = o(n)$ an integer, (x, y_2, y_3) ε -good for some $\varepsilon > 0$, and $m^- := \max(0, y_2 - rz \log n)$, $m^+ := y_2 + rz \log n$. For both phases I and II of the inner loop, the following is true. If, at the beginning of the phase at step (i-1), there are x variables, y_2 2-clauses, and y_3 3-clauses, then the probability that, while dealing with the first z variables in the phase, the number of 2-clauses leaves the interval $[m^-, m^+]$, is $O(n^{-r})$.*

PROOF. For the upper bound m^+ , the probability that the number of 2-clauses exceeds m^+ is bounded from above by the probability that one in a sequence of z independent random variables with $\text{Bin}(m, \frac{3}{\varepsilon n/2})$ -distributions is greater than $r \log n$. Here the factor $1/2$ on the denominator takes care of the $z = o(n)$ variables which are used. For n large enough, this probability is at most

$$zO\left(\binom{m}{r \log n} \left(\frac{6/\varepsilon}{n}\right)^{r \log n}\right) = zO(e^{-r \log n}) = O(n^{-r}).$$

For the lower bound m^- , the probability can be bounded by the same argument, noting that, if $m^- = 0$, the corresponding probability is 0. \square

Let R denote the event that a repair is invoked during this run of Algorithm 4. Moreover, denote by Z_I and Z_{II} the length of the first busy period of the Phase I and Phase II queues, respectively. Note that they depend on A_I and A_{II} , respectively. Further let M_I and M_{II} be the total number of colored clauses which are generated during Phase I and Phase II, respectively; let H_I and H_{II} the event that, in some iteration, in steps (i-8), the current variable is found to be contained in a colored clause (other than the current clause C_j); and by $H_I^{\geq 2}$ the probability that in Phase I the current variable is found to be contained in at least two colored clauses (other than the current clause C_j).

Lemma 10.21. *Suppose that $\mathcal{H}(t)$ is 2ε -good. With the $\delta := \delta(\varepsilon)$ and $C := C(\varepsilon)$ from Lemma 10.18, and $r > 1$, the following is true for all n large enough (depending on ε).*

$$(68a) \quad \mathbf{P}[A_I \geq r \log n \mid \mathcal{H}(t)] = O(n^{-r})$$

$$(68b) \quad \mathbf{P}[Z_I \geq \frac{C}{\delta} r \log n \mid \mathcal{H}(t)] = O(n^{-r})$$

$$(68c) \quad \mathbf{P}[M_I \geq \frac{500C}{\varepsilon\delta} r \log n \mid \mathcal{H}(t)] = O(n^{-r})$$

$$(68d) \quad \mathbf{P}[H_I \mid \mathcal{H}(t)] = O_\varepsilon\left(\frac{\log^2 n}{n}\right)$$

$$(68e) \quad \mathbf{P}[H_I^{\geq 2} \mid \mathcal{H}(t)] = O_\varepsilon\left(\frac{\log^4 n}{n^2}\right)$$

$$(68f) \quad \mathbf{P}[R \mid \mathcal{H}(t)] = O_\varepsilon\left(\frac{\log^2 n}{n}\right)$$

$$(68g) \quad \mathbf{P}[A_{II} \geq \frac{500C}{\varepsilon\delta} (r+1) \log n \mid \mathcal{H}(t) \ \& \ R] = O(n^{-r})$$

$$(68h) \quad \mathbf{P}[Z_{II} \geq \frac{500C^2}{\varepsilon\delta} (r+1) \log n \mid \mathcal{H}(t) \ \& \ R] = O(n^{-r})$$

$$(68i) \quad \mathbf{P}[M_{II} \geq \frac{250000C^2}{\varepsilon^2\delta} (r+1) \log n \mid \mathcal{H}(t) \ \& \ R] = O(n^{-r})$$

$$(68j) \quad \mathbf{P}[H_{II} \mid \mathcal{H}(t) \ \& \ R] = O_\varepsilon\left(\frac{\log^2 n}{n}\right)$$

PROOF. For (68a), if $\mathcal{H}(t)$ is 2ε -good, then the probability that $A_I \geq r \log n$ is bounded from above by the probability that a $\text{Bin}(m, \frac{2}{2\varepsilon n})$ -variable is larger than $r \log n$, which is at most n^{-r} , for n large enough, by (67).

Proof of (68b). We use Lemma 10.18 together with Lemma 10.20 to bound the conditional probability that $Z_I \geq \alpha$. If $\mathcal{H}(t)$ is 2ε -good, then the m^+ from Lemma 10.20, with $x := X(t)$, $y_2 := Y_2(t)$, $y_3 := Y_3(t)$, and the $z = z_r$ from (60), is such that (64) is satisfied if n is large enough depending on ε .

The requirement for the estimate in (65b) is that $A_I \leq a_0 := \min(\alpha/C, z_r/C)$. Thus, for the probabilities conditional on $\mathcal{H}(t)$, we have

$$\begin{aligned} \mathbf{P}[Z_I \geq \alpha] &= \mathbf{P}[Z_I \geq \alpha \mid A_I \leq a_0] \mathbf{P}[A_I \leq a_0] + \mathbf{P}[Z_I \geq \alpha \mid A_I > a_0] \mathbf{P}[A_I > a_0] \\ &\leq O(e^{-\delta\alpha}) + O(n^{-r}) + \mathbf{P}[A_I > a_0]. \end{aligned}$$

With $\alpha := \frac{C}{\delta} r \log n$, using (68a) and (67), the right-hand side is $O(n^{-r})$.

Proof of (68c). For every iteration, a clause is only colored if the current variable of the iteration is contained in the clause. Hence, the number of clauses colored in the first j iterations is upper bounded by the sum of j independent $\text{Bin}(m, \frac{3}{\varepsilon n})$ -variables. Hence, the probability that in the first j iterations, the number of colored clauses exceeds $j\alpha$ is at most $e^{-\alpha}$ by (67), provided that $\alpha \geq \frac{500}{\varepsilon} j \geq 7 \cdot \frac{3m}{\varepsilon n/2} j$. Moreover, we have $M_I \leq m$ with probability one. Thus, conditioning on $\mathcal{H}(t)$ (and keeping in mind that $\mathcal{H}(t)$ is required to be 2ε -good), the probability that M_I is larger than $\frac{500C}{\varepsilon\delta} r \log n$ is at most

$$O(e^{-r \frac{500C}{\varepsilon\delta} \log n}) + m \mathbf{P}[Z_I \geq r \frac{500C}{\varepsilon\delta} \log n \mid \mathcal{H}(t)] = O(n^{-r}) + O(mn^{-500r}) = O(n^{-r}).$$

Proofs of (68d) and (68e). In the first phase, in the j th iteration, the probability that the current variable x_j occurs in a colored clause (other than the current clause C_j) is $O(\frac{M_I}{X(t)-Z_I})$, and the probability that the number of colored clauses containing x_j (other than the current one C_j) is two or more is $O((\frac{M_I}{X(t)-Z_I})^2)$.

By (68b) and (68c), we can bound the probability that this happens in the first Z_I iterations by $O_\varepsilon(\frac{\log^2 n}{n}) + O(n^{-r})$ and $O_\varepsilon(\frac{\log^4 n}{n^2}) + O(n^{-r})$, respectively, where the constant in the $O_\varepsilon(\cdot)$ depends only on ε .

Proof of (68f). Clearly, the probability that a repair occurs is at most the probability that, in some iteration, the current variable x_j occurs in a colored clause (other than the current one C_j). Thus, the inequality follow from (68d).

Proof of (68g). Since $A_{II} \leq M_I$, this inequality follows from (68c) and (68f), with r replaced by $r + 1$, by conditioning on R :

$$\begin{aligned} \mathbf{P}[M_I \geq \frac{500C}{\varepsilon\delta}(r+1)\log n \mid \mathcal{H}(t) \ \& \ R] \\ &\leq \mathbf{P}[M_I \geq \frac{500C}{\varepsilon\delta}(r+1)\log n \mid \mathcal{H}(t)] / \mathbf{P}[R \mid \mathcal{H}(t)] \\ &= O(n^{-r-1} \frac{n}{\log^2 n}) = O(n^{-r}). \end{aligned}$$

Proof of (68h). We now apply Lemmas 10.18 and 10.20 to the Phase-II queue. Let $r' := \frac{500C^2}{\varepsilon\delta}(r+1)$. If $\mathcal{H}(t)$ is 2ε -good, then the m^+ from Lemma 10.20, with $x := X(t)$, $y_2 := Y_2(t)$, $y_3 := Y_3(t)$, and the $z = z_{r'}$ from (60), is such that (64) is satisfied if n is large enough depending on ε .

Again, the requirement for the estimate in (65b) is that $A_{II} \leq a'_0 := \min(\alpha/C, z_{r'}/C)$. Thus, for the probabilities conditional on $\mathcal{H}(t) \ \& \ R$, we have

$$\begin{aligned} \mathbf{P}[Z_{II} \geq \alpha] \\ &= \mathbf{P}[Z_{II} \geq \alpha \mid A_{II} \leq a'_0] \mathbf{P}[A_{II} \leq a'_0] + \mathbf{P}[Z_{II} \geq \alpha \mid A_{II} > a'_0] \mathbf{P}[A_{II} > a'_0] \\ &\leq O(e^{-\delta\alpha}) + O(n^{-r'}) + \mathbf{P}[A_{II} > a'_0] \end{aligned}$$

With $\alpha := \frac{500C^2}{\varepsilon\delta}(r+1)\log n$, we have $a'_0 = \frac{500C}{\varepsilon\delta}(r+1)\log n$, so that, by (68g), the probability that $A_{II} > a'_0$ is $O(n^{-r})$. In total, we obtain an upper bound of $O(n^{-r})$ for the probability that $Z_{II} \geq \frac{500C^2}{\varepsilon\delta}(r+1)\log n$.

Proof of (68i). For every iteration, a clause is only colored if the current variable of the iteration is contained in the clause. Hence, the number of clauses colored in the first j iterations is upper bounded by the sum of j independent $\text{Bin}(m, \frac{3}{\varepsilon n/2})$ -variables. (The factor of $1/2$ in the denominator is to take care of the fact that the number of variables, while starting with at least εn , might drop below εn during the run of Phase I or Phase II.) Hence, the probability that in the first j iterations, the number of colored variables exceeds $j\alpha$ is at most $e^{-\alpha}$ by (67), provided that $\alpha \geq \frac{500}{\varepsilon}j \geq 7 \cdot \frac{3m}{\varepsilon n/2}j$. Moreover, we have $M_{II} \leq m$ with probability one. Thus, conditioning on $\mathcal{H}(t) \ \& \ R$ (and keeping in mind that $\mathcal{H}(t)$ is 2ε -good), the probability that M_{II} is larger than $\frac{500^2C^2}{\varepsilon^2\delta}(r+1)\log n$ is at most

$$\begin{aligned} O(e^{-\frac{500^2C^2}{\varepsilon^2\delta}(r+1)\log n}) + m \mathbf{P}[Z_{II} \geq \frac{500^2C^2}{\varepsilon^2\delta}(r+1)\log n \mid \mathcal{H}(t)] \\ = O(n^{-r}) + O(mn^{-500r}) = O(n^{-r}). \end{aligned}$$

Proof of (68j). In the second phase, in the j th iteration, the probability that the current variable x_j occurs in a colored clause (other than the current one C_j) is $O(\frac{M_I}{X(t)-Z_I})$. By (68h) and (68i), we can bound the probability that this happens in the first Z_{II} iterations by $O_\varepsilon(\frac{\log^2 n}{n}) + O(n^{-r})$, where the constant in the $O_\varepsilon(\cdot)$ depends only on ε . \square

10.6.3. Changes of the parameters $X(t)$, $Y_2(t)$, and $Y_3(t)$. We now move to study the differences between successive values of these parameters, and we start with $X(t+1) - X(t)$. Denote by F_I and F_{II} the number of iterations of the inner loop in the first and second phase, respectively. Clearly, $X(t) - X(t+1) = 1 + F_I + F_{II}$, where the leading 1 accounts for the free variable x_0 . Moreover, we have $F_I \leq Z_I$ and $F_{II} \leq Z_{II}$, and the inequality can be strict

for two reasons: in Phase I, a repair can occur, thus terminating the phase before Q_I drops to zero; in both phases a red clause can vanish (i.e. become black) in (i-9). However, note that

$$(69) \quad \begin{aligned} F_I &= Z_I && \text{with probability } 1 - O_\varepsilon\left(\frac{\log^2 n}{n}\right), \text{ and} \\ F_I \mathbf{I}[\bar{R}] &\geq Z_I \mathbf{I}[\bar{R}] - 1 && \text{with probability } 1 - O_\varepsilon\left(\frac{\log^4 n}{n^2}\right) \end{aligned}$$

by (68d), (68f) and (68e).

Let us abbreviate

$$\Delta X := -1 - \frac{\frac{Y_2(t)}{X(t)}}{1 - \frac{13Y_2(t)}{12X(t)}} = -1 - \frac{12Y_2(t)}{12X(t) - 13Y_2(t)} = -\frac{12X(t) - Y_2(t)}{12X(t) - 13Y_2(t)}.$$

Lemma 10.22. *If $\mathcal{H}(t)$ is 2ε -good and n large enough depending on ε , then*

$$(70a) \quad \left| -1 - \Delta X - \mathbf{E}(Z_I \mid \mathcal{H}(t)) \right| = O_\varepsilon\left(\frac{\log^2 n}{n}\right)$$

$$(70b) \quad \left| \Delta X - \mathbf{E}(X(t+1) - X(t) \mid \mathcal{H}(t)) \right| = O_\varepsilon\left(\frac{\log^4 n}{n}\right)$$

and

$$(70c) \quad \mathbf{P}\left[|X(t+1) - X(t)| \geq \log^2 n \mid \mathcal{H}(t)\right] = O(n^{-10})$$

PROOF. By what we have said above on the relationship between F_I , F_{II} and $X(t+1) - X(t)$, we have $F_I = Z_I \mathbf{I}[\bar{R}] - E_I$ and $F_{II} = Z_{II} - E_{II}$, where E_I and E_{II} are error terms accounting for red clauses vanishing. We have $\mathbf{E}(E_I \mid \mathcal{H}(t)), \mathbf{E}(E_{II} \mid \mathcal{H}(t)) = O_\varepsilon\left(\frac{\log^4 n}{n}\right)$ by (68d) and (68e) (noting that $E_I, E_{II} \leq m$).

We compute the mean of Z_I using Lemma 10.18 with the m^\pm from Lemma 10.20 with $z := \frac{r}{5} \log n$ as in (60). Thus, letting $v := rz \log n$ (the bound from Lemma 10.20), conditional on A_I and $\mathcal{H}(t)$, we have

$$\frac{A_I}{1 - \frac{13Y_2(t)-v}{12X(t)+z}} \leq \mathbf{E}(Z_I \mid A_I \ \& \ \mathcal{H}(t)) \leq \frac{A_I}{1 - \frac{13Y_2(t)+v}{12X(t)-z}},$$

so that

$$\mathbf{E}(Z_I \mid A_I \ \& \ \mathcal{H}(t)) = \frac{A_I}{1 - \frac{13Y_2(t)}{12X(t)}} + O_\varepsilon\left(\frac{A_I \log^2 n}{n}\right),$$

provided that $A_I \leq z/C$, which holds with probability at least $1 - O(n^{-2})$ by (68a) by increasing, if necessary, r beyond $2\delta C$. Since $Z_I = O(n)$ with probability one, we obtain

$$\begin{aligned} \mathbf{E}(Z_I \mid \mathcal{H}(t)) &= \mathbf{E}\left(\frac{A_I}{1 - \frac{13Y_2(t)}{12X(t)}} + O_\varepsilon\left(\frac{A_I \log^2 n}{n}\right) \mid \mathcal{H}(t)\right) \\ &= \frac{\mathbf{E}(A_I \mid \mathcal{H}(t))}{1 - \frac{13Y_2(t)}{12X(t)}} + O_\varepsilon\left(\frac{\log^2 n}{n}\right) = \frac{\frac{Y_2(t)}{X(t)}}{1 - \frac{13Y_2(t)}{12X(t)}} + O_\varepsilon\left(\frac{\log^2 n}{n}\right), \end{aligned}$$

which proves (70a). For F_I , we obtain

$$\begin{aligned} \mathbf{E}(F_I \mid \mathcal{H}(t)) &= \mathbf{E}(Z_I \mid \mathcal{H}(t)) - \mathbf{E}(Z_I \mathbf{I}(R) \mid \mathcal{H}(t)) - \mathbf{E}(E_I \mid \mathcal{H}(t)) \\ &= -1 - \Delta X + O_\varepsilon\left(\frac{\log^2 n}{n}\right) - O_\varepsilon(\log n) \mathbf{P}(R \mid \mathcal{H}(t)) - mO(n^{-r}) - O_\varepsilon\left(\frac{\log^4 n}{n}\right) \\ &= -1 - \Delta X + O_\varepsilon\left(\frac{\log^4 n}{n}\right) \end{aligned}$$

and

$$\mathbf{E}(F_{II} \mid \mathcal{H}(t) \ \& \ R) \leq \mathbf{E}(Z_{II}) = O_\varepsilon(\log n) + O(n^{-r})m,$$

from which (70b) follows.

Since $X(t) - X(t+1) \leq 1 + Z_I + Z_{II}$, the tail inequality (70c) follows immediately from (68b) and (68h). \square

Lemma 10.23. *If $\mathcal{H}(t)$ is 2ε -good, then*

$$(71a) \quad \left| \Delta X \frac{3Y_3(t)}{X(t)} - \mathbf{E}(Y_3(t+1) - Y_3(t) \mid \mathcal{H}(t)) \right| = O_\varepsilon\left(\frac{\log^4 n}{n}\right)$$

and

$$(71b) \quad \mathbf{P}\left[|Y_3(t+1) - Y_3(t)| \geq \log^2 n \mid \mathcal{H}(t)\right] = O(n^{-10})$$

PROOF. Let us denote by $X'(j)$ the number of unused variables after j iterations of the inner loop, i.e., before \mathbf{x}_j is used, for $j = 0, 1, 2, \dots$. In every iteration of the inner loop, regardless of whether in Phase I or Phase II, for every uncolored 3-clause \mathbf{C} , there is a $\frac{3}{X'(j)}$ probability that the current variable \mathbf{x}_j is found to be contained in \mathbf{C} in step (i-10), or (i-3.3), respectively, for the zeroth iteration in Phase I. If that is the case, the 3-clause is colored, and when the inner loop terminates, the clause will no longer be a 3-clause.

If we suppose that, at the beginning of iteration $j = 0, 1, 2, \dots$, before the current variable \mathbf{x}_j is treated, there are $Y_3'(j)$ uncolored 3-clauses and $X'(j)$ unused variables, then the number of 3-clauses which are hit by \mathbf{x}_j is distributed as $\text{Bin}(Y_3'(j), \frac{3}{X'(j)})$. (We have $X'(j) = X(t) - j$ in Phase I, but in Phase II the value of course depends on how Phase I went.)

For (71b), we can just use the fact that the number of 3-clauses which are colored is bounded from above by $M_I + M_{II}$, the total number of colored clauses. Thus, by (68c) and (68i), this number is at most $\log^2 n$ with probability $1 - O(n^{-10})$ for n large enough depending on ε .

For the conditional expectation estimate (71a), we compute, conditional on $\mathcal{H}(t)$,

$$\mathbf{E}(Y_3(t+1) - Y_3(t)) = \mathbf{E}((Y_3(t+1) - Y_3(t)) \mathbf{I}[R]) + \mathbf{E}((Y_3(t+1) - Y_3(t)) \mathbf{I}[\bar{R}]).$$

For the left summand, we have

$$\begin{aligned} \mathbf{E}((Y_3(t+1) - Y_3(t)) \mathbf{I}[R]) &\leq \mathbf{E}(\log^2 n \mathbf{I}[R \ \& \ Y_3(t+1) - Y_3(t) \leq \log^2 n]) \\ &\quad + \mathbf{E}(m \mathbf{I}[R \ \& \ Y_3(t+1) - Y_3(t) \geq \log^2 n]) \\ &\leq \log^2 n \mathbf{P}[R] + m \mathbf{P}[Y_3(t+1) - Y_3(t) \geq \log^2 n] = \log^2 n O_\varepsilon\left(\frac{\log^2 n}{n}\right) + O(n^{-9}) \\ &= O_\varepsilon\left(\frac{\log^4 n}{n}\right), \end{aligned}$$

by (68f) and (71b).

For the right summand, we have

$$\mathbf{E}((Y_3(t+1) - Y_3(t)) \mathbf{I}[\bar{R}]) = \mathbf{E}\left(\sum_{j=1}^{Z_I+1} G(j) \mathbf{I}[\bar{R}]\right) + O_\varepsilon\left(\frac{\log^2 n}{n}\right),$$

where, conditioned on $Y_3'(j)$ as defined above, the $G(j+1)$ are distributed as $\text{Bin}(Y_3'(j), \frac{3}{X(t)-j})$, and the $O(\cdot)$ accounts for the possibility that $F_I < Z_I$, cf. (69). Using (67) and a similar argument as above, we see that

$$\mathbf{E}\left(\sum_{j=1}^{Z_I+1} G(j) \mathbf{I}[\bar{R}]\right) = \mathbf{E}\left(\sum_{j=1}^{Z_I+1} G(j)\right) + O_\varepsilon\left(\frac{\log^4 n}{n}\right).$$

Computing the expectation of the sum can be done in the same way as for classical SAT (e.g. in [Ach00, AS00, Ach01]). Indeed, using the optional stopping theorem ($Z_I + 1$ is a

stopping time for the history of the queue together with all random processes involved; cf. the proof of the next lemma for the details, where the situation is essentially the same, only a bit more complicated), we find that

$$\mathbf{E}\left(\sum_{j=1}^{Z_I+1} G(j)\right) = \mathbf{E}\left(\sum_{j=0}^{Z_I} \frac{3Y_3'(j)}{X(t)-j}\right),$$

where we agree that $0/0 = 0$. By (71b), $Y_3(t) - \log^2 n \leq Y_3'(j) \leq Y_3(t)$ with probability $1 - O(n^{-10})$, and by (68b) we have $Z_I \leq \log^2 n$, implying $X(t) - j \geq \frac{1}{2}X(t)$, with probability $1 - O(n^{-10})$. Thus, we conclude

$$\begin{aligned} & \mathbf{E}\left(\sum_{j=0}^{Z_I} \frac{3Y_3'(j)}{X(t)-j}\right) \\ &= \mathbf{E}\left(\mathbf{I}\left[Y_3(t) - \log^2 n \leq Y_3'(j) \ \& \ X(t) - j \geq \frac{1}{2}X(t)\right] \cdot \sum_{j=0}^{Z_I} \left(\frac{3Y_3(t)}{X(t)} + O\left(\frac{X(t) \log^2 n}{X(t)^2}\right)\right)\right) \\ & \qquad \qquad \qquad + O(n^{-7}) \\ &= (1 + \mathbf{E} Z_I) \left(\frac{3Y_3(t)}{X(t)} + O\left(\frac{\log^2 n}{n}\right) + O(n^{-7})\right) = -\Delta X \frac{3Y_3(t)}{X(t)} + O\left(\frac{\log^2 n}{n}\right), \end{aligned}$$

by (70a). This concludes the proof of (71a). \square

Lemma 10.24. *If $\mathcal{H}(t)$ is 2ε -good, then*

(72a)

$$\left| \frac{3Y_3(t)}{2X(t)} - (\Delta X + 1) \frac{13Y_3(t)}{8X(t)} + \Delta X \frac{2Y_2(t)}{X(t)} - \mathbf{E}(Y_2(t+1) - Y_2(t) \mid \mathcal{H}(t)) \right| = O_\varepsilon\left(\frac{\log^4 n}{n}\right)$$

and

$$(72b) \quad \mathbf{P}\left[|Y_2(t+1) - Y_2(t)| \geq \log^2 n \mid \mathcal{H}(t)\right] = O(n^{-10})$$

PROOF. The tail inequality is obtained by referring to (68c) and (68i) again, since very clause which changes its length has been colored before that can happen.

Let us denote by $X'(j)$ the number of unused variables after j iterations of the inner loop, i.e., before x_j is selected. In every iteration of the inner loop, regardless of whether in Phase I or Phase II, for every uncolored 2-clause C , there is a $\frac{2}{X'(j)}$ probability that the current variable x_j is found to be contained in C in step (i-6.10), or (i-3.3), respectively, for the zeroth iteration in Phase I. If that is the case, the 2-clause is colored, and when the inner loop terminates, the clause will no longer be a 2-clause. The same is true for 3-clauses which have become red in some previous iteration. Denote the total number of 2-clauses and pink 3-clauses which are hit by the current variable in some iteration over the whole run of Algorithm 4 by $L_{2\times}$.

The analysis of the expectation and tail of $L_{2\times}$ is almost identical to the analysis done in the previous lemma for the 3-clauses. Here, too, we have to condition on the number of uncolored 2-clauses and pink 3-clauses not changing too much. The difference is the need to control the number of pink 3-clauses and, after a repair, the number of 3-clauses becoming 2-clauses. The latter two numbers are bounded from above by $Y_3(t+1) - Y_3(t)$, which is at most $\log^2 n$ with probability $1 - O(n^{-10})$. Thus, for $L_{2\times}$, we just note that its expectation accounts for the summand $-\Delta X \frac{2Y_2(t)}{X(t)}$ in (72a).

Now let us denote the number of 3-clauses which become 2-clauses during the two phases of the inner loop by $L_{3\rightarrow 2}$, and let us also focus on the case when no repair occurs.

In this case $L_{3 \rightarrow 2}$ behaves similarly to $Y_3(t+1) - Y_3(t)$, with two differences: The probabilities that a 3-clause is colored pink is different; and the probability in the zeroth iteration differs from the others. Let us first consider the zeroth iteration. The probability that the tentative value $1/2$ of \mathbf{x}_0 makes a 3-clause pink is $1/2$ by Lemma 10.4. Thus, if there is no repair, this contribution is distributed as $\text{Bin}(Y_3(t), \frac{1}{2} \cdot \frac{3}{X(t)})$.

For the other iterations, $j = 1, 2, 3, \dots$, if an uncolored 3-clause \mathcal{C} contains the current variable \mathbf{x}_j , the probability that \mathcal{C} becomes pink in (i-13) depends on the current interval \mathbb{I}_j , and is distributed as P defined in (56). Indeed, if we denote the number of uncolored 3-clauses in iteration j by $Y'_3(j)$ again, then, conditioned on $Y'_3(j)$ and $X'(j)$, the number $G(j+1)$ of uncolored 3-clauses which become pink in iteration j is distributed as $\text{Bin}(Y'_3(j), \frac{3P(j+1)}{X'(j)})$, i.e., binomial with random parameter $P(j+1)$. The $P(j)$ are the iid random variables distributed as P in (56) defined by $\bar{x}(\mathbb{I}_j)$, in other words $P(j+1) = 1 - 2\bar{x}(\mathbb{I}_j)(1 - \bar{x}(\mathbb{I}_j))$.

Let $G(1)$ be distributed as $\text{Bin}(Y_3(t), \frac{3}{2X(t)})$, define $D(j+1) := G(j+1) - \frac{13Y'_3(j)}{8X'(j)}$, where we agree that $0/0 = 0$, and denote by $\mathcal{F}(j)$ the history of the process up to iteration j , i.e., before the variable \mathbf{x}_j is treated. Then $\sum_{j=1}^{\ell} D(j)$, $\ell = 1, 2, 3, \dots$, is a martingale with respect to $\mathcal{F}(j)$, $j = 0, 1, 2, \dots$, and $Z_I + 1$ is a stopping time, because deciding whether $Z_I + 1 \leq \ell$ amounts to checking whether $Q_I(\ell) = 0$.

To estimate the expectation of the contribution of these, we use the optional stopping theorem again; note that the stopping time is finite with probability one, because $Z_I \leq m$. We conclude that $\mathbf{E} \left(\sum_{j=1}^{Z_I+1} D(j) \right) = 0$, which means

$$\mathbf{E} \left(\sum_{j=1}^{Z_I+1} G(j) \right) = \mathbf{E} \left(\sum_{j=0}^{Z_I} \frac{13Y'_3(j)}{8X'(j)} \right).$$

Arguing as we have done a number of times in regard of the possible deviations of $Y'(j)$ from $Y(t)$, we see that the right hand side equals

$$(\mathbf{E} Z_I + 1) \frac{13Y_3(t)}{8X(t)} + O_\varepsilon \left(\frac{\log^4 n}{n} \right).$$

Getting rid of the conditioning on the event that no repair occurs is done in the same way as in the previous lemma, and we leave the details to the reader. \square

10.6.4. Failure probability. We now bound the probability that an empty clause is generated by a run of the inner loop, including, possibly, the repair and following second phase.

Lemma 10.25. *If $\mathcal{H}(t)$ is 2ε -good, then the probability that Algorithm 4 produces an empty clause, is $o(1/n)$.*

PROOF. We use Lemma 10.1. Let us first deal with Phase II. The probability that the algorithm “raises a flag” in Phase II is $O_\varepsilon(\frac{\log^2 n}{n})$ by (68j), conditioned on a repair occurring, so that by the law of total probability, the probability that the algorithm “raises a flag” in Phase II is at most $O_\varepsilon(\frac{\log^4 n}{n^2})$, by (68f).

For Phase I, we need to go through the possible reasons for the algorithm to “raise a flag”. First of all, by (68e), the probability that the current variable \mathbf{x}_j is contained in a colored (red or not) clause other than the current one \mathcal{C}_j is $O_\varepsilon(\frac{\log^4 n}{n^2})$, which takes care of step (i-6.3).

The probability that a fixed clause contains the current variable of a fixed iteration depends only on the number of variables and the number of unexposed atoms in the clause, and so it can always be bounded by $\frac{3}{\varepsilon n}$. In order for a 3-clause to become red or blue (or even black), it must contain the current variable of (at least) two iterations. The probability of this happening is $O_\varepsilon(\frac{\log n}{n^2})$, where we have used (68b). This gives the case of step (i-9.1).

Similarly, for step (i-9.2), a 2-clause must have been hit twice by the current variable of an iteration, the probability of which is again bounded by $O_\varepsilon\left(\frac{\log n}{n^2}\right)$.

In total, the failure probability can be bounded by $O\left(\frac{\text{polylog } n}{n^2}\right)$ \square

10.7. The outer loop

At the heart of analysis of the outer loop is the well-known theorem of Wormald's which, in certain situations, allows to estimate parameters of random processes by solutions to differential equations. Here is the first goal of our analysis.

Lemma 10.26. *For every $c \in]0, 3]$, the initial value problem*

$$(73a) \quad \frac{dy}{dx} = \frac{-18cx^4 + 2y(12x - y)}{x(12x - y)}$$

$$(73b) \quad y(1) = 0$$

has a unique solution y defined on the interval $]0, 1]$.

See Fig 10.1 for a rough sketch of the direction field (73a) with $c = 2$, and a solution to the IVP. Since, ultimately, we will solve the IVP (73) numerically for the right value of c anyway, strictly speaking, this lemma is not needed to complete our argument. However, we would like to reduce our reliance on numerical computations as much as possible.

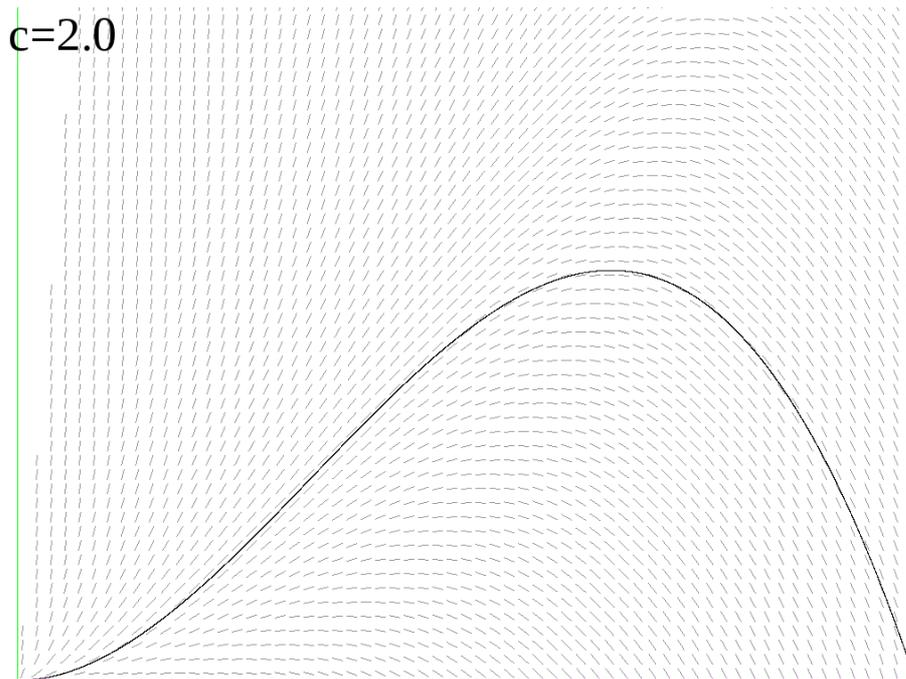


FIGURE 10.1. Direction field and solution for IVP (73)

PROOF OF LEMMA 10.26. To use the known theorems on IVPs, note that the right hand side of (73a), seen as a function of (x, y) , is continuously differentiable on $\{(x, y) \mid x > 0, y < 12x\}$.

We make the following claims:

- (a) For $4/5 \leq x \leq 1$, the solution to the IVP never crosses the line $y = 5(1 - x)$;
- (b) for $0 < x \leq 4/5$, the solution to the IVP never crosses the line $y = 6x$.

Thus, the solution to the IVP does not approach the $y = 12x$, which implies that the solution extends to the whole interval $]0, 1]$.

Let $g(x, y) := \frac{-18cx^4 + 2y(12x-y)}{x(12x-y)}$, the right hand side of the ODE (73a). To prove claim (b), it suffices to show that, with $y(x) := 6x$, whenever $0 < x \leq 4/5$, we have $\frac{dy}{dx} < g(x, y(x))$. The computation is easy but tedious and can be found in the appendix, see 10.8.5. Similarly, for claim (a), with $y(x) := 5(1-x)$, for every $4/5 \leq x \leq 1$, we have $\frac{dy}{dx} < g(x, y(x))$. The computation is in the appendix, too. \square

Lemma 10.27. *Let $c \leq 3$ and y a solution to (73), and let x_0 be the infimum over all $x \geq 3\varepsilon$ for which*

$$(74) \quad 13y(x) < (1 - 3\varepsilon)12x$$

holds. Then there exists a $\tau > 0$ and a strictly decreasing smooth function $x: [0, \tau] \rightarrow \mathbb{R}$ with $x(0) = 1$ and $x(\tau) = x_0$, such that whp for all t with $t/n < \tau$:

$$(75a) \quad X(t) = nx(t/n) + o(n)$$

$$(75b) \quad Y_2(t) = ny(x(t/n)) + o(n)$$

$$(75c) \quad Y_3(t) = ncx(t/n)^3 + o(n).$$

Moreover, we have the relationship

$$(75d) \quad \frac{dx}{dt} = -1 - \frac{y(x)}{x - \frac{13}{12}y(x)} = -\frac{12x - y(x)}{12x - 13y(x)}$$

PROOF. For the proof we use Wormald's well-known theorem, which requires some set up and computations. Using the notation of Theorem 5.1 in [Wor99], let

$$\begin{aligned} D &:= \{(t, x, y_2, y_3) \in]-\varepsilon, c + \varepsilon[^4 \mid (nx, ny_2, ny_3) \text{ is } 2\varepsilon\text{-good}\} \\ C_0 &:= 10 \\ \beta &:= \log^2 n \\ \gamma &:= 3n^{-2} \\ \lambda_1 &:= \frac{\log^5 n}{n} \\ \lambda &:= \frac{\log^{7/3} n}{n^{1/3}}, \end{aligned}$$

Note that $\lambda > \lambda_1 + C_0n\gamma$, and $\lambda = o(1)$, as required in Theorem 5.1 in [Wor99].

Obviously, we have $0 \leq X, Y_2, Y_3 < C_0n$.

- (i) Equations (70c), (72b), and (71b), respectively, show that, if $(t/n, X(t)/n, Y_2(t)/n, Y_3(t)/n) \in D$, then, conditioned on $\mathcal{H}(t)$, the probability that $X(t+1) - X(t) \leq \beta$, $Y_2(t+1) - Y_2(t) \leq \beta$, and $Y_3(t+1) - Y_3(t) \leq \beta$ hold, is at least $1 - \gamma$.
- (ii) The first parts of Lemmas 10.22, 10.24, and 10.23, respectively, show that, if $(t, x, y_2, y_3) := (t/n, X(t)/n, Y_2(t)/n, Y_3(t)/n) \in D$,

$$\begin{aligned} \left| f(t, x, y_2, y_3) - \mathbf{E}(X(t+1) - X(t) \mid \mathcal{H}(t)) \right| &\leq \lambda_1 \\ \left| g_2(t, x, y_2, y_3) - \mathbf{E}(Y_2(t+1) - Y_2(t) \mid \mathcal{H}(t)) \right| &\leq \lambda_1 \\ \left| g_3(t, x, y_2, y_3) - \mathbf{E}(Y_3(t+1) - Y_3(t) \mid \mathcal{H}(t)) \right| &\leq \lambda_1, \end{aligned}$$

where

$$\begin{aligned} f(t, x, y_2, y_3) &:= -1 - \frac{12y_2(t)}{12x(t) - 13y_2(t)} \\ g_2(t, x, y_2, y_3) &:= \frac{3y_3(t)}{2x(t)} + (-1 - f(t, x, y_2, y_3)) \frac{13y_3(t)}{8x(t)} + f(t, x, y_2, y_3) \frac{2y_2(t)}{x(t)} \\ g_3(t, x, y_2, y_3) &:= f(t, x, y_2, y_3) \frac{3y_3(t)}{x(t)}. \end{aligned}$$

(iii) There exists an L depending on ε such that f, g_2, g_3 are L -lipschitz continuous on D .

Let x, y_2, y_3 be the solution to the initial value problem

$$(76a) \quad \frac{dx}{dt} = f(t, x(t), y_2(t), y_3(t))$$

$$(76b) \quad \frac{dy_2}{dt} = g_2(t, x(t), y_2(t), y_3(t))$$

$$(76c) \quad \frac{dy_3}{dt} = g_3(t, x(t), y_2(t), y_3(t))$$

$$(76d) \quad x(0) = 1 \quad y_2(0) = 0 \quad y_3(0) = c.$$

From Wormald's theorem, we conclude that with probability

$$1 - O\left(n\gamma \frac{\beta}{\lambda} e^{-n(\lambda/\beta)^3}\right) = 1 - O\left(\frac{1}{n}\right),$$

it is true that, for all $t = 0, \dots, \sigma n$, we have $X(t) = nx(t/n) + O(\lambda n)$, $Y_2(t) = ny_2(t/n) + O(\lambda n)$, and $Y_3(t) = ny_3(t/n) + O(\lambda n)$, where $\sigma = \sigma(n)$ is the supremum over all s for which the solution to (76) can be extended before reaching within a distance of $C\lambda$ from the boundary of D , for a large constant C .

We now need to study the initial value problem (76). Let us start with the first equation (76a), which we write as

$$\frac{dx}{dt} = -\frac{12x - y_2}{12x - 13y_2},$$

which amounts to

$$(77) \quad -dt = \frac{12x - 13y_2}{12x - y_2} dx = \left(1 - \frac{12y_2}{12x - y_2}\right) dx,$$

The third inequality

$$\frac{dy_3}{dt} = \frac{dx}{dt} \frac{3y_3}{x},$$

is equivalent to

$$\frac{dy_3}{dx} = \frac{3y_3}{x},$$

which immediately integrates to¹

$$y_3 = cx^3,$$

where the constant before the x^3 is derived from the initial value conditions $y_3(0) = c$ and $x(0) = 1$. Finally, we write the second equation as

$$\frac{dy_2}{dt} = -\frac{y_3}{8x} - \frac{dx}{dt} \frac{13y_3}{8x} + \frac{dx}{dt} \frac{2y_2}{x} = -\frac{y_3}{8x} - \frac{13}{8} cx^2 \frac{dx}{dt} + \frac{2y_2}{x} \frac{dx}{dt}$$

from which we obtain

$$\frac{dy_2}{dx} = -\frac{c}{8} x^2 \frac{dt}{dx} - \frac{13}{8} cx^2 + \frac{2y_2}{x},$$

¹It should be noted that this is the same relationship between x and y_3 as in the case of classical 3-SAT (see [Ach01]).

which, by (77), yields

$$\frac{dy_2}{dx} = \frac{c}{8}x^2 \frac{12x - 13y_2}{12x - y_2} - \frac{13}{8}cx^2 + \frac{2y_2}{x} = \frac{-18cx^4 + 2y_2(12x - y_2)}{x(12x - y_2)}$$

which is an ODE of the function y_2 in the variable x . In fact, with $y_2(1) = 0$, we recognize the IVP (73), and thus $y = y_2$ in the interval on which both are defined.

To summarize, we have $y_3 = cx^3$, and $y_2 = y$ as a function of x is a solution to the IVP (73), and x as a function of t solves the ODE (76a) with boundary condition $x(0) = 1$.

From Lemma 10.26, we know that the solution y to (73) can be extended to a solution of the IVP defined on the full interval $]0, 1]$. Moreover, $\frac{dx}{dt} < 0$ whenever $13y(x) < 12x$, so the derivative of x is strictly negative provided that $x \geq x_0$. This implies that the solutions x , y_2 , y_3 to (76) can be extended to the interval $[0, \tau]$, where τ is the unique number satisfying $x(\tau) = x_0$; in particular we have $\sigma < \tau$.

This completes the proof of the lemma. \square

We are now ready to prove Theorem 10.2.

PROOF OF THEOREM 10.2. Lemma 10.27 gives the behavior of the parameters $X(t)$, $Y_2(t)$, and $Y_3(t)$ up to an error with high probability for all $t = 0, \dots, \tau n$. We need to check that

- (a) the algorithm terminates before t grows beyond τn ,
- (b) in this region of t , whp, the algorithm does not produce an empty clause.

For (a), we solve the IVP (73) numerically for $c = 2.3$. The solution is drawn in Fig. 10.2. The figure also shows the line $13y = 12x$. For this value of c , we see that there is an $\varepsilon > 0$ such that the solution $y(x)$ to the IVP (73) satisfies $13y(x) < 12(1 + 2\varepsilon)x$ for all $x > 2\varepsilon$; w.l.o.g., we may assume that $\varepsilon < 1/9$. Consequently, the x_0 from Lemma 10.27 equals 3ε . Algorithm 3 terminates as soon as $Y_2(t) + Y_3(t) \leq c'X(t)$. Thus, by Lemma 10.27, we have an $s < \tau$ such that $x(s) = 1/3 > x_0$, and that, if we let $c' := \frac{50}{39}$, whp, for this $t := \lceil sn \rceil$

$$\begin{aligned} Y_2(t) + Y_3(t) &= ny(1/3) + nc(1/3)^3 + o(n) \\ &\leq n \left((1 - 2\varepsilon) \frac{12}{13} \cdot \frac{1}{3} + \frac{c}{27} \right) + o(n) \leq \frac{49}{39} \cdot \frac{1}{3}n \leq c'X(t) - o(n), \end{aligned}$$

if n is large enough. Thus, the algorithm terminates before the parameters $X(\cdot)$, $Y_2(\cdot)$, $Y_3(\cdot)$ fail to be 2ε -good.

It follows that Lemma 10.25 gives a failure probability of $o(1/n)$ per iteration, so that the total failure probability is $o(1)$. This proves (b) and completes the proof of Theorem 10.2. \square

10.8. Conclusions and outlook

The presented algorithm and its analysis provide a first systematic approach to random iSAT formulas. In the course of the paper, analytical methods for dealing with intervals in CNF formulas have been established, in particular as intervals imply correlation between the variables when choosing a value. These tools will be useful in the study of random algorithms for iSAT as well as in approximating a threshold in random 3-iSAT.

We have given an algorithm for k -iSAT, for $k = 3$, which succeeds with high probability on instances for which $m/n \leq 2.3$. It is conceptually easy to extend the algorithm and the analysis to general k up to the point where the initial value problem has to be solved. For k -iSAT there are $k - 2$ ordinary differential equations to be numerically solved after the transformation in Lemma 10.27, which makes it improbable that a general formula for the maximal ratio can be derived. Solving the system for small values of k , we obtain the results shown in Table 2 (we always rounded down generously).

It is possible to show that, whp, our algorithm fails to produce a satisfying interpretation if $m/n = c$ where c is a constant for which the solution to the IVP (73) crosses the line $13y = 12x$

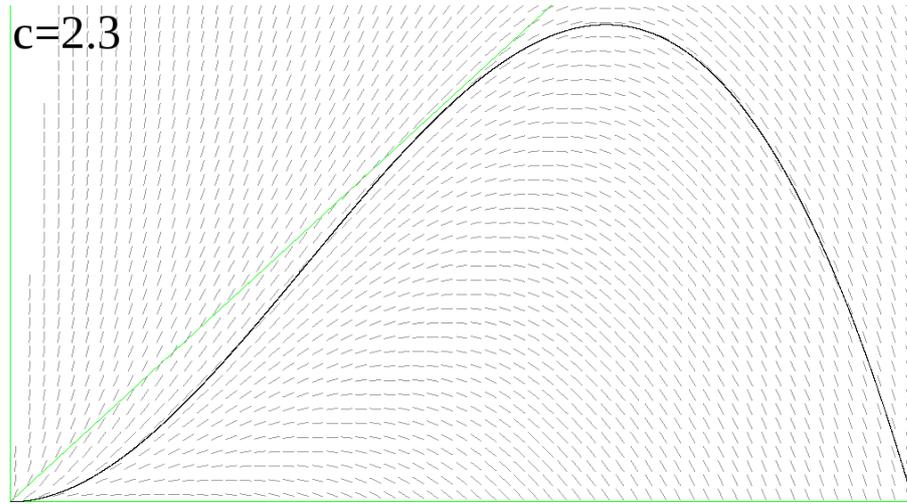


FIGURE 10.2. Solution of IVP (73) with bounding curves

k	3	4	5	6	7	8
max. m/n	2.3	3.75	6.25	10.5	18.5	32.5

TABLE 2. Performance for different values of k

(the green line in Fig 10.2), e.g., $c = 2.4$. This is so because then the inner loop runs for $\Omega(n)$ steps, and thus, whp, the algorithm “raises a flag”. (However, such a result appears futile, given the very limited repair routine which we refer to.)

Some further questions will be of interest. Firstly, the proposed algorithm can be improved in an obvious way: Whenever a variable is set, choose a value which satisfies the maximum number of literals containing the variable. This, however, requires that the following question be answered. Let λ be a nonnegative real number. Suppose that I_0, I_1, I_2, \dots are random intervals drawn independently uar from the sub-intervals of $[0, 1]$, and N is a Poisson random variable with mean λ , independent from the I_j . What is the expectation $\xi(\lambda)$ of the following random variable?

$$\max \left\{ |K| \mid K \subset \{1, \dots, N\}, I_0 \cap \bigcap_{j \in K} I_j \neq \emptyset \right\}?$$

Secondly, a bound for the ratio above which random 3-iSAT formulas are wpp/whp not satisfiable might be interesting and worthwhile to be considered.

Thirdly, there might be a sharp threshold for random 2-iSAT as for classical 2-SAT [CR92, Goe96]. In fact, we conjecture that there is a threshold at $c = 3/2$ (the value from Proposition 10.12). For this it remains to prove that for $c > 3/2$, a random 2-iSAT formula with $m/n = c$ is whp not satisfiable. More generally, it may be of interest whether the results of Friedgut (and Bourgain) [Fri99] (see also [Mol02, Mol03, CD03, CD04, CD09]) can be applied to random iSAT formulas to prove that a threshold (function) exists for k -iSAT for $k \geq 3$.

Fourthly, possibly, a stronger bound for 3-iSAT could be derived by adapting the algorithm of [KKL06] to the interval case. This would pose two problems: First we do are interested in a whp result, which is not offered by the algorithm in [KKL06], so a backtracking routine would have to be devised; Secondly, the rule for the value assignment significantly complicates the computations for random intervals. In their algorithm a randomly chosen variable is assigned the value such that most clauses, in which it occurs, are satisfied, i.e., a variable is assigned a 1

if it mostly occurs not negated, and 0 otherwise. For intervals this translates to assigning a value to a variable that is contained in the non-empty intersection of a maximal number of associated intervals. But the analysis of the probability of this maximal number turns out to be demanding for general intervals.

As a final question, we would like draw attention to the fact that several papers have raised questions concerning the existence and location of a threshold for random regular 3-iSAT [BHM00b, BM99a, BMC⁺07, MBEI98].

We would like to close by thanking the anonymous referees for their very valuable comments!

Appendix: Deferred proofs

10.8.1. Computations for Lemma 10.6. For (a), we compute

$$\begin{aligned}
1 - \mathbf{E}P &= \mathbf{E}(2X(1 - X)) = \int 2t(1 - t) dF(t) \\
&= \int_{[0, 1/2[} 2t(1 - t) \partial_t F(t) dt + 2t(1 - t) \Big|_{t=1/2} \cdot \frac{1}{2} + \int_{]1/2, 1]} 2t(1 - t) \partial_t F(t) dt \\
&= \int_0^{1/2} 2t(1 - t) 2t dt + \frac{1}{4} + \int_{1/2}^1 2t(1 - t) 2(1 - t) dt \\
&= \frac{5}{48} + \frac{1}{4} + \frac{5}{48} = \frac{11}{24}.
\end{aligned}$$

For (b), we compute

$$\begin{aligned}
\mathbf{E}(X^2(1 - X)^2) &= \int t^2(1 - t)^2 dF(t) \\
&= t^2(1 - t)^2 \Big|_{t=1/2} \cdot \frac{1}{2} + \int_{[0, 1/2[} t^2(1 - t)^2 \partial_t F(t) dt + \int_{]1/2, 1]} t^2(1 - t)^2 \partial_t F(t) dt \\
&= \frac{1}{2^5} + \int_{[0, 1/2[} t^2(1 - t)^2 2t dt + \int_{]1/2, 1]} t^2(1 - t)^2 2(1 - t) dt \\
&= \frac{1}{2^5} + 4 \int_0^{1/2} t^3(1 - t)^2 dt \\
&= \frac{1}{2^5} + 4 \left(\frac{1}{4} t^4(1 - t)^2 \Big|_{t=0}^{1/2} + \frac{1}{10} t^5(1 - t) \Big|_{t=0}^{1/2} + \frac{1}{60} t^6 \Big|_{t=0}^{1/2} \right) \\
&= \frac{1}{2^5} + 4 \left(\frac{1}{4} \frac{1}{2^6} + \frac{1}{10} \frac{1}{2^6} + \frac{1}{60} \frac{1}{2^6} \right) = \frac{1}{2^5} + \frac{1}{2^6} \left(1 + \frac{2}{5} + \frac{1}{15} \right) = \frac{1}{2^5} + \frac{22}{15 \cdot 2^6} = \frac{1}{2^5} + \frac{11}{15 \cdot 2^5} = \frac{15+11}{15 \cdot 2^5} = \frac{13}{15 \cdot 2^4}.
\end{aligned}$$

Hence, using (a), we obtain

$$\mathbf{E}(P^2) = 1 - 2(1 - \mathbf{E}P) + 4\mathbf{E}X^2(1 - X)^2 = 1 - \frac{11}{12} + \frac{13}{60} = \frac{18}{60} = \frac{3}{10}.$$

10.8.2. Proof of Lemma 10.13. The proof is taken almost word for word from Grimmett & Stirzaker [GS01], Theorem 11.3.17, with some changes due to the discrete arrival- and servicing points.

We say that the *sons* of a customer Paul are those customers arriving in the time interval in which Paul is serviced. Paul's *family* consists of himself and all of his descendants.

Fix a time interval j in which the queue is not empty and denote by X the size of the family of the customer served at that time interval. We have the relation

$$X = 1 + \sum_{i=1}^{B(j+1)} X_i,$$

where X_i denotes the family size of the i 'th customer arriving in the time interval j .

The important observation now is that the family sizes are iid because the $B(j)$ are iid, and that the X_i are independent of $B(j+1)$. Consequently, for the common probability generating function y of X and the X_i , we have

$$(*) \quad y(x) = x g_B(y(x)).$$

The length of the first busy period coincides with sum of the family sizes of the a customers arriving in the first time interval. Thus, we obtain

$$(**) \quad h(x) = y(x)^a.$$

Solving $(*)$ for x and inserting into $(**)$, we obtain

$$(***) \quad h\left(\frac{y(x)}{g_B(y(x))}\right) = y(x)^a.$$

If $y(0) = 0$, then $B = 0$, and thus $h(y) = y^a$, which coincides with equation (57a). Otherwise, by $(***)$, equation (57a) holds for all y in the interval $[y(0), y(1)]$, and thus for all y for which the power series on both sides of the equality sign converge.

We derive the statement about the mean length of the first busy period by differentiating (57a), and possibly invoking Abel's Theorem to evaluate the power series at the point 1.

Finally, the statement about the tail probability follows directly from the standard exponential moment argument: If $y \geq g_B(y) > 0$, then, with $x := y/g_B(y) \geq 1$, we have

$$\mathbf{P}[Z \geq \alpha] = \mathbf{P}[x^Z \geq x^\alpha] \leq \frac{\mathbf{E} x^Z}{x^\alpha} = \frac{h(x)}{x^\alpha} = \frac{y^a}{(y/g_B(y))^\alpha} = \frac{g_B(y)^\alpha}{y^{\alpha-a}},$$

as claimed.

10.8.3. Computations for Lemma 10.15.

Computations regarding equation $(**)$:

$$\begin{aligned} \alpha r + \frac{12^3}{13^{2.5}} \alpha r^2 u - (\alpha - a) \frac{1}{u+1} &= 0 \\ \alpha r(u+1) + \frac{12^3}{13^{2.5}} \alpha r^2 u(u+1) - (\alpha - a) &= 0 \\ \left(\frac{12^3}{13^{2.5}} \alpha r^2\right) u^2 + \left(\alpha r + \frac{12^3}{13^{2.5}} \alpha r^2\right) u - ((1-r)\alpha - a) &= 0 \\ u = - \frac{\left(\alpha r + \frac{12^3}{13^{2.5}} \alpha r^2\right) \pm \sqrt{\left(\alpha r + \frac{12^3}{13^{2.5}} \alpha r^2\right)^2 + 4((1-r)\alpha - a) \left(\frac{12^3}{13^{2.5}} \alpha r^2\right)}}{2 \cdot \left(\frac{12^3}{13^{2.5}} \alpha r^2\right)} \end{aligned}$$

We need to be close to 0, so we take the “ \pm ” = “ $-$ ”:

$$\begin{aligned}
u_r &:= \frac{-\left(\alpha r + \frac{12^3}{13^{2.5}}\alpha r^2\right) + \sqrt{\left(\alpha r + \frac{12^3}{13^{2.5}}\alpha r^2\right)^2 + 4((1-r)\alpha - a)\frac{12^3}{13^{2.5}}\alpha r^2}}{2 \cdot \frac{12^3}{13^{2.5}}\alpha r^2} \\
&= \frac{-\left(1 + \frac{12^3}{13^{2.5}}r\right) + \sqrt{\left(1 + \frac{12^3}{13^{2.5}}r\right)^2 + 4(1-r-a/\alpha)\frac{12^3}{13^{2.5}}}}{2 \cdot \frac{12^3}{13^{2.5}}r} \\
&= \frac{-\left(1 + \frac{12^3}{13^{2.5}}r\right) + \sqrt{\left(1 + \frac{12^3}{13^{2.5}}r\right)^2 - 4r\frac{12^3}{13^{2.5}} + 4(1-a/\alpha)\frac{12^3}{13^{2.5}}}}{2 \cdot \frac{12^3}{13^{2.5}}r} \\
&= \frac{-\left(1 + \frac{12^3}{13^{2.5}}r\right) + \sqrt{\left(1 - \frac{12^3}{13^{2.5}}r\right)^2 + 4(1-a/\alpha)\frac{12^3}{13^{2.5}}}}{2 \cdot \frac{12^3}{13^{2.5}}r} \\
&= \frac{-\left(1 + \frac{12^3}{5 \cdot 13^2}r\right) + \sqrt{\left(1 - \frac{12^3}{5 \cdot 13^2}r\right)^2 + \frac{4 \cdot 12^3}{5 \cdot 13^2} - \frac{4 \cdot 12^3}{5 \cdot 13^2} \cdot \frac{a}{\alpha}}}{\frac{2 \cdot 12^3}{5 \cdot 13^2}r} \\
&= \frac{-\left(1 + \frac{12^3}{5 \cdot 13^2}r\right) + \sqrt{\left(1 - \frac{12^3}{5 \cdot 13^2}r\right)^2 + \frac{4 \cdot 12^3}{5 \cdot 13^2}}}{\frac{2 \cdot 12^3}{5 \cdot 13^2}r} - O(a/\alpha),
\end{aligned}$$

with an absolute constant in the $O(\cdot)$, because $a \leq \alpha$ and $1/2 \leq r \leq 1$.

Computation regarding equation (***):

$$\begin{aligned}
\frac{(*)}{\alpha}(u_r) &= \frac{\alpha r u + \frac{12^2 \cdot 3 \cdot 2}{13^{2.5}}\alpha r^2 u^2 - (\alpha - a) \log(u + 1)}{\alpha} \Big|_{u:=u_r} \\
&= r u_r + \frac{6 \cdot 12^2}{5 \cdot 13^2} r^2 u_r^2 - (1 - a/\alpha) \log(u_r + 1) \\
&= r u_r + \frac{6 \cdot 12^2}{5 \cdot 13^2} r^2 u_r^2 - \log(u_r + 1) + O(a/\alpha),
\end{aligned}$$

with an absolute constant in the $O(\cdot)$, because $u_r + 1 \leq 2$.

10.8.4. Proof of Lemma 10.18. Suppose that the $B(j)$ are represented as a sum as in (62) above, and define

$$B^\pm(j) := \sum_{i=1}^{m^\pm} \mathbf{I}\left[U(j, i) \leq \frac{P(j)}{n - (\pm z)}\right].$$

Then the $B^+(j)$, $j = 1, 2, 3, \dots$, are iid, so that Lemma 10.17 is applicable. The same is true for the $B^-(j)$, $j = 1, 2, 3, \dots$. We clearly have, with probability $1 - O(n^{-r})$,

$$B^-(j) \leq B(j) \leq B^+(j) \quad \text{for all } j = 1, \dots, z.$$

Defining two queues $Q^\pm(j)$ based on the $B^\pm(j)$ and respective lengths of first busy periods Z^\pm , we obtain, with probability $1 - O(n^{-r})$

$$(*) \quad Z^- \leq Z \leq Z^+,$$

where we have also used that $Z^\pm \leq z$ with probability $1 - O(n^{-r})$ (Lemma 10.17).

Denote by E the event that $(*)$ holds. If $(*)$ does not hold, we still have $Z = O(n)$ by (63), so that we obtain

$$\mathbf{E} Z = \mathbf{E}(Z | E) \mathbf{P}(E) + \mathbf{E}(Z | \bar{E}) \mathbf{P}(\bar{E}) \leq \mathbf{E}(Z^+ | E) \mathbf{P}(E) + O(n^{1-r}) \leq \mathbf{E}(Z^+) + O(n^{1-r}).$$

For the lower bound, we similarly have

$$\mathbf{E} Z \geq \mathbf{E}(\mathbf{I}(E)Z^-) = \mathbf{E}(Z^-) - \mathbf{E}(\mathbf{I}(\overline{E})Z^-)$$

Clearly, $\mathbf{E}(\mathbf{I}(\overline{E})Z^-) \leq z \mathbf{P}(\overline{E}) + \mathbf{E}(\mathbf{I}(\overline{E})Z^- \mathbf{I}[Z^- > z]) = z \mathbf{P}(\overline{E}) + mO(n^{-r}) = O(n^{1-r})$. Thus we conclude that $\mathbf{E} Z \geq \mathbf{E} Z^- - O(n^{1-r})$.

For the tail estimate, we use Z^+ :

$$\begin{aligned} \mathbf{P}[Z \geq \alpha] &\leq \mathbf{P}[Z \geq \alpha \ \& \ Z \leq Z^+] + \mathbf{P}[Z \geq \alpha \ \& \ Z > Z^+] \\ &\leq \mathbf{P}[Z^+ \geq \alpha] + \mathbf{P}[Z > Z^+] \leq e^{-\delta\alpha} + O(n^{-r}) \end{aligned}$$

by Lemma 10.17.

10.8.5. Computations for Lemma 10.26.

For the proof of Claim (b). Let $g(x, y) := \frac{-18cx^4 + 2y(12x - y)}{x(12x - y)}$, the right hand side of the ODE (73a). As mentioned in the proof of the lemma, we show $g(x, y(x)) > 6 = \frac{dy}{dx}$, for $0 < x \leq 4/5$. We compute

$$\begin{aligned} g(x, y(x)) &= \frac{-18cx^4 + 2 \cdot 6x(12x - 6x)}{x(12x - 6x)} = \frac{-18cx^2 + 2 \cdot 6(12 - 6)}{(12 - 6)} = \frac{-18cx^2 + 72}{6} \\ &= -3cx^2 + 12 \underset{c \leq 3}{\geq} -9x^2 + 12 \geq -9(4/5)^2 + 12 = 12 - \frac{9 \cdot 16}{25} = \frac{25 \cdot 12 - 9 \cdot 16}{25} \\ &= \frac{12(25 - 3 \cdot 4)}{25} = \frac{12 \cdot 13}{25} > 6. \end{aligned}$$

For the proof of Claim (a). Let $g(x, y)$ as above. As mentioned in the proof of the lemma, we show $g(x, y(x)) > -5 = \frac{dy}{dx}$, for $4/5 \leq x \leq 1$. To show that

$$g(x, y(x)) = \frac{-18cx^4 + 2 \cdot 5(1 - x)(12x - 5(1 - x))}{x(12x - 5(1 - x))} > -5,$$

we compute

$$\begin{aligned} &-18cx^4 + 2 \cdot 5(1 - x)(12x - 5(1 - x)) + 5x(12x - 5(1 - x)) \\ &= -18cx^4 + 10(1 - x)(17x - 5) + 5x(17x - 5) = -18cx^4 + (10 - 5x)(17x - 5) \\ &= -18cx^4 - 85x^2 + 195x - 50 \underset{c \leq 3}{\geq} -54x^4 - 85x^2 + 195x - 50. \end{aligned}$$

The derivative $-216x^3 - 170x + 195$ of the last polynomial is strictly decreasing, and evaluating it at $4/5$ gives $-216(4/5)^3 - 170 \cdot 4/5 + 195 \approx -51.592 < 0$. Thus, it suffices to check the inequality $-54x^4 - 85x^2 + 195x - 50 > 0$ for $x = 1$: $-54 - 85 + 195 - 50 = 6 > 0$.

Cops & Robber on non-orientable surfaces

Jointly with

Nancy E. Clarke (Acadia University, Wolfville, NS),
 Samuel Fiorini (U.L.B., Brussels),
 and Gwenaël Joret (U.L.B., Brussels)

Abstract. We consider the two-player, complete information game of Cops and Robber played on undirected, finite, reflexive graphs. A number of cops and one robber are positioned on vertices and take turns in sliding along edges. The cops win if, after a move, a cop and the robber are on the same vertex. The minimum number of cops needed to catch the robber on a graph is called the cop number of that graph.

Let $\text{cop}(g)$ be the supremum over all cop numbers of graphs embeddable in a closed orientable surface of genus g , and likewise $\widetilde{\text{cop}}(g)$ for non-orientable surfaces. It is known (Andreae, 1986) that, for a fixed surface, the maximum over all cop numbers of graphs embeddable in this surface is finite. More precisely, Quilliot (1985) showed that $\text{cop}(g) \leq 2g + 3$, and Schröder (2001) sharpened this to $\text{cop}(g) \leq \frac{3}{2}g + 3$. In his paper, Andreae gave the bound $\widetilde{\text{cop}}(g) \in O(g)$ with a weak constant, and posed the question whether a stronger bound can be obtained. Nowakowski & Schröder (1997) obtained $\widetilde{\text{cop}}(g) \leq 2g + 1$.

In this short note, we show $\widetilde{\text{cop}}(g) \leq \text{cop}(g - 1)$, for any $g \geq 1$. As a corollary, using Schröder's results, we obtain the following: the maximum cop number of graphs embeddable in the projective plane is 3; the maximum cop number of graphs embeddable in the Klein Bottle is at most 4, $\widetilde{\text{cop}}(3) \leq 5$, and $\widetilde{\text{cop}}(g) \leq \frac{3}{2}g + 3/2$ for all other g .

For an integer $k \geq 1$, the *Cops and Robber game with k cops* is a pursuit game played on a reflexive graph, i.e. a graph with a loop at every vertex. There are two opposing sides, a set of k cops and a single robber. The cops begin the game by each choosing a (not necessarily distinct) vertex to occupy, and then the robber chooses a vertex. The two sides move alternately, where a move is to slide along an edge or along a loop. The latter is equivalent to passing were the game played on a loopless graph. There is perfect information, and the cops win if any of the cops and the robber occupy the same vertex at the same time, after a finite number of moves. Graphs on which one cop suffices to win are called *copwin* graphs. In general, we say that a graph G is *k -copwin* if k cops can win on G . The minimum number of cops that suffice to win on G is the cop number of G , denoted $c(G)$. The game has been considered on infinite graphs but, here, we only consider finite graphs.

Nowakowski & Winkler [NW83] and Quilliot [Qui78] have characterized the class of copwin graphs. The class of k -copwin graphs, $k > 1$, has been characterized by Clarke and MacGillivray [CM11]. Families of graphs with unbounded cop number have been constructed [AF84], even families of d -regular graphs, for each $d \geq 3$ [And84].

By a surface, we mean a closed surface, i.e. a compact two dimensional topological manifold without boundary. For any non-negative integer g , we denote by $\text{cop}(g)$ the supremum over all $\text{cop}(G)$, with G ranging over all graphs embeddable in an orientable surface of genus g , and we call this the cop number of the surface. Similarly, we define the cop number $\widetilde{\text{cop}}(g)$ of

a non-orientable surface of genus g to be the supremum over all $\text{cop}(G)$, with G ranging over all graphs embeddable in this surface.

Aigner & Fromme [AF84] proved that the cop number of the sphere is equal to three; i.e. $\text{cop}(0) = 3$. Quilliot [Qui85] gave an inductive argument to the effect that the cop number of an orientable surface of genus g is at most $2g + 3$. Schröder [Sch01] was able to sharpen this result to $\text{cop}(g) \leq \frac{3}{2}g + 3$. He also proved that the cop number of the double torus is at most 5.

Andreae [And86] generalized the work of Aigner & Fromme. He proved that, for any graph H satisfying a mild connectivity assumption, the class of graphs which do not contain H as a minor has cop number bounded by a constant depending on H . Using this, and the well known formula for the non-orientable genus of a complete graph, he obtained an upper bound for the cop number of a non-orientable surface of genus g , namely

$$\widetilde{\text{cop}}(g) \leq \left(\frac{\lfloor 7/2 + \sqrt{6g + 1/4} \rfloor}{2} \right).$$

Nowakowski & Schröder [NS] use a series of technically challenging arguments to prove a much stronger bound: $\widetilde{\text{cop}}(g) \leq 2g + 1$.

In this short note, we prove the following.

Theorem 11.1. *For any positive integer g , $\text{cop}(\lfloor g/2 \rfloor) \leq \widetilde{\text{cop}}(g) \leq \text{cop}(g - 1)$.*

This immediately improves the best known upper bound for the non-orientable surface of genus g to $\widetilde{\text{cop}}(g) \leq \frac{3}{2}(g - 1) + 3 = \frac{3}{2}(g + 1)$. The following table gives the new and status quo for the concrete upper bounds.

N/o genus	1	2	3	4	5	6	7
N. & S. [NS]	3	5	7	9	11	13	15
Here	3	4	5 ¹	7	9	10	12

TABLE 1. Comparison of the new and status quo upper bounds for $\widetilde{\text{cop}}(g)$.

We say that a *weak cover* of H by G is a surjective mapping $p: V(G) \rightarrow V(H)$ which maps vertex neighborhoods onto vertex neighborhoods; i.e. for every vertex u of G , we have $p(N(u)) = N(p(u))$. (This terminology lends on the classical definition of a “cover” without weak, where the restriction to the vertex neighborhood $p: N(u) \rightarrow N(p(u))$ is required to be a bijection.) Using the same technique as for the inequality “ \leq ” in the proof of Theorem 11.1, it is possible to show the following:

Lemma 11.2. *If G is a weak cover of H , then $\text{cop}(H) \leq \text{cop}(G)$.*

This is similar in spirit to the seminal result of Berarducci & Ingridila [BI93], saying that if H is a retract of G , then the same inequality holds. Note, however, that neither of the two notions generalizes the other. We will not prove Lemma 11.2; the proof is only slightly more technical than the geometric proof of Theorem 11.1.

11.1. Proof

Familiarity with the classification of combinatorial surfaces is assumed. See any standard textbook on topology. We will make use of the standard representation of surfaces as quotients of polygonal discs with labelled and directed edges. Each label occurs twice, and the two edges with the same label are identified according to their orientations. Reading the labels of the edges in counterclockwise (i.e. *positive*) order and adding an exponent -1 whenever the orientation of the edges is negative (i.e. clockwise) gives the *word* of the surface.

For a graph G , let $\gamma(G)$ denote the smallest integer g such that G can be embedded in an orientable surface of genus g ; similarly define $\tilde{\gamma}(G)$ as the smallest integer g such that G can be embedded in a non-orientable surface of genus g . For the proof of Theorem 1, we use the following well-known fact. (The proof can be found in [MT01]).

Lemma 11.3 (Folklore). *For any graph G , $\tilde{\gamma}(G) \leq 2\gamma(G) + 1$.*

In the proof of the inequality $\text{cop}(g) \leq \text{cop}(g-1)$, we make use of the well-known fact that every manifold X has a 2-sheeted covering $X' \rightarrow X$ by an orientable manifold. If X is a non-orientable surface of genus g , it is a textbook exercise to see that the standard construction yields a surface of genus $g-1$. This is Lemma 11.4. The proof is straightforward (consider Figure 1), and is thus omitted.

Lemma 11.4. *A non-orientable surface of genus g has an orientable surface of genus $g-1$ as a 2-sheeted covering space.*

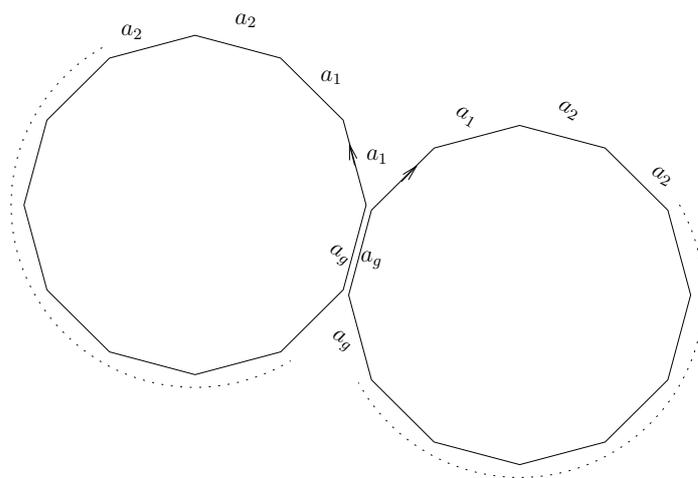


FIGURE 11.1. A figure to accompany Lemma 4.

We are now ready for the proof of our main result.

Proof of Theorem 1. Lemma 3 immediately implies that $\text{cop}(g) \leq \widetilde{\text{cop}}(2g+1)$, and hence $\widetilde{\text{cop}}(g) \geq \text{cop}(\lfloor g/2 \rfloor)$.

For the proof of the remaining inequality $\widetilde{\text{cop}}(g) \leq \text{cop}(g-1)$, let X be the non-orientable surface of genus g on which a graph G is embedded. We identify the graph G with its embedding; i.e. we think of the vertex set $V(G)$ as a set of points of X and the edge set of $E(G)$ as a set of internally disjoint injective curves connecting the respective end vertices of the edge.

By Lemma 11.4, there exists a covering $p: X' \rightarrow X$ of X by an orientable surface X' with genus $g' := g-1$. Consider the graph G' whose vertex set is $\{p^{-1}(V(G))\}$ and whose edge set consists of the curves obtained by lifting the edges of G . By construction, G' is embedded in the orientable surface X' of genus g' .

We now give a strategy for $k := \text{cop}(g')$ cops to win the Cops and Robber game on G , by “simulating” a game on G' and using any winning strategy for k cops on this graph, who chase an “imaginary” robber. In such a strategy, the k cops first choose their starting vertices $u_1, \dots, u_k \in V(G')$. In the strategy for G , we let the starting vertices be $p(u_1), \dots, p(u_k)$. Suppose now that, in the game on G , the robber chooses a starting vertex r . We choose an arbitrary starting vertex for an imaginary robber on G' arbitrarily in the fibre $p^{-1}(r)$.

Throughout the game, the position of each player in G' will be in the fibre $p^{-1}(x)$ of the position x of the corresponding player in G . Moreover, the movements of the players on G describe curves on X , which can be lifted (uniquely, although this is not essential) to curves on X' forming walks in G' .

Now, whenever it be the cops' turn in any game on G , the robber is at a certain vertex s of G' , and the k cops are on vertices v_1, \dots, v_k . The strategy for the cops on G' now prescribes moves for the cops. The corresponding moves in G are then given as images under p .

Since we have a winning strategy, after a finite number of moves, the "imaginary robber" on G' will be on the same vertex as a cop in G' . Consequently, the same holds on G , and thus the cops have won the game on G . \square

11.2. Conclusion

We conclude with a conjecture.

Conjecture. *For a non-negative integer g , $\widetilde{\text{cop}}(g) = \text{cop}(\lfloor g/2 \rfloor)$.*

One might wonder whether it is possible to improve Theorem 11.1 by taking a different covering, or possibly a branched covering. This is impossible: It is a well-known fact that, whenever $p: X' \rightarrow X$ is a (branched) covering with X' orientable and X non-orientable, then p lifts to a (branched) covering $\tilde{p}: X' \rightarrow \tilde{X}$, where \tilde{X} is the orientable double cover constructed in Lemma 11.4.

Cops & Robber on graphs with forbidden (induced) subgraphs

Jointly with
Gwenaël Joret (U.L.B., Brussels)
and Marcin Kamiński (U.L.B., Brussels)

Abstract. The two-player, complete information game of Cops and Robber is played on undirected finite graphs. A number of cops and one robber are positioned on vertices and take turns in sliding along edges. The cops win if, after a move, a cop and the robber are on the same vertex. The minimum number of cops needed to catch the robber on a graph is called the cop number of that graph.

In this paper, we study the cop number in the classes of graphs defined by forbidding one or more graphs as either subgraphs or induced subgraphs. In the case of a single forbidden graph we completely characterize (for both relations) the graphs which force bounded cop number. En passant, we bound the cop number in terms of the tree-width.

12.1. Introduction

Graphs studied in this paper are finite, undirected, without loops and multiple edges. We use standard notation and terminology; for what is not defined here, we refer the reader to Diestel [Die06].

The game of *Cops and Robber* is played on a connected graph by two players – the cops and the robber. The cop player has at her disposal k pieces (cops), for some integer $k \geq 1$, and the robber player has only one piece (the robber). The game begins with the cop player placing her k cops on (not necessarily distinct) vertices of the graph. Next, the robber player chooses a vertex for his piece. Now, starting with the cop player, the two players move their pieces alternately. In the cops’ move, she decides for each of her cops whether it stands still or is moved to an adjacent vertex. In the robber’s move, he can choose to move or not to move the piece. The game ends when a cop and the robber are on the same vertex (that is, the cops catch the robber); in this case the cop player wins. The robber wins if he can never be caught by the cops. Both players have complete information, that is, they know the graph and the positions of all the pieces.

The key problem in this game is to know how many cops are needed to catch a robber on a given graph. For a connected graph G , the smallest integer k such that with k cops, the cop player has a winning strategy is called the *cop number* of G and is denoted by $\text{cop}(G)$. We follow Berarducci and Intrigila [BI93] in defining the cop number of a non-connected graph as the maximum cop number of its connected components. Nowakowski and Winkler [NW83] and Quilliot [Qui78] characterized the class of graphs with cop number 1. Finding a combinatorial characterization of graphs with cop number k (for $k \geq 2$) is a major open problem in the field. On the other hand, algorithmic characterizations of such graphs, which are polynomial in the size of the graph but not in k , do exist [BI93, GR95, HM06]. However, determining the cop number of a graph is a computationally hard problem [FGK08]. For literature review we refer the reader to a recent survey on graph searching [FT08] (see also [Hah07]).

In this paper we study the cop number for different types of graph classes. Our motivation is to learn what structural properties of graphs force the cop number to be bounded. (We say

that the cop number is *bounded* for a class of graphs, if there exists a constant C such that the cop number of every graph from the class is at most C ; otherwise the cop number is *unbounded* for this class.) We consider several containment relations and study the cop number for classes of graphs with a single forbidden graph with respect to these relations.

Families of graphs with unbounded cop number have been constructed [AF84]. For every fixed $d \geq 3$, there even exist families of d -regular graphs with unbounded cop number [And84]. On the other hand, Aigner and Fromme [AF84] proved that the cop number of a planar graph is at most 3. This result has been generalized to the class of graphs with genus g ; Schroeder [Sch01] proved that the cop number of a graph is bounded by $\lfloor 3g/2 \rfloor + 3$ (improving an earlier bound of Quilliot [Qui85]), and conjectured that this bound can be reduced to $g + 3$.

A graph is called *H -minor-free* (*H -topological minor-free*) if it does not contain H as a minor (as a topological minor). Andreae [And86] studied classes of H -minor-free graphs and showed that the cop number of a K_5 -minor-free graph (or $K_{3,3}$ -minor-free graph) is at most 3. Since a planar graph does not have a K_5 or $K_{3,3}$ as a minor this result extends the result on planar graphs. However, for our purposes the most interesting result of Andreae [And86] is that for any graph H the cop number is bounded in the class of H -minor-free graphs. In other words, forbidding a minor is enough to bound the cop number.

Andreae [And86] also observed that excluding a topological minor does not necessarily bound the cop number. In fact, it is an easy corollary of his work that the class of H -topological minor-free graphs has bounded cop number if and only if the maximum degree of H is at most 3.

Inspired by these results we study other containment relations: subgraphs and induced subgraphs. A graph is called *H -subgraph-free* (*H -free*) if it does not contain H as a subgraph (as an induced subgraph). We give necessary and sufficient conditions for the class of H -subgraph-free graphs and H -free graphs to have bounded cop number. First we present our results for induced subgraphs.

Theorem 12.1. *The class of H -free graphs has bounded cop number if and only if every connected component of H is a path.*

Let us remark that a single vertex is considered to be a path. The graph consisting of a path on ℓ ($\ell \geq 1$) vertices is denoted by P_ℓ . The backward implication of Theorem 12.1 is a consequence of the following proposition.

Proposition 12.2. *For every $\ell \geq 3$, every P_ℓ -free graph has cop number at most $\ell - 2$.*

Using the same technique, it is in fact possible to show the following stronger result.

Proposition 12.3. *For every $\ell \geq 3$, every graph with no induced cycle of length at least ℓ has cop number at most $\ell - 2$.*

Notice that it is possible to rephrase the condition of Theorem 12.1 and say that every connected component of H is a tree with at most two leaves. Here is our result for H -subgraph-free graphs.

Theorem 12.4. *The class of H -subgraph-free graphs has bounded cop number if and only if every connected component of H is a tree with at most three leaves.*

It is easy to see that the cop number of a tree is 1. As an intermediate step towards Theorem 12.4, we study how the cop number of a graph G is related to its tree-width, which is denoted by $\text{tw}(G)$.

Proposition 12.5. *The cop number of a graph G is at most $\text{tw}(G)/2 + 1$.*

This bound is sharp for tree-width up to 5. (This is easy to prove for $\text{tw}(G) \leq 3$; for $\text{tw}(G) = 4$ and 5, the Petersen graph and the disjoint union of the Petersen graph and K_6 plus an edge linking them are tight examples, respectively.)

12.2. Forbidding induced subgraphs

Our goal in this section is to prove Theorem 12.1. Notice that a graph whose every connected component is a path, is an induced subgraph of some sufficiently long path. Hence, the following proposition proves the backward implication of Theorem 12.1.

Proposition 12.2. *For every $\ell \geq 3$, every P_ℓ -free graph has cop number at most $\ell - 2$.*

Let us remark that, for $\ell = 1, 2$, the cop number of a P_ℓ -free graph is trivially 1.

PROOF OF PROPOSITION 12.2. Let G be a P_ℓ -free graph and let us also assume, without loss of generality, that G is connected. We will give a winning strategy for $\ell - 2$ cops. Initially all $\ell - 2$ cops are on the same arbitrary vertex. The strategy is divided into stages. The distance between the cops and the robber is the minimum distance from the robber to a cop. The goal of each stage is to decrease the distance between the cops and the robber. Once the distance is decreased we begin the next stage. We will show that a stage lasts a finite number of rounds.

At the beginning of each stage we choose a *lead cop* (for this stage) among the pieces which are at the minimum distance from the robber. All distances in this proof are measured after the robber's and before the cops' move. We route the lead cop and instruct the other pieces to follow the lead cop in single file; the cops should form a path of length $\ell - 2$.

If the distance between the cops and the robber is at most one, then the cops clearly win. Suppose that the distance between the lead cop on vertex x and the robber on vertex y is $d \geq 2$. We order the lead cop to travel along the shortest path from x to y and then follow the exact route the robber took from vertex y . Notice that since the graph is P_ℓ -free the distance between the cops and the robber will decrease after at most $\ell - d - 1$ moves. Once the distance decreased, we move to the next stage. \square

We mention the following result which can be derived using almost the same strategy as in Proposition 12.2.

Proposition 12.3. *For every $\ell \geq 3$, every graph with no induced cycle of length at least ℓ has cop number at most $\ell - 2$.* \square

Before completing the proof of Theorem 12.1, we look at bipartite graphs with no long induced paths. A simple modification of the proof of Proposition 12.2 yields a better bound for the bipartite case. Here is how the cops' strategy needs to be modified: the cops follow the lead cop in such a way that the distance between any two consecutive cops is 2. We leave the details of this proof to the reader.

Proposition 12.4. *For every $\ell \geq 1$, every $P_{2\ell}$ -free bipartite graph has cop number at most ℓ .*

To prove the forward implication of Theorem 12.1, we need to introduce two graph operations which do not decrease the cop number: clique substitution and edge subdivision. Let $N(v)$ be the set of neighbors of a vertex v . A *clique substitution* at a vertex v consists in replacing v with a clique of size $|N(v)|$ and creating a matching between vertices of the clique and the vertices of $N(v)$. The graph obtained from a graph G by substituting a clique at each vertex of G will be denoted by G^+ . More formally, the vertex set of G^+ is $\bigcup_v (\{v\} \times N(v))$ and two vertices (v_1, u_1) and (v_2, u_2) are adjacent if and only if $v_1 = v_2$, or $v_1 = u_2$ and $u_1 = v_2$.

Lemma 12.5. *Clique substitution does not decrease the cop number.*

PROOF. Let G be a graph. To each vertex $v \in V(G)$ there corresponds a clique in G^+ , which we denote by $\phi(v)$. We simultaneously play two games: one on G and another on G^+ . We assume that we have a winning strategy for the cop player on G^+ . We use the same number of cops in G as in G^+ .

At the beginning, the cops are placed on G^+ according to the strategy, and the corresponding cops in G are placed in the obvious way: If a cop in G^+ is on a vertex of the clique $\phi(v)$

for some vertex $v \in V(G)$, then the corresponding cop in G is put on vertex v . Then, we put the robber in G^+ on an arbitrary vertex of the clique $\phi(v)$, where v is the vertex on which the robber is in G . For simplicity of presentation, we do not move the cops at all in G during the first turn. Thus, the robber will move first.

Now, let us consider a robber's move in G , say from vertex u to vertex v . In G^+ , the robber is on some vertex of $\phi(u)$. If $u = v$, we do not move the robber in G^+ . Next, suppose $u \neq v$, and let $u'v'$ be the (unique) edge between $\phi(u)$ and $\phi(v)$, with $u' \in \phi(u)$ and $v' \in \phi(v)$. If in G^+ , the robber is on u' , we move it to v' . Otherwise, the robber is on another vertex of $\phi(u)$, and we move it first to u' , then let the cops react to that move, and finally move the robber to v' (unless it has been caught). Once the robber is in its final position, we let the cops move in G^+ . We refer to this sequence of 1 or 2 turns in G^+ as a *stage*.

Once a stage is finished in G^+ , we translate the moves of the cops back to the graph G : Consider a cop in G^+ . Let u and v be the vertices of G such that the cop was in the clique $\phi(u)$ at the end of the previous stage and in the clique $\phi(v)$ at the end of the current stage. Observe that, either $u = v$, or $uv \in E(G)$. We move the corresponding cop in G from u to v (or let it stay on u if $u = v$).

This describes our strategy for the cops in G . By our assumption, the robber will be caught during some stage in G^+ . At the end of that stage, both a cop and the robber on the clique $\phi(v)$ for some vertex $v \in V(G)$. Hence, when the moves of cops from that stage are translated back to G , the corresponding cop in G will be on the same vertex as the robber. \square

The *claw* is the complete bipartite graph with sides of size 1 and 3. The operation of clique substitution will be used to show that the cop number of claw-free graphs is unbounded.

Lemma 12.6. *The class of claw-free graphs has unbounded cop number.*

PROOF. Let \mathcal{G} be a class of graphs with unbounded cop number and $\mathcal{G}^+ := \{G^+ \mid G \in \mathcal{G}\}$. Notice that all graphs in \mathcal{G}^+ are claw-free. Applying Lemma 12.5, we see that the cop number of graphs in \mathcal{G}^+ is unbounded. \square

The other graph operation needed for the proof of Theorem 12.1 is edge subdivision. Berarducci and Intrigila [BI93] proved the following lemma.

Lemma 12.7 ([BI93]). *Subdividing all edges of a graph an even number of times does not decrease the cop number.*

This leads to the following result. Recall that the *girth* of a graph is the length of its shortest cycle if it has one, $+\infty$ otherwise.

Lemma 12.8. *For every integer $\ell \geq 3$, the class of graphs with girth at least ℓ has unbounded cop number.*

PROOF. Let \mathcal{G} be an arbitrary class of graphs with unbounded cop number. For every $G \in \mathcal{G}$, let G' be a graph with girth at least ℓ obtained from G by subdividing all edges sufficiently often. Let $\mathcal{G}' := \{G' \mid G \in \mathcal{G}\}$. Applying Lemma 12.7, we see that the class \mathcal{G}' has unbounded cop number. \square

Now we are ready to complete the proof of Theorem 12.1.

Theorem 12.1. *The class of H -free graphs has bounded cop number if and only if every connected component of H is a path.*

PROOF. The backward implication of the theorem follows from Proposition 12.2. Indeed, notice that if every connected component of H is a path, then H is a subgraph of the path on $|H| + p - 1$ vertices, where p is the number of connected components of H . Hence, the cop number of an H -free graph is bounded by $\max\{|H| + p - 3, 1\}$.

Now we will prove the forward implication of the theorem. Let H be a graph such that the class of H -free graphs has bounded cop number. Suppose that H contains a cycle and let ℓ be the length of the longest cycle of H . Clearly, the class of graphs with no induced cycle of length at most ℓ is contained in the class of H -free graphs. However, by Lemma 12.8 the class of graphs with no induced cycle of length at most ℓ has unbounded cop number; a contradiction. Hence, H is a forest.

Now suppose that H contains a vertex of degree at least 3. Since H is a forest, it must contain a claw as an induced subgraph. Clearly, the class of claw-free graphs is contained in the class of H -free graphs. However, by Lemma 12.6 the class of claw-free graphs has unbounded cop number; a contradiction. Hence, H is a forest of maximum degree at most 2, that is, H is a disjoint union of paths. \square

We note that in the second part of the proof (removing cycles) we could have used some known constructions which show that graphs simultaneously having an arbitrarily large girth and large cop number do exist; see for instance Andreae [And84] and Frankl [Fra87].

12.2.1. Some remarks about edge subdivisions. Lemma 12.7 by Berarducci and Intrigila [BI93] gives a bound on the cop number of graphs which result by uniformly subdividing all edges an even number of times. By modifying the proof of Lemma 12.5, the following can be shown.

Lemma 12.9. *Subdividing all edges of a graph once does not decrease the cop number.*

Combining Lemmas 12.7 and 12.9 we obtain the general result.

Corollary 12.10. *For every positive integer r , subdividing every edge of a graph r times does not decrease the cop number.*

PROOF. Let G be a graph. The proof is by induction on r . The base case $r = 1$ is given by Lemma 12.9. For the inductive case, assume $r \geq 2$. If r is even, then the claim follows from Lemma 12.7. If r is odd, then by induction subdividing each edge of G $(r - 1)/2$ times does not decrease the cop number. Subdividing once every edge of the resulting graph we obtain a subdivision of G where each edge has been subdivided $(r - 1)/2 + ((r - 1)/2 + 1) = r$ times, and its cop number is at least that of G by Lemma 12.9. \square

Berarducci and Intrigila [BI93] noted that subdividing edges in a non-uniform manner can both increase and decrease the cop number. However, for uniform subdivisions it is possible to give an estimate.

Proposition 12.11. *Subdividing each edge r times increases the cop number by at most one.*

PROOF (SKETCH). Denote by \tilde{G} the graph which results from the graph G by subdividing each edge r times. A winning strategy for $\text{cop}(\tilde{G}) + 1$ cops on \tilde{G} is the following. Let an auxiliary cop pursue the strategy described for the lead cop in the proof of Proposition 12.2. By this we make sure that the robber cannot change his direction or pass in the middle of a subdivided edge except for a finite number of times. The other $\text{cop}(\tilde{G})$ cops simulate their winning strategy for G on \tilde{G} . \square

To further enlighten what happens if edges are subdivided, we propose the following construction. Let G be an arbitrary graph with n vertices and cop number at least 2. We construct a graph \hat{G} by adding paths of length n to G : every pair of non-adjacent vertices of G is joined by such a path. It is not difficult to see that $\text{cop}(\hat{G}) = \text{cop}(G)$. But by subdividing edges of \hat{G} , we can obtain a graph resulting from K_n by subdividing every edge n times. From Proposition 12.11 we know that the cop number of this graph is at most 2.

Considering this construction, it seems natural to propose the following conjecture, which implies the conjecture of Meyniel (see Frankl [Fra87]) that $\text{cop}(G)$ is in $O(\sqrt{|G|})$.

Conjecture 12.12. *For graphs G obtained by subdividing edges of complete graphs K_n we have $\text{cop}(G)$ in $O(\sqrt{n})$.*

12.3. Forbidding (not necessarily induced) subgraphs

We now turn our attention to classes of graphs for which we forbid (not necessarily induced) subgraphs. One key ingredient for the proof of Theorem 12.4 is the fact that families of graphs with bounded circumference have bounded cop number. Although this already follows from Proposition 12.3, in this section we give a better upper bound based on an estimate on the cop number in terms of the tree-width, which we believe to be of interest in its own.

Let us first briefly recall the definition of the tree-width of a graph. A *tree decomposition* of a graph G is a pair $(T, \{W_x \mid x \in V(T)\})$ where T is a tree, and $\{W_x \mid x \in V(T)\}$ a family of subsets of $V(G)$ (called “bags”) such that

- $\bigcup_{x \in V(T)} W_x = V(G)$;
- for every edge $uv \in E(G)$, there exists $x \in V(T)$ with $u, v \in W_x$, and
- for every vertex $u \in V(G)$, the set $\{x \in V(T) \mid u \in W_x\}$ induces a subtree of T .

The *width* of tree decomposition $(T, \{W_x \mid x \in V(T)\})$ is $\max\{|W_x| - 1 \mid x \in V(T)\}$. The *tree-width* $\text{tw}(G)$ of G is the minimum width among all tree decompositions of G . We refer the reader to Diestel’s book [Die06] for an introduction to the theory around tree-width.

Our proof of Proposition 12.5 relies on a well-known strategy for the cops and robber game: guarding a shortest path. Assume that P is a shortest uv -path, for two distinct vertices u, v of a graph G , and that a cop is sitting at the beginning on some vertex of P . The cop’s strategy consists in moving along P in such a way that his distance to u is as close as possible to the robber’s distance to u . It is easily seen that, after a finite number of initial moves, when it is the robber’s turn to play, the cop’s distance to u will be the same as the robber’s distance to u when the latter is no more than $|P|$. This ensures that the robber cannot go on any vertex of P without being caught. (This strategy has been first used by Aigner and Fromme [AF84], in their proof that the cop number of planar graphs is at most 3.)

Proposition 12.5. *The cop number of a graph G is at most $\text{tw}(G)/2 + 1$.*

PROOF. Let us consider an optimal tree decomposition of G . Since the tree-width of G equals the maximum tree-width of its connected components, we may assume without loss of generality that G is connected. For a bag $X \subseteq V(G)$ of the tree decomposition, we denote by t_X the vertex of T corresponding to X .

At the beginning, an arbitrary bag $B \subseteq V(G)$ of the tree decomposition is selected, and all its vertices are guarded in the following way: Letting b_1, b_2, \dots, b_k denote the vertices in B , we let the i th cop ($1 \leq i \leq \lfloor k/2 \rfloor$) guard a shortest $b_{2i-1}b_{2i}$ -path in G , and, if k is odd, we put an additional cop on vertex b_k . This ensures that, after a finite number of moves, the robber cannot go on any vertex in B , and hence is confined to (the subgraph corresponding to) some tree T' of $T \setminus t_B$. (We may assume $B \neq V(G)$, as otherwise the robber is trivially caught.)

Let $B' \subseteq V(G)$ be the unique bag of the tree decomposition that is adjacent to B in T with $B' \cap B \neq \emptyset$. Observe that $B \cap B'$ is a cutset of the graph G . We show that the cops can move in such a way that the vertices of $B \cap B'$ remains guarded, and after a finite number of moves all the vertices of B' (instead of B) are guarded.

Consider each cop. Suppose first that the cop sits on a vertex of $B \setminus B'$ or guards a shortest path between two vertices in $B \setminus B'$. Then we send him to guard a shortest path between two unguarded vertices in $B' \setminus B$ (or to sit on the last unguarded vertex if there is only one such vertex). Assume now that the cop sits on a vertex of $B \cap B'$ or guards a shortest $b_i b_j$ -path with $b_i \in B \cap B'$ and $b_j \in B \setminus B'$. Then the cop first goes to b_i (if he is not already there) along the path he keeps. Then he starts guarding an arbitrary $b_i b'_j$ -path, where b'_j is any unguarded vertex of $B' \setminus B$. Notice that, while it may take some moves before all the vertices of the path

are safely guarded, at least the vertex b_i is guarded at every time. Suppose finally that the cop guards a shortest $b_i b_j$ -path with $b_i, b_j \in B \cap B'$. In this case, the cop does not modify his strategy, and keeps guarding his path.

After a finite number of moves all the vertices in B' are guarded, and the robber did not have, at any time, the opportunity to go on a vertex in $B \cap B'$ without being caught. Moreover, the number of necessary cops is at most $\lceil |B'|/2 \rceil \leq \text{tw}(G)/2 + 1$, and the robber is reduced to stay in (the subgraph corresponding to) some tree of $T \setminus t_{B'}$ which is a proper subtree of T' . Therefore, by repeating this operation a finite number of times the robber will eventually be caught. This completes the proof. \square

We remark that the bound given in Proposition 12.5 is best possible for small values of the tree-width: For every $k = 1, 2, \dots, 5$, there are graphs with tree-width k and cop number $\lfloor k/2 \rfloor + 1$ (this is easily seen for $k = 1, 2, 3$, and the Petersen graph and the graph which is the disjoint union of the Petersen graph and a complete graph on 6 vertices are such examples for $k = 4$ and 5, respectively). On the other hand, we do not know whether there exists a constant $c > 0$ and an infinite family of graphs such that $\text{cop}(G) \geq c \cdot \text{tw}(G)$ holds for every graph G in the family.

Let us recall that the *circumference* of a graph is the length of its longest cycle if it has one, $+\infty$ otherwise.

Corollary 12.13. *The cop number of a graph is less than or equal to half its circumference.*

PROOF. It is a well-known fact that $\text{tw}(G) \leq \text{circum}(G) - 1$ holds for every graph G , where $\text{circum}(G)$ denotes the circumference of G (see for instance Exercise 12.18 in Diestel's book [Die06]). With Proposition 12.5, we conclude $\text{cop}(G) \leq \text{circum}(G)/2$. \square

Theorem 12.4. *The class of H -subgraph-free graphs has bounded cop number if and only if every connected component of H is a tree with at most three leaves.*

PROOF. We first show that the requirements in the statement of the theorem are necessary. Let H be a graph such that the family \mathcal{F} of connected graphs not containing H has bounded cop number.

First, suppose that H contains a cycle, and let ℓ be the length of a longest cycle in H . Then \mathcal{F} contains the family of connected graphs with girth at least $\ell + 1$. However, by Lemma 12.8, the cop number of this family is unbounded. Hence, H is a forest.

Second, suppose that H has a vertex of degree at least 4. This implies that \mathcal{F} contains all connected graphs with maximum degree 3, but Andreae [And84] proved that there exists a family of 3-regular graphs on which the cop number is unbounded. Hence, H has maximum degree at most 3.

Third, suppose that there is a tree in H which has two vertices of degree 3. Let ℓ denote the distance between these two vertices in H . Now \mathcal{F} contains the family of all those connected graphs in which every two vertices of degree 3 or more have distance at least $\ell + 1$. Starting from an arbitrary family of graphs on which the cop number is unbounded, a family with this property can be constructed by subdividing every edge ℓ times, as follows from Corollary 12.10. Thus, each connected component of H contains at most one vertex of degree 3.

We now show that any H meeting the conditions in the theorem yields a family of graphs with bounded cop number. The proof will be by induction on the number of connected components of H . For a single component, by Proposition 12.2, we may assume that a vertex of degree 3 does in fact exist. We will prove the following claim.

Claim. Let H be a tree with maximum degree 3 which has precisely one vertex v of degree 3. Denote by r the maximum distance of a vertex from v . If G does not contain H , then $\text{cop}(G) \leq 2r$.

Before we prove the claim, let us complete the induction. The start of the induction is settled. Let T be a connected component of H , and assume that $\text{cop}(G') \leq k$ for every graph G' not containing $H \setminus V(T)$. Let G be a graph not containing H . If G does not contain T , we are done by the claim and the remark preceding it. Otherwise, let T' be a subgraph of G isomorphic to T . We place $|T|$ cops on the vertices of T' . This corners the robber in a connected component of $G \setminus V(T')$. Noting that $G \setminus V(T')$ does not contain $H \setminus V(T)$, by induction, by restricting to the connected component containing the robber, we can catch the robber in $G \setminus V(T')$ using k cops. This bounds the cop number of G by $k + |T|$ and concludes the induction.

Proof of Claim. We prove the claim in the case when each leaf of H has distance exactly r from v . The general case follows easily from this.

By Proposition 12.2, we may assume that G contains a path P on $2r$ vertices, because otherwise we have $\text{cop}(G) \leq 2r$. We guard the path by placing r cops on every other vertex of P , and show that what remains of G has cop number at most r . Assume that $G \setminus V(P)$ contains a cycle C of length at least $2r + 1$. Then, since G is connected, we can identify a subgraph isomorphic to H choosing v to be a vertex on C which has minimum distance to a vertex in P , while two of the three branches of the tree are wound around C , the other extends to P . Hence, $G \setminus V(P)$ contains no such cycle. By invoking Corollary 12.13 for the connected component of $G \setminus V(P)$ containing the robber, we see that the cop number of $G \setminus V(P)$ is at most r . \square

Bibliography

- [ABCC01] D. Applegate, R. Bixby, V. Chvátal, and W. Cook, *TSP cuts which do not conform to the template paradigm*, Computational Combinatorial Optimization (M. Jünger and D. Naddef, eds.), Springer-Verlag Berlin Heidelberg, 2001, pp. 261–303. [4](#), [21](#), [26](#)
- [ABCC03] ———, *Implementing the Dantzig-Fulkerson-Johnson algorithm for large Traveling Salesman Problems*, Math. Program. Ser. B **97** (2003), no. 1–2, 91–153. [26](#)
- [ABCC06] ———, *The traveling salesman problem – a computational study*, Princeton Series in Applied Mathematics, Princeton, 2006. [25](#), [26](#), [29](#)
- [ABCM04] Carlos Ansotegui, Ramon Béjar, Alba Cabiscol, and Felip Manyà, *The interface between p and np in signed cnf formulas*, Proceedings. 34th International Symposium on Multiple-Valued Logic, 2004, pp. 251–256. [122](#)
- [ACGH09] Júlio Araújo, Nathann Cohen, Frédéric Giroire, and Frédéric Havet, *Good edge-labelling of graphs*, LAGOS’09—V Latin-American Algorithms, Graphs and Optimization Symposium, Electron. Notes Discrete Math., vol. 35, Elsevier Sci. B. V., Amsterdam, 2009, pp. 275–280. MR 2579442 [12](#), [93](#)
- [ACGH12] ———, *Good edge-labelling of graphs*, Discrete Appl. Math. **vol** (2012+), no. no, (in press). [12](#), [13](#), [93](#), [94](#), [95](#), [111](#)
- [Ach00] Dimitris Achlioptas, *Setting 2 variables at a time yields a new lower bound for random 3-SAT (extended abstract)*, Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing (New York), ACM, 2000, pp. 28–37 (electronic). MR 2114514 [15](#), [123](#), [139](#)
- [Ach01] ———, *Lower bounds for random 3-SAT via differential equations*, Theoret. Comput. Sci. **265** (2001), no. 1-2, 159–185, Phase transitions in combinatorial problems (Trieste, 1999). MR 1848217 (2002h:68077) [122](#), [123](#), [139](#), [144](#)
- [AF84] M. Aigner and M. Fromme, *A game of cops and robbers*, Discrete Appl. Math. **8** (1984), 1–12. [17](#), [151](#), [152](#), [156](#), [160](#)
- [AHL02] Noga Alon, Shlomo Hoory, and Nathan Linial, *The Moore bound for irregular graphs*, Graphs Combin. **18** (2002), no. 1, 53–57. MR 1892433 (2003b:05084) [11](#)
- [AL02] Alon Amit and Nathan Linial, *Random graph coverings. I. General theory and graph connectivity*, Combinatorica **22** (2002), no. 1, 1–18. MR 1883559 (2003a:05131) [14](#), [113](#)
- [AL06] ———, *Random lifts of graphs: edge expansion*, Combin. Probab. Comput. **15** (2006), no. 3, 317–332. MR 2216470 (2007a:05125) [14](#), [113](#)
- [AL09] A. R. S. Amaral and A. N. Letchford, *A polyhedral approach to the linear arrangement problem*, Working paper, Dept. of Management Science, Lancaster University, 2009. [8](#), [48](#), [50](#), [51](#), [52](#)
- [ALM02] Alon Amit, Nathan Linial, and Jiří Matoušek, *Random lifts of graphs: independence and chromatic number*, Random Structures Algorithms **20** (2002), no. 1, 1–22. MR 1871947 (2003a:05132) [14](#), [113](#)
- [ALMR01] Alon Amit, Nathan Linial, Jiří Matoušek, and Eyal Rozenman, *Random lifts of graphs*, Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms (Washington, DC, 2001) (Philadelphia, PA), SIAM, 2001, pp. 883–894. MR 1958564 [14](#), [113](#)
- [ALN07] S. Arora, J. R. Lee, and A. Naor, *Fréchet embeddings of negative type metrics*, Discrete Comput. Geom. **38** (2007), 726–739. [47](#)
- [ALN08] ———, *Euclidean distortion and the sparsest cut*, J. of the American Mathematical Society **21** (2008), 1–21. [47](#)
- [AM03] C. Ansotegui and F. Manyà, *New logical and complexity results for signed-sat*, Proceedings of the 33rd International Symposium on Multiple-Valued Logic, 2003, pp. 181–187. [122](#)
- [And84] T. Andreae, *Note on a pursuit game played on graphs*, Discrete Appl. Math. **9** (1984), 111–115. [151](#), [156](#), [159](#), [161](#)
- [And86] ———, *On a pursuit game played on graphs for which a minor is excluded*, J. Combin. Th., Ser. B **41** (1986), 37–47. [17](#), [152](#), [156](#)
- [AP04] Dimitris Achlioptas and Yuval Peres, *The threshold for random k -SAT is $2^k \log 2 - O(k)$* , J. Amer. Math. Soc. **17** (2004), no. 4, 947–973 (electronic). MR 2083472 (2005g:68137) [122](#)

- [APT79] Bengt Aspvall, Michael F. Plass, and Robert Endre Tarjan, *A linear-time algorithm for testing the truth of certain quantified Boolean formulas*, Inform. Process. Lett. **8** (1979), no. 3, 121–123. MR 526451 (80b:68050) [122](#), [129](#)
- [AS00] Dimitris Achlioptas and Gregory B. Sorkin, *Optimal myopic algorithms for random 3-SAT*, 41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000), IEEE Comput. Soc. Press, Los Alamitos, CA, 2000, pp. 590–600. MR 1931856 [15](#), [123](#), [139](#)
- [Ass80] P. Assouad, *Plongements isométriques dans L^1 : aspect analytique*, Initiation Seminar on Analysis: G. Choquet-M. Rogalski-J. Saint-Raymond, 19th Year: 1979/1980, Publ. Math. Univ. Pierre et Marie Curie, vol. 41, Univ. Paris VI, Paris, 1980, pp. Exp. No. 14, 23. [47](#)
- [Bal12] Kathrin Ballerstein, *Logical interaction networks in biology: Theory and application*, Ph.D. thesis, Eidgenössische Technische Hochschule Zürich (ETHZ), 2012. [122](#)
- [BC91] S. C. Boyd and W. H. Cunningham, *Small Travelling Salesman Polytopes*, Math. Oper. Res **16** (1991), no. 2, 259–271. [3](#), [26](#)
- [BCCF06] K. Burgin, P. Chebolu, C. Cooper, and A. M. Frieze, *Hamilton cycles in random lifts of graphs*, European J. Combin. **27** (2006), no. 8, 1282–1293. MR 2260121 (2007e:05081) [14](#), [113](#)
- [BCCP06] J.-C. Bermond, M. Cosnard, D. Coudert, and S. Pérennes, *Optimal solution of the maximum all request path grooming problem*, Proceedings of the Advanced International Conference on Telecommunications, AICT06, Guadeloupe, France, 2006. [12](#), [93](#)
- [BCF⁺01] Ramon Béjar, Alba Cabiscol, Cèsar Fernández, Felip Manà, and Carla P. Gomes, *Capturing structure with satisfiability*, Proceedings of CP-2001, LNCS, vol. 2239, Springer, 2001, pp. 137–152. [122](#)
- [BD92] H.-J. Bandelt and A. W. M. Dress, *A canonical decomposition theory for metrics on a finite set*, Advances in Mathematics **92** (1992), 47–105. [47](#)
- [BFP11] Paul Bonsma, Arthur M. Farley, and Andrzej Proskurowski, *Extremal graphs having no matching cuts*, J. Graph Theory **vol** (2011), no. num, pages. [13](#), [94](#)
- [BFT] Michel Bode, Babak Farzad, and Dirk Oliver Theis, *Good edge labelings and graph with girth at least five*, Submitted to *Discrete Math. Theor. Comput. Sci.* ([arXiv:1109.1125](#)). [12](#)
- [BHM00a] Bernhard Beckert, Reiner Hähnle, and Felip Manyà, *The 2-SAT problem of regular signed CNF formulas*, 30th IEEE International Symposium on Multiple-Valued Logic (ISMVL 2000) (Portland, OR), IEEE Computer Soc., Los Alamitos, CA, 2000, pp. 331–336. MR MR1896571 [122](#)
- [BHM00b] Bernhard Beckert, Reiner Hähnle, and Felip Manyà, *The SAT problem of signed CNF formulas*, Labelled Deduction (M. D’Agostino, D. Basin, D. Gabbay, S. Matthews, and L. Vigano, eds.), Kluwer, Dordrecht, 2000, pp. 59–80. [121](#), [122](#), [147](#)
- [BI93] A. Berarducci and B. Intrigila, *On the cop number of a graph*, Adv. Appl. Math. **14** (1993), 389–403. [17](#), [18](#), [152](#), [155](#), [158](#), [159](#)
- [BL06] Yonatan Bilu and Nathan Linial, *Lifts, discrepancy and nearly optimal spectral gap*, Combinatorica **26** (2006), no. 5, 495–519. MR 2279667 (2008a:05160) [14](#), [113](#)
- [BM86] F. Barahona and A. R. Mahjoub, *On the cut polytope*, Math. Program. **36** (1986), 157–173. [49](#), [52](#)
- [BM99a] Ramon Béjar and Felip Manyà, *Phase transitions in the regular random 3-SAT problem.*, ISMIS’99, 1999, pp. M292–300. [15](#), [122](#), [147](#)
- [BM99b] ———, *Solving combinatorial problems with regular local search algorithms*, Proceedings of the 6th International Conference on Logic for Programming and Automated Reasoning (LPAR), LNCS, vol. 1705, Springer, 1999, pp. 33–43. [122](#)
- [BMC⁺07] Ramón Béjar, Felip Manyà, Alba Cabiscol, Cèsar Fernández, and Carla Gomes, *Regular-SAT: a many-valued approach to solving combinatorial problems*, Discrete Appl. Math. **155** (2007), no. 12, 1613–1626. MR 2341632 (2008f:68135) [15](#), [122](#), [147](#)
- [BN11] Anthony Bonato and Richard J. Nowakowski, *The game of cops and robbers on graphs*, American Mathematical Society, Providence, RI, 2011. [17](#)
- [Bod11] Michel Bode, *Good edge-labelings of graphs*, Diploma thesis, Otto von Guericke University Magdeburg, 2011, in preparation (will be on arXiv). [95](#)
- [Bon05] Paul Bonsma, *A characterization of extremal graphs with no matching-cut*, Proceedings of the European Conference on Combinatorics, Graph Theory and Applications (EuroComb’05), Discrete Mathematics and Theoretical Computer Science Proceedings AE, 2005, pp. 135–138. [13](#), [94](#)
- [Bor89] Oleg V. Borodin, *On the total coloring of planar graphs*, J. Reine Angew. Math. **394** (1989), 180–185. [88](#)
- [Bri04] Katarina Britz, *Interval satisfiability*, Proceedings of the 2004 annual research conference of the South African institute of computer scientists and information technologists on IT research in developing countries (Republic of South Africa), SAICSIT ’04, South African Institute for Computer Scientists and Information Technologists, 2004, pp. 244–248. [122](#)
- [BS96] L. J. Billera and A. Sarangarajan, *All 0-1 polytopes are traveling salesman polytopes*, Combinatorica **16** (1996), 175–188. [3](#), [26](#)

- [BT] Kathrin Ballerstein and Dirk Oliver Theis, *An algorithm for random signed 3-SAT with intervals*, Submitted to *Random Struct. Algorithms* (arXiv:1105.2525). 15, 16
- [Car04] R. Carr, *Separation algorithms for classes of STSP inequalities arising from a new STSP relaxation*, *Math. Oper. Res.* **29** (2004), no. 1, 80–91. 4, 21, 26
- [CCHS10] Victor Chepoi, Nadia Creignou, Miki Hermann, and Gernot Salzer, *The Helly property and satisfiability of Boolean formulas defined on set families*, *European J. Combin.* **31** (2010), no. 2, 502–516. MR 2565343 (2011c:68070) 121, 122, 125, 129, 130
- [CD03] Nadia Creignou and Hervé Daudé, *Generalized satisfiability problems: minimal elements and phase transitions*, *Theoret. Comput. Sci.* **302** (2003), no. 1-3, 417–430. MR 1981958 (2004d:68055) 146
- [CD04] ———, *Combinatorial sharpness criterion and phase transition classification for random CSPs*, *Inform. and Comput.* **190** (2004), no. 2, 220–238. MR 2049708 (2005a:82033) 146
- [CD09] ———, *The SAT-UNSAT transition for random constraint satisfaction problems*, *Discrete Math.* **309** (2009), no. 8, 2085–2099. MR 2510334 (2010d:68062) 146
- [CF86] Ming-Te Chao and John Franco, *Probabilistic analysis of two heuristics for the 3-satisfiability problem*, *SIAM J. Comput.* **15** (1986), no. 4, 1106–1118. MR 861375 (88b:68079) 15, 122, 123
- [CF90] ———, *Probabilistic analysis of a generalization of the unit-clause literal selection heuristics for the k -satisfiability problem*, *Inform. Sci.* **51** (1990), no. 3, 289–314. MR 1072035 (91g:68076) 15, 123
- [CFJT] Nancy E. Clarke, Samuel Fiorini, Gwenaél Joret, and Dirk Oliver Theis, *A note on the cops & robber game on graphs embedded in non-orientable surfaces*, Submitted to *Graphs Comb.* (arXiv:0803.0538). 17
- [CFN85] G. Cornuéjols, J. Fonlupt, and D. Naddef, *The Traveling Salesman Problem on a Graph and some related Integer Polyhedra*, *Math. Program.* **33** (1985), 1–27. 28, 29, 30
- [CGR08] S. Chawla, A. Gupta, and H. Räcke, *Embeddings of negative-type metrics and an improved approximation to generalized sparsest cut*, *ACM Trans. on Algorithms* **4** (2008). 47
- [Cha03] L. Sunil Chandran, *A high girth graph construction*, *SIAM J. Discrete Math.* **16** (2003), no. 3, 366–370. 74, 87, 88
- [CHKR08] M. Charikar, M. T. Hajiaghayi, H. Karloff, and S. Rao, *l_2^2 spreading metrics for vertex ordering problems*, *Algorithmica*, to appear (2008). 48
- [CJR91] T. Christof, M. Jünger, and G. Reinelt, *A complete description of the Traveling Salesman Polytope on 8 nodes*, *Oper. Res. Lett.* **10** (1991), 497–500. 3, 26
- [CLSG09] A. Caprara, A. N. Letchford, and J. J. Salazar-Gonzalez, *Decorous lower bounds for minimum linear arrangement*, Working paper, DEIS, University of Bologna, 2009. 48
- [CM11] N. E. Clarke and G. MacGillivray, *Characterizations of k -copwin graphs*, *Discrete Math.* **vol** (2011), pgs. 17, 151
- [CO10] Amin Coja-Oghlan, *A better algorithm for random k -SAT*, *SIAM J. Comput.* **39** (2010), no. 7, 2823–2864. MR 2645890 15, 123
- [COF] Armin Coja-Oghlan and Alan Frieze, *Analyzing Walksat on random formulas*, (arXiv:1106.0120). 15
- [CR92] Vašek Chvátal and Bruce Reed, *Mick gets some (the odds are on his side)*, *FOCS*, 1992, pp. 620–627. 129, 130, 146
- [CR96] T. Christof and G. Reinelt, *Combinatorial optimization and small polytopes*, *Top* **4** (1996), no. 1, 1–53. 3, 26
- [CR01] ———, *Decomposition and parallelization techniques for enumerating the facets of combinatorial polytopes*, *Int. J. Comput. Geom. Appl.* **11** (2001), 423–437. 3, 26
- [Dez61] M. Deza, *On hamming geometry of unitary cubes*, *Cybernetics and Control Theory* **134** (1961), 940–943. 47
- [DFJ54] G. Dantzig, R. Fulkerson, and S. Johnson, *Solution of a large-scale Traveling Salesman Problem*, *Oper. Res.* **2** (1954), 393–410. 3, 25
- [DGL93] M. Deza, V. P. Grishukhin, and M. Laurent, *The hypermetric cone is polyhedral*, *Combinatorica* **13** (1993), 397–411. 47
- [Die05] Reinhard Diestel, *Graph theory*, 3rd ed., Springer-Verlag, Berlin, 2005. 75, 79
- [Die06] R. Diestel, *Graph theory*, GTM, Springer-Verlag New York, 2006. 93, 155, 160, 161
- [DL97] M. M. Deza and M. Laurent, *Geometry of cuts and metrics*, New York: Springer-Verlag, 1997. 7, 45, 46
- [DL06] Yotam Drier and Nathan Linial, *Minors in lifts of graphs*, *Random Structures Algorithms* **29** (2006), no. 2, 208–225. MR 2245501 (2007e:05074) 14, 113
- [Dow01] J. Downward, *The ins and outs of signalling*, *Nature* **411** (2001), no. 6839, 759–762. 122
- [DPS02] J. Díaz, J. Petit, and M. Serna, *A survey of graph layout problems*, *ACM Computing Surveys* **34** (2002), 313–356. 8, 47

- [EGRS08] F. Eisenbrand, F. Grandoni, T. Rothvoß, and G. Schäfer, *Approximating connected facility location problems via random facility sampling and core detouring*, Proceedings of the XIX Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2008), 2008, pp. 1174–1183. [65](#), [66](#), [71](#)
- [EIM94] Gonzalo Escalada-Imaz and Felip Manyà, *The satisfiability problem for multiple-valued horn formulae*, Proceedings of the IEEE International Symposium on Multiple-Valued Logics, 1994, pp. 250–256. [122](#)
- [ENRS00] G. Even, J. Naor, S. Rao, and B. Schieber, *Divide-and-conquer approximation algorithms via spreading metrics*, J. of the ACM **47** (2000), 585–616. [8](#), [48](#)
- [ER04] T. Erlebach and M. Rüegg, *Optimal bandwidth reservation in hose-model vpns with multi-path routing*, Proceedings of the Twenty-Third Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2004), 2004, pp. 2275–2282. [65](#)
- [FGK08] Fedor V. Fomin, Petr A. Golovach, and Jan Kratochvíl, *On tractability of cops and robbers game*, Fifth IFIP International Conference on Theoretical Computer Science—TCS 2008, IFIP Int. Fed. Inf. Process., vol. 273, Springer, New York, 2008, pp. 171–185. MR 2757374 [17](#), [155](#)
- [FJP10] Samuel Fiorini, Gwenaël Joret, and Ugo Pietropaoli, *Hitting diamonds and growing cacti*, Integer programming and combinatorial optimization, Lecture Notes in Comput. Sci., vol. 6080, Springer, Berlin, 2010, pp. 191–204. MR 2661098 (2011j:05321) [11](#), [91](#)
- [FJTW] Samuel Fiorini, Gwenaël Joret, Dirk Oliver Theis, and David Wood, *Small minors in dense graphs*, To appear in *Eur. J. Comb.* ([arXiv:1005.0895](#)). [11](#)
- [FL07] U. Feige and J. R. Lee, *An improved approximation ratio for the minimum linear arrangement problem*, Inf. Proc. Lett. **101** (2007), 26–29. [48](#)
- [FN92] J. Fonlupt and D. Naddef, *The Traveling Salesman Problem in graphs with some excluded minors*, Math. Program. **53** (1992), no. 2, 147–172. [3](#), [26](#)
- [FOST10] Samuel Fiorini, Gianpaolo Oriolo, Laura Sanità, and Dirk Oliver Theis, *The VPN problem with concave costs*, SIAM J. Discrete Math. **24** (2010), no. 3, 1080–1090. MR 2736009 [9](#)
- [FP84] Arthur M. Farley and Andrzej Proskurowski, *Extremal graphs with no disconnecting matching*, Proceedings of the second West Coast conference on combinatorics, graph theory, and computing (Eugene, Ore., 1983), vol. 41, 1984, pp. 153–165. MR 749610 (85g:05089) [13](#), [94](#)
- [FP01] A.M. Frisch and T.J. Peugniez, *Solving bon-boolean satisfiability problems with stochastic local search*, Proceedings of IJCAI-2001, 2001, pp. 282–288. [122](#)
- [Fra87] P. Frankl, *Cops and Robbers in graphs with large girth and Cayley graphs*, Discrete Appl. Math. **17** (1987), 301–305. [159](#)
- [Fri99] Ehud Friedgut, *Sharp thresholds of graph properties, and the k -sat problem*, J. Amer. Math. Soc. **12** (1999), no. 4, 1017–1054, With an appendix by Jean Bourgain. MR 1678031 (2000a:05183) [15](#), [146](#)
- [FS96] Alan Frieze and Stephen Suen, *Analysis of two simple heuristics on a random instance of k -SAT*, J. Algorithms **20** (1996), no. 2, 312–355. MR 1379227 (97c:68062) [15](#), [16](#), [122](#), [123](#), [126](#)
- [FST97] J.A. Fingerhut, S. Suri, and J.S. Turner, *Designing least-cost nonblocking broadband networks*, Journal of Algorithms **24** (1997), 287–309. [9](#), [65](#)
- [FT] Babak Farzad and Dirk Oliver Theis, *Random lifts of $K_5 \setminus e$ are 3-colourable.*, Accepted for *SIAM J. Discrete Math.* [14](#)
- [FT08] Fedor V. Fomin and Dimitrios M. Thilikos, *An annotated bibliography on guaranteed graph searching*, Theoret. Comput. Sci. **399** (2008), no. 3, 236–245. MR 2419780 (2010a:05176) [17](#), [155](#)
- [GB93] M. X. Goemans and D. J. Bertsimas, *Survivable networks, linear programming relaxations and the parsimonious property*, Math. Program. **60** (1993), no. 2, 145–166. [28](#), [33](#)
- [GI06] F. Grandoni and G. Italiano, *Improved approximation for single-sink buy-at-bulk*, Proceedings of the 17th International Symposium on Algorithms and Computation (ISAAC 2006), 2006, pp. 111–120. [65](#), [66](#), [69](#)
- [GJR10] Catherine Greenhill, Svante Janson, and Andrzej Ruciński, *On the number of perfect matchings in random lifts*, Combin. Probab. Comput. **19** (2010), no. 5-6, 791–817. MR 2726080 (2011h:05229) [14](#)
- [GK02] A. Gupta and A. Kumar, *Personal communication*, 2002. [65](#)
- [GKK⁺01] A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, and B. Yener, *Provisioning a virtual private network: A network design problem for multicommodity flow*, Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC 2001), 2001, pp. 389–398. [9](#), [64](#), [65](#), [66](#), [67](#), [70](#)
- [GKOS08] F. Grandoni, V. Kaibel, G. Oriolo, and M. Skutella, *A short proof of the VPN tree routing conjecture on ring networks*, Operation Research Letters **36** (2008), 361–365. [64](#), [65](#), [66](#), [67](#)
- [GKPR07] A. Gupta, A. Kumar, M. Pål, and T. Roughgarden, *Approximation via Cost Sharing: Simpler and better approximation algorithms for network design*, Journal of the ACM (JACM) **54** (2007), no. 3, Article No. 11. [65](#), [66](#), [69](#), [71](#)
- [Goe95] M. X. Goemans, *Worst-case comparison of valid inequalities for the TSP*, Math. Program. **69** (1995), no. 2, 335–349. [28](#), [29](#), [33](#)

- [Goe96] Andreas Goerdt, *A threshold for unsatisfiability*, J. Comput. System Sci. **53** (1996), no. 3, 469–486, 1994 ACM Symposium on Parallel Algorithms and Architectures (Cape May, NJ, 1994). MR 1423858 (98i:03012) [146](#)
- [GOS08a] N. Goyal, N. Olver, and B. Shepherd, *Personal communication*, 2008. [72](#)
- [GOS08b] ———, *The VPN conjecture is true*, Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC 2008), 2008, pp. 443–450. [9](#), [65](#), [66](#), [67](#), [68](#), [70](#), [71](#)
- [GP79a] M. Grötschel and M. W. Padberg, *On the Symmetric Travelling Salesman Problem I: inequalities*, Math. Program. **16** (1979), 265–280. [33](#), [34](#)
- [GP79b] M. Grötschel and M. W. Padberg, *On the Symmetric Travelling Salesman Problem II: lifting theorems and facets*, Math. Program. **16** (1979), 281–302. [33](#), [34](#)
- [GP85] M. Grötschel and M.W. Padberg, *Polyhedral theory*, The Travelling Salesman Problem (E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan, and D. Shmoys, eds.), Wiley, New York, 1985, pp. 251–305. [25](#), [26](#)
- [GP90] G.M. Guisewite and P.M. Pardalos, *Minimum concave-cost network flow problems: Applications, complexity, and algorithms*, Annals of Operations Research **25** (1990), 75–99. [64](#)
- [GR95] A. S. Goldstein and E. M. Reingold, *The complexity of pursuit on a graph*, Theoret. Comput. Sci. **143** (1995), 93–112. [17](#), [155](#)
- [Gra70] R. L. Graham, *On primitive graphs and optimal vertex assignments*, Ann. New York Acad. Sci. **175** (1970), 170–186. MR 0269533 (42 #4428) [12](#), [93](#)
- [Grü03] B. Grünbaum, *Convex polytopes*, 2nd ed., Springer-Verlag New York, 2003. [28](#)
- [GS01] Geoffrey R. Grimmett and David R. Stirzaker, *Probability and random processes*, third ed., Oxford University Press, New York, 2001. MR 2059709 (2004m:60002) [147](#)
- [H91] Reiner Hähnle, *Towards an efficient tableau proof procedure for multiple-valued logics*, Lecture Notes in Computer Science **533/1991** (1991), 248–260, DOI: 10.1007/3-540-54487-9_62. [121](#)
- [Hah07] Geňa Hahn, *Cops, robbers and graphs*, Tatra Mt. Math. Publ. **36** (2007), 163–176. MR 2378748 (2009b:05254) [155](#)
- [Hel55a] I. Heller, *Geometric characterization of cyclic permutations [abstract]*, Bull. Am. Math. Soc. **61** (1955), 227. [3](#)
- [Hel55b] ———, *Neighbor relations on the convex hull of cyclic permutations [abstract]*, Bull. Am. Math. Soc. **61** (1955), 440. [3](#)
- [Hel56] ———, *Neighbor relations on the convex hull of cyclic permutations*, Pacific J. Math. **6** (1956), 467–477. [3](#)
- [HKS07] C.A.J. Hurkens, J.C.M. Keijsper, and L. Stougie, *Virtual private network design: A proof of the tree routing conjecture on ring networks*, SIAM Journal on Discrete Mathematics **21** (2007), 482–503. [65](#)
- [HM06] G. Hahn and G. MacGillivray, *A note on k -cop, l -robber games on graphs*, Discrete Math **306** (2006), 2492–2497. [17](#), [155](#)
- [HNTW09] Utz-Uwe Haus, Kathrin Niermann, Klaus Truemper, and Robert Weismantel, *Logic integer programming models for signaling networks*, Journal of Computational Biology **16** (2009), no. 5, 725–743. [122](#)
- [HS03] MohammadTaghi Hajiaghayi and Gregory B. Sorkin, *The satisfiability threshold of random 3-SAT is at least 3.52*, Tech. Report MIT-CSAIL-TR-2003-029, Massachusetts Institute of Technology, 2003. [123](#)
- [ILO06] G. Italiano, S. Leonardi, and G. Oriolo, *Design of trees in the hose model: the balanced case*, Operation Research Letters **34** (2006), 601–606. [64](#), [65](#), [70](#)
- [JCBP09] M. Cosnard J.-C. Bermond and S. Pérennes, *Directed acyclic graphs with unique path property*, Tech. Report 6932, INRIA, May 2009. [12](#), [93](#)
- [Jia11] Tao Jiang, *Compact topological minors in graphs*, J. Graph Theory **67** (2011), no. 2, 139–152. [74](#)
- [JKT10] Gwenaél Joret, Marcin Kamiński, and Dirk Oliver Theis, *The cops and robber game on graphs with forbidden (induced) subgraphs*, Contrib. Discrete Math. **5** (2010), no. 2, 40–51. MR 2791289 [18](#)
- [JLR00] Svante Janson, Tomasz Łuczak, and Andrzej Rucinski, *Random graphs*, Wiley-Interscience Series in Discrete Mathematics and Optimization, Wiley-Interscience, New York, 2000. MR 1782847 (2001k:05180) [135](#)
- [JM96] Stanislav Jendroľ and Tomáš Madaras, *On light subgraphs in plane graphs of minimum degree five*, Discuss. Math. Graph Theory **16** (1996), no. 2, 207–217. [88](#)
- [JRR95] M. Jünger, G. Reinelt, and G. Rinaldi, *The Traveling Salesman Problem*, Handbooks in Operations Research and Management Science (M. O. Ball, T. L. Magnanti, C. L. Monma, and G. L. Nemhauser, eds.), vol. 7, Elsevier Science B. V., 1995, pp. 225–330. [25](#)
- [JSW90] Joyce Justicz, Edward R. Scheinerman, and Peter M. Winkler, *Random intervals*, Amer. Math. Monthly **97** (1990), no. 10, 881–889. MR MR1079974 (91m:60023) [121](#), [128](#)

- [JV05] Stanislav Jendroř and Heinz-Jürgen Voss, *Light subgraphs of order at most 3 in large maps of minimum degree 5 on compact 2-manifolds*, *European J. Combin.* **26** (2005), no. 3–4, 457–471. [88](#)
- [Kau69] Stuart A. Kauffman, *Metabolic stability and epigenesis in randomly constructed genetic nets*, *J. Theor. Biol.* **22** (1969), no. 3, 437–467. [122](#)
- [KKL06] Alexis C. Kaporis, Lefteris M. Kirousis, and Efthimios G. Lalas, *The probabilistic analysis of a greedy satisfiability algorithm*, *Random Structures Algorithms* **28** (2006), no. 4, 444–480. MR 2225702 (2007g:68067) [123](#), [146](#)
- [KM00] D.R. Karger and M. Minkoff, *Building steiner trees with incomplete global knowledge*, *Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS 2000)*, 2000, pp. 613–623. [71](#)
- [KO03] Daniela Kühn and Deryk Osthus, *Minors in graphs of large girth*, *Random Structures Algorithms* **22** (2003), no. 2, 213–225. [79](#)
- [KO05] ———, *Forcing unbalanced complete bipartite minors*, *European J. Combin.* **26** (2005), no. 1, 75–81. [91](#)
- [KO06] ———, *Improved bounds for topological cliques in graphs of large girth*, *SIAM J. Discrete Math.* **20** (2006), no. 1, 62–78. [75](#)
- [Kos82] Alexandr V. Kostochka, *The minimum Hadwiger number for graphs with a given mean degree of vertices*, *Metody Diskret. Analiz.* **38** (1982), 37–58. [73](#), [75](#)
- [Kos84] ———, *Lower bound of the Hadwiger number of graphs by their average degree*, *Combinatorica* **4** (1984), no. 4, 307–316. [73](#), [75](#)
- [Kot55] Anton Kotzig, *Contribution to the theory of Eulerian polyhedra*, *Mat.-Fyz. Časopis. Slovensk. Akad. Vied* **5** (1955), 101–113. [88](#)
- [KP88] A. Kostochka and L. Pyber, *Small topological complete subgraphs of “dense” graphs*, *Combinatorica* **8** (1988), no. 1, 83–86. MR 951996 (89e:05079) [11](#), [74](#), [75](#), [91](#)
- [KP08] Alexandr Kostochka and Noah Prince, *On $K_{s,t}$ -minors in graphs with given average degree*, *Discrete Math.* **308** (2008), no. 19, 4435–4445. [91](#)
- [KR06] R. Krauthgamer and Y. Rabani, *Improved lower bounds for embeddings into l_1* , In Proc. of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2006. [47](#)
- [KS09] Michael Krivelevich and Benjamin Sudakov, *Minors in expanding graphs*, *Geom. Funct. Anal.* **19** (2009), no. 1, 294–331. [75](#)
- [KSRL⁺06] Steffen Klamt, Julio Saez-Rodriguez, Jonathan A. Lindquist, Luca Simeoni, and Ernst Dieter Gilles, *A methodology for the structural and functional analysis of signaling and regulatory networks.*, *BMC Bioinformatics* **7** (2006), no. 56, available from <http://www.biomedcentral.com/1471-2105/7/56>. [122](#)
- [Kuh55] H. W. Kuhn, *On certain convex polyhedra [abstract]*, *Bull. Am. Math. Soc.* **61** (1955), 557–558. [3](#)
- [KV05] S. A. Khot and N. K. Vishnoi, *The unique games conjecture, integrality gaps for cut problems and embeddability of negative type metrics into l_1* , In Proc. of the 2005 IEEE Symposium on Foundations of Computer Science (FOCS), 2005. [47](#)
- [Lee05] J. R. Lee, *On distance scales, embeddings, and efficient relaxations of the cut cone*, In Proc. of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2005. [47](#)
- [Let05] A. N. Letchford, *Personal communication*, 2005. [33](#)
- [Lov07] László Lovász, *Combinatorial problems and exercises*, second ed., AMS Chelsea Publishing, Providence, RI, 2007. MR 2321240 [114](#)
- [LP10] Nati Linial and Doron Puder, *Word maps and spectra of random graph lifts*, *Random Structures Algorithms* **37** (2010), no. 1, 100–135. MR 2674623 (2011k:60026) [14](#)
- [LR05] Nathan Linial and Eyal Rozenman, *Random lifts of graphs: perfect matchings*, *Combinatorica* **25** (2005), no. 4, 407–424. MR 2143247 (2006c:05110) [14](#), [113](#)
- [LRST10] Adam N. Letchford, Gerhard Reinelt, Hanna Seitz, and Dirk Oliver Theis, *On a class of metrics related to graph layout problems*, *Linear Algebra Appl.* **433** (2010), no. 11–12, 1760–1777. MR 2736096 [8](#), [9](#)
- [Lu20] Jan Łukasiewicz, *O logice trjwartociowej*, *Ruch Filozoficzny* (1920), 169–171, In polish. [121](#)
- [LV95] W. Liu and A. Vannelli, *Generating lower bounds for the linear arrangement problem*, *Discr. Appl. Math.* **59** (1995), 137–151. [47](#)
- [Mad67] Wolfgang Mader, *Homomorphieigenschaften und mittlere Kantendichte von Graphen*, *Math. Ann.* **174** (1967), 265–268. [73](#)
- [Mad68] Wolfgang Mader, *Homomorphiesätze für Graphen*, *Math. Ann.* **178** (1968), 154–168. [73](#)
- [Mad98] ———, *Topological subgraphs in graphs of large girth*, *Combinatorica* **18** (1998), no. 3, 405–412. [79](#)
- [Man00] Felip Manyà, *The 2-SAT problem in signed CNF formulas*, *Multiple-Valued Logic* **5** (2000), no. 4, 307–325. [122](#)

- [MBEI98] Felip Manyà, Ramon Béjar, and Gonzalo Escalada-Imaz, *The satisfiability problem in regular CNF-formulas*, Soft Computing: A Fusion of Foundations, Methodologies and Applications **2** (1998), no. 3, 116–123. [15](#), [122](#), [147](#)
- [Meh12] Abbas Mehrabian, *On the density of nearly regular graphs with a good edge labelling*, SIAM J. Discrete Math. **vol** (2012+), no. no. ?–?, arXiv:1110.2391. [94](#)
- [Mol02] Michael Molloy, *Models and thresholds for random constraint satisfaction problems*, Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing (New York), ACM, 2002, pp. 209–217 (electronic). MR 2121144 [146](#)
- [Mol03] ———, *Models for random constraint satisfaction problems*, SIAM J. Comput. **32** (2003), no. 4, 935–949 (electronic). MR 2001891 (2004i:68197) [146](#)
- [MT01] Bojan Mohar and Carsten Thomassen, *Graphs on surfaces*, Johns Hopkins University Press, Baltimore, 2001. [75](#), [88](#), [152](#)
- [MT05] Joseph Samuel Myers and Andrew Thomason, *The extremal function for noncomplete minors*, Combinatorica **25** (2005), no. 6, 725–753. [91](#)
- [Mye03] Joseph Samuel Myers, *The extremal function for unbalanced bipartite minors*, Discrete Math. **271** (2003), no. 1–3, 209–222. [91](#)
- [Nad02] D. Naddef, *Polyhedral Theory and Branch-and-Cut Algorithms for the Symmetric TSP*, The Traveling Salesman Problem and Its Variations (Gregory Gutin and Abraham P. Punnen, eds.), Kluwer Academic Publishers, 2002, pp. 29–116. [4](#), [21](#), [25](#), [29](#)
- [Nic94] Alexandru Nica, *On the number of cycles of given length of a free word in several random permutations*, Random Structures Algorithms **5** (1994), no. 5, 703–730. MR 1300595 (95m:60017) [14](#)
- [Nor55] R. Z. Norman, *On the convex polyhedra of the Symmetric Traveling Salesman Problem [abstract]*, Bull. Am. Math. Soc. **61** (1955), 559. [3](#), [26](#), [29](#)
- [NP01] Denis Naddef and Yves Pochet, *The symmetric traveling salesman polytope revisited*, Math. Oper. Res. **26** (2001), no. 4, 700–722. MR 1870741 (2002i:90127) [29](#)
- [NR91] Denis Naddef and Giovanni Rinaldi, *The symmetric traveling salesman polytope and its graphical relaxation: composition of valid inequalities*, Math. Programming **51** (1991), no. 3, (Ser. A), 359–400. MR 1130332 (92i:90065) [4](#), [5](#), [6](#), [21](#), [26](#), [29](#), [33](#), [34](#)
- [NR93] ———, *The graphical relaxation: a new framework for the symmetric traveling salesman polytope*, Math. Programming **58** (1993), no. 1, Ser. A, 53–88. MR 1214035 (94c:90081) [4](#), [5](#), [6](#), [21](#), [22](#), [26](#), [29](#), [31](#), [34](#), [35](#)
- [NR07] ———, *The symmetric traveling salesman polytope: new facets from the graphical relaxation*, Math. Oper. Res. **32** (2007), no. 1, 233–256. MR 2292509 (2007m:90142) [29](#)
- [NS] R. J. Nowakowski and B. S. W. Schröder, *Bounding the cop number using the crosscap number*, Preprint. [17](#), [152](#)
- [NW83] R. J. Nowakowski and P. Winkler, *Vertex to vertex pursuit in a graph*, Discrete Math. **43** (1983), 23–29. [17](#), [151](#), [155](#)
- [ORT05] Marcus Oswald, Gerhard Reinelt, and Dirk Oliver Theis, *Not every GTSP facet induces an STSP facet*, Integer programming and combinatorial optimization, Lecture Notes in Comput. Sci., vol. 3509, Springer, Berlin, 2005, pp. 468–482. MR 2210041 (2006k:90147) [5](#), [26](#), [33](#)
- [ORT07] ———, *On the graphical relaxation of the symmetric traveling salesman polytope*, Math. Program. **110** (2007), no. 1, Ser. B, 175–193. MR 2306135 (2008j:52023) [3](#), [4](#), [5](#), [6](#), [7](#), [21](#), [26](#), [28](#), [33](#), [43](#)
- [Qui78] A. Quilliot, *Jeux et points fixes sur les graphes*, Thèse de 3ème cycle, Université de Paris VI, 1978. [17](#), [151](#), [155](#)
- [Qui85] A. Quilliot, *A short note about pursuit games played on a graph with a given genus*, J. Combin. Theory Ser. B **38** (1985), 89–92. [17](#), [152](#), [156](#)
- [QW93] M. Queyranne and Y. Wang, *Hamiltonian path and Symmetric Travelling Salesman polytopes*, Math. Program. **58** (1993), no. 1, 89–110. [7](#)
- [RC98] F. J. Rispoli and S. Cosares, *A bound of 4 for the diameter of the Symmetric Traveling Salesman Polytope*, SIAM J. Discrete Math. **11** (1998), 343–380. [26](#)
- [Ris98] F. J. Rispoli, *The monotonic diameter of traveling salesman polytopes*, Oper. Res. Lett. **22** (1998), 69–73. [26](#)
- [RR05] S. Rao and A. W. Richa, *New approximation techniques for some linear ordering problems*, SIAM Journal on Computing **34** (2005), 388–404. [8](#), [48](#)
- [Sch35] I. J. Schoenberg, *Remarks to maurice fréchet’s article “sur la définition axiomatique d’une classe d’espace distanciér vectoriellement applicable sur l’espace de hilbert”*, Annals of Mathematics **36** (1935), 724–732. [47](#)
- [Sch88] E. R. Scheinerman, *Random interval graphs*, Combinatorica **8** (1988), no. 4, 357–371. MR MR981893 (90a:05171) [121](#), [123](#), [124](#), [126](#), [128](#), [130](#)

- [Sch01] B. S. W. Schröder, *The copnumber of a graph is bounded by $\lfloor \frac{3}{2} \text{genus}(G) \rfloor + 3$* , “Categorical Perspectives” — Proceedings of the Conference in Honor of George Strecker’s 60th Birthday, Birkhäuser, 2001, pp. 243–263. [17](#), [152](#), [156](#)
- [Sch03] Alexander Schrijver, *Combinatorial optimization. Polyhedra and efficiency.*, Algorithms and Combinatorics, vol. 24, Springer-Verlag, Berlin, 2003. MR 1956926 [25](#)
- [Sei09] H. Seitz, *The linear arrangement problem*, Ph.D. thesis, University of Heidelberg, 2009, In preparation. [58](#)
- [Sie98] G Sierksma, *Interchange graphs and the Hamiltonian cycle polytope*, Discrete Appl. Math. **81** (1998), 217–224. [26](#)
- [ST92] G. Sierksma and G. A. Tjissen, *Faces with large diameter on the Symmetric Traveling Salesman Polytope*, Oper. Res. Lett. **12** (1992), 73–77. [26](#)
- [STT95] G. Sierksma, R. H. Teunter, and G. A. Tjissen, *Faces of diameter two on the Hamiltonian cycle polytope*, Oper. Res. Lett. **18** (1995), 59–64. [26](#)
- [SW07] Lingsheng Shi and Nicholas Wormald, *Colouring random 4-regular graphs*, Combin. Probab. Comput. **16** (2007), no. 2, 309–344. MR 2298817 (2008f:05180) [14](#), [119](#)
- [The] Dirk Oliver Theis, *On the facial structure of symmetric and graphical traveling salesman polyhedra*, Submitted to *Discrete Opt.* ([arXiv:0712.1269](#)). [5](#), [6](#)
- [The05] D. O. Theis, *Polyhedra and algorithms for the general routing problem*, Ph.D. thesis, University of Heidelberg, Germany, 2005. [5](#), [6](#), [7](#)
- [The10] Dirk Oliver Theis, *A note on the relationship between the graphical traveling salesman polyhedron, the symmetric traveling salesman polytope, and the metric cone*, Discrete Appl. Math. **158** (2010), no. 10, 1118–1120. MR 2609617 (2011c:90116) [4](#), [27](#)
- [Tho83] Carsten Thomassen, *Girth in graphs*, J. Combin. Theory Ser. B **35** (1983), no. 2, 129–141. [79](#)
- [Tho84] Andrew Thomason, *An extremal function for contractions of graphs*, Math. Proc. Cambridge Philos. Soc. **95** (1984), no. 2, 261–265. [73](#), [75](#)
- [Tho01] ———, *The extremal function for complete minors*, J. Combin. Theory Ser. B **81** (2001), no. 2, 318–338. [73](#)
- [Tho06] ———, *Extremal functions for graph minors*, More sets, graphs and numbers, Bolyai Soc. Math. Stud., vol. 15, Springer, Berlin, 2006, pp. 359–380. [73](#), [91](#)
- [Wag37] Klaus Wagner, *Über eine Eigenschaft der ebene Komplexe*, Math. Ann. **114** (1937), 570–590. [91](#)
- [Wer04] Paul Wernicke, *Über den kartographischen Vierfarbensatz*, Math. Ann. **58** (1904), 413–426. [88](#), [89](#)
- [Wik10] Contributors Wikipedia, *Snub dodecahedron*, 2010, Wikipedia, the free encyclopedia. [81](#)
- [Wor95] Nicholas C. Wormald, *Differential equations for random processes and random graphs*, Ann. Appl. Probab. **5** (1995), no. 4, 1217–1235. MR MR1384372 (97c:05139) [14](#), [123](#)
- [Wor99] ———, *The differential equation method for random graph processes and greedy algorithms*, Lectures on Approximation and Randomized Algorithms (M. Karonski and H. J. Proemel, eds.), PWN, Warsaw, 1999, pp. 73–155. [123](#), [143](#)
- [Zie98] G. M. Ziegler, *Lectures on Polytopes*, Springer-Verlag New York, 1998. [28](#)