Odd cycles in planar graphs

Nadia Hardy

Department of Mathematics and Statistics
McGill University, Montreal, Canada
hardy@math.mcgill.ca

JUNE 20, 2005

A thesis submitted to McGill University in partial fulfilment of the requirements of the degree of Master of Science.

Your file Votre référence
ISBN: 978-0-494-22729-9
Our file Notre référence
ISBN: 978-0-494-22729-9

Abstract

Given a graph $G = (V, E)$, an *odd cycle cover* is a subset of the vertices whose removal makes the graph bipartite, that is, it meets all odd cycles in G. A *packing* in G is a collection of vertex disjoint odd cycles. This thesis addresses algorithmic and structural problems concerning odd cycle covers and packings. In particular, we consider the two NP-hard problems of finding a maximum packing and a minimum covering. In 1994 Brass [53] conjectured that τ, the minimum size of an odd cycle cover, is at most twice ν, the maximum size of a packing. The conjecture is known to be false in general [11, 41]. We prove here that $\tau \leq 10\nu$ for planar graphs. Our structural results leads to the first constant approximation algorithm for the packing problem. The covering problem was shown to be tractable for graphs of constant sized solutions [42]. We give a linear time algorithm for the covering problem restricted to the case where the graphs have constant sized solutions and are planar.

Résumé

Étant donné un graphe $G = (V, E)$, une *couverture des cycles impairs* est un sous ensemble de sommets qui, s'ils sont enlevés, laissent le graphe biparti. C'est-à-dire que chacun des cycles de longueur impaire contient un des ces sommets. Un *empaquetage* dans G est un ensemble de cycles impairs dont les sommets sont disjoints. Cette thèse traite de problèmes algorithmiques et structuraux ayant trait aux couvertures et aux empaquetage des cycles impairs. En particulier, nous considérons la recherche de la couverture minimum et de l'empaquetage maximum, tous deux connus pour être NP-difficiles. En 1994, Brass [53] fit la conjecture que τ, la taille d'une couverture minimum des cycles impairs, était au plus deux fois ν, la taille d'un empaquetage maximum. On sait maintenant que c'est faux dans le cas général [11, 41]. Nous prouvons que $\tau \leq 10\nu$ pour les graphes planaires. Nos résultats structuraux conduisent à la première approximation à facteur constant pour le problème de l'empaquetage. Il a été montré par [42] que le problème de la couverture était polynomial lorsque la taille de la solution est bornée. Nous donnons un algorithme qui trouve la couverture en temps linéaire pour les graphes planaires dont la solution est de taille bornée.

Acknowledgements

I am deeply grateful to my supervisor Professor Adrian R. Vetta. I feel very lucky to have had the opportunity to be his student. He has been incredibly generous with his time, his knowledge and his resources.

It has also been a privilege to study and work with Professor Bruce A. Reed. Like Prof. Vetta, Prof. Reed is an amazingly generous person. Furthermore, studying and working with him involved a perfect balance between seriousness and fun.

I am in great debt to Samuel Fiorini. He was, together with Prof. Vetta and Prof. Reed, my co-author of the results presented in this thesis. I thank him for sharing his ideas and his time with me.

The work presented in this thesis was supported by NSERC Canada and by the Department of Mathematics and Statistics of McGill University.

I am also in debt to the students of the Discrete Mathematics and Optimization group of the School of Computer Science and the Department of Mathematics and Statistics of McGill University. I learned a lot from them about graph theory, combinatorial optimization, and having fun in Montreal. In particular, I am indebted to Nicolas Broutin for translating

I am grateful to my mother for her unconditional support. To my friends in Argentina for keeping me company in cyberspace. To my friends in Canada for helping me to get used to this country.

Finally, I would like to thank Laura for encouraging me and for always being there. This thesis is dedicated to her.

Contents

1 Introduction

Given a graph $G = (V, E)$, an *odd cycle vertex cover* (or *odd cycle vertex transversal*) is a subset of the vertices whose removal makes the graph bipartite, that is, a subset of the vertices that intersects every odd cycle in the graph. A *packing of vertex disjoint odd cycles* is a collection of vertex disjoint odd cycles. *Odd cycle edge cover* and *packing of edge disjoint odd cycles* are defined similarly. For brevity, we will often use the term *cover* to refer to an odd cycle cover, and we will use *packing* to refer to a packing of odd cycles. Unless stated otherwise we always mean odd cycle vertex cover and packing of vertex disjoint odd cycles, respectively.

In this thesis, we consider the problem of finding a minimum odd cycle cover (the *covering problem* or the *graph bipartization problem*) and the problem of finding a maximum packing of odd cycles (the *odd cycle packing problem*). In particular, for both these problems, we will tackle the case in which the underlying graph is restricted to be planar.

The graph bipartization problem is $\mathcal{APX}$-hard (see [37] and Section 4.1). However, it was recently shown [42] that the problem becomes tractable if the size of the optimal cover is bounded by a fixed constant k. Here, we present a linear time algorithm to find minimum odd cycle covers in planar graphs where the size of the cover is at most k (see Section 4). For the odd cycle packing problem, we show in Section 3.4 that the problem is NP-hard in planar graphs. We give a 10-approximation algorithm for the packing problem. This algorithm follows from the following structural result: the cardinality of a minimum odd cycle cover is at most ten times the cardinality of a maximum odd cycle packing.

Both the exact algorithm for the covering problem and the structural theorem for the packing problem are the result of a collaboration with Samuel Fiorini, Bruce Reed and Adrian Vetta. The papers will appear in [19] and [18], respectively.

We begin with an overview of the thesis. Our results rely on important connections between odd cycle covers and T-joins[1] in specific auxiliary graphs. In particular, let T correspond to the set of odd faces of G. Then, in the vertex version of the problems, the auxiliary graph we need to consider is the face-vertex incidence graph G^+. Specifically, we show that minimum odd cycle covers correspond to T-joins in G^+ covering the least number of vertices of G (namely, the vertices in a minimum cover). For the edge version, the appropriated auxiliary graph is the planar dual graph G^*; this relationship was first used by Hadlock [26] to derive a polynomial time algorithm for the maximum cut problem in planar graphs. In Section 2 we present the necessary background on T-joins and detail these connections between odd cycle covers and T-joins. These relationships will play a vital role in our proofs of both the theoretical and algorithmic results.

Section 3 contains the structural result that $\tau \leq 10\nu$ for planar graphs; here τ denotes the size of a minimum odd cycle cover and ν denotes the size of a maximum odd cycle packing. The history of this problem began, in 1994, with a conjecture of Brass [53] stating that $\tau \leq 2\nu$. This was shown to be false for general graphs; see the discussion below (and [11, 41]). We obtain our result by combining a local and a global approach to the problem (see Section 3.2). First we prove that for planar graphs in which all odd faces are pairwise (vertex) intersecting, it is true that $\tau \leq 2\nu$. Then we prove that this also holds for 4-connected planar graphs in which there is a face (possibly an even face) intersecting every odd face. Finally, we show that the same bound holds

[1] See Section 2 for definitions.

if all the odd faces are pairwise "far" apart. The proof for the general class of planar graphs combines the techniques developed for these special classes. We let G be a minimum counterexample to the claim that $\tau \leq 10\nu$. We take a minimum collection of faces of G, termed *centers*, with the property that every odd face of G intersects some face of the collection. Then we show that G must be 4-connected with all its centers far apart. Next, using the techniques developed for our second special class of planar graphs we find a *local* cover around each center. At the same time, we find a packing of odd faces around each center. Since the centers are far apart the union of these packings is also a packing. We call this union the *local* packing. We then extend the *local* covers to a cover of the whole graph. Associated with this cover we find a different packing of odd cycles, which we call the *global* packing. The size of the cover we obtain is within a constant factor of the size of the greater of the local packing and the global packing. Our main result then follows.

The factor 2 bound is tight for the three special cases considered. It is our belief that the factor 10 bound for planar graphs is not tight, and that Brass conjecture is, in fact, true for the class of planar graphs.

In Section 3.4 we detail our hardness result and present an approximation algorithm for the packing problem. Our structural result implies that the greater of the cardinality of the local packing and the the cardinality of the global packing is at least one tenth τ, the size of the minimum cover. On the other hand, it is always true that $\tau \geq \nu$. Therefore, since our structural result gives a constructive method to find those packings, we have a 10-approximation algorithm. The hardness result follows from a result by Caprara and Rizzi [12] showing that packing vertex disjoint triangles is NP-hard for planar graphs.

Finally, in Section 4, we proffer a hardness result showing that the covering problem is NP-hard in planar graphs in general. In contrast, Reed, Smith and Vetta [42] presented a quadratic exact algorithm that determines whether a graph G has an odd cycle cover of size at most k, an exhibits such cover, thus proving that the problem is tractable if restricted to graphs with constant sized solutions. We continue this line of work by giving a linear time algorithm for finding odd cycle covers in planar graphs for instances with constant sized optimal solutions.

Our main technique consists of proving that, given a planar graph G, a vast number of vertices are irrelevant when looking for covers of size bounded by a fixed constant k. We show that vertices that are "far" from every odd face in the face-vertex incidence are irrelevant. Our algorithm finds and removes those irrelevant vertices. We are then left with at most k subgraphs of G (defined by the vertices that are "close" to some odd face) which have the following property: if G has optimal cover of size at most k then this cover is the union of optimal covers in each of the subgraphs. Moreover, the subgraphs are shown to have constant tree-width. In Section 4.4, we show how to find in linear time a minimum odd cycle cover in graphs of bounded tree-width. Consequently, we have a linear time algorithm for planar graphs. As mentioned, the relationship between T-joins and odd cycle covers is key in proving the validity of the algorithm.

As we pointed out above, there are corresponding edge versions of the covering and packing problems. That is, finding a minimum set of edges whose removal makes the graph bipartite and finding a maximum collection of edge disjoint odd cycles. In Section 4.1 we show that the first problem is equivalent to the maximum cut problem, and therefore becomes polynomial when restricted to planar graphs [26].

For the packing problem in its edge version, Král and Voss [30] recently proved that $\tau_e \leq 2\nu_e$, where τ_e and ν_e denote the size of a minimum odd cycle edge cover and the size of a maximum packing of edge disjoint odd cycles respectively. In Section 3.1 we present a very short proof of their result.

This proof, again, relies upon the relationship between T-joins and minimum edge covers. As stated, the auxiliary graph here is the dual graph G^* and a minimum cover corresponds to a T-join in G^* with the fewest number of edges. This correspondence was also exploited by Král and Voss in their paper [30]. A further similarity between the proofs is that we also require the use of the Four Colour Theorem. One major difference between their proof and ours, is the fact that we apply a result of Lovász linking T-joins and 2-packings of T-cuts. We observe that the Four Colour Theorem implies that any laminar 2-packing of k odd cycles in G contains $\frac{1}{4}k$ edge disjoint odd cycles. The result by Lovász then guarantees the existence of a 2-packing of T-cuts whose size is twice the minimum size of an odd cycle edge cover. This gives the result.

We close this introductory section with a discussion on related work. Brandt [11] disproved Brass' conjecture for general graphs. The question remained as to whether the minimum size of a cover is bounded by a function of the maximum size of a packing. Such a correspondence is known as the *Erdös-Pósa property.* A class of graphs $\mathcal{F}$ is said to have the Erdös-Pósa property if for every integer k there is an integer $f(k, \mathcal{F})$ such that every graph G either contains k vertex disjoint subgraphs each isomorphic to a graph in $\mathcal{F}$ or a set W of at most $f(k, \mathcal{F})$ vertices such that $G - W$ has no subgraph isomorphic to a graph in $\mathcal{F}$.

Reed [41] proved that the *Erdös-Pósa* property does not hold in general for the family of odd cycles. He showed that for each positive integer k, there exists a projectiveplanar graph G with size of a minimum odd cycle cover τ equal k and size of a maximum packing ν equal 1. That is, τ is not bounded by any function of ν. In addition, the only obstruction for τ to be bounded is the existence of *Escher walls*[2]. In his work Reed considered only the vertex version of the problem, but as he considered graphs which are cubic, his result is valid for the edge version.

Interestingly, as a corollary of his proof, odd cycles in planar graphs do have the Erdös-Pósa property. This follows as planar graphs do not contain Escher walls. Therefore, the size of a minimum cover is in fact bounded by a function of the size of a maximum packing. The results in [41], though, lead to a super-exponential function of ν. In [7], Berge and Reed gave a super exponential function of ν that bounds τ in planar graphs. Previously, Thomassen [52] had shown that $\tau \leq 2\nu$ for $2^{3^{9\nu}}$-connected graphs. A similar result involving large connectivity was proved by Rautenbach and Reed [40]. They showed that each $2000(k+1)$-connected graph contains either $k+1$ vertex disjoint odd cycles or $2k$ vertices hitting all the odd cycles.

The idea of finding irrelevant vertices to reduce a problem on a given graph to subproblems in particular subgraphs appears in the work of Robertson and Seymour concerning *Realization problems* in planar graphs [46].

As mentioned, the relation between T-joins and T-cuts was first used by Hadlock [26] to obtain a polynomial time algorithm for the maximum cut problem in planar graphs. Král and Voss [30] exploited these relationships in their result. Other works exploring these relationships are due to Lovász [33] and Seymour [49]. They gave min-max relations between T-cuts and T-joins. For a survey on T-joins and their relation with T-cuts, see [20].

[2] Reed wrote in [41] that this was a personal communication by Lóvasz and Schrijver.

We remark that the minimum odd cycle cover problem and the maximum odd cycle packing problem can be formulated as integer programs whose linear program (LP) relaxations are duals. Letting V denote the vertex set of G and $\mathcal{O}$ denote the set of odd cycles in G, then these dual LPs are:

COVERING LP

$$\min \sum_{v \in V} y_v$$

$$\sum_{v : v \in C} y_v \geq 1 \quad \forall C \in \mathcal{O}$$

$$y_v \geq 0 \quad \forall v \in V$$

PACKING LP

$$\max \sum_{C \in \mathcal{O}} x_C$$

$$\sum_{C : v \in C} x_C \leq 1 \quad \forall v \in V$$

$$x_C \geq 0 \quad \forall C \in \mathcal{O}$$

Our structural result, leading to a 10-approximation algorithm for the packing problem in planar graphs, also shows that the integrality gap of the Packing LP is bounded by a constant. Previously, in [25], Goemans and Williamson presented a primal-dual $\frac{9}{4}$-approximation algorithm for the covering problem restricted to planar graphs. In their paper they approach the more general problem of finding a minimum vertex cover for a class of cycles in planar graphs. The class of odd cycles is one of the families they consider. They proved that the integrality gap of the Covering LP is at most $\frac{9}{4}$ and conjectured that it is actually $\frac{3}{2}$. For general graphs, Garg, Vazirani and Yannakakis [23] gave a $O(\log n)$-approximation algorithm for the covering problem.

The covering problem we tackle in this thesis has its origins in *feedback set problems.* The most general feedback set problem consists in finding a minimum-weight set of vertices (or edges, or arcs) that hits all the cycles in a specific subset $\mathcal{C}$ of the cycles of a graph G. These problems originated in combinatorial circuit design but now have a vast number of applications. For example, deadlock prevention in operating systems, constraint satisfaction and Bayesian inference in artificial intelligence, and graph theory. Further applications for the bipartization problem arise in biology. In particular, in [43] a bipartization routine is used to develop an exact algorithm for the *single individual SNP haplotyping* problem in instances where the number of holes in the input sequence is bounded by a fixed constant k. The single individual SNP haplotyping problem comes from combinatorial biology. A *haplotype* is the complete sequence of single nucleotide polymorphisms (SNPs) of the two copies of a given chromosome in a diploid genome. These polymorphisms are the most frequent form of human genetic variation and of foremost importance for a variety of applications including medical diagnosis, phylogenies and drug design [43]. The haplotyping problem is, given a set of fragments from one individual's DNA, to find a maximally consistent pair of SNP's haplotypes. The fragments can overlap or repeat themselves, but they are usually incomplete in the sense that they do not cover the complete information of the DNA sequence. The parameter k corresponds to the number of holes in the input data.

When C is the set of all cycles in G, then our problem is known also as the *hitting cycle problem*, which has been well studied from an algorithmic point of view. See [17] for a detailed survey on the subject.

For a large variety of classes of graphs, polynomial time algorithms were found for different versions of these hitting problems. Some examples are the minimum feedback vertex set problem in reducible flow graphs [51], chordal graphs [13, 55], interval graphs [38], permutation graphs [10, 9]. For classes of graphs in which the problems are not known to be polynomial time solvable, approximation algorithms have been developed [16, 39, 4, 6, 3, 25]. On a more theoretical perspective, several bounds were found on the size of the minimum feedback vertex (edge) set for different classes of graphs and different versions of the problem ($\mathcal{C}$ being all the cycles in the directed or undirected case [16, 36], $\mathcal{C}$ being a subset of the cycles, in particular the set of odd cycles [41, 52, 40, 30, 18]).

In the Conclusion, open problems and possible approaches for further work are discussed.

2 Definitions and background

Here, we define the problems considered in this thesis. We also present the necessary definitions and background on T-joins. In particular, we describe the relationship between T-joins and odd cycle covers.

Definition 2.1. Given a graph $G = (V, E)$, $W \subseteq V$ is an *odd cycle vertex cover* (or an *odd cycle vertex transversal*) if $G - W$ is bipartite, that is, if W intersects all the odd cycles in G.

Definition 2.2. Given a graph $G = (V, E)$, a *packing of vertex disjoint odd cycles* is a collection of vertex disjoint odd cycles of G.

There are similar definitions when we consider edge subsets and collections of edge disjoint odd cycles. We often simply write *cover* to refer to odd cycle covers. We also use the term *packing* to refer to a packing of odd cycles. Unless otherwise stated, we always mean odd cycle vertex cover and packing of vertex disjoint odd cycles, respectively.

Covering problem (*or Bipartization problem*)**:** Given a graph $G = (V, E)$, find an odd cycle cover of minimum cardinality.

Packing problem: Given a graph $G = (V, E)$, find a packing of odd cycles of maximum cardinality.

Our focus is upon the restrictions of both these problems to planar graphs. For general graphs, the covering and the packing problem are NP-hard in their edge and vertex versions. However, as we show in Section 4.1, the edge version of the covering problem is equivalent to the maximum cut problem and hence it is polynomial when restricted to planar graphs [32]. On the other hand, the covering problem in its vertex version and the packing problem in its both versions, are NP-hard even for planar graphs (see Sections 4.1 and 3.4).

We start by stating a useful relation between the odd cycles and the odd faces in a planar graph. We consider an embedding of the planar graph G. When a vertex v is deleted from G, all the faces incident to v are merged together in a new face F_v. The other faces are unchanged. We denote the new face by a capital letter to stress the fact that it determines a set of faces of G, namely, the faces of G included in it. The *parity* of a face of G is defined as the parity of the edge set of its boundary, counting bridges twice. By induction on the number of faces incident to v, the parity F_v equals the sum mod 2 of the parities of the faces of G it contains. As a corollary, we obtain that G is bipartite if and only if all its faces are even.

Lemma 2.3. *Given a planar graph G and a vertex $v \in V(G)$, removing v from G creates a new face F_v which contains all faces of G incident to v. Let $par(f)$ denote the parity of face f, then $par(F_v) = \sum_{f \text{ incident to } v} par(f) \pmod 2$.* □

Corollary 2.4. *A planar graph G is bipartite if and only if all its faces are even.*

Proof. Clearly, if a planar graph is bipartite, all its faces are even. On the other hand, assume all faces are even and fix an embedding of G. Any cycle C divides the embedding of G into two regions of the plane, G_1 and G_2. By removing all vertices of G in one of these two regions but not the vertices in C, say in G_1, the boundary of C becomes the boundary of a face. This new face contains all faces of G in region

G_1. But from the previous lemma, the parity of this new face must be even. Hence C is even in G. I.e., all cycles in G are even and therefore, G is bipartite. □

Next we present the definition of a T-join and state some useful facts concerning T-joins, T-cuts and odd cycle covers.

Definition 2.5. Given a graph $G = (V, E)$ and $T \subseteq V$, a *T-join* is a subset J of the edges of G such that T is exactly the set of odd degree vertices in the subgraph of G induced by J.

There exists a T-join in G if and only if each connected component of G contains an even number of vertices of T. In particular, if G has a T-join then $|T|$ is even.

Definition 2.6. Given a graph $G = (V, E)$, a *T-cut* in G is a subset $\delta(S) \subseteq E$ such that $S \subseteq V$ and both $|T \cap S|$ and $|T \cap (V - S)|$ are odd.

Here $\delta(S)$ is the subset of the edges of G with one endpoint in S and the other in $V - S$. Consequently, a T-cut is simply a cut having an odd number of vertices of T on each side of the cut. Clearly, for G to have a T-cut, $|T|$ must be even.

Definition 2.7. Two sets X and Y of vertices of a graph G are said to be *laminar* if either $X \subseteq Y$ or $Y \subseteq X$ or $X \cap Y = \varnothing$. The sets X and Y are *cross-free* when they are laminar or $X \cup Y = V$.

Definition 2.8. A collection of subsets of the vertices of a graph G is said to be *laminar* (resp. *cross-free*) if any two of its members are laminar (resp. cross-free).

In particular, when we consider a collection of cuts in G, i.e., a family $\delta(\mathcal{F}) = \{\delta(X) : X \in \mathcal{F}\}$, where $\mathcal{F}$ is a family of subsets of $V(G)$, the collection of cuts $\delta(\mathcal{F})$ is said to be *laminar* (resp. *cross-free*) whenever $\mathcal{F}$ is.

The *size* or the *length* of a T-join is the number of edges it contains. A *packing* of T-cuts is a collection of edge disjoint T-cuts. Because every T-join intersects every T-cut, the minimum length of a T-join in G is at least the maximum size of a packing of T-cuts in G. In fact, equality holds for bipartite graphs.

Proposition 2.9 (Seymour [49]). *Let H be a bipartite graph and let T be an even set of vertices of H. The minimum length of a T-join in H equals the maximum size of a packing of T-cuts in H. Moreover, the maximum is attained by a cross-free collection of T-cuts.* □

The last proposition implies the next, where a *2-packing* of T-cuts is a collection of T-cuts such that each edge is contained in at most two T-cuts of the collection.

Proposition 2.10 (Lovász [33]). *Let H be a graph and T be an even set of vertices of H. The minimum length of a T-join in G equals half the maximum cardinality of a 2-packing of T-cuts in H. The maximum is attained by a cross-free collection of T-cuts.* □

The following observation will be useful.

Observation 2.11. *In Propositions 2.9 and 2.10, there exists an optimal collection of T-cuts which is laminar and consists on inclusion-wise minimal T-cuts, with respect to $\delta(X)$.*

Proof. Let $\mathcal{F}$ denote a collection of subsets of $V(G)$ such that $\delta(\mathcal{F})$ is optimal. Now assume that $\mathcal{F}$ is chosen in such a way that the total length of $\delta(\mathcal{F})$, that is $\sum_{\delta(X)\in\delta(\mathcal{F})} |\delta(X)|$, is minimum. Then each T-cut in $\delta(\mathcal{F})$ is inclusion-wise minimal. Otherwise, we could replace any non-minimal T-cut by a smaller T-cut and uncross the resulting collection of T-cuts by standard uncrossing techniques (see, e.g., Proposition 3.4 in [20] or Section 80.7b in [48]) and obtain a new cross-free packing (resp. 2-packing) of T-cuts with the same size and a shorter total length, a contradiction. Now let t denote any element of T. Whenever some member X of $\mathcal{F}$ contains t, we replace it by its complement $\bar{X}$. Because two sets X and Y are cross-free if and only if $\bar{X}$ and Y are cross-free, the resulting collection $\mathcal{F}$ is cross-free. Moreover, $\mathcal{F}$ is laminar because none of its members contains t. $\square$

Now we point out the relationship between edge covers and T-joins in the dual graph. It will be used to give a short proof of a result by Král and Voss [30] (see Sections 1 and 3.1). This relationship was first pointed out by Hadlock [26]. Let G be a planar graph and consider $G^* = (V^*, E^*)$ the dual graph of G. Let $T \subseteq V^*$ be the set of odd degree vertices in G^*. Hence, $|T|$ is even and corresponds to the set of odd faces of G. Note that given any subset J of E, $G - J$ is bipartite if and only if $G - J$ has only even faces if and only if $G^* - J^*$ has only vertices of even degree. However, this last fact implies that J^* is a T-join in G^*, since T is exactly the set of odd vertices of G^*. Therefore, we have proved

Proposition 2.12. *Let G be a planar graph and let G^* denote its face-dual graph. Let T be the set of odd degree vertices in G^* (that is, T corresponds to the set of odd faces of G). A subset of edges F of G is an odd cycle edge cover if and only if $F^* = \{e^* \;:\; e \in F\}$ is a T-join in G^*. Hence, the minimum size of a T-join in G^* is exactly the minimum size of an edge cover in G.* $\square$

We now show how odd cycle vertex covers relate to T-joins. This relationship will be used to prove our main theoretical result, a linear bound of the minimum size of a vertex cover in terms of the size of a maximum packing in planar graphs (see Section

3.2). It will also be used to prove the correctness of the linear time algorithm to find a minimum odd cycle cover in planar graphs where the size of the covering is known to be small (see Section 4).

Let G be a planar graph and let T denote the set of odd faces of G. Hence $|T|$ is even. The *parity* of a face equals the parity of its boundary, counting bridges twice. The *face-vertex incidence graph* of G is the graph G^+ on the faces and vertices of G whose edges are the pairs fv, where f is a face of G and v is a vertex of G incident to f. Note that G^+ is bipartite. Moreover, it is planar because it can be drawn in the plane as follows. Keep all vertices of G as vertices of G^+ and add a new vertex v_f in each face f of G. Then link each new vertex v_f to the vertices of G which are incident to f by an arc whose interior is contained in f. Do this in such a way that two arcs never have a common interior point. The resulting drawing of G^+ is referred to as a *standard drawing*. Below, and henceforth, $F(G)$ denotes the face set of G.

Observation 2.13. *Let $\delta(X)$ be a T-cut in the face-vertex incidence graph G^+ and let R denote the subgraph of G determined by the edges incident to a face in X and to a face in $\bar{X}$. Then R is Eulerian. Furthermore, R contains an odd cycle.*

Proof. Pick some vertex v of G. Let $f_1, \ldots, f_d$ denote the faces of G incident to v listed in clockwise order. Each face f_i belongs either to X or to $\bar{X}$. Because there is an even number of switches between X and $\bar{X}$ when one goes clockwise around v, the degree of v in R is even. In other words, R is Eulerian. So it can be decomposed into edge disjoint cycles. Consider each of these cycles. Since X contains an odd number of vertices of T, that is, an odd number of odd faces of G, at least one of these cycles surrounds an odd number of odd faces of G. Hence, by Lemma 2.3, this cycle is in fact an odd cycle. $\square$

Lemma 2.14. *A subset W of $V(G)$ is an odd cycle cover of G if and only if the subgraph of the face-vertex incidence graph G^+ induced by $W \cup F(G)$ contains a T-join, that is, every component of the subgraph has an even number of vertices of T.*

Proof. We first prove the forward direction. Suppose, by contradiction, that some connected component X of the subgraph of G^+ induced on $W \cup F(G)$ contains an odd number of vertices of T. Then $\delta(X)$ is a T-cut in G^+. Consider the edges of G incident to a face in X and to a face in $\bar{X}$. These edges determine a subgraph R of G. Let e be an edge of R. None of the endpoints of e belongs to W because otherwise all the faces incident to this endpoint would be in X and e would not belong to R, a contradiction. Therefore, R is vertex disjoint from W. By Observation 2.13, we know that R contains an odd cycle. So W is not a cover, a contradiction.

To prove the backward direction, consider an odd cycle C and a T-join J in G^+ covering some vertices of W and no vertex of $G - W$. Let Y be the set of faces of G contained in C and let $X = Y \cup W$. Because C is odd, X contains an odd number of odd faces, that is, an odd number of elements of T. Because $|T|$ is even, there is an odd number of elements of T in $\bar{X}$ too. It follows from Definition 2.5 that J contains a path P from an element of T in X to an element of T in $\bar{X}$. Let v be any vertex of G on P incident to a face in X and to a face in $\bar{X}$. Then v is a vertex of C covered by J. In other words, W intersects C. Therefore, W is a cover. $\square$

3 Approximate Min-Max relations

We begin this Section with our short proof of Král and Voss' result [30]. That is, we show that $\tau_e \leq 2\nu_e$ for planar graphs, where τ_e denotes the size of a minimum odd cycle edge cover and ν_e denotes the size of a maximum packing of edge disjoint odd cycles. The rest of the section is devoted to a proof of the fact that $\tau \leq 10\nu$ in planar graphs, where τ denotes the minimum size of an odd cycle vertex cover and ν denotes the maximum cardinality of a packing of vertex disjoint odd cycles. To do this, we first consider some special subclasses of planar graphs for which we can prove that $\tau \leq 2\nu$. The techniques we develop for these special cases will then be combined to give our main result for general planar graphs.

3.1 A Min-Max relation in the edge case

The following result was proved by Král and Voss [30]. Here we present an alternative (and shorter) proof. Our proof relies, as does the one given in [30], upon the relationship between T-joins and odd cycle covers (see Section 2) and upon the Four Color Theorem [2, 44].

Theorem 3.1. *[30] Let $G = (V, E)$ be a planar graph. Let ν_e denote the maximum number of edge disjoint odd cycles in G and let τ_e denote the minimum size of an odd cycle edge cover. Then $\tau_e \leq 2\nu_e$.*

Proof. [18] The theorem trivially holds if $\nu_e = 0$. Assume that $\nu_e > 0$ and fix a planar embedding of G such that the outer face is odd. Consider $G^* = (V^*, E^*)$ the dual graph of G. Let $T \subseteq V^*$ be the set of odd degree vertices in G^*. Hence, $|T|$ is even and it corresponds to the set of odd faces of G. By Proposition 2.12, the minimum size of a T-join in G^* is exactly τ_e. Then, by Proposition 2.10 and Observation 2.11, there is a laminar family $\mathcal{F}$ of $2\tau_e$ subsets of $V(G^*)$ such that $\delta(\mathcal{F}) = \{\delta(X) : X \in \mathcal{F}\}$ is a 2-packing of inclusion-wise minimal T-cuts in G^*. Without loss of generality, we can assume that no member of $\mathcal{F}$ contains the outer face of G (from the proof of Observation 2.11).

Now let H denote the graph with vertex set $\mathcal{F}$ in which X and Y are adjacent whenever the corresponding T-cuts intersect. We claim that H is planar. Clearly, the claim implies the theorem because, by the Four Color Theorem, H has a stable set of size at least $|V(H)|/4 = 2\tau_e/4 = \tau_e/2$. As this stable set corresponds to a set of edge disjoint T-cuts, any cover must have at least $\tau_e/2$ edges. Hence we have $\tau_e \leq 2\nu_e$. In order to show that H is planar, it suffices to show that every block K of H is planar. Let $\mathcal{F}'$ denote the vertex set of K. Since $\mathcal{F}$ is laminar, $\mathcal{F}'$ is also laminar and the set $\mathcal{F}'$ partially ordered by inclusion is a forest, i.e., every point is covered by at most one point. Let X, Y and Z be three distinct elements of $\mathcal{F}'$. The following cannot occur: (i) $X \subsetneq Y \subsetneq Z$, (ii) $X \subsetneq Y$ and $Y \cap Z = \varnothing$. Indeed, if (i) or (ii) holds then every X–Z path in K intersects Y because $\delta(\mathcal{F})$ is a 2-packing. This contradicts our assumption that K is a block of H. Then $\mathcal{F}'$ partially ordered by inclusion is either a forest of height 0 (that is, an antichain) or a tree of height 1. In both cases, it is easy to construct a planar drawing for K from G. Each element of $\mathcal{F}'$ determines a cycle in the plane graph G. In the first case, we pick any point in the bounded face of each of these cycles and connect the points by an arc whenever there is an edge in K between the two corresponding elements of $\mathcal{F}'$. This can be done in such a way that the resulting graph is planar. The second case is simpler. For if $\mathcal{F}'$ is a tree of height 1, all T-cuts that correspond to children are disjoint (as

$\mathcal{F}$ is a 2-packing of T-cuts). Hence, K itself is a tree of height 1. $\square$

The following example shows that the bound in Theorem 3.1 is tight. Consider the hexagonal grid with at least k vertices and replace k of the vertices with a gadget as shown in Figure 2. (The gadget we use here is the same as the one used in [30]).

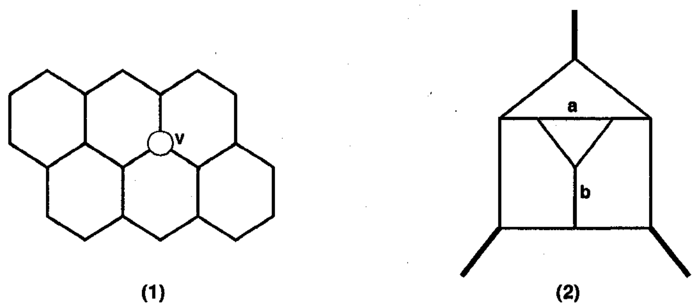

Figure 1: Vertex v is replaced by the gadget shown in (2). The heavy edges in (2) are the original edges adjacent to vertex v. Exactly k of these replacements are performed.

Trivially, $\nu_e \geq k$ since the k gadgets are edge disjoint and each of them contains an odd cycle. Notice that to remove all the odd cycles in a single gadget it is necessary to remove at least two edges, therefore $\tau_e \geq 2k$. On the other hand, if we remove edges a and b for every one of the k gadgets, every face is even and by Corollary 2.4, the resulting graph is bipartite. This implies that $\tau_e \leq 2k$. It also implies that every odd cycle in the graph contains an edge of type a or b. Thus, every odd cycle in the graph contains at least two vertices of the inner triangle in at least one gadget. Hence, if there were $k+1$ disjoint odd cycles, by pigeon-hole principle, two of them must contain two vertices of the inner triangle of the same gadget, so they share a vertex. Since the graph is cubic this means they share an edge. Therefore, $\nu_e \leq k$. So we have $\tau_e = 2\nu_e$.

Corollary 3.2. *Let ν_v and τ_v denote the maximum cardinality of a packing of vertex disjoint odd cycles and the minimum size of an odd cycle vertex cover, respectively. Then $\tau_v \leq 2\nu_v$ for all planar graphs of maximum degree at most 3.*

Proof. For any graph G we have: (i) $\tau_v \leq \tau_e$ and (ii) $\nu_v \leq \nu_e$. Inequality (i) follows immediately from the fact that for any edge cover we can build a vertex-cover of the same size by just choosing one vertex adjacent to each edge in the edge-cover; (ii) holds since vertex disjoint implies edge disjoint. In the case that G has maximum degree at most 3, any pair of edge disjoint cycles is also vertex disjoint (since any vertex has at most 3 incident edges). Hence, $\nu_v = \nu_e$. Combining this last equality with (i) and Theorem 3.1, we have the result. □

Moreover, the result stated in Corollary 3.2 is tight since the example showed before is cubic and it can be easily proved (following the same arguments for the edge case) that $\tau_v = 2\nu_v$.

3.2 A Min-Max relation in the vertex case

In what follows, *cover* stands for odd cycle vertex cover and *packing* stands for packing of vertex disjoint odd cycles. We show that the minimum size of a cover is at most twice the maximum size of a packing for a collection of special classes of planar graphs. For technical reasons it will be useful to assume that G is signed. A *signed* graph is a graph whose edges are labeled *odd* ('$-$') or *even* ('$+$'). In a signed graph, a cycle (or more generally a subgraph) is said to be *odd* if it contains an odd number of odd edges and *even* otherwise. Similarly, a face of a plane signed graph is said to be *odd* if its boundary has an odd number of odd edges, counting bridges twice. Otherwise, the face is said to be *even.* A signed graph is said to be *balanced* if it has no odd cycles. Odd cycle covers and odd cycle packings are defined as in the unsigned case. We denote by $\tau(G)$ and $\nu(G)$ the minimum size of a cover of G and the maximum size of a packing in G, respectively. We assume that G has no loops and no multiple edges, with one exception. We allow *odd digons*, that is, subgraphs with two vertices and two edges between them, one of which is odd and the other even. Clearly, unsigned graphs are a special case of signed graphs where every edge is odd.

3.2.1 When all odd faces mutually intersect

We begin by stating two technical lemmas. Our first lemma is Lemma 4.1.2 in Diestel's book [14]. By arc we mean a continuous curve in $\mathbb{R}^2$. If P is an arc, by $\mathring{P}$ we mean P minus its endpoints.

Lemma 3.3. *Let P_1, P_2 and P_3 be three arcs, between the same two endpoints but otherwise disjoint. Then (i) $\mathbb{R}^2 \setminus (P_1 \cup P_2 \cup P_3)$ has exactly three regions, with frontiers $P_1 \cup P_2$, $P_2 \cup P_3$ and $P_1 \cup P_3$, and (ii) if P is an arc between a point in $\mathring{P}_1$ and a point in $\mathring{P}_3$ whose interior lies in the region of $\mathbb{R}^2 \setminus (P_1 \cup P_3)$ that contains $\mathring{P}_2$, then $\mathring{P} \cap \mathring{P}_2 \neq \varnothing$.* □

To prove our next lemma, we implicitly use the two following facts: in a 2-connected plane graph, every face is bounded by a cycle and every edge is in two distinct faces (one for each side of the edge).

Lemma 3.4. *Let G be a 3-connected plane graph. Then every face of G is bounded by a cycle. For any two faces f and f' of G whose respective boundaries C and C' intersect, the following holds. Either C and C' share exactly one vertex, or two adjacent vertices and the edge between them.*

Proof. Suppose otherwise. Then C and C' share at least two vertices. Pick any vertex u in their intersection. Now walk along C' from u in some direction until some other vertex of C, say v, is hit. Let P_2 be the u–v path in C' that we have walked through, and let P_1 and P_3 be the two u–v paths contained in C. Because G is 3-connected, $G - u - v$ is connected. In particular, if P_1 and P_3 both contain an internal vertex, then there is a path P in $G - u - v$ from an internal vertex of P_1 to an internal vertex of P_3. By shortening P if necessary, we can assume that no internal vertex of P is on P_1 or P_3. The path P satisfies the hypotheses of Lemma 3.3, so it must intersect P_2 in some internal vertex. But then f' is not a face, a contradiction. So, without loss of generality, we can assume that P_3 has no internal vertex. Now exchange the roles of C and C'. Let $P_1' = P_2$, $P_2' = P_3$ and P_3' be the u–v path in C' that is distinct from P_1'. By the same arguments as above, we infer that P_1' or P_3' has no internal vertex. In other words, uv is an edge of both C and C'. Then C and C' have to share at least three vertices. Let w denote the first vertex of C we hit when we walk along C' from v in the same direction as before (away from u). As before, we have that vw is an edge of both C and C'. Hence C and C' share the path of length 2 with vertex sequence u, v, w. It follows that v has degree 2 in G, contradicting the fact that G is 3-connected. □

The following result will be used as a base case to prove our result for general planar graphs.

Proposition 3.5. *If every two odd faces of G have intersecting boundaries, then G has an odd cycle cover of size at most* 2.

Proof. Note that the boundaries of every pair of odd faces of G intersect if and only if every two odd cycles of G intersect, that is, if and only if $\nu(G) \leq 1$. We prove by induction on the number of vertices of G that $\nu(G) = 1$ implies $\tau(G) \leq 2$. This clearly implies the proposition. If G has at most three vertices or has an odd digon, then the proposition trivially holds. Now assume that G is simple and has at least four vertices. We claim that we can also assume that G is 3-connected.

If G is not 3-connected, then it has a cutset consisting of two vertices u and v. Let U be a connected component of $G - u - v$, let G_1 denote the subgraph of G induced on $U \cup \{u, v\}$ and let $G_2 = G - U$. If both G_1 and G_2 are unbalanced then $\{u, v\}$ is a cover because we have $\nu(G) = 1$. By symmetry, we can assume that G_1 is balanced. Then all u–v paths in G_1 have the same parity. Let G' be the graph obtained from G_2 by adding an edge e with endpoints u and v that is labeled odd if all u–v paths in G_1 are odd, and even otherwise. We don't add edge e if there is already an edge between u and v or if there is no u–v path in G_1. Because G' has less vertices than G and $\nu(G') = 1$, there is a cover of cardinality 2 in G'. The same two vertices form a cover in G. This concludes the proof of our first claim.

From now on, we assume that G is 3-connected. We claim that if every vertex of G is incident to at most three odd faces, then G has a cover of size 2. Indeed, if it is the case then consider the *intersection graph* of the odd faces, that is, the graph whose vertices are the odd faces of G and whose edges are the pairs ff' where f and f' have a common incident vertex. The intersection graph is complete because any two odd cycles in G have a common vertex. Any standard drawing of the face-vertex incidence graph G^+ can be modified to obtain a drawing of the intersection graph, so the latter is planar. It follows that G has either 2 or 4 odd faces. If there are 2 odd faces f_1 and f_2, let v be a vertex incident to both f_1 and f_2. By Lemma 2.14, $\{v\}$ is a cover. If there are 4 odd faces f_1, f_2, f_3 and f_4, let v be a vertex incident to both f_1 and f_2 and w be a vertex incident to both f_3 and f_4. By Lemma 2.14, $\{v, w\}$ is a cover. This concludes the proof of our second claim.

Let now v be a vertex which is incident to at least four odd faces, say, f_1, f_2, f_3 and f_4, in counterclockwise order. By Lemma 3.4, the boundaries of f_i and f_{i+2} share exactly one vertex, namely v, for $i = 1, 2$. If all the odd faces are incident to v, then $\{v\}$ is a cover by Lemma 2.14 and we are done. So we can assume there is an odd face f that is not incident to v. Let f' be an odd face distinct from f which is also not incident to v. If there is no such face, then $\{v, w\}$ is a cover (again by Lemma 2.14), where w is any vertex incident to both f and f_1.

Consider the subgraph H of G^+ obtained by adding to the subgraph of G^+ induced on v, f_1, f_2, f_3 and f_4 four paths of length two from f to f_1, f_2, f_3 and f_4 respectively. Let f, u_i, f_i be the vertex sequence of the i-th path, and let $U = \{u_1, u_2, u_3, u_4\}$. (See examples in Figure 2.) We chose the paths in such a way that the number of vertices in U is minimum (we maximize the intersections between them). In other words, we ask that H be an induced subgraph of G^+. Now consider any standard drawing of G^+. Where is the vertex $v_{f'}$ corresponding to the odd face f'? It has to lie in one of the faces of H. Moreover, there is a path of length 2 in G^+ from f' to each of the f_i's. Each path avoids v because f' is not incident to v. Because the f_i's are arranged around v in counterclockwise order, Lemma 3.4 implies $u_1 \neq u_3$ and $u_2 \neq u_4$. So it suffices to consider the following three cases (see Figure 2).

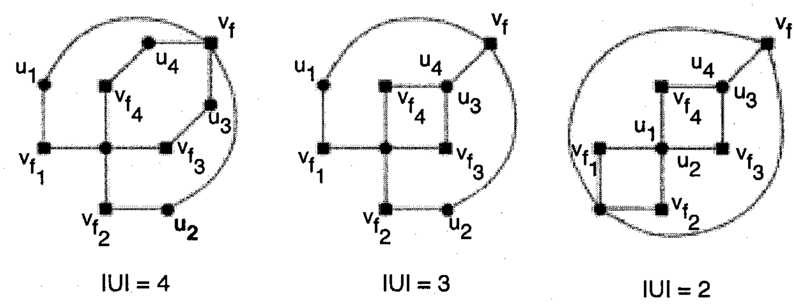

Figure 2: Subgraph H in each of the three cases.

Case 1. $|U| = 4$. Without loss of generality, $v_{f'}$ lies in the face of H bounded by the cycle with vertex sequence v, v_{f_1}, u_1, v_f, u_2, v_{f_2}, v. Then there cannot be a path of length 2 avoiding v from f' to f_3 in G^+, a contradiction.

Case 2. $|U| = 3$. Without loss of generality, we assume that $u_3 = u_4$. As in Case 1, we see that there is nowhere $v_{f'}$ could be. For instance, if it lies in the face bounded by the cycle with vertex sequence v, v_{f_1}, u_1, v_f, u_4, v_{f_4}, v, then there cannot be a path of length 2 avoiding v from f' to f_2 in G^+, a contradiction.

Case 3. $|U| = 2$. Without loss of generality, we assume that $u_1 = u_2$ and $u_3 = u_4$. Vertex $v_{f'}$ can be in the face bounded by the cycle with vertex sequence v, v_{f_2}, u_2, v_f, u_3, v_{f_3}, v or in the face bounded by the cycle with vertex sequence v, v_{f_1}, u_1, v_f, u_4, v_{f_4}, v. In both cases, it must be adjacent to both u_1 and u_3. By Lemma 2.14, we see that $\{v, u_1\}$ is a cover of G. □

3.2.2 When some face intersects every odd face

Now we consider the graphs that have some face whose boundary intersects the boundary of every odd face.

Proposition 3.6. *Assume G is 4-connected, simple, has at least five vertices and is such that the boundary of the outer face intersects the boundary of every odd face. Then the minimum size of an odd cycle cover of G is at most twice the maximum size of a packing in G.*

Proof. We assume that G is not balanced. Otherwise, the result trivially holds. The hypotheses severely restrict the way face boundaries intersect each other. Consider a vertex y not incident to the outer face. If there are two distinct vertices x, z such that each of them is incident to the outer face and to a face incident to y, then x and z have to be adjacent. Suppose not. Then there exists a polygon P in $\mathbb{R}^2$ intersecting G exactly in x, y and z. By the Jordan Curve Theorem, we know that all paths from a vertex of G in the bounded region of $\mathbb{R}^2 \setminus P$ to a vertex of G in the unbounded region of $\mathbb{R}^2 \setminus P$ go through x, y or z. The situation is depicted in Figure 3. If x and z are not adjacent, then the two neighbours of x on the boundary of the outer face lie in a different region of $\mathbb{R}^2 \setminus P$. Hence $X = \{x, y, z\}$ is a cutset of size 3 in G, a contradiction.

Now consider two distinct and intersecting odd faces f and g different from the outer face. By Lemma 3.4, the boundaries of faces f and g intersect in a vertex or in a common edge. If the boundaries of f and g share a unique vertex y, then y is incident to the outer face (see Figure 4.a), unless the following occurs. There are vertices x and z incident to the outer face such that xy, xz and yz are edges, x is incident to f

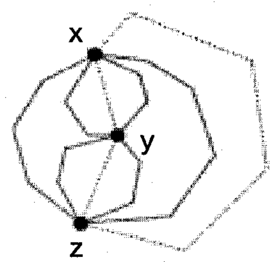

Figure 3: $X = \{x, y, z\}$ is a cutset.

and z is incident to g (see Figure 4.b). Moreover, x is the only vertex incident to both f and the outer face, and z is the only vertex incident to both g and the outer face. We refer to the triangle on x, y and z as a *junctional triangle.* Note that junctional triangles can be even because G is signed. If the boundaries of f and g intersect in a common edge e (see Figure 4.c), then one of the endpoints of e is on the outer face and the other is not. Moreover, in that case f and g cannot both have a common incident edge with the outer face because otherwise $G = K_4$, contradicting the fact that G has at least five vertices.

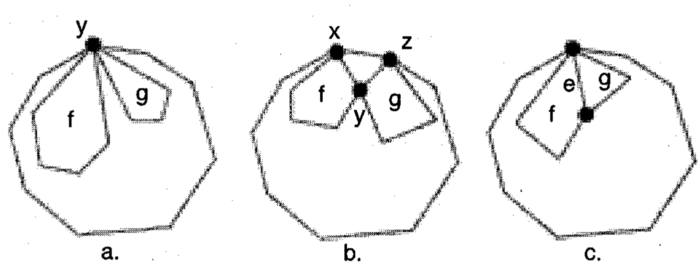

Figure 4: The three ways the boundaries of f and g can intersect.

Enumerate the vertices of the outer face in clockwise order as v_1, v_2, $\ldots$, v_n. For the sake of simplicity, let $v_0 = v_n$ and $v_{n+1} = v_1$. Let I be the set of indices i such that there is a junctional triangle containing the edge v_iv_{i+1}. For each $i \in I$, we let

u_i be the vertex of the junctional triangle not incident to the edge $e = v_iv_{i+1}$ and opposite to e, and we let w_i be any point in the interior of the edge e. For each odd face f different from the outer face, we define an arc A_f contained in the frontier of the outer face, as follows. If the boundary of f intersects the boundary of the outer face in an edge v_iv_{i+1}, then we let A_f be the edge v_iv_{i+1}. Otherwise, the boundary of f intersects the boundary of the outer face in a vertex v_i. If f is incident neither to u_{i-1} nor to u_i then we let A_f be the point $\{v_i\}$. If f is incident to u_{i-1} and not to u_i then we let A_f be the part of the edge $v_{i-1}v_i$ between w_{i-1} and v_i. If f is incident to u_i and not to u_{i-1} then we let A_f be the part of the edge v_iv_{i+1} between v_i and w_i. Finally, if f is incident to both u_i and u_{i-1} then we let A_f be the arc linking w_{i-1} and w_i on the outer face and containing v_i. By construction, two odd faces f and g different from the outer face are incident to some common vertex if and only if their corresponding arcs A_f and A_g have a nonempty intersection.

Let H denote the graph whose vertices are the odd faces different from the outer face and whose edges are the pairs fg such that $A_f \cap A_g \neq \varnothing$. Then H is a circular arc graph. The maximum size of a packing in G is precisely equal to the maximum size of a stable set in H, that is, we have $\nu(G) = \alpha(H)$. We consider the following two cases.

Case 1. There is some point x on the boundary of the outer face that is not in any arc A_f. In this case, H is an interval graph. Let W be a minimum cardinality subset of $\{v_i : 1 \leq i \leq n\} \cup \{w_i : i \in I\}$ meeting all the arcs. By Dilworth's chain partitioning theorem [15], the complement of an interval graph H is perfect, hence we have $|W| = \alpha(H) = \nu(G)$. Now replace each $w_i \in W$ by v_i and v_{i+1}. Let W' be the resulting set of vertices of G. Then W' is a cover of G of cardinality at most

$2|W|$. In other words, we have $\tau(G) \leq 2\nu(G)$.

Case 2. The arcs A_f cover the whole boundary of the outer face. It follows that for each edge $e = v_i v_{i+1}$, the face f_i incident to e and different from the outer face is either an odd face or an even junctional triangle. If f_i is odd, then we let $g_i = f_i$. Otherwise, we let g_i be the odd face incident to $v_i u_i$. As above, u_i denotes the vertex of the junctional triangle incident to e which is opposite to e. Thus, every edge of the outer face has a corresponding odd face. Note that if the boundaries of g_i and g_j intersect then $i \in \{j-1, j+1\}$ or $i = j$. So if n is even, then we have $\tau(G) \leq 2\nu(G)$ because $\{v_1, \ldots, v_n\}$ is a cover of size n and $\{f_1, f_3, \ldots, f_{n-1}\}$ yields a packing of size $n/2$. Now assume that n is odd. Let H' be the graph whose vertices are the faces g_i and whose edges are the pairs $g_i g_j$ such that the boundary of g_i intersects that of g_j and $i \neq j$. By what precedes, we know that the graph H' is a subgraph of the odd cycle with vertex sequence $v_1, \ldots, v_n, v_1$. If H' is not connected, then it has a stable set of size $(n+1)/2$ and we get $\tau(G) \leq 2\nu(G)$ as before. For the rest of the proof, we assume that H' is connected. We claim that either all f_i's are odd faces or all f_i's are even junctional triangles. Otherwise, there is some index i such that f_i is an even junctional triangle and f_{i+1} is an odd face. By what precedes, the boundaries of g_i and g_{i+1} cannot intersect. So our claim holds.

If all f_i's are odd faces then consider vertex v_1. If $\{v_1, \ldots, v_n\} \setminus \{v_1\}$ is a cover then we have $\tau(G) \leq n - 1 \leq 2\nu(G)$ because H' has a stable set of size $(n-1)/2$. Otherwise, there is some odd face f incident to v_1 and to no other v_i. Then $\{f\} \cup \{f_2, f_4, \ldots, f_{n-1}\}$ yields a packing of size $(n+1)/2$. Hence, we have $\tau(G) \leq 2\nu(G)$. If all f_i's are even junctional triangles, then the odd faces of G are exactly the outer face and the faces g_i for $i = 1, \ldots, n$. It is easy to see that $\{u_1, v_3, v_4, \ldots, v_n\}$ is a cover of size $n - 1$.

Because H' has a stable set of size $(n-1)/2$, we have $\tau(G) \leq 2\nu(G)$. This concludes the proof. $\square$

We need a slight generalization of Proposition 3.6. Consider some face f of G, which we refer to as a *center*. The odd faces of G whose boundary intersects the boundary of the center are called the *targets* (around f). In particular, if f is odd then f is itself a target. A *local cover* is a set W of vertices of G satisfying the following properties:

(i) every target is incident to some vertex of W;

(ii) at most one vertex of W is not incident to the center;

(iii) if $u \in W$ is not incident to the center, then u is incident to exactly two targets.

The proof of Proposition 3.6 in fact shows:

Lemma 3.7. *Assume G is 4-connected, simple and has at least five vertices. Let f be a face of G acting as center. Then the minimum size of a local cover of G is at most twice the maximum number of boundary-disjoint targets in G.* $\square$

3.2.3 When odd faces are disjoint

We begin this section by recasting the minimum cover and the maximum packing problems entirely in terms of T-joins and T-cuts in the face-vertex incidence graph. This slight change of terminology simplifies the proofs and enables us to state our results with more generality. Let H denote any bipartite graph with bipartition $\{A, B\}$, and let T be any even subset of B. The *width* of a T-join in H is the number of vertices of A it covers. The *fringe* of a T-cut $\delta(X)$ in H is the set of vertices of A which have a neighbour in X and a neighbour in $\bar{X}$. Note that the minimum width of a T-join in H is at least the maximum number of fringe disjoint T-cuts in H. This is due to the fact that every T-join covers some element in the fringe of every T-cut.

We now relate the above definitions to odd cycle vertex cover and packing in plane signed graphs. Consider the case where H is the face-vertex incidence graph G^+ of the plane signed graph G, set A is the vertex set of G, set B is the face set of G, and set T is, as before, the set of odd faces of G. By Lemma 2.14, every T-join in H defines a cover of G, namely, the vertices of A it covers. Reciprocally, to every cover W there corresponds a T-join in H which covers some vertices of W and no vertex of $A \setminus W$. So the minimum width of a T-join in H equals the minimum size of a cover of G. Furthermore, there is a correspondence between T-cuts in H and odd cycles in G. By Observation 2.13, every T-cut in H determines an Eulerian subgraph of G with an odd number of odd edges. This subgraph contains an odd cycle. The vertex set of the subgraph is the fringe of the T-cut. Reciprocally, every odd cycle in G determines a T-cut in H whose fringe is the vertex set of the cycle. Hence the maximum size of a collection of fringe disjoint T-cuts in H equals the maximum size of a packing in G. We use the following notation: let ν denote the maximum size of a collection of fringe disjoint T-cuts in H, let τ denote the minimum width of a T-join in H and let ℓ denote the minimum length of a T-join in H.

Proposition 3.8. *Let H be any bipartite graph with bipartition $\{A, B\}$ and let T denote any even subset of B. Assume that the shortest path distance $d_H(t, t')$ between any two distinct elements t and t' of T is at least $2c$ for some $c \geq 1$. Then we have*

$$\nu \geq \frac{1}{2}(\ell - |T| + 1) \geq \left(1 - \frac{1}{c}\right)\tau.$$

Proof. Let $\mathcal{F}$ denote a laminar collection of ℓ sets of faces and vertices of G such that $\delta(\mathcal{F}) = \{\delta(X) : X \in \mathcal{F}\}$ is a collection of edge-disjoint T-cuts in H. Such a laminar collection of T-cuts is guaranteed to exist by Proposition 2.9 and Observation 2.11.

We claim that whenever X, Y and Z are three distinct elements of $\mathcal{F}$ such that $X \subsetneq Y \subsetneq Z$ or $X \subsetneq Y$ and $Y \cap Z = \varnothing$, then T-cuts $\delta(X)$ and $\delta(Z)$ are fringe disjoint. It suffices to consider the first case. Suppose there exists an element $a \in A$ which belongs to the fringes of X and Z. In particular, a has a neighbor b in X and a neighbor b' in $\bar{Z}$. If $a \in Y$ then $ab' \in \delta(Y) \cap \delta(Z)$, a contradiction. If $a \in \bar{Y}$ then $ab \in \delta(X) \cap \delta(Y)$, a contradiction. So our claim holds.

The set $\mathcal{F}$ partially ordered by inclusion is a forest. Without loss of generality, we can assume that the leaves of this forest are singletons of the form $\{b\}$ for some $b \in T$. It follows that $\mathcal{F}$ has at most $|T|$ leaves. Note that the claim above implies that two nodes of the forest, X and Y, are fringe disjoint unless X is the parent of Y, Y is the parent of X, X and Y are siblings or X and Y are roots.

Rank the children of each node of the forest $\mathcal{F}$ arbitrarily and order its roots arbitrarily also. Let $\mathcal{F}'$ denote the subset of $\mathcal{F}$ formed by all nodes which are ranked first in their respective ordering. Letting λ denote the number of leaves of $\mathcal{F}$, we claim that $\mathcal{F}'$ contains at least $|\mathcal{F}| - \lambda + 1 \geq |\mathcal{F}| - |T| + 1$ elements. Let X be any node. We define a function $f : \mathcal{F} \to \mathcal{F}$ as follows. If X is a leaf then we set $f(X) = X$. Otherwise, we set $f(X) = f(Y)$, where Y is the first child of X. So $f(X)$ is the "first" leaf amongst the descendants of X. Observe that, for any leaf node Z the preimage of Z, under f, either contains exactly one node in $\mathcal{F}'$ or contains the highest ranked root node. Moreover, by construction, the preimages of any pair of leaves are disjoint. Thus, $\lambda = |\mathcal{F}| - |\mathcal{F}'| + 1$. Our second claim follows.

To obtain a packing of fringe disjoint T-cuts, colour the elements of $\mathcal{F}'$ black or white in such a way that no parent and child have the same colour, i.e., whenever X is the parent of Y then X and Y have different colours. In other words, colour the subgraph of the Hasse diagram of $\mathcal{F}$ induced on $\mathcal{F}'$ with two colours. Let $\mathcal{F}''$ denote the biggest of the two colour classes. Then $\delta(\mathcal{F}'')$ is a collection of fringe disjoint T-cuts of size at least $\frac{1}{2}(\ell - |T| + 1)$.

Note that every minimum length T-join in G can be thought of as a perfect matching on T whose edges have become edge-disjoint shortest paths in H. Hence ℓ is at least $\frac{|T|}{2}$ times $2c$. Note also that τ is at most $\frac{\ell}{2}$ because the width of any T-join is at most half its length. It follows that we have

$$\nu \geq \frac{1}{2}(\ell - |T| + 1) \geq \frac{1}{2}(\ell - |T|) \geq \left(1 - \frac{1}{c}\right)\frac{\ell}{2} \geq \left(1 - \frac{1}{c}\right)\tau.$$

□

Corollary 3.9. *If the boundaries of the odd faces of G are pairwise disjoint then the minimum size of an odd cycle cover of G is at most twice the maximum size of a packing in G.*

Proof. This follows directly from Proposition 3.8 with $H = G^+$, $A = V(G)$, $B = F(G)$ and $c = 2$. □

3.2.4 Combining the local and global approaches

We are now ready to prove our result for general planar graphs. We will combine the local and global approaches we have described to give our main approximate min-max result. Towards this end, let $\rho(G)$ denote the minimum size of a collection of faces of G such that the boundary of every odd face of G intersects the boundary of some face in the collection. The following two lemmas are simple and we omit their proofs. Lemma 3.10 implies that $\rho(G') \leq \rho(G)$ for any subgraph G' of G.

Lemma 3.10. *Let G be a plane signed graph with $\rho(G) = r$, and let $f_1, \ldots, f_r$ be a collection of faces such that for every odd face f there is an index i such that the boundary of f_i intersects the boundary of f. Then we have $\rho(G-e) \leq r$ for each edge e. Moreover, we have $\rho(G-e) \leq r-1$ if edge e is incident to f_i and f_j for some distinct indices i and j.* □

Lemma 3.11. *Let G be a plane graph and X be a cutset of G with at most three vertices (we allow the case $X = \varnothing$). If G has no cutset with fewer than $|X|$ elements, then there exists a polygon $P \subset \mathbb{R}^2$ intersecting G only in vertices and such that X is precisely the intersection of P and G and each region of $\mathbb{R}^2 \setminus P$ contains a vertex of G.* □

The next lemma, combined with Lemma 3.11, will allow us to focus on 4-connected graphs G.

Lemma 3.12. *Let G be a plane signed graph, let $P \subset \mathbb{R}^2$ be a polygon intersecting G only in vertices, and let $X = P \cap V(G)$. Assume that each region R_1 and R_2 of $\mathbb{R}^2 \setminus P$ contains at least one vertex of G. Then X is a cutset in G. For $i = 1, 2$, let G_i be the part of G contained in the closure of region R_i. Then we have $\rho(G_1) + \rho(G_2) \leq \rho(G) + 2$.*

Proof. Let $\{f_1, \ldots, f_r\}$ denote a collection of $r = \rho(G)$ faces of G such that the boundary of every odd face of G intersects the boundary of some face of the collection.

Without loss of generality, we can assume that there are some indices r_1 and r_2 with $r_1 \leq r_2$ such that $f_1, \ldots, f_{r_1}$ are contained in R_1 and incident to no vertex of X, and $f_{r_2}, \ldots, f_r$ are contained in R_2 and incident to no vertex of X. For $i = 1, 2$, let g_i denote the face of G_i containing R_{2-i+1}. Then the boundary of every odd face of G_1 intersects the boundary of some face in $\{f_1, \ldots, f_{r_1}\} \cup \{g_1\}$. Similarly, the boundary of every odd face of G_2 intersects the boundary of some face in $\{f_{r_2}, \ldots, f_r\} \cup \{g_2\}$. The lemma follows. □

Theorem 3.13. *For every unbalanced plane signed graph G, we have $\tau(G) \leq 7\nu(G) + 3\rho(G) - 8$.*

Proof. Let G be a counterexample with $|V(G)|$ as small as possible, and let $\{f_1, \ldots, f_r\}$ denote any minimum collection of faces of G such that the boundary of every odd face of G intersects the boundary of some face in the collection. Note that we have $r = \rho(G) \geq 1$. We claim: (1) G has a packing of size 2, no cover of size at most 9, and G is simple; (2) G is 4-connected; (3) the shortest path distance $d_{G^+}(f_i, f_j)$ between f_i and f_j is at least 8 whenever $i \neq j$.

Proof of Claim (1). If G has no packing of size 2, then by Proposition 3.5, we have

$$\tau(G) \leq 2 = 7 + 3 - 8 \leq 7\nu(G) + 3\rho(G) - 8,$$

a contradiction. So G has a packing of size 2, that is, we have $\nu(G) \geq 2$. Now a similar argument shows that G has no cover of size at most 9, that is, we have $\tau(G) > 9$. If G is not simple, then it has an odd digon. Letting x and y be the vertices of the digon and $X = \{x, y\}$, we have $\nu(G - X) \leq \nu(G) - 1$. By Lemma 3.10, we have also $\rho(G - X) \leq \rho(G)$. Therefore, because $G - X$ is unbalanced and has less vertices than

G, we have

$$
\begin{aligned}
\tau(G) &\leq 2+\tau(G-X) \\
&\leq 2+7\nu(G-X)+3\rho(G-X)-8 \\
&\leq 7\nu(G)+3\rho(G)+2-7-8 \leq 7\nu(G)+3\rho(G)-8,
\end{aligned}
$$

a contradiction. So Claim (1) holds.

Proof of Claim (2). By the previous claim, G has at least 10 vertices. Therefore, to prove the present claim, it suffices to prove that G has no cutset of size 3. However, in order to use Lemma 3.11 we need to show that G is 3-connected; we leave this straightforward task to the reader. Now assume that G is 3-connected. Suppose that G has a cutset X consisting of three vertices x, y and z. Let $Y = \{y, z\}$. By Lemma 3.11, there exist induced subgraphs G_1 and G_2 of G and a polygon $P \subset \mathbb{R}^2$ determining two regions R_1 and R_2 in the plane such that P intersects G precisely in x, y and z, and G_i equals the restriction of G to the closure of region R_i, for $i = 1, 2$. It suffices to consider the following two cases. Indeed, if $G_1 - X$ and $G_2 - X$ are both balanced then G has a cover of size at most 3, contradicting Claim (1).

Case 1. Neither $G_1 - X$ nor $G_2 - X$ is balanced. It follows that neither $G_1 - Y$ nor $G_2 - Y$ is balanced. If we have $\nu(G) \geq \nu(G_1 - Y) + \nu(G_2 - Y)$ then Lemma 3.12 implies

$$
\begin{aligned}
\tau(G) &\leq 2+\tau(G_1-Y)+\tau(G_2-Y) \\
&\leq 2+7\nu(G_1-Y)+3\rho(G_1-Y)-8+7\nu(G_2-Y)+3\rho(G_2-Y)-8 \\
&\leq 7\nu(G)+3\rho(G)+2+6-16 = 7\nu(G)+3\rho(G)-8,
\end{aligned}
$$

a contradiction. Else, we have $\nu(G) = \nu(G_1 - Y) + \nu(G_2 - Y) - 1$. It follows that every maximum packing of $G_1 - Y$ and every maximum packing of $G_2 - Y$ hit the vertex x. So we have $\nu(G_1 - X) = \nu(G_1 - Y) - 1$, $\nu(G_2 - X) = \nu(G_2 - Y) - 1$ and $\nu(G) = \nu(G_1 - X) + \nu(G_2 - X) + 1$. Therefore, we have

$$\begin{aligned}
\tau(G) &\leq 3 + \tau(G_1 - X) + \tau(G_2 - X) \\
&\leq 3 + 7\nu(G_1 - X) + 3\rho(G_1 - X) - 8 + 7\nu(G_2 - X) + 3\rho(G_2 - X) - 8 \\
&\leq 7\nu(G) + 3\rho(G) + 3 - 7 + 6 - 16 \leq 7\nu(G) + 3\rho(G) - 8,
\end{aligned}$$

a contradiction.

Case 2. $G_1 - X$ is balanced and $G_2 - X$ is not balanced. If G_1 is not balanced, then we have $\nu(G) \geq \nu(G_2 - X) + 1$ and hence

$$\begin{aligned}
\tau(G) &\leq 3 + \tau(G_2 - X) \leq 3 + 7\nu(G_2 - X) + 3\rho(G_2 - X) - 8 \\
&\leq 7\nu(G) + 3\rho(G) + 3 - 7 - 8 \leq 7\nu(G) + 3\rho(G) - 8,
\end{aligned}$$

a contradiction. Otherwise, G_1 is balanced. Consider the graph G_2' obtained from G_2 by adding a triangle on x, y and z to G_2'. We do not add an edge if it is already present in G_2. Since we can easily modify G to get a drawing of G_2', we can regard G_2' as a plane graph. Consider any two distinct vertices u, v in $X = \{x, y, z\}$. Because G_1 is balanced, all u–v paths in G_1 have the same parity. We let the parity of the edge uv in G_2' be the parity of all u–v paths in G_1. Note that we have $\tau(G) \leq \tau(G_2')$ and $\nu(G_2') \leq \nu(G)$. Moreover, we have $\rho(G_2') \leq \rho(G)$, as we now prove. Since G is 3-connected, there is a vertex t in $G_1 - X$ sending three independent paths to x, y and z in G_1. By Lemma 3.10, if we delete from G all edges which are contained in G_1

except those which belong to one of the three paths, the resulting graph G' satisfies $\rho(G') \leq \rho(G)$. Since the triangle on x, y, z determines an even face in G_2', we have $\rho(G_2') \leq \rho(G')$. Hence, we have $\rho(G_2') \leq \rho(G)$, as claimed. It follows that we have

$$\tau(G) \leq \tau(G_2') \leq 7\nu(G_2') + 3\rho(G_2') - 8 \leq 7\nu(G) + 3\rho(G) - 8,$$

a contradiction. In conclusion, Claim (2) holds.

Proof of Claim (3). Suppose that $d_{G^+}(f_i, f_j) \leq 6$ for some distinct indices i and j. Let X denote the set of vertices of G on a shortest path between f_i and f_j in G^+. So X contains at most three vertices. By Lemma 3.10, we have $\rho(G - X) \leq \rho(G) - 1$. Because $G - X$ is not balanced, we have

$$\tau(G) \leq 3 + \tau(G - X) \leq 3 + 7\nu(G - X) + 3\rho(G - X) - 8 \leq 7\nu(G) + 3\rho(G) - 8,$$

a contradiction. So Claim (3) holds.

Now we would like to apply Lemma 3.7 around each face in the collection $\{f_1, \ldots, f_r\}$. So each face f_i will perform as a center. The targets around f_i are the odd faces of G whose boundary intersects the boundary of f_i. By Claim (3), whenever g is a target around f_i and g' is a target around f_j with $i \neq j$, the boundaries of g and g' are disjoint. By Lemma 3.7, for each center f_i there exists a packing of odd cycles $\mathcal{C}_i$ formed by target boundaries, and a local cover W_i whose size is at most twice the size of packing $\mathcal{C}_i$. Let $\mathcal{C}_{\text{local}}$ denote the union of packings $\mathcal{C}_1, \ldots, \mathcal{C}_r$. Then $\mathcal{C}_{\text{local}}$ is a packing.

Now let $H = G^+$, let $A = V(G)$ and let $B = F(G)$. Consider the graph $\tilde{H}$ obtained from H by contracting, for $1 \le i \le r$, all vertices of H at distance at most 2 from f_i to a single vertex $\tilde{f}_i$. Note that $\tilde{H}$ is still bipartite, with bipartition $\{\tilde{A}, \tilde{B}\}$, where

$$\begin{aligned}\tilde{A} &= A \setminus \{a \in A : d_H(a, f_i) \le 2 \text{ for some } i \text{ with } 1 \le i \le r\}, \\ \tilde{B} &= B \setminus \{b \in B : d_H(b, f_i) \le 2 \text{ for some } i \text{ with } 1 \le i \le r\} \cup \{\tilde{f}_i : 1 \le i \le r\}.\end{aligned}$$

Let $\tilde{T}$ denote the set of those $\tilde{f}_i$'s that correspond to centers f_i which have an odd number of targets around them. So $\tilde{T}$ is an even subset of $\tilde{B}$. Let $\tilde{J}$ denote a minimum length $\tilde{T}$-join in $\tilde{H}$. Then, by Proposition 3.8, there is a collection of fringe disjoint $\tilde{T}$-cuts $\delta(\tilde{\mathcal{F}})$ in $\tilde{H}$ such that

$$|\delta(\tilde{\mathcal{F}})| \ge \frac{1}{2}(|\tilde{J}| - |\tilde{T}| + 1) \ge \frac{1}{2}(|\tilde{J}| - r + 1) \Rightarrow |\tilde{J}| \le 2|\delta(\tilde{\mathcal{F}})| + r - 1.$$

This collection of fringe disjoint $\tilde{T}$-cuts yields a packing of odd cycles $\mathcal{C}_{\text{global}}$ in G, of the same size. The $\tilde{T}$-join $\tilde{J}$ defines a set of edges J_{global} in $H = G^+$, as follows. Every edge of $\tilde{J}$ that belongs to H is kept as it is. Every other edge of $\tilde{J}$ is of the form $v\tilde{f}_i$ and is replaced by any shortest path between v and f_i in H. Because we have $d_{\tilde{H}}(\tilde{f}_i, \tilde{f}_j) \ge 4$ whenever $i \ne j$ and because $\tilde{J}$ is the edge-disjoint union of shortest paths between pairs of vertices of $\tilde{T}$, the length of J_{global} is at most twice the length of $\tilde{J}$.

For each local cover W_i, let J_i denote the set of edges vf of the face-vertex incidence graph G^+ such that $v \in W_i$ and f is a target around f_i incident to v. Let J_{local} denote the union of $J_1, \ldots, J_r$. The union of J_{local} and J_{global} contains a T-join, say J. Because the width of J is at most the width of J_{local} plus the width of J_{global} and because the width of a T-join is at most half of its length, the width of J is at most

$$\sum_{i=1}^{r} |W_i| + \frac{1}{2}|J_{\text{global}}| \leq 2|\mathcal{C}_{\text{local}}| + 2|\mathcal{C}_{\text{global}}| + r - 1 \leq 4\nu(G) + \rho(G) - 1.$$

By Claim (1), we have $\nu(G) \geq 2$. Therefore, we have

$$\tau(G) \leq 4\nu(G) + \rho(G) - 1 \leq 7\nu(G) + 3\rho(G) - 8,$$

a contradiction. This concludes the proof of the theorem. □

Because $\rho(G)$ is at most the size of any inclusion-wise maximal collection of boundary disjoint odd faces in G, which is in turn at most $\nu(G)$, we obtain our main result from Theorem 3.13.

Corollary 3.14. *For every plane signed graph G, we have $\tau(G) \leq 10\nu(G)$.* □

3.3 Hardness result for the packing problem

In this Section we prove that packing vertex disjoint odd cycles in planar graphs is NP-hard. This follows from a result by Caprara and Rizzi [12] showing that packing vertex disjoint triangles is NP-hard in planar graphs. They obtain this result from a reduction of Planar 3-SAT.

PLANAR 3-SAT: Let X be a set of n variables and let ϕ be a boolean formula in conjunctive normal form on m clauses $c_1, \ldots, c_m$ over the variables in X, where each clause c_j has exactly 3 literals. Consider the bipartite graph B with vertex set $X \cup \{c_1, \ldots, c_m\}$ and edges xc if variable x occurs in clause c (in its affirmative or negative form). The formula ϕ is called planar if the graph B is planar. The problem Planar 3-SAT consists of finding, if it exists, a truth assignment for the variables in X that satisfies ϕ, i.e., that satisfies all clauses in ϕ. It is known that Planar 3-SAT is NP-complete [32].

In the reduction of [12], given an instance ϕ of Planar 3-SAT, a graph $G(\phi)$ is constructed with the following structure:

1) For each variable x_i there is a gadget which we call a *sun*. It consists on an even cycle of length $2m_i$, where m_i is the number of clauses in which x_i occurs. We call this cycle the *inside* cycle of the *sun*. From each edge in the cycle there is a triangle pointing outside the cycle. The vertices of the $2m_i$ triangles which are not incident to the even cycle are labeled a_1^i, b_1^i, a_2^i, b_2^i, ..., $a_{m_i}^i$, $b_{m_i}^i$

An example of the gadget built for a variable with 4 occurrences is shown in Figure 5(a).

2) For each clause c_j there is a gadget which we call a *starred-pentagon*. It consists of a pentagon with five triangles, one per edge, pointing to the outside. The vertices of the triangles which are not incident to the pentagon are labeled t_1^j, t_2^j, t_3^j and q_1^j and q_2^j. An example is shown in Figure 5(b). By t-triangle (q-triangle) we mean a triangle on a *starred-pentagon* whose vertex is labeled t (q).

3) Connecting *starred-pentagons* to *suns*: If variable x_i appears in clause c_j in its affirmative (respectively negative) form, then an unused vertex of the type t_r^j, $r = 1, 2, 3$, is identified with an unused vertex of the type a_s^i (respectively b_s^i), $s = 1, 2, \ldots, m_i$.

As each clause has exactly 3 literals, every *starred-pentagon* has each of its three t-

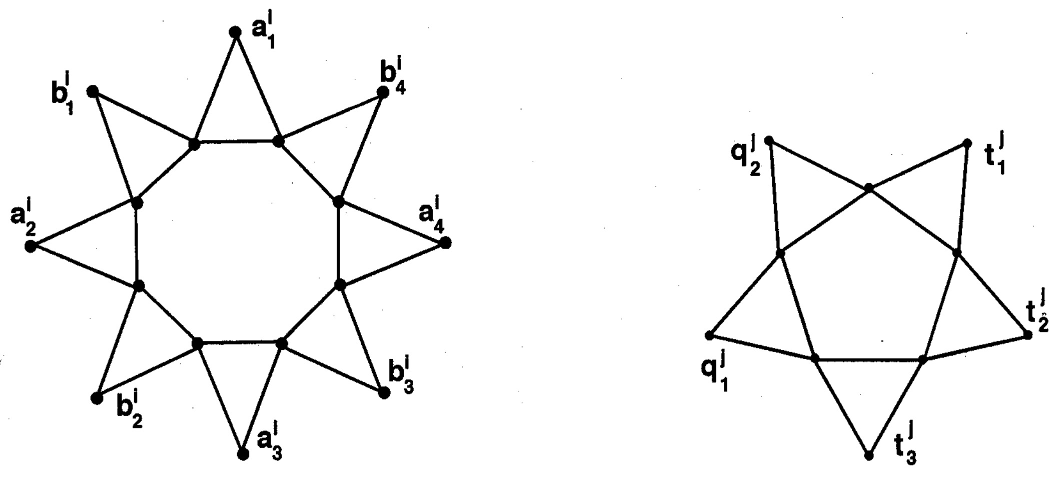

(a) (b)

Figure 5: (a) The *starred-pentagon* corresponding to clause c_j and (b) the *sun* for a variable x_i with $m_i = 4$ occurrences.

vertices (vertices labeled t) identified with a vertex of a *sun*. Each *sun* i (corresponding to variable x_i) has exactly half of its $2m_i$ triangles incident to a triangle on a *starred-pentagon*.

A planar embedding for $G(\phi)$ can be constructed from a planar embedding of the bipartite graph used in the definition of Planar 3-SAT. As before, let n be the number of variables and let m_i be the number of occurrences of variable x_i, counting affirmative and negative occurrences. Let m be the number of clauses in ϕ. Then $\sum_{i=1}^{n} m_i + 2m$ is an upper bound for the size of a packing of triangles in $G(\phi)$. To see this recall that the only triangles in $G(\phi)$ are those in the *suns* and in the *starred-pentagons*. In each *sun* i we can pack at most m_i triangles (all triangles with vertex labeled a or all triangles with vertex labeled b). Hence, the *suns* contribute to the packing with at most $\sum_{i=1}^{n} m_i$ triangles. In each *starred-pentagon* we can pack at most 2 triangles.

Therefore, *starred-pentagons* contribute to the packing with at most $2m$ triangles. The following theorem [12] states that the bound $\sum_{i=1}^{n} m_i + 2m$ is reached if and only if ϕ is satisfiable.

Theorem 3.15. *[12] $G(\phi)$ has a packing of triangles of size $\sum_{i=1}^{n} m_i + 2m$ if and only if ϕ is satisfiable.* □

On what follows, we show how this last theorem implies our hardness result:

Theorem 3.16. *Unless P=NP, finding a maximum packing of vertex disjoint odd cycles is NP-hard for planar graphs.*

Before proving the theorem we need some definitions. Let H be a plane graph. Then, for each cycle C in H we can define two *sides.* Two faces of H lie on opposite sides with respect to C if and only if every path in the face-vertex incidence graph H^+ that links the two faces intersects C in at least one vertex. Recall from Section 2 that the parity of a cycle equals the sum mod 2 of the parities of the faces on one of this sides (see Section 2, Lemma 2.3). Hence, if cycle C is odd, there must be exactly an odd number of odd faces in each of the two sides defined by C.

Now consider a planar embedding of $G(\phi)$. Notice that the only odd faces in $G(\phi)$ are triangles in the *suns* and triangles and pentagons in the *starred-pentagons.* The other faces in $G(\phi)$ are the even cycles *inside* the *suns* and faces made up with paths on *starred-pentagons* and *suns.* Note that these paths must alternate from a path on a *starred-pentagons* to a path on *suns*, connected through identified vertices (the vertices mentioned in 3 above). Furthermore, each of these paths must be of even length since they have to link two t-vertices (in the case of the *starred-pentagons* and

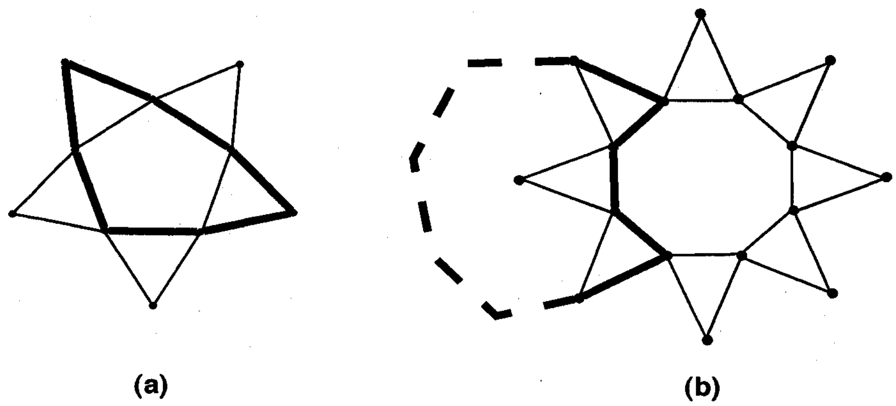

Figure 6: Two examples of odd cycles (heavy edges) separating (a) a *starred-pentagon* and (b) a *sun.*

two vertices labeled either a or b (in the case of the *suns*) without using any edge on the *inside* cycle. Hence the faces made up this way are even.

It follows from the last two paragraphs that for every odd cycle C in $G(\phi)$ there is a gadget (a *sun* or a *starred-pentagon*) such that exactly an odd number of odd faces of the gadget lie on one of the sides defined by C while the rest of the odd faces lie on the opposite one. We abbreviate this by saying that the odd cycle C *separates* the gadget (see Figure 6).

Proof of Theorem 3.16. We show that an optimal packing of vertex disjoint odd cycles in $G(\phi)$ has size $\sum_{i=1}^{n} m_i + 2m$ if and only if ϕ is satisfiable. More precisely, we show that if an optimal packing of vertex disjoint odd cycles in $G(\phi)$ has size $\sum_{i=1}^{n} m_i + 2m$, then it can be transformed into a packing of vertex disjoint triangles of the same size. Otherwise, it has size strictly less than $\sum_{i=1}^{n} m_i + 2m$. This implies that any packing of vertex disjoint triangles has size strictly less than $\sum_{i=1}^{n} m_i + 2m$ and therefore, by Theorem 3.15 ϕ is not satisfiable.

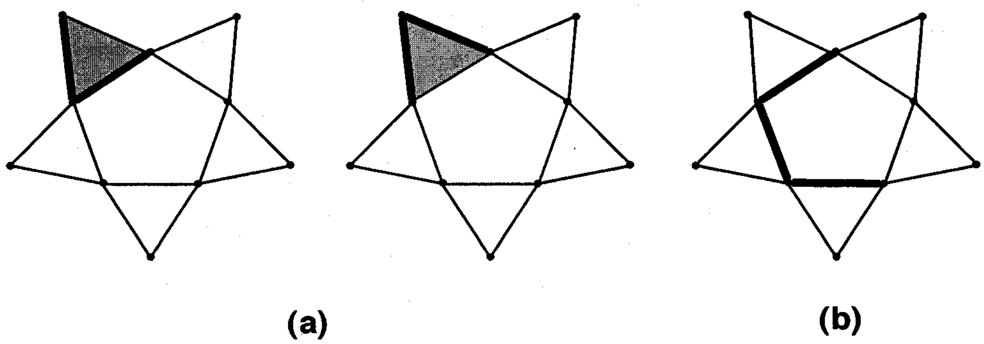

Figure 7: The heavy lines represent part of the intersection between the odd cycle C and the gadget. The shadowed triangles in (a) can replace C. In (b), any q-triangle can replace C.

Let $\mathcal{P}$ be an optimal packing of vertex disjoint odd cycles in $G(\phi)$. If an odd cycle $C \in \mathcal{P}$ separates some *starred-pentagon* then it can be replaced by a triangle of that *starred-pentagon.* This is easy to see if the C is contained in a *starred-pentagon.* So suppose not. Then there are two cases. Either C contains two consecutive edges of a triangle of some *starred-pentagon* that it separates (see Figure 7(a)), or it has at most one edge of each triangle of every *starred-pentagon* that it separates. In the first case, cycle C can be replaced by a triangle of which it contains two consecutive edges (see Figure 7(a)). In the second case assume first that C contains three or more edges of a pentagon on some *starred-pentagon* that it separates (see Figure 7(b)). Hence, it can be replaced by any q-triangle on that *starred-pentagon.* If this is not the case, C contains at most 2 edges of the pentagon corresponding to each *starred-pentagon* that it separates. The only two possibilities are shown in Figure 8. If C is as shown in Figure 8(a), at most one of the two q-triangles can be in the packing and therefore, the triangle adjacent to the q triangle that is not in the packing can be chosen to replace C. If C is as shown in Figure 8(b), one of the q-triangles can be used to replace C.

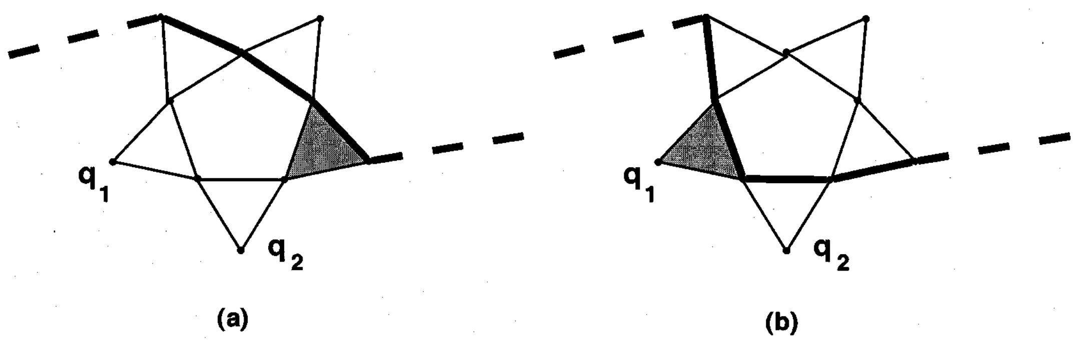

Figure 8: The heavy lines indicate the intersection between the odd cycle C and the gadget. In (a), without loss of generality, the q_1-triangle belongs to $\mathcal{P}$ and therefore the shadowed triangle can replace C. In (b) any of the two q-triangles can be used to replace C.

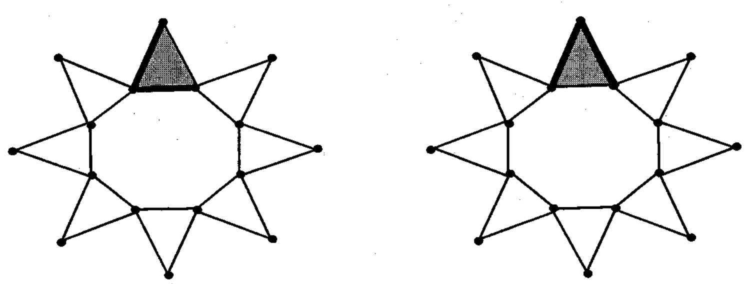

Figure 9: The heavy lines represent part of the intersection between the odd cycle C and the gadget. The shadowed triangles can replace cycle C.

Now suppose an odd cycle $C \in \mathcal{P}$ does not separates any *starred-pentagon*. Then C separates at least one *sun*. As before, it is easy to see that if C is contained in a *sun* it can be replaced by one of the triangles in that *sun*. Hence, let C separate some *sun* but in such a way that C is not completely contained in it. If C contains two consecutive edges of a triangle in some of the *suns* that it separates (see Figure 9), then it can be replaced by that triangle. Otherwise, C contains at most one edge of each triangle in the *suns* that it separates. Without loss of generality, the intersection of C and each of these *suns* is as shown in Figure 10. Let i be one of the

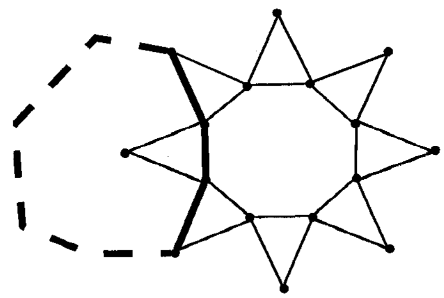

Figure 10: An example where C separates a *sun* and it is not completely contained in it.

suns separated by C.

Since *suns* are only connected to *starred-pentagons* and these are only connected to *suns*, and C is not completely contained within any *sun*, there must be some other *sun* j with which C shares at least two edges. Then, if we do not count cycle C, there are at most $m_j - 1$ odd cycles separating *sun* j and at most $m_i - 1$ odd cycles separating *sun* i. Hence, there are at most $\sum_{k=1}^{n} m_k + 2m - 2 + 1$ cycles in $\mathcal{P}$. (Cycle C is in $\mathcal{P}$ but for every *sun* that C intersects we lose at least one cycle; as C intersects at least two *suns*, the size of $\mathcal{P}$ is strictly less than $\sum_{k=1}^{n} m_k + 2m$.) The argument above also tells us that $\sum_{i=1}^{n} m_i + 2m$ is an upper bound for the size of $\mathcal{P}$. This concludes the proof. □

3.4 A 10-approximation algorithm for packing odd cycles in planar graphs

We briefly discuss how the results proved in the previous sections imply a 10-approximation algorithm for the packing problem restricted to planar graphs.

Note that the proof of Theorem 3.13 implies that τ is within a constant factor of the maximum between the size of the global and the local packing. In particular, we proved that $\tau \leq 10 \max\{\mathcal{C}_{local}, \mathcal{C}_{global}\}$, and this immediately implies that $\nu \leq 10 \max\{\mathcal{C}_{local}, \mathcal{C}_{global}\}$. Hence, the largest of the two packings is a 10-approximation for the packing problem. We follow the proof of Theorem 3.13. We start by checking if graph G is 4-connected. If not, we split the problem into the subgraphs defined by a cutset of size at most 3. If the graph is 4-connected, find a maximal collection of vertex disjoint faces hitting all odd faces in G. This can be done in linear time. As in Lemma 3.7 we call these faces *centers.* If these centers are such that the distance between any pair of them is at least 8, then we can easily find a *local packing* around each center and the union of these is $\mathcal{C}_{local}$. In Section 3.2.2 we showed that this is essentially the same as finding a maximum stable set in a circular arc graph (see Proposition 3.6). Hence, it can be done in polynomial time [24].

If it is the case that two centers are close to each other, remove the vertices in a shortest path linking them (distances and paths considered in the face-vertex incidence graph). Check whether the obtained graph is 4-connected and recurse.

Finally, we need to find a *global packing.* If the graph is 4-connected we find a *global packing* using the ideas in the proof of Theorem 3.13. Otherwise, we split the problem into the subgraphs defined by a cut set of size at most 3. In Section 3.2.4 we showed how a packing of T-cuts in a particular subgraph of the face-vertex incidence graph yields a packing of odd cycles in G of the same size. This packing is the one we call $\mathcal{C}_{global}$ (see Theorem 3.13). Then, as the face-vertex incidence graph is bipartite and optimal packings of T-cuts can be found in polynomial time in bipartite graphs [49] (see also [5]), we have a 10-approximation for the packing problem.

4 Planar graph bipartization in linear time

In this section we present the linear time algorithm given in [19] for finding minimum odd cycle vertex covers in planar graphs. We start with the hardness result showing that this problem is NP-hard. Lund and Yannakakis [37] proved that the problem in its both versions (edge and vertex) is $\mathcal{APX}$-hard.

4.1 Hardness results

The edge version of the bipartization problem is equivalent to MAX CUT and therefore, it is polynomial when restricted to planar graphs. In contrast, the vertex version remains NP-hard when considering the restriction to planar graphs [26].

In what follows we prove the equivalence between MAX CUT and the problem of finding a minimum odd cycle edge cover (MOC-EC). Then we present a reduction from VERTEX COVER (VC) to the problem of finding a minimum odd cycle vertex cover (MOC-VC). It is clear that both reductions preserve planarity, proving that, for planar graphs, MOC-EC is easy and MOC-VC is hard.

We use *odd cycle cover* to refer to an *odd cycle edge cover* or to an *odd cycle vertex cover*. It will be clear from the context to which of them we are making reference.

MAX CUT

- Input: Graph $G = (V, E)$ and a non-negative weight function over the edge set.
- Goal: Find a partition of V, S, $V - S$ that maximizes $\sum_{e\in\delta(S)} w(e)$.

It was shown in [22] that this problem remains NP-hard if all weights equal 1.

Theorem 4.1. *Unless P=NP, the problem of finding a minimum odd cycle edge cover (MOC-EC) is NP-hard.*

Proof. We show a reduction from MAX CUT with all weights equal 1 to MOC-EC. Let $G = (V, E)$ be an instance of MAX CUT. Suppose G has a minimum odd cycle cover of size τ. This odd cycle cover gives a partition of V into two subsets such that the number of edges going across them is exactly $|E| - \tau$. Therefore, if $S \subseteq V$ is optimal for MAX CUT then $\delta(S) \geq |E| - \tau$. On the other hand, given any subset S' of V, the edges with both endpoints in S' plus the edges with both endpoints in $V - S'$ give an odd cycle cover of size exactly $|E| - \delta(S')$. Then, $\tau \leq |E| - \delta(S')$ for any $S' \subseteq V$. In particular, $\tau \leq |E| - \delta(S)$. It follows that $\tau = |E| - \delta(S)$. Hence, if we can find a minimum odd cycle cover then we also obtain a maximum cut. $\square$

VERTEX COVER

- Input: Graph $G = (V, E)$.
- Goal: Find a subset V' of V of minimum cardinality such that each edge of G has at leas one of its endpoints in V'.

The problem is known to be NP-hard [28]. We use it here to show that finding a minimum odd cycle vertex cover is also NP-hard.

Theorem 4.2. *Unless P=NP, the problem of finding a minimum odd cycle vertex cover (MOC-VC) is NP-hard.*

Proof. Let $G = (V, E)$ be an instance of VC. We construct an instance $G' = (V', E')$ of MOC-VC as follows. Let $V' = V \cup \{e \text{ for every edge } e \in E\}$. That is, we add one vertex for each edge in E. We let $E' = E \cup \{(u, e)\text{: } u \text{ is an endpoint of } e\}$. Suppose we have a vertex cover $W \subseteq V$ for G. Look at W as a subset of V' and let C be any odd cycle in G'. Take any edge g of C. If $g = (u, v)$ for some vertices u and v originally members of V, then we know either u or v are in W and therefore cycle C is hit. If $g = (u, e)$ for some vertex $u \in V$ and some edge $e \in E$, then, since e has degree 2 in G', C must also contain the other endpoint of e, and, hence C is hit, either by u or by the other endpoint of e. It follows that the size of a minimum odd cycle cover for G' is at most the size of a minimum vertex cover for G. On the other hand, suppose $W' \subseteq V'$ is an odd cycle cover for G'. Consider a subset W of V of the following form. First, W includes every vertex of W' which is also a vertex in V. Second, for each vertex in W' of the form $e = (u, v)$ place in W an endpoint already chosen for W, or an arbitrary endpoint if none of the two was previously chosen. We claim that W is a vertex cover for G. This is clear from the fact that every edge $e = (u, v)$ of G belongs to an odd cycle in G' of the form u, v, e. Therefore, at least one of u, or v or e is in W' and, thus, at least one of u or v is in W. It follows that the size of a minimum vertex cover for G is at most the size of a minimum odd cycle cover for G'. Hence, we have shown that the size of a minimum odd cycle cover G' equals the size of a minimum vertex cover for G. □

Clearly, this reduction preserves planarity as we can obtain an embedding of G' from the embedding of G by placing vertex $e = (u, v)$ with edges (u, e) and (v, e) as close to the embedding of edge e as necessary.

Corollary 4.3. *The problem of finding a minimum odd cycle vertex cover restricted to planar graphs is NP-hard.* □

4.2 Preliminaries

Through all this section by *cover* and *packing* we mean odd cycle vertex cover and packing of vertex disjoint odd cycles respectively. As it was shown in Section 4.1, the problem of finding a minimum odd cycle cover, or the bipartization problem, is NP-hard, even when restricted to planar graphs. Garg, Vazirani and Yanakakis [23] gave an $O(\log n)$-approximation algorithm for general graphs and Goemans and Williamson [25] gave a (9/4)-approximation for planar graphs. In [42] Reed, Smith and Vetta presented an $O(n^2)$ exact algorithm to determine whether a given graph G has a cover of order at most k. In what follows we present the algorithm given in [19] that determines in $O(n)$ time whether a given planar graph G has a cover of order at most k.

Theorem 4.4. *Given a planar graph G and an integer k, there exists a linear time algorithm that either finds a minimum odd cycle cover of size at most k or gives a certificate that no such cover exists.*

Essentially, the algorithm consists in finding ℓ subgraphs of G with the property that the vertices which are not in any of these subgraphs are irrelevant when looking for a minimum cover of size at most k. That is, the vertices of G which are vertices of these ℓ subgraphs form a cover for G. If $\ell > k$ we can affirm that every odd cycle cover of G has size $> k$. If $\ell \leq k$, we prove that if the union of minimum odd cycle covers for each of these subgraphs has size $\leq k$, then it is a minimum odd cycle cover for the whole graph. Otherwise, every cover of G has size $> k$. Finally we show that the mentioned subgraphs have tree-width bounded by a function of k. In Section 4.4, we show that when this is the case, a minimum cover can be found in linear time, using standard dynamic programming on the tree decomposition.

We recall the relation between the odd faces and the odd cycles of a planar graph (see Section 2, Lemma 2.3 and Corollary 2.4). Consider an embedding of a planar graph G. When a vertex v is deleted from G, all the faces incident to v are merged together in a new face F_v. The other faces are unchanged. We denote the new face by a capital letter to stress the fact that it determines a set of faces of G, namely, the faces of G included in it. The *parity* of a face of G is defined as the parity of the edge set of its boundary, counting bridges twice. By induction on the number of faces incident to v, the parity F_v equals the sum mod 2 of the parities of the faces of G it contains. As a corollary, we obtain that G is bipartite if and only if all its faces are even.

Lemma 2.3. Given a planar graph G and a vertex $v \in V(G)$, removing v from G creates a new face F_v which contains all faces of G incident to v. Let $par(f)$ denote the parity of face f, then $par(F_v) = \sum_{f \text{ incident to } v} par(f) \pmod 2$. □

Corollary 2.4. *A planar graph G is bipartite if and only if all its faces are even.* □

Let now W denote any set of vertices in G. By deleting from G the vertices in W one after the other in some order, we see that each face of $G - W$ corresponds to a set of faces of G. This set is a singleton if the corresponding face is a face of G that survived in $G - W$. Furthermore, form the lemma above, a face of $G - W$ is odd precisely if it contains an odd number of odd faces of G. Hence, since a subset W is an odd cycle cover if and only if $G - W$ is bipartite, we obtain our

Key Fact: *A set W of vertices is an odd cycle cover of G precisely if every face of $G - W$ contains an even number of odd faces of G.* □

Recall that given an embedding of G, the *face-vertex incidence graph* of G is the bipartite graph G^+ on the vertices and faces of G whose edges are the pairs fv, where f is a face of G and v is a vertex of G incident to f. (The face-vertex incidence graph was first defined in Section 2.)

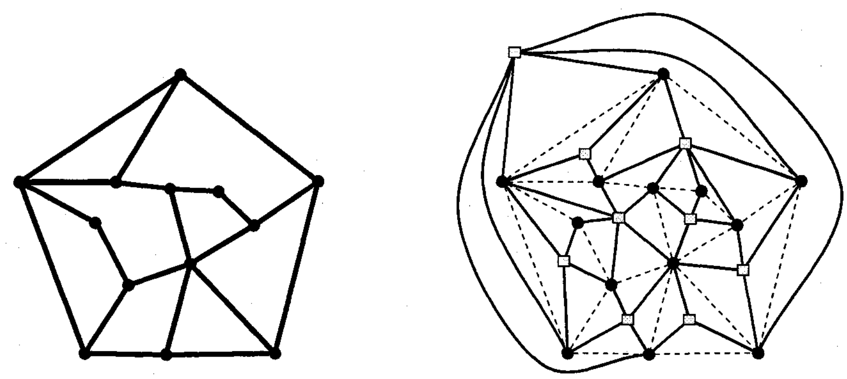

Figure 11: A graph G (left) with its face-vertex incidence graph G^+ (right). Black dots in G^+ correspond to vertices of G, white squares correspond to faces of G.

Corollary 4.5. *No vertex v of G is in an inclusion-wise minimal odd cycle cover of size less than $d_{\min}(v)$, the minimum length of a path from v to an odd face in the face-vertex incidence graph.* □

To prove this corollary we will need to deal with T-joins, which were described in Section 2 but which we redefine here. Before that, we give the intuition behind Corollary 4.5.

Suppose G is a planar graph with exactly two odd faces which are far apart from each other in the face-vertex incidence graph, as shown in Figure 12(a). At least one vertex of each of these faces must be in any cover. Let these vertices be u and v. When we remove them from G, we obtain two new faces F_v and F_u, see Figure

12(b). Both these faces are odd since they contain only one odd face from the original graph G. Then, any cover must contain at least one vertex of each of these new faces (see Figure 12(b)). If we recurse on this process we see that we are constructing two paths in the face-vertex incidence graph of G (see Figure 12(c)). We refer to the vertices in these paths that are vertices of G, as G-vertices, and to the vertices which represent faces, as face-vertices (represented in Figure 12 by white squares). Furthermore, if considering minimal inclusion-wise covers, and since our graph G contained originally only two odd faces, intuition tells us that these two paths must intersect at some vertex w (see Figure 12(d)). Moreover, note that all G-vertices in this path are vertices in the cover. Then, the length of this path is related to the size of the cover. Using T-joins, we will show that this is always the case: having a minimum cover is equivalent to having a forest on the face-vertex incidence graph that connects an even number of odd faces of G. All G-vertices in this forest must be in the cover, therefore, if we are looking for a minimum cover of size at most k, the vertices in it cannot be more than k away of any odd face, distances counted on the face-vertex incidence graph.

We now go back to the definition of T-joins and state a useful reformulation of the Key Fact in terms of T-joins in the face-vertex incidence graph.

Consider any graph H and set of vertices T in H. A *T-join* in H is a subset J of edges of H such that T equals the set of odd degree vertices in the subgraph of H induced by J. There exists a T-join in H if and only if each connected component of H contains an even number of vertices of T. In particular, if H has a T-join then $|T|$ is even. Now let T be the set of odd faces of the planar graph G. So T is an even set of vertices in the face-vertex incidence graph G^+. Let $F(G)$ denote the set of faces of

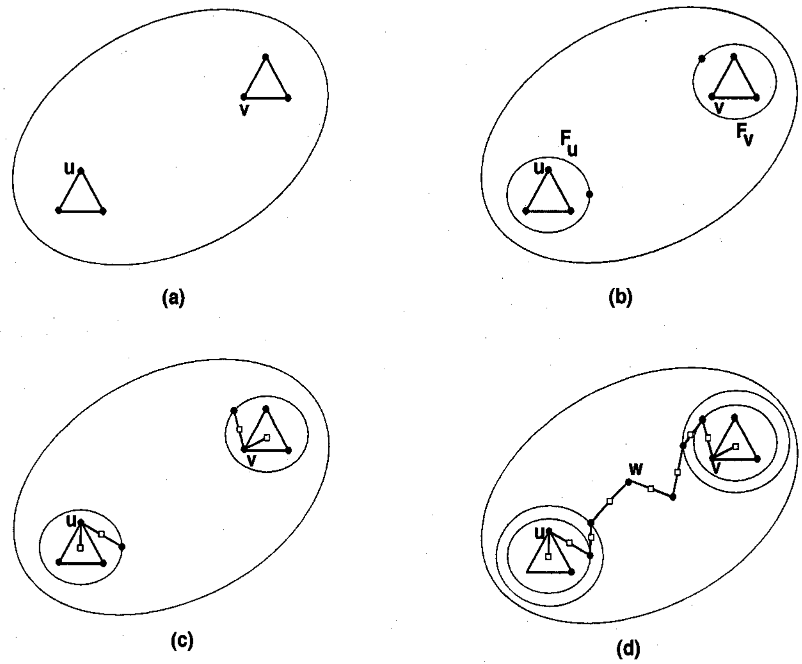

Figure 12: Distances in the face-vertex incidence graph are related to the size of inclusion-wise minimal covers.

G. The relation between odd cycle covers and T-joins was shown in Lemma 2.14.

Lemma 2.14. *A subset W of $V(G)$ is an odd cycle cover of G if and only if the subgraph of the face-vertex incidence graph G^+ induced by $W \cup F(G)$ contains a T-join, that is, every component of the subgraph has an even number of vertices of T.* □

Recall that the lemma follows easily from the fact that deleting a vertex v of G corresponds to contracting the edges incident to v in the face-vertex incident graph. Every new face of $G - W$ corresponds to a connected component of a T-join in the face-vertex incidence graph.

The above lemma is useful because it enables us to visualize an odd cycle cover of G as forest in the face-vertex incidence graph G^+ such that each tree of the forest contains an even number of vertices of T. Indeed, consider an inclusion-wise minimal odd cycle cover W of G. By Lemma 2.14, there is a T-join J in G^+ covering each vertex of W and no vertex of $G - W$. Without loss of generality, we can assume that J is inclusion-wise minimal. Then, J contains no cycles, because if there is a cycle, we can remove from J all edges in the cycle and we will not change the parity of the degree of any vertex with respect to J, i.e., T will still be the set of odd-degree vertices in the graph induced by J. Hence, if J is inclusion-wise minimal, J is a forest and every leaf of J is in T. Note that some vertices of T can be internal nodes of J. For every vertex v of W, there are two internally disjoint paths in J between v and T. So, letting $d_{\min}(v)$ be the minimum length of a path from v to an odd face in the face-vertex incidence graph, we see that the Key Fact implies Corollary 4.5. For suppose $v \in V$ is in an odd cycle cover W of size $< d_{min}(v)$, then, v is in the path between two leaves of a minimal T-join in the subgraph of G^+ induced by $W \cup F(G)$, where every G-vertex in that path is also in the cover. There are at least $d_{min}(v)$ G-vertices in the path, contradicting the fact that the odd cycle cover has size $< d_{min}(v)$.

Thus, letting G' be the subgraph of G induced by $\{v | d_{\min}(v) > k\}$, we see that if G has a cover of size at most k then G' must be bipartite. So applying the Key Fact to the embedding of G' which appears as a sub-embedding of our embedding of G, we obtain:

Corollary 4.6. *If G has an odd cycle cover of size at most k then every face of G' contains an even number of odd faces of G.* □

We note further that the boundary, $bd(F)$, of every face F of G' is disjoint from the boundaries of the odd faces of G within it by the definition of G' (except for the trivial case $k = 0$). Thus we have:

Observation 4.7. *If G has an odd cycle cover of size at most k then there are at most k faces of G' which contain an odd face of G.* □

For some $r \leq k$, we let $\{F_1, ..., F_r\}$ be the set of faces of G' containing an odd face of G and let $G_i = G \cap (F_i \cup bd(F_i))$. Applying Corollary 4.5 again, it is easy to show:

Corollary 4.8. *If G has an odd cycle cover of order at most k then W is a minimum odd cycle cover of G precisely if $W_i = W \cap G_i$ is a minimum odd cycle cover of G_i for every i between 1 and r.*

Proof. Consider a cover W of G of order at most k. By Corollary 4.5, W_i is disjoint from the boundary of F_i, and each face of $G - W$ which is not a face of G' is a face of $G_i - W_i$ for some i. Thus, applying the Key Fact to $G - W$ and $G_i - W_i$ for each i we see that W is a cover of G if and only if W_i is a cover of G_i for each i. Since W_i is disjoint from the boundary of F_i, the W_i's are disjoint and the result follows. □

We will show later that the face-vertex incidence graph of each G_i has radius $O(k^2)$. Hence each G_i has tree-width (defined below) which is $O(k^2)$. We show in Section 4.4 that we can find minimum covers in linear time in graphs of bounded tree-width. So if we could find all the G_i's in linear time then we could compute a minimum cover for each G_i in linear time and by taking their union, find a minimum cover of G (or determine that G has no odd cycle cover of order at most k).

In the next section we show how to transform this into an algorithm. Note that our argument relies in the fact that vertices far apart from all odd faces of G are irrelevant. To find these irrelevant vertices we have to deal with inclusion-wise maximal sets of vertex disjoint odd faces. Finally, we prove that each subgraph G_i of the face-vertex incidence graph has bounded tree-width.

4.3 The algorithm

Our algorithm works as follows. First obtain an embedding of G in linear time [27], and construct the face-vertex incidence graph G^+. Then find a collection $\mathcal{F} = \{f_1, \dots, f_s\}$ of boundary-disjoint odd faces of G which either has $k+1$ faces or is inclusion-wise maximal. This part of the algorithm can be implemented to run in $O(kn)$ time which is linear as k is fixed.

If $s > k$ then return the information that G has no odd cycle cover of size at most k and stop. Otherwise, let B_i denote the set of faces and vertices of G whose distance to f_i in G^+ is at most $k+3$. Determine the sets B_i for all $i = 1, \dots, s$ via a breadth first search in G^+. Let G'' be the subgraph of G obtained by deleting all the vertices of G in each B_i.

Determine the set $F_1, ..., F_r$ of faces of the embedding of G'' which contain an odd face of G. Note that $r \leq s \leq k$ as each F_i contains some $f_j \in \mathcal{F}$. We let D_i be the subgraph of G contained in the union of F_i and its boundary. We refer to these graphs as *discs*. Now find a minimum odd cycle cover W_i in each disc D_i. Since, as we show below, each disc has bounded tree-width, this can be done in linear time using the techniques described in Section 4.4. Let W be the union of $W_1, \dots, W_r$. If W has size at most k, then W is a minimum odd cycle cover of G; output W. Otherwise, return the information that G has no cover of size at most k. This concludes the description of the algorithm. The correctness follows immediately from the fact that if v is at distance at least $k+3$ from every face in $\mathcal{F}$, then $d_{min}(v)$ is at least $k+1$ and therefore v is irrelevant when looking for covers of size at most k.

Proposition 4.9. *The algorithm finds a minimum cardinality odd cycle cover if G has an odd cycle cover of size at most k or otherwise detects that no such cover exists.*

Proof. The proof of this proposition mimics exactly the proof of Corollary 4.8 with G' replaced by G'' and G_i replaced by D_i. □

In Section 4.4, we will describe how to find minimum odd cycle covers in graphs of bounded tree-width in linear time. Since all of the steps described in this section can be carried out in linear time, Proposition 4.9 tells us that we will obtain a linear time algorithm for general planar graphs if we can show that each disc has bounded tree-width. This, though, follows simply from the following result.

Lemma 4.10. *([1], for a more general result see [45, 47]) If a planar graph contains no $h \times h$ grid minor, then its tree-width is at most $8h$.* □

Because the radius of the face-vertex incidence graph of any planar graph containing a $h \times h$ grid minor is at least h, the preceding lemma has the following immediate corollary:

Corollary 4.11. *Let G be a planar graph. If the radius of the face-vertex incidence graph of G is less than h, then the tree-width of G is at most $8h$.* □

Lemma 4.12. *The tree-width of each disc is $O(k^2)$.*

Proof. By Corollary 4.11, it suffices to show that the radius of each disc is $O(k^2)$. Consider any disc D_i. Let I be the set of indices ℓ such that $f_\ell \in \mathcal{F}$ is a face of D_i. Let H denote the graph whose vertex set is I and whose edges are the pairs $\ell\ell'$ of indices such that some vertex of G in B_ℓ and some vertex of G in $B_{\ell'}$ are incident to some common face of G. We know that H is connected and has at most k vertices, so its radius is at most $k/2$. Let j be a vertex of H such that the distance in H between j and any vertex of H is at most $k/2$. The distance in D_i^+ between any f_ℓ with $\ell \in I$ and f_j is at most the distance in H between ℓ and j times $2(k+1)+2 = 2k+4$. Moreover, for every face or vertex of D_i there is an index $\ell \in I$ such that the distance in D_i^+ between the considered face or vertex and f_ℓ is at most $k+2$. So the distance in D_i^+ between any face or vertex of D_i and f_j is at most $(k/2)(2k+4)+k+2 = k^2+3k+2$. So the radius of D_i^+ is indeed $O(k^2)$. □

4.4 Linear time bipartization of graphs of bounded tree-width

As we have seen, it suffices to have a linear time algorithm for graph bipartization for graphs of bounded tree-width. This can be done using standard techniques; we present such an algorithm below. We begin with the required technical definitions.

A *tree-decomposition* of G is a pair $(T, \mathcal{V})$, where T is a tree and $\mathcal{V} = (V_t \subseteq V(G) : t \in V(T))$ is a family of subsets of $V(G)$ with the following properties:

1. $\bigcup(V_t : t \in V(T)) = V(G)$.

2. For each edge $e \in E(G)$ there is a node $t \in V(T)$ such that both endpoints of e are in V_t.

3. For t_0, t_1 and t_2 in $\in V(T)$, if t_0 is on the path of T between t_1 and t_2, then $V_{t_1} \cap V_{t_2} \subseteq V_{t_0}$.

The *width* of the tree-decomposition $(T, \mathcal{V})$ is defined as $\max_{t \in V(T)}(|V_t| - 1)$. The *tree-width* of a graph G is the minimum ω such that G has a tree-decomposition of width ω. It is well known that there are minimum tree decompositions of G that use at most n nodes. Moreover, we can easily convert a tree decomposition $(T, \mathcal{V})$ to another $(T', \mathcal{V}')$ of the same width, such that T' is a binary tree with at most twice as many nodes as T.

Let G be a graph of bounded tree-width $\omega - 1$ and let $(T, \mathcal{V})$ be a binary minimum tree-decomposition of G. We denote by t the nodes of T and by V_t the subset of $V(G)$ assigned to t. We have that $|V_t| \leq \omega$ for all $t \in T$. Pick an arbitrary root node $t^* \in T$. Then, given a node $t \in T$ we let S_t be the subtree of T rooted at t. From (2) we may assign to each edge $e = (u, v)$ of G a specific node $t(e) \in T$ for which $u, v \in V_t$. Thus, for each $t \in T$ there is an associated edge set $E_t \subseteq E(G)$. Hence, we may define the graphs $G(t) = (V_t, E_t)$ and $G(S_t) = (\bigcup_{t' \in S_t} V_{t'}, \bigcup_{t' \in S_t} E_{t'})$. We associate with each node $t \in T$ a set $\mathcal{A}_t$ of all the ordered triplets $\Pi_t = (L_t, R_t, W_t)$ where L_t, R_t and W_t form a vertex partition of V_t. Clearly $|\mathcal{A}_t|$ is at most 3^ω. Our algorithm will work up from the leaves maintaining the property that for each partition Π_t we (implicitly) store a minimum odd cycle cover $\hat{W}_t$ in $G(S_t)$ that is *accordant* with the partition. That is, $W_t \subseteq \hat{W}_t$ and L_t and R_t are on opposites sides of the bipartition in $G(S_t) - \hat{W}_t$. If such a cover exists then we will set $f(\Pi_t) = |\hat{W}_t|$; otherwise if there is no such accordant cover then we set $f(\Pi_t) = \infty$. Hence, for a leaf $t \in T$ we have $f(\Pi_t) = |W_t|$ if L_t and R_t both induce stable sets in E_t. Otherwise $f(\Pi_t) = \infty$. Now take a non-leaf node $t \in T$ with children r and s. If L_t or R_t induce an edge in E_t then we set $f(\Pi_t) = \infty$. So suppose not. We say that a partition $\Pi_r = (L_r, R_r, W_r)$ in $\mathcal{A}_r$ is *consistent* with a partition $\Pi_t = (L_t, R_t, W_t)$ in $\mathcal{A}_t$ if $W_t \cap V(S_r) \subseteq W_r$, $L_t \cap V(S_r) \subseteq L_r$ and $R_t \cap V(S_r) \subseteq R_r$. We use the notation $\Pi_r \sim \Pi_t$ to denote consistency. Note, by property (3), that if Π_r and Π_s are both consistent with Π_t then they are consistent with each other. Then set

$$f(\Pi_t) = \min_{\Pi_r \sim \Pi_t, \Pi_s \sim \Pi_t} f(\Pi_r) + f(\Pi_s) + |W_t - (W_r \cup W_s)| - |W_r \cap W_s|$$

Note that it may still be the case that $f(\Pi_t) = \infty$. We repeat this process up the tree. Observe that, by storing pointers from a partition Π_t to the partitions Π'_r and Π'_s in its children that produced the minimum value $f(\Pi_t)$, we may implicitly store the set $\hat{W}_t$. We then obtain the following result.

Lemma 4.13. *For each Π_t, either $f(\Pi_t)$ is the size of the minimum odd cycle cover in $G(S_t)$ accordant with the partition Π_t, or $f(\Pi_t) = \infty$ and no such a cover exists.*

Proof. This is clearly true if t is a leaf. So let $t \in T$ be a non-leaf with children r and s. Take Π_t and assume first that $f(\Pi_t)$ is finite. Next take consistent partitions Π_r and Π_s with optimal covers $\hat{W}_r$ and $\hat{W}_s$, respectively. Then, since $\hat{W}_r$ and $\hat{W}_s$ are accordant with Π_r and Π_s, by property (3) we have that $W_t - (\hat{W}_r \cup \hat{W}_s) = W_t - (W_r \cup W_s)$. Thus, in obtaining $\hat{W}_t$ we only need to add the vertices in $W_t - (W_r \cup W_s)$. Moreover any vertex in $W_r \cap W_s$ is double counted by $f(\Pi_r) + f(\Pi_s)$. Thus $f(\Pi_t)$ is in fact the size of a cover in $G(t)$ accordant with Π_t. Therefore, since we are examining all consistent pairs of partitions for the children, it is clear that $f(\Pi_t)$ is the size of a minimum odd cycle cover $\hat{W}_t$ in $G(S_t)$ accordant with the partition Π_t. Now suppose $f(\Pi_t) = \infty$ and that there is a cover W for $G(S_t)$ accordant with Π_t. Then, for all pairs of partitions Π_r and Π_s that are consistent with Π_t, at least one of $f(\Pi_r)$ or $f(\Pi_s)$ is infinite. We obtain a contradiction as the restrictions of W to $G(S_r)$ and $G(S_s)$ give odd cycle covers for these subgraphs that are accordant with Π_r and Π_s, respectively. $\square$

It immediately follows that the minimum odd cycle cover can be found by considering the partition Π_{t^*} with the minimum f value. We may obtain a binary tree-decomposition in linear time [8]. For each node in the tree we have $O(3^\omega)$ partitions. It takes $O(|E_t|)$ time to check whether L_t or R_t induce stable sets in $G(t)$. There are then $O(9^\omega)$ possible pairs of partitions for the children. Thus it takes $O(\omega 9^\omega)$ time to check for consistencies and to calculate $f(\Pi_t)$. In total, therefore the algorithm runs in time $O(\omega 3^{3\omega} n)$. Thus we have proven Theorem 4.14 and Theorem 4.4.

Theorem 4.14. *Let G be a graph with bounded tree-width. Then there is an linear time algorithm to find a minimum odd cycle cover in G.* $\square$

5 Conclusions

We have shown that Brass conjecture [53], i.e., that the size of a minimum odd cycle cover (τ) is at most twice the size of a maximum packing (ν) is true for several subclasses of planar graphs, namely, when all odd faces are disjoint, when one face intersects every odd face and when all odd faces mutually intersect. Then, using these results, we proved that for the general class of planar graphs, $\tau \leq 10\nu$. As a corollary of our proofs, we obtained a 10-approximation algorithm for the packing problem.

Secondly, we have shown a linear time algorithm to find minimum odd cycle covers for planar graphs with constant sized covers. In doing so, we proved that odd cycle covers can be found in linear time in graphs of bounded tree-width.

The immediate question that arises is whether there exists a linear time exact algorithm to find minimum odd cycle covers in general graphs where the size of the cover is known to be small. In [42], Reed, Smith and Vetta presented a quadratic algorithm for such problem, and here we showed there is a linear algorithm for the problem restricted to planar graphs.

Concerning approximation algorithms for the covering problem, improvements can be done on showing the exact integrality gap for the Covering LP. Goemans and Williamson proved this gap to be at most $9/4$ and conjectured it to be $3/2$.

With respect to our theoretical result that bounds the size of a minimum odd cycle cover τ by a linear function of the maximum size of a packing ν, we ask the question whether the constant can be reduced to 2. That is, to show that $\tau \leq 2\nu$. In doing so, it would be of interest to derive a better approximation algorithm for the packing problem, than the one presented in this thesis.

On the other hand, Král and Voss in [30], asked whether a relation between the minimum size of an odd cycle edge cover and the maximum size of a packing of edge disjoint odd cycles exists for graphs that can be embedded in orientable surfaces. The results given in [41] imply that τ is bounded by a function of ν and g for any graph that can be embedded on an orientable surface of genus g, for both the edge and the vertex versions of the problem. It will be of interest to find such a function.

Also, it seems reasonable to think that, using the techniques presented here, a fast algorithm can be derived for the packing problem (vertex version) for planar graphs where the size of the packing is known to be small.

References

[1] N. Alon, P. Seymour, and R. Thomas, "Planar separators", *SIAM J. Discrete Math.*, **7(2)**, pp184-193, 1994.

[2] K. Appel, and W. Haken, "A proof of the four color theorem", *Discrete Mathematics*, **16**, pp179-180, 1976.

[3] V. Bafna, P. Berman, and T. Fujito, "Constant ratio approximations of the weighted feedback vertex set problem for undirected graphs", *ISAAC95, Algorithms and Computation*, Springer LNCS, **1004**, pp142-151, 1995.

[4] R. Bar-Yehuda, D. Geiger, J. Naor, and R.M. Roth, "Approximation algorithms for the vertex feedback set problem with applications to constraint satisfaction and Bayesian inference", *SIAM Journal on Computing* **27(4)**, pp942-959, 1998.

[5] F. Barahona, "Planar multicommodity flows, max cut, and the Chinese postman problem", *Proceedings DIMACS Workshop on Polyhedral combinatorics*, DIMACS Series in Discrete Mathematics and Theoretical Computer Sciences, **1**, pp189-202, 1990.

[6] A. Becker, and D. Geiger, "Approximation algorithm for the loop cutset problem" *Proceedings of the 10th Conference on Uncertainty in Artificial Intelligence Morgan Kaufman*, pp60-68, 1994.

[7] C. Berge, and B. Reed, "Optimal packings of edge-disjoint odd cycles", *Discrete Mathematics*, **211**, pp197-202, 2000.

[8] H. Bodlaender, "A linear time algorithm for finding tree-decompositions of small treewidth", *SIAM Journal on Computing*, **25**, pp1305-1317, 1996.

[9] A. Brandstädt, "On improved time bounds for permutation graph problems", *Proceedings of the 18th Workshop on Graph-theoretic concepts in computer science*, Springer LNCS, **657**, pp1-10, 1993.

[10] A. Brandstädt, and D. Kratsch, "On the restriction of some NP-complete graph problems to permutation graphs", *Proceedings of Fundamentals of Computing Theory*, Springer LNCS, **199**, pp53-62, 1985.

[11] S. Brandt, "Lösung zu Problem 66", *Mathematische Semesterberichte*, Springer, **44**, pp95-96, 1997.

[12] A. Caprara, and R. Rizzi, "Packing triangles in bounded degree graphs", *Information Processing Letters*, **84(4)**, pp175-180, 2002.

[13] D.G. Corneil, and J. Fonlupt, "The complexity of generalized clique covering", *Discrete Applied Mathematics*, **22**, pp109-118.

[14] R. Diestel, "Graph Theory" *(Second Edition), Number 173 in Graduate Texts in Mathematics*, Springer Verlag, New York, 2000.

[15] R. Dilworth, "A decomposition theorem for partially ordered sets" *Annals of Mathematics*, **51**, pp161-166, 1950.

[16] P. Erdös, and L. Pósa, "On independent circuits contained in a graph", *Canadian Journal of Mathematics*, **17**, pp347-352, 1965.

[17] P. Festa, P. Pardalos, and M. Resende, "Feedback set problems", *Handbook of Combinatorial Optimization*, **4**, Kluwer Academic Publishers, 1999.

[18] S. Fiorini, N. Hardy, B. Reed, and A. Vetta, "Approximate min-max relations for odd cycles in planar graphs", *Proceedings of the11th International Integer Pro-*

gramming and Combinatorial Optimization Conference, Springer LNCS, **3509**, pp35-50, 2005.

[19] S. Fiorini, N. Hardy, B. Reed, and A. Vetta, "Planar graph bipartization in linear time", *Proceedings of the 2nd Brazilian Symposium on Graph Algorithms and Combinatorics*, pp226-232, 2005.

[20] A. Frank, "A survey on T-joins, T-cuts, and conservative weightings", *Combinatorics: Paul Erdös is Eighty*, **2**, pp213-252, 1994.

[21] M. R. Garey, and D. S. Johnson, *Computers and intractability, a guide to the Theory of NP-completness*, New York, 1979.

[22] M. R. Garey, D. S. Johnson, and L. Stockmeyer, "Some simplified NP-complete problems", *Theory of Computer Science*, **1**, pp237-267, 1976.

[23] N. Garg, V. Vazirani, and M. Yannakakis, "Approximate max-flow min-(multi)cut theorems and their applications", *SIAM Journal on Computing*, **25**, pp235-251, 1996.

[24] F. Gavril, "Algorithms on circular arc graphs", *Networks*, **4**, pp357-369, 1974.

[25] M. Goemans, and D. Williamson, "Primal-dual approximation algorithms for feedback problems in planar graphs", *Combinatorica*, **18**, pp37-59, 1998.

[26] F. Hadlock, "Finding a maximum cut of a planar graph in polynomial time", *SIAM Journal on Computing*, **4(3)**, pp221-225, 1975.

[27] J. Hopcroft, and R. Tarjan, "Efficient planarity testing", *Journal of the Association for Computing Machinery*, **21**, pp549-568, 1974.

[28] R. Karp, "Reducibility among combinatorial problems", in *Complexity of Computer Computations*, R. Miller, J. Thatcher, eds., Plenum Press, New York, pp85-103, 1972.

[29] J. Kleinberg, and A. Kumar, "Wavelength conversion in optical networks", *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms*, Baltimore, MD, pp566-575, 1999.

[30] D. Král, and H. Voss, "Edge-Disjoint Odd Cycles in Planar Graphs", *Journal of Combinatorial Theory, Series B*, **90(1)**, pp107-120, 2004.

[31] Y.D. Liang, "On the feedback vertex set problem in permutation graphs", *Information Processing Letters*, **52**, pp123-129, 1994.

[32] D. Lichtenstein, "Planar formulae and their uses", *SIAM Journal on Computing*, **11**, pp329-343, 1982.

[33] L. Lovasz, "2-matchings and 2-covers of hypergraphs", *Acta Mathematica Academiae Scientiarum Hungaricae*, **26**, pp433-444, 1975.

[34] Chin Lung Lu, and Chuan Yi Tang, "A linear-time algorithm for the weighted feedback vertex problem on interval graphs", *Information Processing Letters*, **61**, pp107-111, 1997.

[35] C. Lucchesi, and D. Younger, "A minimax theorem for directed graphs", *Journal of The London Mathematical Society*, **17**, pp369-374, 1978.

[36] F.L. Luccio, "Almost exact minimum feedback vertex set in meshes and butterflies", *Information Processing Letters*, **66**, pp59-64, 1998.

[37] C. Lund, and M. Yannakakis, "The approximation of maximum subgraphs problems", *Proceedings of 20th International Colloquium on Automata, Languages and Programming*, Springer, pp40-51, 1994.

[38] M.V. Marathe, C. Pandu Rangan, and R. Ravi, "Efficient algorithms for generalized clique covering on interval graphs", *Discrete Applied Mathematics*, **39**, pp87-93, 1992.

[39] B. Monien, and R. Schultz, "Four approximation algorithms for the feedback vertex set problems", *Proceedings of the 7th Conference on Graph Theoretic Concepts of Computer Science*, Hanser-Verlag, Münich, pp315-326, 1981.

[40] D. Rautenbach, and B. Reed, "The Erdos-Possa property for odd cycles in highly connected graphs", *Combinatorica* **21(2)**, pp267-278, 2001.

[41] B. Reed, "Mangoes and blueberries", *Combinatorica*, **19**(2), pp267-296, 1999.

[42] B. Reed, K. Smith, and A. Vetta, "Finding odd cycle transversals", *Operations Research Letters*, **32**, pp299-301, 2004.

[43] R. Rizzi, V. Bafna, S. Istrail, and G. Lancia, "Practical algorithms and fixed parameter tractability for the single individual SNP haplotyping problem", *Algorithms in Bioinformatics*, LNCS **2452**, Springer, pp29-43, 2002.

[44] N. Robertson, D. Sanders, P. Seymour and R. Thomas, "The four color theorem", *Journal of combinatorial Theory, Series B*, **70**, pp2-44, 1997.

[45] N. Robertson, and P. Seymour, "Graph Minors V. Excluding a planar graph", *Journal Combinatorial Theory, Series B*, **36**, pp92-114, 1986.

[46] N. Robertson, and P. Seymour, "Graph Minors XXII: irrelevant vertices in linkage problems", preprint, 1992.

[47] N. Robertson, P. Seymour and R. Thomas, "Quickly excluding a planar graph", *Journal Combinatorial Theory, Series B*, **62**, pp323-348, 1994.

[48] A. Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency*, Springer, 2003.

[49] P. Seymour, "On odd cuts and plane multicomodity flows", *Proceedings of the London Mathematical Society*, (Third Series), **42**, pp178-192, 1981.

[50] P. Seymour, "Packing directed circuits fractionally", *Combinatorica*, **15**, pp281-288, 1995.

[51] A. Shamir, "A linear time algorithm for finding minimum cutsets in reduced graphs", *SIAM Journal on Computing*, **8(4)**, pp645-655, 1979.

[52] C. Thomasen, "The Erdös-Pósa property for odd cycles in graphs of large connectivity", *MAT-report*, **8**, Technical University of Denmark, 1999.

[53] D. West, "Open problems Column", *SIAM Activity Group Newsletter in Discrete Mathematics*, **14**, Spring 1994.

[54] M. Yannakakis, "Node and edge-deletion NP-complete problems", *Proceedings of the 10th Annual ACM Symposium on Theory of Computing*, pp253-264, 1978.

[55] M. Yannakakis, and F. Gavril, "The maximum k-colorable subgraph problem for chordal graphs", Information Processing Letters, **24**, pp133-137, 1987.