On the Edge Density of Bipartite 3-Planar and Bipartite Gap-Planar Graphs

Aaron Büngener
Universität Tübingen, Germany

Maximilian Pfister
Universität Tübingen, Germany

Abstract

We show that if a bipartite graph G with $n \geq 3$ vertices can be drawn in the plane such that (i) each edge is involved in at most three crossings per edge or (ii) each crossing is assigned to one of the two involved edges and each edge is assigned at most one crossing, then G has at most $4n - 8$ edges. In both cases, this bound is tight up to an additive constant as witnessed by lower-bound constructions. The former result can be used to improve the leading constant for the crossing lemma for bipartite graphs which in turn improves various results such as the biplanar crossing number or the maximum number of edges a bipartite k-planar graph can have.

2012 ACM Subject Classification Mathematics of computing → Combinatorics; Mathematics of computing → Graph theory

Keywords and phrases Edge Density, Beyond Planarity, bipartite Graphs, Discharging Method

Digital Object Identifier 10.4230/LIPIcs.GD.2024.28

Funding *Maximilian Pfister*: This research was supported by the DFG grant SCHL 2331/1-1.

Acknowledgements We thank the anonymous reviewers for their valuable feedback.

1 Introduction

There is a rich literature about planar graphs which includes both algorithmic and combinatorial results. Since minimizing the number of edge crossings in a drawing improves its readability [18], a drawing should be planar whenever possible. Using Euler's Formula, one can easily derive that a planar graph with n vertices can have at most $3n - 6$ edges, which, however, implies that most graphs are in fact not planar. Empirical studies showed that not only the number of crossings but also their topological/geometrical properties play a crucial role regarding the readability of a drawing [9, 12]. This gave rise to the research area of *beyond planarity*, where graph classes are defined in terms of forbidden crossing configurations – see [8] for a survey of the area. While there is a plethora of beyond-planar classes, we are here concerned with k-planar and k-gap-planar graphs. A graph G is k-planar if there exists a drawing of G in the plane such that each of its edges has at most k crossings. A graph G is k-gap-planar, if G can be drawn in the plane such that there exists an assignment of every crossing to one of the involved edges such that each edge is assigned at most k crossings.
While the research of k-planar graphs (in particular, 1-planar graphs) started already more than half a century ago [15], k-gap-planar graphs were introduced quite recently [5] and can, in some sense, be interpreted as an asymmetric version of k-planar graphs. The authors of [5] showed that every $2k$-planar graph is k-gap-planar, but, for any fixed choice of k, there exists a 1-gap-planar graph which is not k-planar.
One of the most studied questions regarding a beyond-planar graph class is to determine its *edge density*, i.e., the maximum number of edges an n-vertex graph that belongs to this class can have, see e.g. [2, 7, 11] for some work in this direction. There is an additional motivation to study the edge density of k-planar graphs in particular: Improved results

32nd International Symposium on Graph Drawing and Network Visualization (GD 2024).
Editors: Stefan Felsner and Karsten Klein; Article No. 28; pp. 28:1–28:21
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Table 1 Overview of related edge density bounds and summary of our results (red angle brackets).

Graph class	Bipartite Upper bound Previous	Bipartite Upper bound Ours	Bipartite Lower bound
1-planar:	$3n-8$ [10]	–	$3n-8$ [10]
2-planar:	$3.5n-7$ [3]	–	$3.5n-12$ [3]
3-planar:	$5.205n$ [3]	$4n-8$ ⟨6⟩	$4n-12$ ⟨7⟩
k-planar:	$3.005\sqrt{k}n$ [3]	$2.871\sqrt{k}n$ ⟨17⟩	–
1-gap-planar:	$4.25n$ [16]	$4n-8$ ⟨5⟩	$4n-16$ [16]
k-gap-planar:	$4.25\sqrt{k}n$ [16]	$4.06n$ ⟨18⟩	–

in turn improve the leading constant of the celebrated *Crossing Lemma* which has various applications, see [1] – currently there are tight (up to a constant number of edges) bounds for $k \in \{1,2,3,4\}$ [15, 14, 13, 1]. The edge density was also studied when one imposes additional restrictions on (i) the drawings (e.g., the drawing is *outer* [6]) or on (ii) the graphs, where the most common restriction is to consider bipartite graphs, i.e., graphs which do not contain any odd cycle. In Table 1, we highlight the important past results regarding the edge density of bipartite graphs for our work.

Our contribution

In Section 3 and Section 4, we show that both bipartite 1-gap-planar graphs as well as bipartite 3-planar graphs on n vertices have at most $4n-8$ edges, thus (partially) answering an open problem posed in [3] and proving a conjecture posed in [16]. In Section 5, we use the result of Section 4 to improve the constant of the bipartite crossing lemma from $\frac{1}{18.1}$ to $\frac{1}{16.5}$ which in turn improves the upper bound on the edge density of bipartite k-planar graphs from $\approx 3.005\sqrt{k}n$ to $\approx 2.871\sqrt{k}n$.

2 Preliminaries

Throughout the paper, we will assume that all drawings are *simple* in the sense that no two adjacent edges cross and no edge crosses itself[1]. We consider bipartite graphs $G = (A, B, E)$ with $A \cap B = \emptyset$ and let $n = |A \cup B|$ and $m = |E|$. We will usually denote a vertex of A (of B) by a_i (by b_i). We do not require that G is simple, i.e., several edges between the same two endpoints are allowed if they are non-homotopic in the corresponding drawing Γ, i.e., for any such pair of edges, both the interior and the exterior of the closed region (defined by the pair) contains at least one vertex of G.

Let $\mathcal{B}$ be a beyond-planar graph class. Denote by $\mathcal{G}$ the set of all tuples (G, Γ) where $G \in \mathcal{B}$ is a bipartite graph of n vertices and Γ is a drawing of G (satisfying the constraints of $\mathcal{B}$) where any two copies of an edge are non-homotopic. Let $\mathcal{G}' \subset \mathcal{G}$ consist of all elements (G, Γ) such that G has the maximum number of edges among all graphs contained in $\mathcal{G}$. Finally, let $\mathcal{G}'' \subset \mathcal{G}'$ consist of all elements (G, Γ) such that Γ has the minimum number of crossings among all drawings contained in $\mathcal{G}'$. In the remainder, we will refer to such a tuple $(G, \Gamma) \in \mathcal{G}''$ as a *MaxMin* tuple (since G has the maximum number of edges and since Γ has the minimum number of crossings).

[1] we explicitly allow that two edges cross more than once

3 Bipartite 1-gap-planar graphs

Since we only consider 1-gap-planar graphs, we will abbreviate it henceforth by *gap-planar*. The *crossing graph* of a drawing Γ of graph G has a vertex v_e for any edge e of G and an edge $(v_e, v_{e'})$ if and only if the edges e and e' intersect in Γ. The following lemma is directly derived from the definition of gap-planarity.

▶ **Lemma 1** ([5]). *Let G be a gap-planar graph. Then G admits a gap-planar drawing Γ such that the crossing graph of Γ is a pseudoforest.*

For the remainder of this section, we will fix (G, Γ) as a MaxMin tuple (regarding gap-planarity). In order to show that G has at most $4n - 8$ edges, we want to find a set of edges of G, denoted by $\mathcal{X}$, such that no two edges of $\mathcal{X}$ cross in Γ and such that $|\mathcal{X}| \geq \frac{m}{2}$. The result then follows immediately since $\mathcal{X}$ induces a bipartite planar subgraph of G, hence $|\mathcal{X}| \leq 2n - 4$ and thus $m \leq 4n - 8$ follows. To define $\mathcal{X}$, we will consider the components of the crossing graph $\mathcal{I}$ of Γ. Recall that by definition, no two edges that belong to different components of $\mathcal{I}$ intersect, thus we can consider the components separately.

▶ **Lemma 2.** *Let X be an arbitrary component of $\mathcal{I}$. If X (i) is a tree, (ii) contains an even-length cycle or (iii) contains an odd-length cycle C and at least one rooted tree at a vertex of C (that is edge-disjoint from C) is a path of odd length, then X contains an independent set of size at least $\frac{|X|}{2}$.*

Proof. If X is a tree or contains an even-length cycle, then by definition and Lemma 1 X is bipartite and thus its vertices can be colored using two colors. Each color induces an independent set, from which one has size at least $\frac{|X|}{2}$. For the third case, let $C = (v_1, v_2, \ldots, v_k)$ be the unique odd cycle of X and w.l.o.g. assume that the tree rooted at v_1 is in fact a path $(u_1, u_2, \ldots, u_{2j+1})$ such that $(v_1, u_1) \in \mathcal{I}$. Coloring u_i with i odd implies that we have colored $j + 1$ vertices in one color, which is enough to accommodate for all vertices of the path in addition to v_1. Clearly, the only vertex that is not on the path which is influenced by the coloring is v_1. Removing v_1 (and the whole path) from X yields a tree which has an independent set of the desired size as shown in the previous case. Combining both independent sets then concludes the proof. ◀

Components of $\mathcal{I}$ which do not meet the criteria of Lemma 2 are called *critical*. By Lemma 2, any such component X is a pseudotree whose unique odd cycle C is of odd length and none of the trees rooted at the vertices of C are odd-length paths. For these critical components, we cannot directly find an independent set in $\mathcal{I}$ of appropriate size. To be more precise, we can only find an independent set of size $\left\lceil \frac{|X|-1}{2} \right\rceil$. To overcome this issue, we will show that for any such component X, there exists an uncrossed edge in Γ (i.e., a singleton in $\mathcal{I}$) which we can uniquely assign to X. The next lemma follows by our choice of (G, Γ).

▶ **Lemma 3.** *Let $e_1 = (a_1, b_1)$ and $e_2 = (a_2, b_2)$ be two edges of G that intersect in Γ. Then, (a_1, b_2) or (a_2, b_1) drawn along the curves of e_1 and e_2 is present in Γ.*

Proof. Suppose for a contradiction that neither (a_1, b_2) nor (a_2, b_1) drawn along e_1 and e_2 exist in Γ . Denote by x the intersection point between e_1 and e_2. W.l.o.g. assume that the crossing between e_1 and e_2 was assigned to e_1 in Γ. This implies that no crossing on the (open) segments a_1x and xb_1 can be assigned to e_1. Moreover, at most one crossing is assigned to e_2 by definition – w.l.o.g. assume this crossing is due to an edge that intersects e_2 on the segment a_2x. Now, consider the graph $G' = G \setminus \{(a_1, b_1), (a_2, b_2)\} \cup \{(a_1, b_2), (a_2, b_1)\}$

with corresponding drawing Γ' where the drawing of all edges but (a_1, b_2) and (a_2, b_1) is inherited from Γ, while the edges (a_1, b_2) and (a_2, b_1) are drawn along (the original curves of) e_1 and e_2, refer to Fig. 1a. First observe that G' is a non-homotopic multigraph, as neither (a_1, b_2) nor (a_2, b_1) drawn along e_1 and e_2 were present by assumption. Further, Γ' is a valid gap-planar drawing as we do not need to assign (a_1, b_2) any crossing, while (a_2, b_1) is assigned at most one crossing. But now we have a contradiction to our choice of G and Γ, as G' contains the same number of edges as G, but Γ' contains less crossings than Γ. ◀

▶ **Lemma 4.** *Let S be the set of singletons in $\mathcal{I}$ and let Z be the set of critical components of $\mathcal{I}$. Then $|S| \geq |Z|$.*

Proof. As we will argue about graph G and its crossing graph $\mathcal{I}$ simultaneously, we will assume in the following that an edge $e_i = (a_i, b_i)$ of G corresponds to a vertex v_i of $\mathcal{I}$. Let $X \in Z$ be a critical component in $\mathcal{I}$ and let $C = (v_1, v_2, \ldots, v_{2j+1})$ be its unique odd cycle in $\mathcal{I}$. Pick two adjacent vertices v_1 and v_2 of C with corresponding edges $e_1 = (a_1, b_1)$ and $e_2 = (a_2, b_2)$ of G. Lemma 3 ensures that at least one of (a_1, b_2) or (a_2, b_1) exists in G such that its curve follows e_1 and e_2 in Γ. W.l.o.g. assume that (a_1, b_2) exists and denote by v the corresponding vertex of (a_1, b_2) in $\mathcal{I}$. We distinguish between the following two cases based on whether v is adjacent to a vertex of C or not.

Assume first that v is adjacent to some vertex of C, i.e., (a_1, b_2) intersects an edge of G associated with a vertex of C in Γ. As (a_1, b_2) is drawn along e_1 and e_2, this edge is either e_{2j+1} or e_3 by construction. We first observe that (a_1, b_2) cannot intersect both e_3 and e_{2j+1}, as otherwise X is not a pseudoforest (this also holds in the case where $e_3 = e_{2j+1}$, i.e., C is a 3-cycle, in which case (a_1, b_2) intersects this edge at most once). W.l.o.g assume that (a_1, b_2) crosses e_3, the other case is symmetric. If there is an additional edge e' besides e_3 that is crossing (a_1, b_2), then e' also crosses either e_1 or e_2 as (a_1, b_2) is drawn along e_1 and e_2, but then X is not a pseudoforest as this crossing would close another cycle. In particular, if it crosses e_1, then we obtain the cycle (v_1, v_2, v_3, v, v'), and otherwise we obtain (v_2, v_3, v, v'), where v' is the corresponding vertex of e' in $\mathcal{I}$. Hence, (a_1, b_2) only crosses e_3 – but then we have an odd-length path rooted at v_3 in $\mathcal{I}$ (that only contains vertex v), in which case X is not critical, a contradiction.

Thus we can assume from now on that v is not adjacent to any vertex of C. This means that either (a_1, b_2) is crossing free in Γ, or there is an edge e' that intersects (a_1, b_2) and thus either e_1 or e_2. We keep the former case in mind and consider the latter case. W.l.o.g. assume that e' intersects e_1, the other case is symmetric. Denote by v' the corresponding vertex of e' in $\mathcal{I}$. Now, in $\mathcal{I}$, we have a tree $T \subset \mathcal{I}[X]$ rooted at v_1 such that $(v_1, v') \in T$ and $(v', v) \in T$. Let t be the depth of T, let u_k be a leaf of depth t and let u_{k-1} be the unique parent of u_k. Denote by $u_k^1, \ldots, u_k^r$ the children of u_k with $u_k = u_k^1$ and let (α, β) be the corresponding edge to u_{k-1} in G. By traversing (α, β) starting from α, the first intersection we encounter is either with an edge that corresponds to a leaf u_k^i or with the edge that corresponds to the unique parent of u_{k-1}. If the latter case occurs, observe that by traversing (α, β) starting at β we encounter a leaf first. Denote by (α', β') the corresponding edge to u_k^i in G. Hence, w.l.o.g. assume that the crossing x with (α', β') is the first one that we encounter when traversing (α, β) starting at α. This implies that the segment αx is crossing free. Moreover, since u_k^i is a leaf, both segments $\alpha' x$ and $x\beta'$ are crossing free – but then the edge (α, β') exists (crossing free) in Γ by maximality.

In both cases, we found an uncrossed edge for a fixed pair of consecutive vertices. By repeating this argumentation for any two consecutive vertices of C, we can find a set of edges P with $|P| = |C| \geq 3$ such that any edge in P is uncrossed in Γ, i.e., belongs to S.

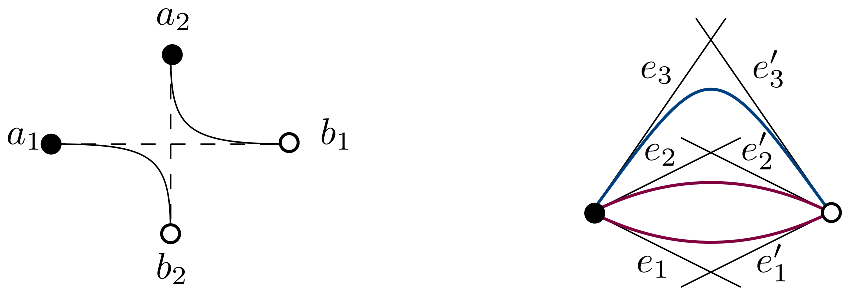

(a) **(b)**

Figure 1 (a) Illustration for Lemma 3. (b) Subgraph formed by edges e_i, e'_1 for $i \leq 3$. The blue edge is a non-homotopic copy of the red ones.

It remains to consider the interaction of different components of $\mathcal{I}$. Given two critical components X and Y of $\mathcal{I}$, it is possible that an edge occurs in both P_X and P_Y. We claim the following: any non-homotopic copy of such an edge occurs in at most two such sets. Assuming we have this claim at hand, the total number of uncrossed edges in Γ is at least

$$\frac{1}{2}\sum_{i=1}^{|Z|} X_i \geq \frac{1}{2}\sum_{i=1}^{|Z|} 3 = \frac{3|Z|}{2} \geq |Z|$$

since any odd cycle of $\mathcal{I}$ has size at least three. It remains to prove the claim. Suppose for a contradiction that one copy of an edge (a, b) belongs to at least three (critical) components X_1, X_2 and X_3. Denote by e_i and e'_i the two edges of X_i where (a, b) is drawn along. By definition, no edge of X_i crosses an edge of X_j for $i \neq j$. But then e_1 and e'_1, e_2 and e'_2 and e_3 and e'_3 need to bound the same cell of Γ which is impossible, as e_1, e_2 and e_3 are all incident to a, while e'_1, e'_2 and e'_3 are all incident to b, see Fig.1b. ◀

By Lemma 4, there exists a bijective mapping from S to Z. We will call a critical component with an additional (singleton) edge an *augmented* component. Observe that every augmented component X has an independent set of size $\left\lceil \frac{|X|-1}{2} \right\rceil + 1 \geq \frac{|X|}{2}$.

▶ **Theorem 5.** *An n-vertex bipartite gap-planar multigraph without homotopic parallel edges has at most $4n - 8$ edges.*

Proof. Let $\mathcal{X}$ be the union of the maximum independent sets of every (augmented) component X of $\mathcal{I}$. Since we established that every component of $\mathcal{I}$, in particular also the augmented critical components, has an independent set of size at least half its size, it follows that $|\mathcal{X}| \geq \frac{m}{2}$. Clearly, no two edges of $\mathcal{X}$ intersect in $\mathcal{I}$, hence the edges of $\mathcal{X}$ induce a planar bipartite multi-graph $G_{\mathcal{X}}$. Since $G_{\mathcal{X}}$ does not contain any homotopic multiedges by construction, it still holds that any face of $G_{\mathcal{X}}$ has length at least four. Since Euler's formula can also be applied to non-simple graphs, we have that $G_{\mathcal{X}}$ has at most $2n - 4$ edges, and thus $m \leq 4n - 8$ which concludes the proof. ◀

The lower-bound construction in [16], which yields n-vertex bipartite gap-planar graphs with $4n - 16$ edges, asserts that our bound is tight up to an additive constant.

4 Bipartite 3-planar graphs

In this section, we will establish an upper bound on the number of edges a bipartite 3-planar graph can have.

▶ **Theorem 6.** *A bipartite* 3*-planar graph* G *with* n *vertices has at most* $4n - 8$ *edges.*

Before we prove Theorem 6, we will provide a lower-bound construction to show its tightness (which has been suggested in [3]).

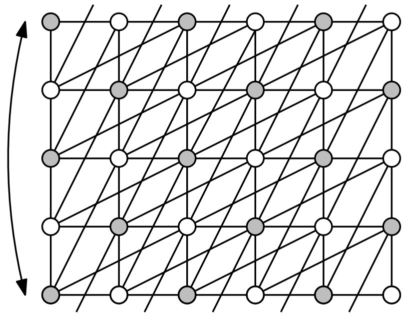

Figure 2 Lower bound construction for bipartite 3-planar graphs. The vertices of the first and last row coincide.

▶ **Theorem 7.** *For infinitely many* n*, there exists a bipartite* 3*-planar graph* G *with* n *vertices and* $4n - 12$ *edges.*

Proof. Assume n is divisible by four. We arrange the vertices equally in four rows (i.e., every row contains $\frac{n}{4}$ vertices) and wrap the construction around a cylinder; i.e., the topmost and the bottommost row in Figure 2 coincide. Clearly, the drawing is 3-planar. In order to determine the number of edges, let us count the degrees of the vertices. The vertices in the first and last column have degree five, while all other vertices have degree eight. Hence, $2m = 8(n - 8) + 8 \cdot 5$ and thus $m = 4n - 12$. ◀

Similar to the previous section, we will fix (G, Γ) as a MaxMin tuple (w.r.t. 3-planarity) for the remainder of the section. In order to prove Theorem 6, we will use the *Discharging Method*, which was initially introduced in order to prove structural properties of planar graphs, e.g., for the proof of the Four Color Theorem [4]. Our proof will reuse parts of the notation and ideas from [1] where the author proved an upper bound on the number of edges of 4-planar graphs. We denote by $P(\Gamma)$ the so called *planarization* of Γ, i.e, the vertices and crossing points of Γ are the vertices of $P(\Gamma)$, while the edges of $P(\Gamma)$ are the crossing-free segments in Γ which are bounded by vertices and crossing points. We will refer to the vertices of $P(\Gamma) \cap G$ as *original*. We will denote by $e = (a, b)$ an edge of G while the segment of e restricted to a face f of $P(\Gamma)$ is denoted by $f[e]$, or, if the two endpoints x and y of the segment are known, we might also refer to the segment as xy. We will also use (B) and $(3P)$ to abbreviate the bipartite and 3-planar property, respectively.

We will prove Theorem 6 by induction on the number of vertices of G. Clearly, if $n \leq 6$, we have $4n - 8 > \binom{n}{2}$ and the theorem holds. Thus, we assume that $n \geq 7$. Moreover, we can assume that every vertex in G has degree at least 5, as otherwise the theorem follows by removing a vertex of small degree and applying the induction hypothesis. We begin with the following important observation for $P(\Gamma)$; its proof is analogous to the one in [1].

▶ **Property 8.** *If* $P(\Gamma)$ *is not* 2*-connected, then* G *has at most* $4n - 8$ *edges.*

Proof. Assume that $P(\Gamma)$ has a vertex x such that $P(\Gamma) \setminus \{x\}$ is not connected. The vertex x is either a vertex of G or a crossing point of two of its edges. Suppose first that x is vertex of G. Then, $G \setminus \{x\}$ is also not connected. Let $G_1, \ldots, G_k$ be the connected components of $G \setminus \{x\}$, let G' be the graph induced by $V(G_1) \cup \{x\}$ and let G'' be the graph induced by $V(G_2) \cup \ldots \cup V(G_k) \cup \{x\}$. Note that $4 \leq |V(G')|, |V(G'')| < n$, since we established earlier that the minimum degree of a vertex of G is at least 5. Therefore, it follows from the induction hypothesis that $|E(G)| \leq 4|V(G')| - 8 + 4|V(G'')| - 8 = 4(n+1) - 16 < 4n - 8$.

Suppose now that x is a crossing point of two edges e_1 and e_2. Let $\hat{G}$ be the graph we obtain by adding x as a vertex to G. Therefore, $|V(\hat{G})| = n + 1$ and $|E(\hat{G})| = |E(G)| + 2$. Let $G_1, \ldots, G_k$ be the connected components of $\hat{G} \setminus \{x\}$, let G' be the graph induced by $V(G_1) \cup \{x\}$ and let G'' be the graph induced by $V(G_2) \cup \ldots \cup V(G_k) \cup \{x\}$. Again, note that $4 \leq |V(G')|, |V(G'')| < n$ by our observation about the minimum degree. It follows from the induction hypothesis that $|E(G)| \leq 4|V(G')| - 8 + 4|V(G'')| - 8 - 2 = 4(n+2) - 18 < 4n - 8$. ◀

Property 8 allows us to assume that $P(\Gamma)$ is 2-connected. The *boundary* δf of a face f in $P(\Gamma)$ consists of all the edges of $P(\Gamma)$ that are incident to f. Since $P(\Gamma)$ is 2-connected, the boundary of every face in $P(\Gamma)$ is a simple cycle. Thus, we can define the *size* of a face f, $|f|$, as the number of edges of $P(\Gamma)$ on its boundary. We will denote by $V(f)$ the set of original vertices on the boundary of f.

▶ **Observation 9.** *The boundary of every face in $P(\Gamma)$ is a simple cycle.*

Similar to [1], we begin by assigning a *charge* to every face of $P(\Gamma)$ such that the total charge is $4n - 8$. Then, we redistribute the charge in several steps such that the charge of every face is non-negative and the charge of every original vertex v is at least $\deg(v)/2$. Hence, $|E(G)| = \sum_{v \in V(G)} \deg(v)/2 \leq 4n - 8$ and we get the desired bound on $|E(G)|$.

Let V', E', and F' denote the vertex, edge, and face sets of $P(\Gamma)$, respectively. Clearly, $\sum_{f \in F'} |V(f)| = \sum_{v \in V(G)} \deg(v)$ and $\sum_{f \in F'} |f| = 2|E'| = \sum_{u \in V'} \deg(u)$ holds. Every vertex in $V' \setminus V(G)$ is a crossing point in G and therefore its degree in $P(\Gamma)$ is four. Hence,

$$\sum_{f \in F'} |V(f)| = \sum_{v \in V(G)} \deg(v) = \sum_{u \in V'} \deg(u) - \sum_{u \in V' \setminus V(G)} \deg(u) = 2|E'| - 4\left(|V'| - n\right).$$

Assigning every face $f \in F'$ a charge of $|f| + |V(f)| - 4$, we get a total charge of

$$\sum_{f \in F'} \left(|f| + |V(f)| - 4\right) = 2|E'| + 2|E'| - 4\left(|V'| - n\right) - 4|F'| = 4n - 8,$$

Recall that we will redistribute the initial charge s.t. the charge of every face of F' is non-negative, while every original vertex v has charge at least $deg(v)/2$. An equivalent precondition is that

$$ch(f) \geq 0.5 \cdot |V(f)| \text{ for all } f \in F' \tag{1}$$

as we can then redistribute the excess charge from the faces to the original vertices in a final step. Let f be face of F' with $|V(f)| = x$ and $|f| = y$. We will then refer to f as an x-y face. To ease the notation, we will use the terms *triangles*, *quadrangles*, *pentagons*, *hexagons* and *heptagons* to refer to faces of size 3,4,5,6 and 7, respectively. For example, a 2-triangle is a face of size 3 whose boundary contains two original vertices. Since Γ is a non-homotopic drawing with the minimum number of crossings, there are no faces of size 2 in F'. Therefore, initially, the only faces which do not satisfy Equation (1) are 0-triangles and 1-triangles. In order to ensure that a face f still satisfies Equation (1) after it redistributed some of its charge to another face, we will introduce the notion of *local charge* for faces that contain sufficiently many original vertices. Let $f \in F'$ be a face of $P(\Gamma)$ with $|V(f)| \geq 2$. Let

$v_1, v_2, \dots v_k$ be the (ordered) sequence of vertices of $P(\Gamma)$ that are contained in δf. Let v_i and v_j be two consecutive original vertices of δf with $i < j$, i.e., all vertices v_k with $i < k < j$ are crossing points in Γ.
If $|j - i| > 1 \,(\mathrm{mod}(k-1))$ holds, then we split f along the hypothetical edge (v_i, v_j) and obtain a so called *subdivision face* $f_a = (v_i, v_{i+1}, \dots, v_j)$. Observe that $|f_a| \geq 3$ and $|V(f_a)| = 2$.

For example, let f be a 2-pentagon with vertices v_1, v_2, v_3, v_4, v_5 such that v_1 and v_3 are original vertices. By splitting along the pair $\{v_1, v_3\}$ we obtain the subdivision face $f_a = (v_1, v_2, v_3)$ and, as a remainder, $f_b = (v_3, v_4, v_5, v_1)$. Observe that in this particular case, f_b is also a subdivision face as $|f_b| \geq 3$ and $|V(f_b)| = 2$, but this is not the case in general. Let us now consider the charge distribution of such a subdivision. Since a subdivision always occurs at an edge between two original vertices, we have $|V(f_a)| + |V(f_b)| = |V(f)| + 2$ and $|f_a| + |f_b| = |f| + 2$. Hence, $ch(f) = ch(f_a) + ch(f_b)$ after the initial assignment holds. Since both f_a and f_b contain the two original vertices on their boundary which defined the subdivision edge and since the subdivision edges do not contribute to the degree of these two vertices, f_a and f_b have to retain less charge. In particular, we require that every subdivision face f_a satisfies

$$ch(f_a) \geq 0.5 \tag{2}$$

while the remainder f_b has to satisfy

$$ch(f_b) \geq 0.5(|V(f_b)| - s(f_b)) \tag{3}$$

where $s(f_b)$ is the number of subdivision edges on the boundary of f_b. Observe that Equation (2) is a special case of Equation (3) when $|V(f)| = 2$ and $s(f) = 1$. If all subdivision faces f_x of a face f (and its possible remainder) satisfy Equation (2) and Equation (3), it is immediate that f satisfies Equation (1), which allows us to argue about the charge in a more local way. In order to describe the way the charging of $\{0,1\}$-triangles works we will need the following definitions. Let f be a face, let e be one of its edges, and let f' be the other face that shares e with f. We say that f' is the *immediate neighbor* of f at e. Note that $f' \neq f$ since $P(\Gamma)$ is 2-connected. The following two definitions are taken from [1].

Wedge-neighbors. Let f_0 be a $\{0,1\}$-triangle in $P(\Gamma)$ and let x_1 and y_1 be two vertices of f_0 that are crossing points in Γ. Denote by e_x (resp., e_y) the edge of G that contains x_1 (resp., y_1) and does not contain y_1 (resp., x_1). Note that e_x and e_y intersect at the other vertex of f_0. Let f_1 be the immediate neighbor of f_0 at x_1y_1. For $i \geq 1$, if f_i is a 0-quadrangle, then denote by $x_{i+1}y_{i+1}$ the edge of $P(\Gamma)$ opposite to x_iy_i in f_i, such that e_x contains x_{i+1} and e_y contains y_{i+1}, and let f_{i+1} be the immediate neighbor of f_i at $x_{i+1}y_{i+1}$. Observe that $f_i \neq f_j$ for $i < j$, for otherwise x_j coincides with one of x_i and x_{i+1} (which implies that e_x crosses itself) or with one of y_i and y_{i+1} (which implies that e_x and e_y intersect more than once). Let j be the maximum index for which f_j is defined. We then call f_j the *wedge-neighbor* of f_0 at x_1y_1 (note that f_j is uniquely defined). We also say that f_0 is the wedge-neighbor of f_j at x_jy_j. Observe that since the relations being an immediate neighbor at a certain edge of $P(\Gamma)$ and being an opposite edge in a 0-quadrangle are both one-to-one, it follows that indeed there cannot be another triangle but f_0 that is a wedge-neighbor of f_j at x_jy_j. Note also that since e_x and e_y already intersect at a vertex of f_0, and by definition f_j cannot be a 0-quadrangle, either $|f_j| \geq 5$ or $|f_j| = 4$ and $|V(f_j)| \geq 1$.

▶ **Observation 10** ([1]). *Let f be a face and let e be one of its edges. Then there is at most one triangle t such that t is a wedge-neighbor of f at e. If such a triangle exists, then either $|f| \geq 5$ or $|f| = 4$ and $|V(f)| \geq 1$.*

Vertex-neighbors. Let x be a crossing point in Γ and let f_0, f_1, f_2 and f_3 be the four faces that are incident to x in clockwise order around x. Note that these faces are distinct, since $P(\Gamma)$ is 2-connected. We say that f_0 and f_2 (resp., f_1 and f_3) are vertex-neighbors at x.

We will introduce one additional kind of neighbor relation.

Rich immediate neighbor. Let f be a face with edges $e_0, e_1, \ldots e_k$. We call an immediate neighbor f' of f at edge e_i *rich* if, in the facial walk of f', we have e_{j-1}, e_j, e_{j+1} such that $e_i = e_j$, the common endpoint of e_{j-1} and e_j (e_{j+1} and e_j) is a crossing point in Γ, while the other endpoint of e_{j-1} (e_{j+1}) is an original vertex.

4.1 Step 1: Charging the 0-triangles

Consider a 0-triangle f and its immediate neighbors f_0,f_1 and f_2. If one such f_i is a rich immediate neighbor, it charges one unit to f. Otherwise, f obtains $\frac{1}{3}$ units of charge from each of its three wedge-neighbors. Hence, in every case, $ch(f) = 0$ for all 0-triangles f.

▶ **Property 11.** *Let f be a 0-triangle. If one immediate neighbor f_i of f is a 0-x face, then f has a rich immediate neighbor f_j with $f_j \neq f_i$.*

Proof. Assume that the edges that define f are denoted by e_0, e_1 and e_2 such that face f_i shares the edge e_i with f. W.l.o.g. assume that f_0 is a 0-x face. Then, by definition, e_1 and e_2 have a crossing in $\delta f_0 \setminus \delta f$ and are thus crossed three times each. Since e_0 is crossed by e_1 and e_2 already, it can be crossed at most once more, w.l.o.g. assume that e_0 has a crossing in $\delta f_1 \setminus \delta f$. But then $f_2[e_0]$ and $f_2[e_1]$ each contain an original vertex by (3P) and hence f_2 is a rich immediate neighbor. ◀

▶ **Observation 12.** *No 0-x face has to pay in the initial charging step.*

Obviously, a sufficiently large face can be a wedge-neighbor and an immediate neighbor to several other faces. Fix an edge e. Since both the immediate neighbor relation and the wedge-neighbor relation is unique, e can be assigned to at most one wedge-neighbor (to at most one immediate neighbor). Further, if e is used for the rich immediate neighbor relation, it substitutes the wedge-neighbor relation. Hence, a face f is either a wedge-neighbor or a rich immediate neighbor over each of its edges. Since every rich immediate neighbor relation introduces three new edges and an original vertex (every wedge-neighbor relation introduces one new edge), our face gets an additional charge of four (one) units which clearly accommodates for the discharge if our face is sufficiently large. For example, a 2-pentagon that is wedge-neighbor to one rich immediate neighbor has, after discharging, still two units of charge left, which clearly satisfies Equation (1). We conclude with the following observation:

▶ **Observation 13.** *After the initial charging step, only the* 1*-triangles and (possibly) the* 1*-quadrangles do not satisfy Equation* (1)*.*

Observe that the charge of every 1-triangle is 0, while the charge of a 1-quadrangle can be as low as $\frac{1}{3}$ (this occurs when a 1-quadrangle is a wedge-neighbor to two 0-triangles).

4.2 Step 2: Charging the 1-triangles

Every 1-triangle obtains 0.5 units of charge from its unique wedge-neighbor. Since a 1-triangle is not a wedge-neighbor to a 0-triangle by Observation 10, it did not loose charge in the previous step, hence we have $ch(f) = 0.5$ for any 1-triangle $f \in F'$. Observe that after these two rounds, every triangle satisfies Equation (1) (while this is explicit for $\{0, 1\}$-triangles, Observation 10 guarantees that no triangle is a wedge-neighbor and, by definition, they are also not a rich immediate neighbor, hence the charge of 2-triangles still satisfies Equation (1)). Moreover, by definition, 0-quadrangles are not a wedge-neighbor nor a rich neighbor. Since any x-y face can be wedge-neighbor to at most $y - x$ triangles, they loose at most $\frac{1}{2}(y - x)$ charge in the second step and thus always satisfy Equation (1) if $y \geq 8$. Recall that wedge-neighbor relations and rich immediate neighbor relations cannot interfere and rich immediate neighbors give a vast surplus of charge to our face. Thus, the only faces which potentially do not satisfy Equation (1) are 0-heptagons, 0-hexagons, 0-pentagons and 1-quadrangles. We will now establish that the first two in fact satisfy Equation (1):

1. Face f is a 0-heptagon: By Observation 12, f did not discharge to a 0-triangle. Since its initial charge of three is sufficient if f is a wedge-neighbor to at most six triangles, f needs to a be a wedge-neighbor to seven 1-triangles. But this is impossible as the corresponding edges would form an odd cycle in G, a contradiction to (B).
2. Face f is a 0-hexagon: Again, we have by Observation 12 that f did not discharge to a 0-triangle. Its initial charge is sufficient if f is wedge-neighbor to at most four 1-triangles. If f is a wedge-neighbor to five or six 1-triangles, we again have an induced cycle of odd length in G, a contradiction to (B).

4.3 Step 3: Recharging 0-pentagons and 1-quadrangles

In order for a 0-pentagon (1-quadrangle) to not satisfy Equation (1), it requires a quite limited local structure which we will exploit to locally redistribute charges.

Throughout the analysis, we will denote by $e_0, e_1 \ldots e_{k-1}$ the original edges (i.e., the edges in G) whose pairwise crossing points (and the corresponding segments) define face f. Let $e_i = (a_i, b_i)$ where $a_i \in A$ and $b_i \in B$ are original vertices. We will refer to the immediate neighbor to f at the edge e_i as f_i. We will denote by x_i the common endpoint of $f[e_i]$ and $f[e_{i+1}]$ which is part of the boundary δf of f. We will further denote by f'_i the vertex neighbor of f at x_i. Finally, we will denote by t_i the 1-triangle which is the wedge-neighbor to f at edge e_i (if it exists). Note that if t_i exists, then $t_i = f_i$ unless f_i is a 0-quadrangle. We denote the unique real vertex of t_i as v_i. Throughout the proof, our charging scheme will maintain the following invariant.

▶ **Invariant 14.** *Let x be an intersection point that belongs to the boundary δf of a face f and let e_1 and e_2 be the edges which define x. If neither $f[e_1]$ nor $f[e_2]$ contains an original vertex, then f does not discharge over x.*

4.3.1 f is a 1-quadrangle

With a slight abuse of notation, let $v = x_3$ be the real vertex of f. W.l.o.g. assume that $v \in A$.

1. f is wedge-neighbor to two 0-triangles. It follows that e_1 and e_2 are crossed three times each. Consider f'_0. By assumption, $\delta f'_0$ contains $f'_0[e_0], f'_0[e_1]$ and $f'_0[e_2]$ such that the endpoints of $f'_0[e_1] \setminus f'_0[e_0]$ and $f'_0[e_2] \setminus f'_0[e_0]$ are original vertices by $(3P)$, see Fig. 3a. Hence, f'_0 is a rich immediate neighbor to f_1. A similar observation holds for f'_2 and f_2. Thus, f did not loose any charge in the initial step and thus satisfies Equation (1), see Fig. 3a.
2. f is wedge-neighbor to two 1-triangles. Consider the edges e_1 and e_2. They both have exactly two crossings which belong to δf.
   a. $f_0[e_1]$ contains an original vertex. In particular, since f_2 is a 1-triangle, we have by (B) that the endpoint of $f_0[e_1]$ is of the same partition as v and thus it is a_1. This implies that the segment of e_1 delimited by a_1 and x_0 is not crossed (in which case

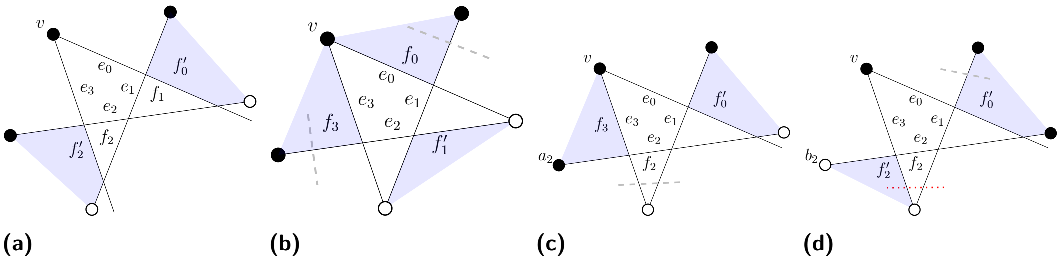

Figure 3 Illustrations used in the recharging of 1-quadrangles.

the gray-dashed edge in Fig. 3b that intersects e_0 is not present) and the sequence (a_1, x_0, v) of f_0 defines a subdivision face, which contains 0.5 excess charge which will be redistributed to f (over its vertex neighbor f_1 at x_0).

b. $f_3[e_2]$ contains an original vertex. Following an analogous argument as before, this original vertex is a_2 and thus the sequence (v, x_2, a_2) forms a subdivision face and distributes its excess charge to f via f_2.

c. Neither $f_0[e_1]$ nor $f_3[e_2]$ contains an original vertex. This setting can be observed in Fig. 3b if we assume that both gray-dashed edges are present. Consider the vertex neighbor f'_1 to f at x_1. Since e_1 and e_2 are both crossed thrice, it follows that both $f'_1[e_1]$ and $f'_1[e_2]$ contain an original vertex. Thus (b_1, x_1, b_2) forms a subdivision face which will transfer its excess charge to its vertex-neighbor f at x_1.

3. f is wedge-neighbor to one 0-triangle and one 1-triangle. W.l.o.g. assume that the wedge-neighbor to f via edge e_1 is a 0-triangle and hence the wedge-neighbor to f via edge e_2 is a 1-triangle. By assumption, edge e_2 has three crossings. Assume first that, $a_2 \in f_3[e_2]$, see Fig. 3c. Suppose first that $f_2 \neq t_2$, i.e., the gray-dashed edge in Fig. 3c is present. Observe that in this case, we have that e_1 and e_2 have three crossings each. But then f'_0 is a rich immediate neighbor to f_1 (as witnessed by the edges $f'_0[e_2], f'_0[e_0]$ and $f'_0[e_1]$) and hence f did not charge f_1 to begin with, a contradiction to our assumption. If $f_2 = t_2$ holds, then we observe that (v, x_2, a_2) defines a subdivision face of f_3, which will distribute its excess charge to its vertex-neighbor f_2 at x_2, which will then propagate it to its unique wedge-neighbor f. Otherwise, we have $b_2 \in f_3[e_2]$ and thus $a_2 \in f'_0[e_2]$. If also $a_1 \in f'_0[e_1]$ holds, then f'_0 is a rich immediate neighbor to f_1 – hence f did not redistribute charge to f_1 and thus its charge was sufficient to begin with. Hence assume that $a_1 \notin f'_0[e_1]$, which implies e_1 has an additional crossing which belongs to $\delta f'_0$. In this case, consider f'_2 and observe that (b_2, x_2, b_3) lie consecutively on the boundary of f'_3, as the dotted red edge in Fig. 3d cannot be present due to $(3P)$, hence the sequence defines a subdivision face which distributes its excess charge to its vertex neighbor f at x_2.

4.3.2 f is a 0-pentagon

Observation 12 establishes that f is not a (discharging) wedge-neighbor to any 0-triangle. By (B), f cannot be a wedge-neighbor to five 1-triangles.

f is a wedge-neighbor to exactly three 1-triangles

Observe that in this case, it is sufficient to distribute 0.5 units of charge to f. Assume first that the three 1-triangles appear consecutively. W.l.o.g. assume that t_0, t_1 and t_2 exist and that the common endpoint of e_0 and e_2 (i.e., v_1) is of the same partition as v_0.

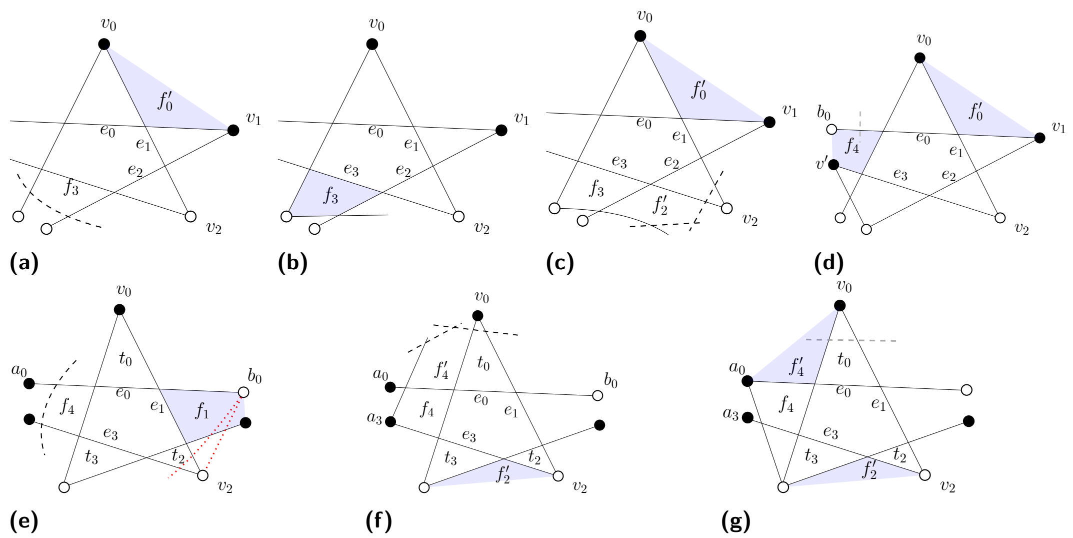

Figure 4 Illustrations used in the recharging of 0-pentagons that are wedge-neighbors to three 1-triangles.

1. f_3 is a 0-x face. In this case, both e_2 and e_4 have a crossing in $\delta f_4 \setminus \delta f$ and are thus crossed three times each – this implies that $f_0 = t_0$ and $f_1 = t_1$ and thus the sequence (v_1, x_0, v_0) forms a subdivision face of f_0' which can distribute its excess charge of 0.5 units to its vertex-neighbor f at x_0, see Fig. 4a.
2. f_3 is a 1-quadrangle, hence either e_2 or e_4 has a crossing in $\delta f_3 \setminus \delta f$.
   - e_2 has a crossing in $\delta f_3 \setminus \delta f$. Again we observe that f_3 cannot be a wedge-neighbor at e_2 by (B). Thus, if f_2' is not a 0-pentagon, f_3 has sufficient excess charge which can be distributed to its immediate neighbor f at e_3, see Fig. 4b. If f_2' is a 0-pentagon, then we necessarily have an additional edge that intersects both e_1 and e_3. But then, by assumption both e_1 and e_2 have three intersections each – hence it holds that $f_0 = t_0$ and $f_1 = t_1$ and thus (v_1, x_0, v_0) forms a subdivision face of f_0' which distributes its excess charge to its vertex-neighbor f at x_0, see Fig. 4c.
   - e_4 has a crossing in $\delta f_3 \setminus \delta f$. In this case, the immediate neighbor of f_3 at e_4 can be a 1-triangle f', see Fig. 4d (if this is not the case, we proceed as in the previous case). Let us denote the unique original vertex of f' by v'. If $b_0 \in f_4[e_0]$ then the sequence (b_0, x_4, x_3, v') either defines f_4 or forms a subdivision face – in both cases it has sufficient excess charge to distribute 0.5 units each to its vertex-neighbor f_0 at x_4, which then propagates the charge to its unique wedge-neighbor f. Otherwise, $b_0 \notin f_4[e_0]$. Since e_4 and e_0 have three crossings each, it follows that $f_0 = t_0$ and $f_1 = t_1$ and hence again (v_1, x_0, v_0) forms a subdivision face of f_0' which distributes its excess charge to its vertex-neighbor f at x_0, again see Fig. 4d.
3. In the remaining cases, f_3 always has sufficient charge which it can distribute to f over edge e_3. The crucial observation is that the edge e_3 is neither part of a wedge-neighbor relation nor of a rich immediate neighbor relation. Since we discharge 0.5 units over e_3, which is the same quantity as a discharge over an edge that defines a wedge-neighbor relation in the second step, the analysis at the start of Step 2 still holds.

Assume now that the three 1-triangles do not appear consecutively. W.l.o.g. assume that t_0, t_2 and t_3 exist and that $v_0 \in A$ (which implies $v_2 \in B$ and $v_3 \in B$). Further, assume that, when traversing e_0 starting at a_0, we encounter x_4 before x_0, see Fig. 4e.

1. f_4 is a 0-x face. This implies that both e_0 and e_3 have a crossing in $\delta f_4 \setminus \delta f$. Consider f_1. By construction $b_0 \in f_1[e_0]$, thus f_1 is not a 0-x face. By assumption, f_1 is not a 1-triangle. If f_1 is a 1-quadrangle, it is not a wedge-neighbor to any 1-triangle by (B) and it cannot be a wedge-neighbor to a 0-triangle due to $(3P)$ (as it would necessarily have to cross e_1 and hence also e_3, which already has three crossings), see Fig. 4e. The crucial observation here is that f_1' cannot be a 0-pentagon, as this would imply $f_2 \neq t_2$ and thus an additional crossing of e_3, a contradiction to $(3P)$. Hence, f_1 does not have to distribute charge over the edge e_2 and hence has sufficient excess charge to distribute 0.5 units of charge to its immediate neighbor f. Finally, if f_1 is any larger face, it again has sufficient excess charge which can be distributed to f (the argument is analogous to the previous case).
2. f_4 is a 1-quadrangle. Assume first that $a_0 \notin f_4[e_0]$. Hence, a_3 is the unique vertex of f_4, see Fig. 4f. By (B), f_4' cannot be a 1-triangle. If f_4' is a 0-x face, then there is an edge e' that intersects e_4 and e_1 such that $f_0 \neq t_0$. But then both e_1 and e_4 have three crossings each, and hence the sequence (v_3, x_2, v_2) forms a subdivision face of f_2' which distributes 0.5 units of charge to its vertex-neighbor f, see Fig. 4f. If f_4' is not a 0-x face (in particular, not a 0-pentagon), then f_4 has sufficient charge which it distributes to its immediate neighbor f at e_4.
   Assume now that $a_3 \notin f_4[e_3]$. Hence, a_0 is the unique vertex of f_4. If $f_0 = t_0$, then the sequence (v_0, x_4, a_0) defines a subdivision face of f_4' which distributes its charge to its vertex-neighbor f at x_4. This situation is depicted in Fig. 4g when the dashed-gray edge does not exist. Otherwise, $f_0 \neq t_0$ and thus f_0 is a 0-quadrangle. This implies that $t_2 = f_2$ and $t_3 = f_3$ by $(3P)$. But then the sequence (v_3, x_2, v_2) defines a subdivision face of f_2' which distributes its excess charge to its vertex-neighbor f at x_2, again see Fig. 4g.
3. f_4 is a larger face. Again, f_4 contains sufficient charge which it can distribute to f.

f is a wedge-neighbor to four 1-triangles

W.l.o.g. we assume that f_4 is not a 1-triangle and that $v_0 \in A$, which implies that $v_1 \in A, v_2 \in B$ and $v_3 \in B$.

1. f_4 is a 0-x face. It follows that both e_0 and e_3 have three crossings each. Let us first consider the case where f_4 is a 0-quadrangle. Let e' be the edge of f_4 opposite to e_4. Further, denote by x (y) the intersection of e' with e_0 (e_3). Let f' be the immediate neighbor of f_4 at e'. By (B) and $(3P)$ we have $b_0 \in \delta f'$ and $a_3 \in \delta f'$. But then the sequence $f_a = (b_0, x, y, a_3)$ forms a subdivision of f' (note that f_a could coincide with f') which has an excess of at least one charge. Since it is not possible that both immediate neighbors of f_4 different from f' and f are 0-pentagons by $(3P)$, f' looses at most 0.5 charge over a vertex-neighbor relation – thus it can distribute 0.5 charge to f through f_4. In order to determine who supplies the remaining 0.5 charge, consider e_0. If e_0 has no crossing in $\delta f_0 \setminus f$, then the sequence (v_1, x_0, v_0) forms a subdivision which distributes its excess charge to f; otherwise, if e_0 has a crossing in $\delta f_0 \setminus f$, then by $(3P)$ e_0 has no crossing in $\delta f_3 \setminus f$ and hence (v_3, x_2, v_2) forms the desired subdivision, see Fig. 5a. Henceforth, we can thus assume that f_4 is not a 0-quadrangle.
   a. e_4 has no crossing outside of δf. Observe that this implies that also e_1 and e_2 do not have any additional crossings outside of δf, as an edge which is crossing e_1 or e_2 would necessarily also cross either e_0, e_3 or e_4. We can thus identify two pairs of consecutive vertices of the same partition, namely (v_0, v_1) and (v_2, v_3). Observe that $\delta f_0'$ ($\delta f_2'$) contains both v_0 and v_1 (v_2 and v_3). Hence, (v_1, x_0, v_0) and (v_3, x_2, v_2) each form a subdivision face which can distribute its excess charge to its vertex-neighbor f at x_0 and x_2, respectively, see Fig. 5b.

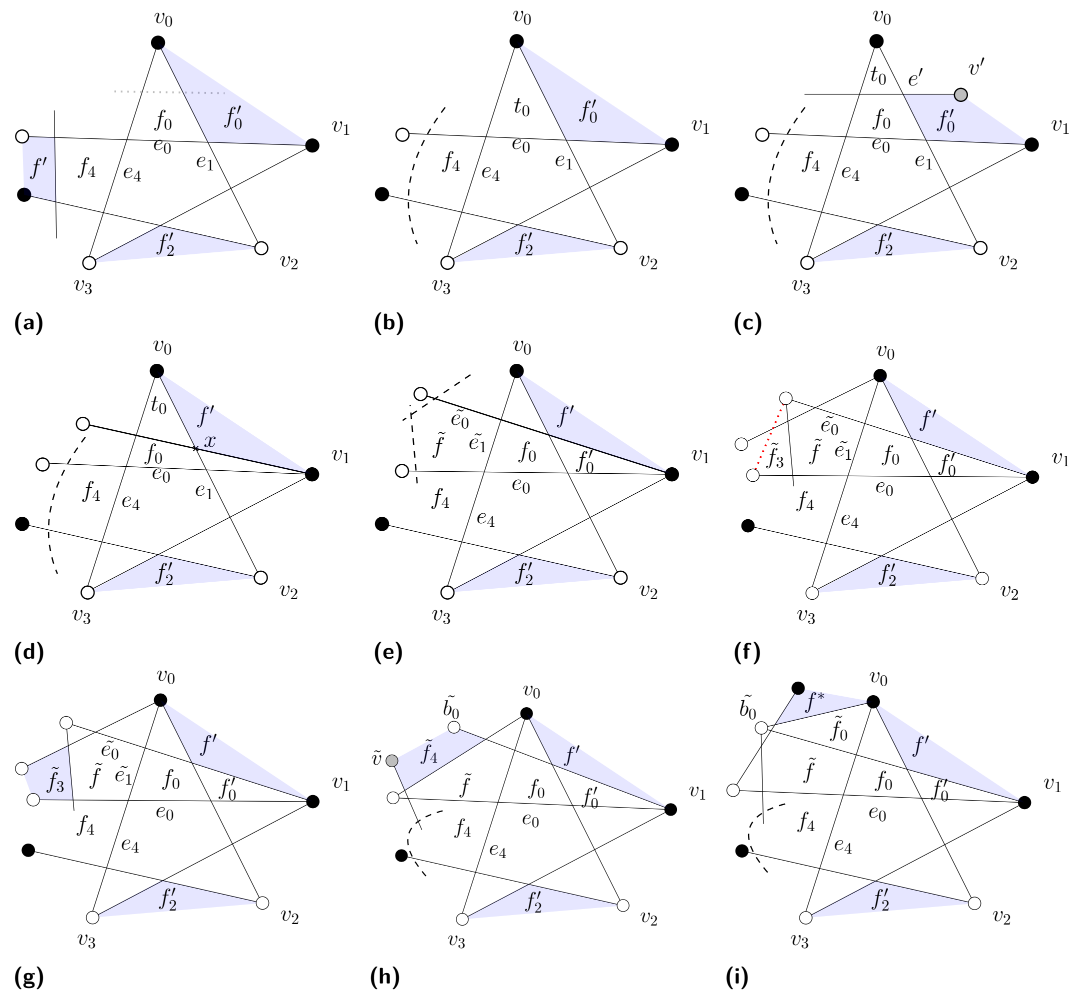

Figure 5 Illustrations used for the first part of the case analysis where face f is a 0-pentagons that is a wedge-neighbor to four 1-triangles.

b. e_4 has a crossing in $\delta f_0 \setminus \delta f$ with an edge $e' = (a', b')$, the case where e_4 has a crossing in $\delta f_3 \setminus f$ is symmetric. This implies that $v_0 \notin f_0[e_4]$ and thus $f_0 \neq t_0$. We can then again observe that $f_3 = t_3$ and $f_2 = t_2$ holds and thus (v_3, x_2, v_2) forms a subdivision face of f_2', which distributes its excess charge of 0.5 to its vertex-neighbor f at x_2. For the remaining 0.5 charge, we consider the following cases.

i. $a' \in f_0'[e']$ and $a' \neq v_1$ or $b' \in f_0'[e']$, see Fig. 5c. Let v' be the original vertex of $f_0'[e']$. Hence, by assumption $v' = b'$ or $v' = a'$, in both cases it holds that $v' \neq v_1$ (by (B) or by assumption). The sequence (v_1, x_0, x, v'), where x is the intersection point of e' and e_1, then defines a subdivision face f_a (in the case of $v' = b'$, it is possible that f_a and f_0' coincide). Observe that the vertex-neighbor of f_a at x is t_0, which does not require any charge. Note that f_a is potentially a wedge-neighbor via the edge e_1. However, since the initial charge of f_a is at least 2 (the extremal case occurs if $f_a = f_0'$ holds). In every case, f_a has an excess charge of at least one unit, it has sufficient charge such that it can distribute 0.5 to its wedge-neighbor and 0.5 to its vertex-neighbor f at x.

ii. $a' \in f_0'[e']$ and $a' = v_1$. Again, let x be the intersection between e' and e_1 and let f' be the vertex-neighbor of f_0 at x. Since $a' \in f_0'[e']$ and since e_1 already has three crossings, it follows that (v_1, x, v_0) forms a subdivision of f', see Fig. 5d, which charges 0.5 units to its vertex-neighbor f_0 at x. Before we propagate this charge from f_0 to f, we first have to consider f_4'. If f_4' is not a 0-pentagon, f_4' has sufficient charge even before the third step, unless f_4' is a 1-quadrangle. However, in the previous case analysis, we never required a 0-quadrangle to propagate the charge to a 1-quadrangle – hence, we can safely propagate 0.5 units of charge from f_0 to f. Hence, consider now the case that f_4' is a 0-pentagon. Note that this implies that f_4 cannot be a 0-triangle by $(3P)$. Since the following analysis is quite detailed, we set $\tilde{f} = f_4'$ and reuse the adjusted notation of f. In particular, we get $\tilde{e_0} = e', \tilde{e_1} = e_4$ and $\tilde{e_2} = e_0$, see Fig. 5e. Observe that $\tilde{t_1} = f_0'$ and since$\tilde{f_2} = f_4$ and since we covered the case where f_4 is a 0-quadrangle already, it follows that $\tilde{t_2}$ does not exist.

A. $\tilde{f}$ is wedge-neighbor to at most two 1-triangles. In this case, $\tilde{f}$ does not require any charge and hence f_0 can propagate 0.5 units to f.

B. $\tilde{f}$ is wedge-neighbor to three 1-triangles.

- $\tilde{t_3}$ is missing. Assume first that $\tilde{f_3}$ is a 1-quadrangle. By (B), it is not a wedge-neighbor at $\tilde{e_4}$ to a 1-triangle, see Fig. 5f. Moreover, the immediate face at $\tilde{e_4}$ is not a 0-x face by $(3P)$. Thus, $\tilde{f_3}$ does not loose charge over $\tilde{e_4}$ and we can therefore distribute the excess charge to $\tilde{f}$. If $\tilde{f_3}$ is not a 1-quadrangle, then the sequence $(\tilde{b_4}, \tilde{x_3}, \tilde{x_2}, \tilde{b_2})$ forms a subdivision of face $\tilde{f_3}$, see Fig. 5g. This subdivision has 1.5 units of excess charge and can therefore charge 0.5 to its immediate neighbor $\tilde{f}$ at $\tilde{e_3}$ as well as 0.5 to its vertex-neighbors at $\tilde{x_2}$ and $\tilde{x_3}$ (if required).
- $\tilde{t_4}$ is missing. By $(3P)$, face $\tilde{f_4}$ is not a 1-quadrangle as it contains an original vertex of $\tilde{e_3}$, which we denote by $\tilde{v}$ as well as $\tilde{b_0}$. The sequence $(\tilde{b_0}, \tilde{x_4}, \tilde{x_3}, \tilde{v})$ forms a subdivision face f_a of face $\tilde{f_4}$ (observe that $f_a = \tilde{f_4}$ is possible) – hence f_a has at least one unit of excess charge, see Fig. 5h. Since its vertex-neighbor at $\tilde{x_4}$ is a 1-triangle, it does not charge over $\tilde{x_4}$ (the same holds for $\tilde{x_3}$, but this is not necessary), and thus its excess charge is sufficient to distribute 0.5 units to its immediate neighbor $\tilde{f}$ at $\tilde{f_4}$.
- $\tilde{t_0}$ is missing. Assume first that $\tilde{f_0}$ is a 1-quadrangle. Note that in this case, it is possible that $\tilde{f_0}$ is a wedge-neighbor at $\tilde{e_4}$ to a 1-triangle t^*. If this is not the case, we proceed as in the first subcase, i.e., $\tilde{f_0}$ charges its excess of 0.5 units to its immediate neighbor $\tilde{f}$ at $\tilde{e_0}$. If $\tilde{f_0}$ is a wedge-neighbor at $\tilde{e_4}$, then we consider the immediate neighbor of $\tilde{f_0}$ at the edge $e^* = (\tilde{a_1}, \tilde{b_0})$, see Fig. 5i, which we denote by f^*. If we denote by x^* the intersection point between e^* and $\tilde{e_4}$, then the sequence $(\tilde{a_1}, x^*, \tilde{a_3})$ defines a subdivision face which distributes its excess charge to its vertex-neighbor t^* at x^*. Thus, $\tilde{f_0}$ can again distribute its excess charge to $\tilde{f}$ as desired. The case where $\tilde{f_0}$ is not a 1-quadrangle is analogous to the previous ones, i.e., we can identify a suitable subdivision face.

C. $\tilde{f}$ is wedge-neighbor to four 1-triangles. In this extremal case, we have the setting depicted in Fig. 6a. We can again observe that the sequence $(\tilde{b_0}, \tilde{x_3}, \tilde{b_2})$ forms a subdivision face of $\tilde{f_3'}$ which can distribute 0.5 units of charge to $\tilde{f}$. In this case, we will evenly split the charge that f_0 obtained earlier such that each of $\tilde{f}$ and f obtain 0.25 units. Note that both $\tilde{f}$ and f need an additional charge of 0.25 units. Consider face f_4. Recall that f_4 is a 0-x face by assumption, but

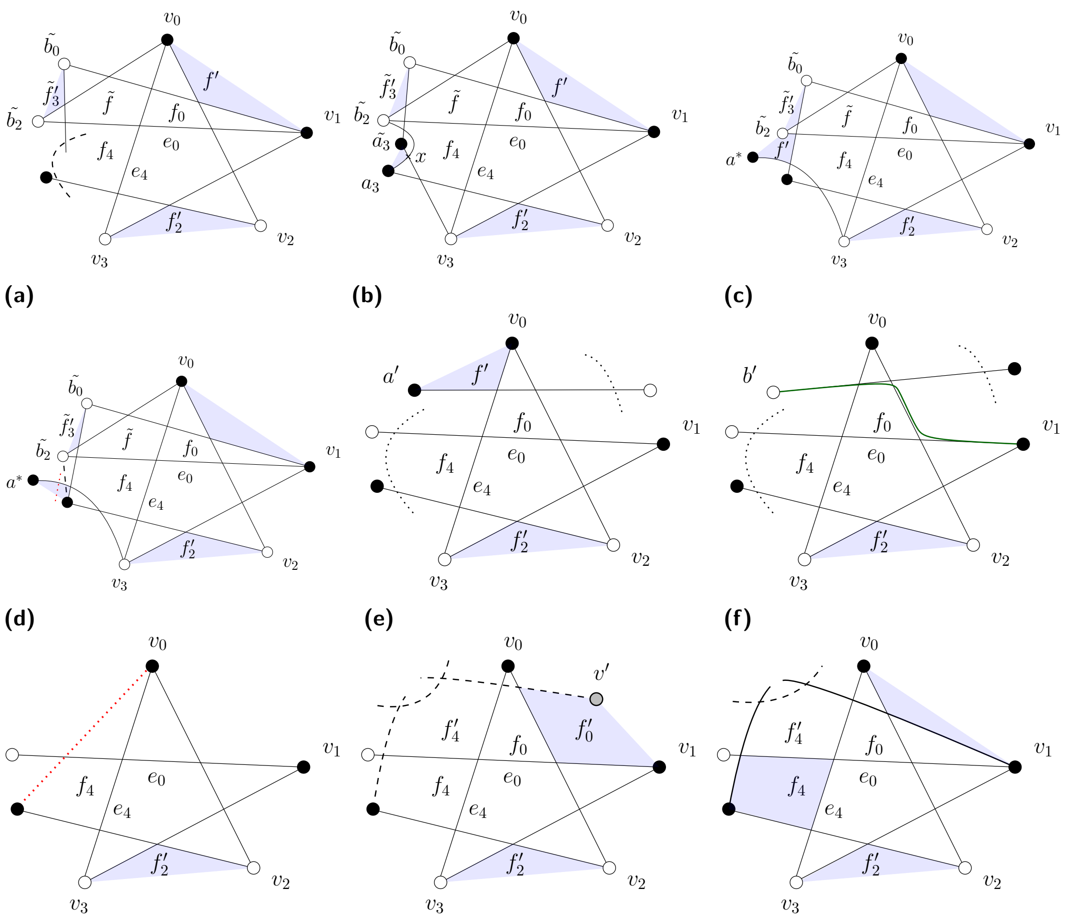

(a) **(b)** **(c)** **(d)** **(e)** **(f)** **(g)** **(h)** **(i)**

Figure 6 Illustrations used for the second part of the case analysis where face f is a 0-pentagons that is a wedge-neighbor to four 1-triangles.

it cannot be a 0-triangle as observed earlier. If f_4 has size at least seven, it has sufficient charge to distribute 0.25 to each of $\tilde{f}$ and f. To see this, let $x = |f_4|$ and observe that f_4 has an excess charge of $x - 4$. Since f_4 does not contain any original vertices, it cannot be a rich immediate neighbor nor a discharging vertex-neighbor by Inv. 14. Further, since f looses a combined charge of 0.5 over the edges e_0 and e_4, it follows that it discharges at most $0.5(x - 1) \leq x - 4$ for $x \geq 7$. If f_4 is a 0-6 and does not have sufficient charge, then we have exactly the setting depicted in Fig. 6b. But then the sequence $(\tilde{a_3}, x, a_3)$ is a subdivision face which distributes 0.5 charge to its vertex-neighbor f_4 at x, which can then be used to charge $\tilde{f}$ and f. Hence, assume that f_4 is a 0-5 face, see Fig. 6c for the extremal case where f_4 is an immediate neighbor to two 0-pentagons $\tilde{f}$ and f as well as a wedge-neighbor to two 1-triangles (the other possible arrangements of the two 1-triangles surrounding f_4 are symmetric). Let f' be the immediate neighbor of f_4 at the edge $\tilde{e_3}$. By $(3P)$ we have $\tilde{b_2} \in \delta f'$. Now, consider the edge $e^* = (a^*, b^*)$ that is incident to v_3 (i.e., $v_3 = b^*$) and intersects the edge $\tilde{e_3}$. If $a^* \in \delta f'$, then we can find a subdivision face, see Fig. 6c, which charges its immediate neighbor f_4 at $\tilde{e_3}$. Otherwise, if $a^* \notin \delta f'$, then e^* is intersected by an additional edge e''. If $e'' = (a_3, \tilde{b_2})$, then a_3 and a^* (together with an intersection point) again define a subdivision face, see Fig. 6d. We remark here that the local configuration which can be observed in Fig. 6d is dense, i.e., we do not have any excess charge left – this then yields, after an appropriate concatenation of the configurations, an alternative construction of bipartite 3-planar graphs on n vertices and $4n - \mathcal{O}(1)$ edges.

If $e'' \neq (a_3, \tilde{b_2})$, we claim that $(G, \Gamma) \notin \mathcal{G}''$. Indeed, by substituting the edge $\tilde{e_3}$ by the edge $(a_3, \tilde{b_2})$ (which will be drawn along the old curve of $\tilde{e_3}$ and e_0) and inheriting the remainder of the drawing, we can construct a graph G' and a corresponding 3-planar drawing Γ' with strictly less crossings. Observe that by our construction, any copy of $(a_3, \tilde{b_2})$, if it exists, is not homotopic to our new curve. Hence, we obtain a contradiction to our choice of (G, Γ).

iii. $a' \notin f_0'[e']$ and $b' \notin f_0'[e']$. Suppose first that, when traversing e' starting at a', we encounter its intersection with e_4 before its intersection with e_0, see Fig. 6e. Let x be the intersection between e' and e_4 and let f' be the vertex-neighbor of f_0 at x. The sequence (v_0, x, a') defines a subdivision face which charges 0.5 units to its vertex-neighbor f_0 at x. The crucial observation is that by $(3P)$, a' is necessarily part of f_4' – by a similar argument as in the previous case, f_4' has sufficient charge and thus f_0 can propagate 0.5 to f.

Suppose now that, when traversing e' starting at b', we encounter its intersection with e_4 before its intersection with e_0, see Fig. 6f. But then our choice of (G, Γ) is a contradiction to our assumption unless b' and b_0 coincide, as the drawing Γ' of $G \setminus e' \cup \{(b', v_1)\}$, which inherits the curve of all edges of Γ and adds the edge (b', v_1) such that its curve is drawn along e', e_1 and e_0, see the green edge in Fig. 6f, has strictly less crossings. Suppose therefore that b' and b_0 would coincide. But since f_4 is by assumption a 0-x face, we have that e_0 contains a crossing in $\delta f_4 \setminus f$ – but then by $(3P)$ we necessarily have $b_0 = b' \in \delta f_0'$, a contradiction to our assumption.

2. f_4 is a x-y face with $x \geq 1$. By assumption, f_4 is not a 1-triangle and by construction it cannot be a 2-triangle.

- f_4 is a quadrangle. Observe that it cannot be a wedge-neighbor via edges e_3 or e_0 by (B), see the red dotted edge in Fig. 6g for the case of e_0, the other is symmetric. W.l.o.g. we assume that $f_4[e_0]$ does not contain an original vertex, i.e., the unique original vertex of f_4 is an endpoint of e_3. Now, consider face f_4'. If f_4' is not a 0-pentagon, then f_4' does not loose any charge over e_0 and hence can distribute its excess charge to its immediate neighbor f at e_4. If f_4' is a 0-pentagon, we necessarily have that $f_0 \neq t_0$, see Fig 6h. But then e_1 and e_4 are crossed by an edge e' which has, by definition, an additional crossing in f_4'. Consequently, by $(3P)$, f_0' contains an endpoint of e'. Let v' be this endpoint and denote by x' the intersection of e' with e_1. If $v' \neq v_1$, we either have $f_0' = (v_1, x_0, x', v')$ or (v_1, x_0, x', v') forms a subdivision face of f_0', see Fig. 6h. In either case, f_0' has sufficient charge such that it can distribute 0.5 units to its vertex-neighbor f at x_0. Otherwise, v' and v_1 coincide, see Fig. 6i and we are in the same setting as in Case 1(b)ii with the difference that f_4 is a 1-quadrangle instead of a 0-x face. Again, we can observe that (v_1, x, v_0), where x is the intersection between e' and e_1, forms a subdivision face and it charges 0.5 units to its vertex-neighbor f_0 at x_0. Similar to the previous case, f_0 distributes 0.5 units to f unless f_4' is a wedge-neighbor to exactly four 1-triangles. If f_4' is a wedge-neighbor to four 1-triangles, then the remaining charge for f_4' and f (we will cover the missing 0.5 charge for f right after) is provided by f_4 – as it has 0.5 excess charge, it can distribute 0.25 to each of its wedge-neighbors, see Fig. 6i.

- f_4 is a larger face. Similar to previous cases, f_4 then always has sufficient excess charge to distribute 0.5 to its immediate neighbor f at edge e_4.

In order to determine who distributes the missing 0.5 charge to f, we consider face f_4. By assumption, f_4 contains at least one original vertex. Assume first that both b_0 and a_3 belong to δf_4. In this case f_4 is a 2-quadrangle (or a larger face) and has an excess charge of at least one. Since the vertex-neighbor at x_3 is a 1-triangle, it does not loose charge over x_3 – hence, it has sufficient charge to (potentially) distribute 0.5 via x_4 and to distribute 0.5 to its immediate neighbor f. Hence, assume that exactly one of b_0 or a_3 belongs to δf_4. Assume w.l.o.g. that a_3 belongs to δf_4 which implies that e_0 has and additional crossing in $\delta f_4 \setminus f$. Now, if $f_0 = t_0$ holds, then (v_1, x_0, v_0) defines a subdivision face which distributes its excess charge to its vertex-neighbor f at x_0. Otherwise, if $f_0 \neq t_0$, then e_1 and e_4 have an additional crossing. But then neither e_2 nor e_3 can have an additional crossing by $(3P)$, hence $t_2 = f_2$ and $t_3 = f_3$ and hence (v_3, x_2, v_2) forms a subdivision face which distributes 0.5 units of charge to its vertex-neighbor f at x_2.

After this final step, every face of F' satisfies Equation (1) which concludes the proof.

5 Implications

We will now use the main result of Section 4 to improve the lower-bound for the number of crossings, which consequently improves various other results. Note that, besides some numerical differences, the proof strategies for Sections 5.1 and 5.2 are identical to the ones of [3], while the proofs of Section 5.3 are identical to the ones of [17].

5.1 Crossing Lemma and Edge Density bounds

▶ **Theorem 15.** *Let G be a simple bipartite graph with $n \geq 3$ vertices and m edges. Then, the crossing number $cr(G)$ satisfies the following:*

$$cr(G) \geq 4m - \frac{25}{2}n + 27$$

Proof. The statements clearly holds when $m \leq 2n - 4$. Hence, we may assume w.l.o.g. that $m > 2n - 4$. It follows from [10] that if $m > 3n - 8$, then G has an edge that is crossed by at least two other edges. Also, by [3], we know that if $m > \frac{7}{2}n - 7$, then G has an edge that is crossed by at least three other edges. Finally, if $m > 4n - 8$, then Theorem 6 establishes that G has an edge that is crossed by at least four other edges. Hence we obtain by induction on the number of edges of G that the crossing number $cr(G)$ is at least:

$$cr(G) \geq (m-(2n-4))+(m-(3n-8))+(m-(\frac{7}{2}n-7))+(m-(4n-8)) = 4m - \frac{25}{2}n + 27$$

◀

▶ **Theorem 16.** *Let G be a simple bipartite graph with n vertices and m edges, where $m \geq \frac{75}{16}n$. Then, the crossing number $cr(G)$ satisfies the following:*

$$cr(G) \geq \frac{1024}{16875} \cdot \frac{m^3}{n^2} \approx \frac{1}{16.5} \frac{m^3}{n^2}$$

Proof. Assume that G admits a drawing on the plane with $cr(G)$ crossings and let $p = \frac{75n}{16m} \leq 1$. Choose independently every vertex of G with probability p, and denote by G_p the graph induced by the vertices chosen in G_p. Let also n_p, m_p and c_p be the random variables corresponding to the number of vertices, of edges and of crossings of G_p. Taking expectations on the relationship $c_p \geq 4m_p - \frac{25}{2}n_p + 27$, which holds by Theorem 15, we obtain:

$$p^4 cr(G) \geq 4p^2 m - \frac{25}{2}np \quad \Rightarrow \quad cr(G) \geq \frac{4m}{p^2} - \frac{25n}{2p^3}$$

The proof of the theorem follows by plugging $p = \frac{75n}{16m}$ (which is at most 1 by our assumption on m) to the inequality above. ◀

▶ **Theorem 17.** *Let G be a simple bipartite k-planar graph with n vertices and m edges, for some $k \geq 2$. Then:*

$$m \leq \sqrt{\frac{16875}{2048}kn} \approx 2.871\sqrt{kn}$$

Proof. For $k = 2$ and $k = 3$, the bounds of this theorem are weaker than the corresponding ones of [3], and of Theorem 6, respectively. So, we may assume w.l.o.g. that $k > 3$. We may also assume that $m \geq \frac{75}{16}n$, as otherwise there is nothing to prove. Combining the fact that G is k-planar with the bound of Theorem 16 we obtain:

$$\frac{1024}{16875} \cdot \frac{m^3}{n^2} \leq cr(G) \leq \frac{1}{2}mk$$

which implies:

$$m \leq \sqrt{\frac{16875}{2048}kn} \approx 2.871\sqrt{kn}$$ ◀

▶ **Theorem 18.** *Let G be a simple bipartite k-gap-planar graph with n vertices and $m \geq \frac{75}{16}n$ edges. Then:*

$$m \leq \sqrt{\frac{16875}{1024}}\sqrt{kn} \approx 4.06\sqrt{kn}$$

Proof. By definition, we have that

$$cr(G) \leq k \cdot m$$

for any k-gap-planar graph G with m edges. On the other hand, Theorem 16 gives us

$$cr(G) \geq \frac{1024}{16875}\frac{m^3}{n^2}$$

since G is bipartite. Thus

$$\frac{1024}{16875}\frac{m^3}{n^2} \leq cr(G) \leq k \cdot m$$

and the result follows. ◀

5.2 Exclusion of complete bipartite graphs

▶ **Theorem 19.** *Let $K_{n,m}$ be a complete bipartite graph and let $n \leq m$. Then, K_{5,m_1} with $m_1 \geq 13$, K_{6,m_2} with $m_2 \geq 9$ and K_{7,m_3} with $m_3 \geq 7$ are not 3-planar and not gap-planar.*

Proof. $K_{5,13}$ has $5 \cdot 13 = 65$ edges, but any bipartite 3-planar (gap-planar) graph on 18 vertices has at most $4 \cdot 18 - 8 = 64$ edges, a contradiction. Similarly, we have that $K_{6,9}$ has $54 > 4 \cdot 15 - 8 = 52$ and $K_{7,7}$ has $49 > 4 \cdot 14 - 8 = 48$ edges. ◀

5.3 Biplanar crossing number

The biplanar (k-planar) crossing number of a graph G, denoted by $cr_2(G)$ ($cr_k(G)$), is the minimum number of crossings over all possible drawings of the edges of G in two (k) disjoint planes.

▶ **Theorem 20.** *Let $K_{p,q}$ be a complete bipartite graph with $p, q \geq 30$. Then*

$$cr_2(K_{p,q}) \geq \frac{p(p-1)q(q-1)}{204}$$

Proof. Applying [17, Lemma 1] together with Lemma 15 yields $cr_2(G) \geq 4m - (\frac{25}{2}n - 27) \cdot 2$ and thus $cr_2(K_{17,17}) \geq 360$. Using the recurrence relation

$$cr_2(K_{p+1,p+1}) \geq \left\lceil \frac{p+1}{p-1} \left\lceil \frac{p+1}{p-1} cr_2(K_{p,q}) \right\rceil \right\rceil$$

repeatedly as in [17], we obtain $cr_2(K_{30,30}) \geq 3723$ and thus

$$cr_2(K_{p,q}) \geq \frac{p(p-1)q(q-1)}{30 \times 29 \times 30 \times 29} cr_2(K_{30,30})$$

which yields the desired result. ◀

▶ **Theorem 21.** *For all $p, q \geq 9k + 2$,*

$$cr_k(K_{p,q}) \geq \frac{p(p-1)q(q-1)}{66.3k^2}$$

Proof. Using Lemma 15, we obtain

$$cr_k(K_{9k+2,9k+2}) \geq 99k^2 + 121k + 16$$

Following the proof of Theorem 7 in [17] we then obtain the desired result. ◀

6 Conclusions and Open Problems

We have established tight upper bounds on the number of edges of bipartite gap-planar and bipartite 3-planar graphs. The following questions follow naturally:

- What is the density of bipartite k-planar graphs, in particular for $k = 4$? One could most likely apply the discharging method in a similar way for any fixed k – the issue that arises for larger k is just the sheer number of cases one has to consider. Hence, we ask as an open problem if one can (partially) automate such a charging proof in a similar way to [4]. This is of course also an interesting question in the normal (non-bipartite) setting.
- A graph is quasi-planar if there is a drawing in which no three edges mutually cross. What is the edge density of bipartite quasi-planar graphs?

References

1 Eyal Ackerman. On topological graphs with at most four crossings per edge. *Comput. Geom.*, 85:101574, 2019. doi:10.1016/j.comgeo.2019.101574.

2 Eyal Ackerman and Gábor Tardos. On the maximum number of edges in quasi-planar graphs. *J. Comb. Theory A*, 114(3):563–571, 2007. doi:10.1016/j.jcta.2006.08.002.

3 Patrizio Angelini, Michael A. Bekos, Michael Kaufmann, Maximilian Pfister, and Torsten Ueckerdt. Beyond-planarity: Turán-type results for non-planar bipartite graphs. In *ISAAC*, volume 123 of *LIPIcs*, pages 28:1–28:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.ISAAC.2018.28.
4 K. Appel and W. Haken. Every planar map is four colorable. Part I: Discharging. *Illinois Journal of Mathematics*, 21(3):429–490, 1977. doi:10.1215/ijm/1256049011.
5 Sang Won Bae, Jean-François Baffier, Jinhee Chun, Peter Eades, Kord Eickmeyer, Luca Grilli, Seok-Hee Hong, Matias Korman, Fabrizio Montecchiani, Ignaz Rutter, and Csaba D. Tóth. Gap-planar graphs. *Theor. Comput. Sci.*, 745:36–52, 2018. doi:10.1016/j.tcs.2018.05.029.
6 Steven Chaplick, Myroslav Kryven, Giuseppe Liotta, Andre Löffler, and Alexander Wolff. Beyond outerplanarity. In *GD*, volume 10692 of *Lecture Notes in Computer Science*, pages 546–559. Springer, 2017. doi:10.1007/978-3-319-73915-1_42.
7 Otfried Cheong, Sariel Har-Peled, Heuna Kim, and Hyo-Sil Kim. On the number of edges of fan-crossing free graphs. *Algorithmica*, 73(4):673–695, December 2015. doi:10.1007/s00453-014-9935-z.
8 Walter Didimo, Giuseppe Liotta, and Fabrizio Montecchiani. A survey on graph drawing beyond planarity. *ACM Comput. Surv.*, 52(1), February 2019. doi:10.1145/3301281.
9 Weidong Huang, Peter Eades, and Seok-Hee Hong. Larger crossing angles make graphs easier to read. *J. Vis. Lang. Comput.*, 25(4):452–465, 2014. doi:10.1016/j.jvlc.2014.03.001.
10 Dmitri V Karpov. An upper bound on the number of edges in an almost planar bipartite graph. *Journal of Mathematical Sciences*, 196:737–746, 2014.
11 Michael Kaufmann and Torsten Ueckerdt. The density of fan-planar graphs. *Electron. J. Comb.*, 29(1):P1–29, 2022. doi:10.37236/10521.
12 Petra Mutzel. An alternative method to crossing minimization on hierarchical graphs. In Stephen C. North, editor, *Graph Drawing, Symposium on Graph Drawing, GD '96, Berkeley, California, USA, September 18-20, Proceedings*, volume 1190 of *Lecture Notes in Computer Science*, pages 318–333, Berlin, Heidelberg, 1996. Springer. doi:10.1007/3-540-62495-3_57.
13 János Pach, Rados Radoicic, Gábor Tardos, and Géza Tóth. Improving the crossing lemma by finding more crossings in sparse graphs. *Discret. Comput. Geom.*, 36(4):527–552, 2006. doi:10.1007/s00454-006-1264-9.
14 János Pach and Géza Tóth. Graphs drawn with few crossings per edge. *Combinatorica*, 17(3):427–439, 1997. doi:10.1007/BF01215922.
15 Gerhard Ringel. Ein Sechsfarbenproblem auf der Kugel. *Abh. Math. Sem. Univ. Hamb.*, 29:107–117, 1965.
16 Thomas Schneck. *New Parameters for Beyond-Planar Graphs*. dissertation, University Tübingen, 2020. URL: https://publikationen.uni-tuebingen.de/xmlui/handle/10900/109680.
17 Alireza Shavali and Hamid Zarrabi-Zadeh. On the biplanar and k-planar crossing numbers. In *CCCG*, pages 293–297, 2022.
18 Colin Ware, Helen Purchase, Linda Colpoys, and Matthew McGill. Cognitive measurements of graph aesthetics. *Information Visualization*, 1:103–110, June 2002. doi:10.1057/palgrave.ivs.9500013.