Towards Better Approximation of Graph Crossing Number

Julia Chuzhoy
Toyota Technological Institute at Chicago
Email: cjulia@ttic.edu.

Sepideh Mahabadi
Toyota Technological Institute at Chicago
Email: mahabadi@ttic.edu.

Zihan Tan
University of Chicago
Email: zihantan@uchicago.edu

***Abstract*—Graph Crossing Number is a fundamental and extensively studied problem with wide ranging applications. In this problem, the goal is to draw an input graph G in the plane so as to minimize the number of crossings between the images of its edges. The problem is notoriously difficult, and despite extensive work, non-trivial approximation algorithms are only known for bounded-degree graphs. Even for this special case, the best current algorithm achieves a $\tilde{O}(\sqrt{n})$-approximation, while the best current negative results do not rule out constant-factor approximation. All current approximation algorithms for the problem build on the same paradigm, which is also used in practice: compute a set E' of edges (called a *planarizing set*) such that $G \setminus E'$ is planar; compute a planar drawing of $G \setminus E'$; then add the drawings of the edges of E' to the resulting drawing. Unfortunately, there are examples of graphs G, in which any implementation of this method must incur $\Omega(\mathsf{OPT}^2)$ crossings, where OPT is the value of the optimal solution. This barrier seems to doom the only currently known approach to designing approximation algorithms for the problem, and to prevent it from yielding a better than $O(\sqrt{n})$-approximation.**

In this paper we propose a new paradigm that allows us to overcome this barrier. We show an algorithm, that, given a bounded-degree graph G and a planarizing set E' of its edges, computes another planarizing edge set E'' with $E' \subseteq E''$, such that $|E''|$ is relatively small, and there exists a near-optimal drawing of G in which no edges of $G \setminus E''$ participate in crossings. This allows us to reduce the Crossing Number problem to *Crossing Number with Rotation System* – a variant of the Crossing Number problem, in which the ordering of the edges incident to every vertex is fixed as part of input. In our reduction, we obtain an instance G' of this problem, where $|E(G')|$ is roughly bounded by the crossing number of the original graph G. We show a randomized algorithm for this new problem, that allows us to obtain an $O(n^{1/2-\epsilon})$-approximation for Graph Crossing Number on bounded-degree graphs, for some constant $\epsilon > 0$.

***Keywords*-approximation algorithm; crossing number**

I. Introduction

Given a graph $G = (V, E)$, a *drawing* of G is an embedding of the graph into the plane, that maps every vertex to a point in the plane, and every edge to a continuous curve connecting the images of its endpoints. We require that the image of an edge may not contain images of vertices of G other than its two endpoints, and no three curves representing drawings of edges of G intersect at a single point, unless that point is the image of their shared endpoint. We say that two edges *cross* in a given drawing of G iff their images share a point p other than the image of their common endpoint; such a point p is called a *crossing*. In the Minimum Crossing Number problem, the goal is to compute a drawing of the input n-vertex graph G with minimum number of crossings. We denote the value of the optimal solution to this problem, also called the *crossing number* of G, by $\mathsf{OPT}_{\mathsf{cr}}(G)$.

The Minimum Crossing Number problem naturally arises in several areas of Computer Science and Mathematics. The problem was initially introduced by Turán [Tur77], who considered the question of computing the crossing number of complete bipartite graphs, and since then it has been a subject of extensive studies (see, e.g., [Tur77], [Chu11], [CMS11], [CH11], [CS13], [KS17], [KS19], and also [RS09], [PT00], [Mat02], [Vrt] for surveys on this topic). The problem is known to be NP-hard [GJ83], and it remains NP-hard, and APX-hard, even on cubic graphs [Hli06], [Cab13]. The Minimum Crossing Number problem appears to be notoriously difficult from the approximation perspective. All currently known algorithms achieve approximation factors that depend polynomially on Δ – the maximum vertex degree of the input graph, and, to the best of our knowlgedge, no non-trivial approximation algorithms are known for graphs with arbitrary vertex degrees. We note that the famous Crossing Number Inequality [ACNS82], [Lei83] shows that, for every graph G with $|E(G)| \geq 4n$, the crossing number of G is $\Omega(|E(G)|^3/n^2)$. Since the problem is most interesting when the crossing number of the input graph is low, and since our understanding of the problem is still extremely poor, it is reasonable to focus on designing algorithms for low-degree graphs. Throughout, we denote by Δ the maximum vertex degree of the input graph. While we do not make this assumption explicitly, it may be convenient to think of Δ as being bounded by a constant or by $\operatorname{poly}\log n$.

The first non-trivial algorithm for Minimum Crossing Number, due to Leighton and Rao [LR99], combined their algorithm for balanced separators with the framework of [BL84], to compute a drawing of input graph G with $O((n + \mathsf{OPT}_{\mathsf{cr}}(G)) \cdot \Delta^{O(1)} \cdot \log^4 n)$ crossings. This bound was later improved to $O((n+\mathsf{OPT}_{\mathsf{cr}}(G)) \cdot \Delta^{O(1)} \cdot \log^3 n)$ by [EGS02], and then to $O((n+\mathsf{OPT}_{\mathsf{cr}}(G)) \cdot \Delta^{O(1)} \cdot \log^2 n)$ as a consequence of the improved algorithm of [ARV09] for Balanced Cut. All these results provide an $O(n \operatorname{poly}(\Delta \log n))$-approximation for Minimum Crossing Number (but perform much better when $\mathsf{OPT}_{\mathsf{cr}}(G)$ is high). For a long time, this remained the best approximation algorithm for Minimum Crossing Number, while the best inapproximability result, to this day, only rules out the existence of a PTAS, unless P=NP [GJ83], [AMS07], [Cab13]. We note that it is highly unusual that achieving an $O(n)$-approximation for an unweighted graph optimization problem is so challenging. However, unlike many other unweighted graph optimization problems, the value of the optimal solution to Minimum Crossing Number may be as large[1] as $\Omega(n^4)$.

A sequence of papers [CMS11], [Chu11] was the first to break the barrier of $\tilde{\Theta}(n)$-approximation, providing a $\tilde{O}\left(n^{9/10} \cdot \Delta^{O(1)}\right)$-approximation algorithm. Recently, a breakthrough sequence of works [KS17], [KS19] has led to an improved $\tilde{O}\left(\sqrt{n} \cdot \Delta^{O(1)}\right)$-approximation for Minimum Crossing Number. All the above-mentioned algorithms exploit the same algorithmic paradigm, that builds on the connection of Minimum Crossing Number to the Minimum Planarization problem, that we discuss next.

Minimum Planarization. In the Minimum Planarization problem, the input is an n-vertex graph $G = (V, E)$, and the goal is to compute a minimum-cardinality subset E^* of its edges (called a *planarizing set*), such that graph $G \setminus E^*$ is planar. This problem and its close variant Minimum Vertex Planarization (where we need to compute a minimum-cardinality subset V' of vertices such that $G \setminus V'$ is planar) are of independent interest and have been studied extensively (see, e.g., [CMS11], [KS17], [KS19]). It is immediate to see that, for every graph G, $\mathsf{OPT}_{\mathsf{mvp}}(G) \leq \mathsf{OPT}_{\mathsf{mep}}(G) \leq \mathsf{OPT}_{\mathsf{cr}}(G)$, where $\mathsf{OPT}_{\mathsf{mvp}}(G)$ and $\mathsf{OPT}_{\mathsf{mep}}(G)$ are the optimal solution values of the Minimum Vertex Planarization and the Minimum Planarization problems on G, respectively. A simple application of the Planar Separator Theorem of [LT79] was shown to give an $O(\sqrt{n \log n} \cdot \Delta)$-approximation algorithm for both problems [CMS11]. Further, [CS13] provided an $O(k^{15} \cdot \mathrm{poly}(\Delta \log n)$-approximation algorithm for Minimum Planarization and Minimum Vertex Planarization, where k is the value of the optimal solution. The more recent breakthrough result of Kawarabayashi and Sidiropoulus [KS17], [KS19] provides an $O(\Delta^3 \cdot \log^{3.5} n)$-approximation algorithm for Minimum Vertex Planarization, and an $O(\Delta^4 \cdot \log^{3.5} n)$-approximation algorithm for Minimum Planarization.

Returning to the Minimum Crossing Number problem, all currently known approximation algorithms for the problem rely on the same paradigm, which is also used in heuristics (see, e.g. [BCG$^+$13]). For convenience of notation, we call it Paradigm Π.

[1]This can be seen by applying the Crossing Number Inequality to the complete n-vertex graph.

PARADIGM Π

1) compute a planarizing set E' of edges for G;
2) compute a planar drawing of $G \setminus E'$;
3) add the images of the edges of E' to the resulting drawing.

We note that graph $G \setminus E'$ may not be 3-connected and thus it may have several planar drawings; there are also many ways in which the edges of E' can be added to the drawing. It is therefore important to understand the following questions:

Can this paradigm be implemented in a way that provides a near-optimal drawing of G? What is the best approximation factor that can be achieved when using paradigm Π?

These questions were partially answered in previous work. Specifically, [CMS11] provided an efficient algorithm, that, given an input graph G, and a planarizing set E' of k edges for G, draws the graph with $O\left(\Delta^3 \cdot k \cdot (\mathsf{OPT}_{\mathsf{cr}}(G) + k)\right)$ crossings. Later, Chimani and Hliněný [CH11] improved this bound via a different efficient algorithm to $O\left(\Delta \cdot k \cdot (\mathsf{OPT}_{\mathsf{cr}}(G) + \log k) + k^2\right)$. Both works can be viewed as an implementation of the above paradigm. Combining these results with the $O(\mathrm{poly}(\Delta \log n))$-approximation algorithm for Minimum Planarization of [KS19] in order to compute the initial planarizing edge set E' with $|E'| \leq O(\mathsf{OPT}_{\mathsf{cr}}(G) \cdot \mathrm{poly}(\Delta \log n))$, yields an implementation of Paradigm Π that produces a drawing of the input graph G with $O\left((\mathsf{OPT}_{\mathsf{cr}}(G))^2 \cdot \mathrm{poly}(\Delta \log n)\right)$ crossings. Lastly, combining this with the $O(n\, \mathrm{poly}(\Delta \log n))$-approximation algorithm of [LR99], [BL84], [EGS02] leads to the best currently known $O(\sqrt{n} \cdot \mathrm{poly}(\Delta \log n))$-approximation algorithm of [KS19] for Minimum Crossing Number.

The bottleneck in using this approach in order to obtain a better than $O(\sqrt{n})$-approximation for Minimum Crossing Number is the bounds of [CMS11] and [CH11], whose algorithms produce a drawing of the graph G with $O\left(k \cdot \mathsf{OPT}_{\mathsf{cr}}(G) + k^2\right)$ crossings when $\Delta = O(1)$, where k is the size of the given planarizing set. The quadratic dependence of this bound on k and the linear dependence on $k \cdot \mathsf{OPT}_{\mathsf{cr}}(G)$ are unacceptable if our goal is to obtain better approximation using this technique. A natural question therefore is:

Can we obtain a stronger bound, that is near-linear in $(\mathsf{OPT}_{\mathsf{cr}}(G) + |E'|)$, using Paradigm Π?

Unfortunately, Chuzhoy, Madan and Mahabadi [CMM16] have answered the latter question in the negative, by showing that the bounds of [CMS11] and [CH11] are almost tight; see the full version of this paper for details. This seems to doom the only currently known approach for designing approximation algorithms for the problem.

In this paper, we propose a new paradigm towards overcoming this barrier, and show that it leads to a better approximation of Minimum Crossing Number. Specifically, we show an efficient algorithm, that, given a planarizing set E' of edges, augments E' in order to obtain a new planarizing set E'', whose cardinality is $O\left((|E'|+\mathsf{OPT}_{\mathsf{cr}}(G))\operatorname{poly}(\Delta\log n)\right)$. Moreover, we show that there exists a drawing φ of the graph G, with at most $O\left((|E'|+\mathsf{OPT}_{\mathsf{cr}}(G))\operatorname{poly}(\Delta\log n)\right)$ crossings, where the edges of $G\setminus E''$ do not participate in any crossings. In other words, the drawing φ of G can be obtained by first computing a planar drawing of $G\setminus E''$, and then inserting the images of the edges of E'' into this drawing. This new paradigm can be summarized as follows:

PARADIGM Π'

1) compute a planarizing set E' of edges for G;
2) compute an augmented planarizing edge set E'' with $E'\subseteq E''$ that has some additional useful properties;
3) compute a planar drawing of $G\setminus E''$;
4) add the images of the edges in E'' to the resulting drawing.

Our result, combined with the $O(\operatorname{poly}(\Delta\log n))$-approximation algorithm for Minimum Planarization of [KS19], provides an efficient implementation of Steps (1) and (2) of Paradigm Π', such that there exists an implementation of Steps (3) and (4), that produces a drawing of G with $O(\mathsf{OPT}_{\mathsf{cr}}(G)\cdot\operatorname{poly}(\Delta\log n))$ crossings. This still leaves open the following question:

Can we implement Steps (3) and (4) of Paradigm Π' efficiently in near-optimal fashion?

One way to address this question is by designing algorithms for the following problem: given a graph G and a planarizing set E^* of its edges, compute a drawing of G, such that the induced drawing of $G\setminus E^*$ is planar (in other words, every crossing in the drawing involves an edge of E^*), while minimizing the number of crossings in the resulting drawing. This problem was considered by Chimani and Hliněný [CH11] who showed an efficient algorithm, that computes a drawing of G with $\mathsf{OPT}_{\mathsf{cr}}^{E^*}(G)+O(\Delta k\log k+k^2)$ crossings, where $k=|E^*|$, and $\mathsf{OPT}_{\mathsf{cr}}^{E^*}(G)$ is the optimal solution value for this problem. Unfortunately, if our goal is to break the $\Theta(\sqrt{n})$ barrier on the approximation factor for Minimum Crossing Number, the quadratic dependence of this bound on k is prohibitive.

We propose a different approach in order to implement Steps (3) and (4) of Paradigm Π'. We provide an efficient algorithm that exploits our algorithm for Steps (1) and (2) in order to reduce the Minimum Crossing Number problem to another problem, called Minimum Crossing Number with Rotation System (MCNwRS). In this problem, the input is a multi-graph G with arbitrary vertex degrees. Additionally, for every vertex v of G, we are given a circular ordering $\mathcal{O}_v$ of the edges that are incident to v in G. The goal is to compute a drawing of G in the plane with minimum number of crossings, that respects the orderings $\{\mathcal{O}_v\}_{v\in V(G)}$ of the edges incident to each vertex (but we may choose whether the ordering is clockwise or counter-clock-wise in the drawing). We denote $\Sigma=\{\mathcal{O}_v\}_{v\in V}$, and we call Σ a *rotation system for graph* G. Given an instance (G,Σ) of MCNwRS, we denote by $\mathsf{OPT}_{\mathsf{cnwrs}}(G,\Sigma)$ the value of the optimal solution for this instance. We show a reduction, that, given an instance G of Minimum Crossing Number with maximum vertex degree Δ, produces an instance (G',Σ) of MCNwRS, such that $|E(G')|\leq O\left(\mathsf{OPT}_{\mathsf{cr}}(G)\cdot\operatorname{poly}(\Delta\log n)\right)$, and moreover, $\mathsf{OPT}_{\mathsf{cnwrs}}(G',\Sigma)\leq O\left(\mathsf{OPT}_{\mathsf{cr}}(G)\cdot\operatorname{poly}(\Delta\log n)\right)$. In particular, our reduction shows that, in order to obtain an $O(\alpha\operatorname{poly}(\Delta\log n))$-approximation for Minimum Crossing Number, it is sufficient to obtain an α-approximation algorithm for MCNwRS. We also show an efficient randomized algorithm, that, given an instance (G,Σ) of MCNwRS, with high probability computes a solution to the problem with at most $\tilde{O}\left((\mathsf{OPT}_{\mathsf{cnwrs}}(G,\Sigma)+|E(G)|)^{2-\epsilon}\right)$ crossings, for $\epsilon=1/20$. Combining this result with our reduction, we obtain an efficient algorithm that computes a drawing of an input graph G with maximum vertex degree Δ with at most $\tilde{O}((\mathsf{OPT}_{\mathsf{cr}}(G))^{2-\epsilon}\cdot\operatorname{poly}(\Delta\log n))$ crossings. We note that this algorithm can be viewed as implementing Steps (3) and (4) of Paradigm Π'. The resulting algorithm, in turn, leads to a $\tilde{O}(n^{1/2-\epsilon'}\operatorname{poly}(\Delta))$-approximation algorithm for Minimum Crossing Number, for some fixed constant $\epsilon'>0$. While this only provides a modest improvement in the approximation factor, we view this as a proof of concept that our new method can lead to improved approximation algorithms, and in particular, this result breaks the barrier that the previously known methods could not overcome.

Other related work. There has been a large body of work on FPT algorithms for several variants and special cases of the Minimum Crossing Number problem (see, e.g. [Gro04], [KR07], [PSŠ07], [BE14], [HD15], [DLM19]). In particular, Grohe [Gro04] obtained an algorithm for solving the problem exactly, in time $f(\mathsf{OPT}_{\mathsf{cr}}(G))\cdot n^2$, where function f is at least doubly exponential. In his paper, he conjectures that there exists an FPT algorithm for the problem with running time $2^{O(\mathsf{OPT}_{\mathsf{cr}}(G))}\cdot n$. The dependency on n in the algorithm of [Gro04] was later improved by [KR07] from n^2 to n.

For cubic graphs (3-regular graphs), the MCNwRS problem is equivalent to the Minimum Crossing Number problem. Hliněný [Hli06] proved that Minimum Crossing Number is NP-hard for cubic graphs, and Cabello [Cab13] proved that Minimum Crossing Number is APX-hard for cubic graphs. Therefore, the MCNwRS problem is also APX-hard on cubic graphs. Pelsmajer et al. [PSŠ11] studied a variant of the MCNwRS problem, where for each vertex, the orientation of the ordering $\mathcal{O}_v$ (clockwise or counterclockwise) is also fixed. They showed that this variant is also NP-hard, and provided approximation algorithms for some special cases. Additionally, they provide an $O(n^4)$-approximation algorithm for this variant with running time $O(m^n \log m)$, where m is the number of edges in the graph[2]. We now proceed to describe our results more formally.

A. Our Results

Our main technical result is summarized in the following theorem.

Theorem 1 *There is an efficient algorithm, that, given an n-vertex graph G with maximum vertex degree Δ and a planarizing set E' of its edges, computes another planarizing edge set E'' for G, with $E' \subseteq E''$, such that $|E''| \leq O\left((|E'| + \mathsf{OPT}_{\mathsf{cr}}(G)) \cdot \mathrm{poly}(\Delta \log n)\right)$, and, moreover, there is a drawing φ of G with at most $O\left((|E'| + \mathsf{OPT}_{\mathsf{cr}}(G)) \cdot \mathrm{poly}(\Delta \log n)\right)$ crossings, such that the edges of $G \setminus E''$ do not participate in any crossings in φ.*

Recall that Kawarabayashi and Sidiropoulos [KS19] provide an efficient $O(\mathrm{poly}(\Delta \log n))$-approximation algorithm for the Minimum Planarization problem. Since, for every graph G, there is a planarizing set E^* containing at most $\mathsf{OPT}_{\mathsf{cr}}(G)$ edges, we can use their algorithm in order to compute, for an input graph G, a planarizing edge set of cardinality $O(\mathsf{OPT}_{\mathsf{cr}}(G) \cdot \mathrm{poly}(\Delta \log n))$. Combining this with Theorem 1, we obtain the following immediate corollary.

Corollary 2 *There is an efficient algorithm, that, given an n-vertex graph G with maximum vertex degree Δ, computes a planarizing set $E' \subseteq E(G)$ of $O(\mathsf{OPT}_{\mathsf{cr}}(G) \cdot \mathrm{poly}(\Delta \log n))$ edges, such that there is a drawing φ of G with $O(\mathsf{OPT}_{\mathsf{cr}}(G) \cdot \mathrm{poly}(\Delta \log n))$ crossings, and the edges of $E(G) \setminus E'$ do not participate in any crossings in φ.*

Next, we show a reduction from Minimum Crossing Number to the MCNwRS problem.

Theorem 3 *There is an efficient algorithm, that, given an n-vertex graph G with maximum vertex degree Δ, computes an instance (G', Σ) of the MCNwRS problem, such that the number of edges in G' is at most $O\left(\mathsf{OPT}_{\mathsf{cr}}(G) \cdot \mathrm{poly}(\Delta \log n)\right)$, and $\mathsf{OPT}_{\mathsf{cnwrs}}(G', \Sigma) \leq O\left(\mathsf{OPT}_{\mathsf{cr}}(G) \cdot \mathrm{poly}(\Delta \log n)\right)$. Moreover, there is an efficient algorithm that, given any solution to instance (G', Σ) of MCNwRS of value X, computes a drawing of G with at most $O\left((X + \mathsf{OPT}_{\mathsf{cr}}(G)) \cdot \mathrm{poly}(\Delta \log n)\right)$ crossings.*

Notice that the above theorem shows that an α-approximation algorithm for MCNwRS immediately gives an $O(\alpha\, \mathrm{poly}(\Delta \log n))$-approximation algorithm for the Minimum Crossing Number problem. In fact, even much weaker guarantees for MCNwRS suffice: if there is an algorithm that, given an instance (G, Σ) of MCNwRS, computes a solution of value at most $\alpha(\mathsf{OPT}_{\mathsf{cnwrs}}(G, \Sigma) + |E(G)|)$, then there is an $O(\alpha\, \mathrm{poly}(\Delta \log n))$-approximation algorithm for Minimum Crossing Number. Recall that [LR99], [EGS02] provide an algorithm for the Minimum Crossing Number problem that draws an input graph G with $\tilde{O}((n + \mathsf{OPT}_{\mathsf{cr}}(G)) \cdot \Delta^{O(1)})$ crossings. While it is conceivable that this algorithm can be adapted to the MCNwRS problem, it only gives meaningful guarantees when the maximum vertex degree in G is low, while in the instances of MCNwRS produced by our reduction this may not be the case, even if the initial instance of Minimum Crossing Number had bounded vertex degrees. Our next result provides an algorithm for the MCNwRS problem.

Theorem 4 *There is an efficient randomized algorithm, that, given an instance (G, Σ) of MCNwRS, with high probability computes a solution of value at most $\tilde{O}\left((\mathsf{OPT}_{\mathsf{cnwrs}}(G, \Sigma) + |E(G)|)^{2-\epsilon}\right)$ for this instance, for $\epsilon = 1/20$.*

By combining Theorem 3 with Theorem 4, we obtain the following immediate corollary.

Corollary 5 *There is an efficient randomized algorithm, that, given any n-vertex graph G with maximum vertex degree Δ, with high probability computes a drawing of G in the plane with at most $O\left((\mathsf{OPT}_{\mathsf{cr}}(G))^{2-\epsilon} \cdot \mathrm{poly}(\Delta \log n)\right)$ crossings, for $\epsilon = 1/20$.*

Lastly, by combining the algorithm from Corollary 5 with the algorithm of Even et al. [EGS02], we obtain the following corollary, whose proof is deferred to the full version of the paper.

Corollary 6 *There is a randomized efficient $O(n^{1/2-\epsilon'} \cdot \mathrm{poly}(\Delta))$-approximation algorithm for Minimum Crossing Number, for some universal constant ϵ'.*

B. Our Techniques

In this subsection, we provide a high-level intuitive overview of the main technical result of our paper – the proof of Theorem 1. As in much of previous work, we focus on the special case of Minimum Crossing Number, where the input graph G is 3-connected. This special case seems to capture the main difficulty of the problem, and the extension of our algorithm to the general case is somewhat standard. We use the standard graph-theoretic notion of well-linkedness: we say that a set Γ of vertices of an n-vertex graph G is *well-linked* in G iff for every pair Γ', Γ'' of disjoint equal-cardinality subsets of Γ, there is a collection of $|\Gamma'|$ paths in G, connecting every vertex of Γ' to a distinct vertex of Γ'', such that every edge of G participates in at most $\mathrm{poly} \log n$ such paths. Given a sub-graph C of

[2]Note that, since the input graph G in both MCNwRS and this variant is allowed to be a multi-graph, it is possible that $m \gg n^2$, and the optimal solution value may be much higher than n^4.

G, we let $\Gamma(C)$ be the set of its *boundary vertices*: all vertices of C that are incident to edges of $E(G) \setminus E(C)$. The following observation can be shown to follow from arguments in [Chu11]: Suppose we are given a collection $\mathcal{C} = \{C_1, \ldots, C_r\}$ of disjoint sub-graphs of G, such that each subgraph C_i has the following properties:

- **3-Connectivity:** graph C_i is 3-connected;
- **Well-Linkedness:** the vertex set $\Gamma(C_i)$ is well-linked in C_i; and
- **Strong Planarity:** graph C_i is planar, and the vertices of $\Gamma(C_i)$ lie on the boundary of a single face in the unique planar drawing of C_i.

Arguments from [Chu11] can then be used to show that there is a near-optimal drawing of G, in which the edges of $\bigcup_{i=1}^{r} E(C_i)$ do not participate in any crossings. Therefore, in order to prove Theorem 1, it is enough to show that there is a small planarizing set E'' of edges in G, with $E' \subseteq E''$, such that every connected component C of $G \setminus E''$ has the above three properties. We note that the Well-Linkedness property is easy to achieve: there are standard algorithms (usually referred to as "well-linked decomposition"), that allow us to compute a set E'' of $O(|E'|)$ edges with $E' \subseteq E''$, such that every connected component of $G \setminus E''$ has the Well-Linkedness property. Typically, such algorithms start with $E'' = E'$, and then iterate, as long as some connected component C of the current graph $G \setminus E''$ does not have the Well-Linkedness property. The algorithm then computes a cut (X, Y) of C that is sparse with respect to the set $\Gamma(C)$ of vertices, adds the edges of $E(X, Y)$ to the set E'' and continues to the next iteration. One can show, using standard arguments, that, once the algorithm terminates, $|E''| \leq O(|E'|)$ holds.

The remaining two properties, however, seem impossible to achieve, if we insist that the set E'' of edges is relatively small. For example, consider a cylindrical grid C, that consists of a large number N of disjoint cycles, and a number of additional paths that intersect each cycle in a way that forms a cylindrical grid (see Figure 1(a)). Consider two additional graphs X and X'. Let G be a graph obtained from the union of C, X, X', by adding two sets of edges: set E_1, connecting some vertices of X to some vertices of the innermost cycle of C (the set of endpoints of these edges that lie in C is denoted by Γ_1), and set E_2, connecting some vertices of X' to some vertices of the outermost cycle of C (the set of endpoints of these edges that lie in C is denoted by Γ_2). Assume that the given planarizing set of edges for G is $E' = E_1 \cup E_2$, so that $\Gamma(C) = \Gamma_1 \cup \Gamma_2$. Then, in order to achieve the strong planarity property for the graph C, we are forced to add N edges to set E'': one edge from every cycle of the cylindrical grid, in order to ensure that the vertices of $\Gamma_1 \cup \Gamma_2$ lie on the boundary of a single face (see Figure 1(b)).

In order to overcome this difficulty, we weaken the Strong

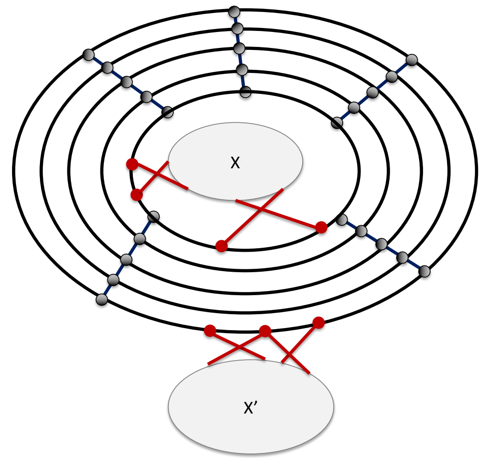

(a) Bad example for the Strong Planarity property. The vertices of $\Gamma_1 \cup \Gamma_2$ and the edges of $E_1 \cup E_2$ are shown in red.

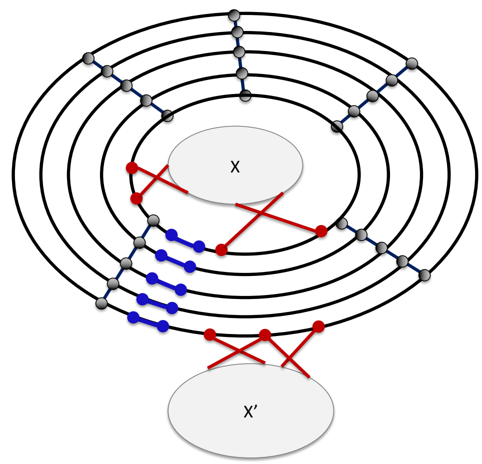

(b) Deleting the blue edges in this figure ensures that all vertices of $\Gamma_1 \cup \Gamma_2$ lie on the boundary a single face. Note that this deletion may cause C to violate the 3-connectivity property.

Figure 1. Obtaining the Strong Planarity property

Planarity requirement, as follows. Informally, given a graph G and a sub-graph C of G, a *bridge* for C is a connected component of $G \setminus C$. If R is a bridge for C, then we let $L(R)$ be the set of *legs* of this bridge, that contains all vertices of C that have a neighbor in R. We replace the Strong Planarity requirement with the following weaker property:

- **Bridge Property:** for every bridge R for C_i, all vertices of $L(R)$ lie on the boundary of a single face in the unique planar drawing of C_i.

We illustrate the intuition for why this weaker property is sufficient on the example from Figure 1(a). Let $G_1 = C \cup E_1 \cup X$, and let $G_2 = C \cup E_2 \cup X'$. Notice that C has the strong planarity property in both G_1 and G_2. Therefore, from the results of [Chu11], there is a drawing φ_1 of G_1, and a drawing φ_2 of G_2, such that the edges of C do not participate in any crossings in either drawing. In particular, the drawings of the graph C induced by φ_1 and φ_2 must be identical to the unique planar drawing of C, and hence to each other. Moreover, in drawing φ_1 of G_1, all edges and vertices of X must be drawn inside a disc whose boundary is the innermost cycle of C, and in the drawing φ_2 of G_2, all edges and vertices of X' must be drawn outside of the disc whose boundary is the outermost cycle of C. Therefore, we can "glue" the drawings φ_1 and φ_2 to each other via the drawing of C in order to obtain the final drawing φ of G, such that the edges of C do not participate in crossings in φ.

Assume now that the Bridge Property does not hold for C (for example, assume that X and X' are connected by an edge). Then we can show that in any drawing of G, at least $\Omega(N)$ edges of C must participate in crossings. We can then add N edges of C to E'' – one edge per cycle, as shown in Figure 1(b), in order to ensure that all vertices of $\Gamma_1 \cup \Gamma_2$ lie on the boundary of a single face of the resulting drawing. This increase in the cardinality of E'' can be charged to the crossings of the optimal drawing of G in which the edges of C participate.

Intuitively, if our goal were to only ensure the Well-Linkedness and the Bridge properties, we could start with $E'' = E'$, and then gradually add edges to E'', until every connected component of $G \setminus E''$ has the bridge property, using reasonings that are similar to the above. After that we could employ the well-linked decomposition in order to ensure the Well-Linkedness property of the resulting clusters. One can show that, once the Bridge Property is achieved, it continues to hold over the course of the algorithm that computes the well-linked decomposition.

Unfortunately, it is also critical that we ensure the 3-connectivity property. Assume that a connected component C of $G \setminus E''$ is 3-connected, let ψ be its unique planar drawing, and let φ^* be the optimal drawing of G. Then one can show that drawing ψ of C is "close" to the drawing φ^*_C of C induced by φ^*. We measure the "closeness" between two drawings using the notion of irregular vertices and edges, that was introduced in [CMS11]. A vertex of C is irregular, if the ordering of the edges incident to C, as they enter C, is different in the drawings ψ and φ^*_C (ignoring the orientation). The notion of irregular edges is somewhat more technical and since we do not need it, we will not define it here. It was shown in [CMS11], that, if C is 3-connected, then the total number of irregular vertices and edges of C is roughly bounded by the number of crossings in which the edges of C participate in φ^*. Therefore, if we think of the number of crossings in φ^* as being low, then the two drawings ψ and φ^*_C are close to each other. However, if graph C is not 3-connected, and, for example, is only 2-connected, then the number of irregular vertices in C may be much higher. Let $S_2(C)$ denote the set of all vertices of C that participate in 2-separators, that is, a vertex v belongs to $S_2(C)$ iff there is another vertex $v' \in V(C)$, such that graph $C \setminus \{v, v'\}$ is not connected. It is easy to see that, even if the drawing φ^*_C is planar, it is possible that there are as many as $|S_2(C)|$ irregular vertices in C. Since we do not know what the optimal drawing φ^* looks like, it seems impossible to fix a planar drawing of C that is close to φ^*_C.

Ensuring the 3-connectivity property for the connected components of $G \setminus E''$, however, seems a daunting task. As edges are added to E'', some components C may no longer be 3-connected. Even if we somehow manage to decompose them into 3-connected subgraphs, while only adding few edges to E'', the addition of these new edges may cause the well-linkedness property to be violated. Then we need to perform the well-linked decomposition from scratch, which in turn can lead to the violation of the 3-connectivity property, and it is not clear that this process will terminate while $|E''|$ is still sufficiently small.

In order to get around this problem, we slightly weaken the 3-connectivity property. We first observe, that, even if graph C is not 3-connected, and is instead 2-connected, but the number of vertices participating in 2-separators (vertices in set $S_2(C)$) is low, then this is sufficient for us. Intuitively, the reason is that, the results of [CMS11] show that the number of irregular vertices in such a graph C is roughly bounded by $|S_2(C)|$ plus the number of crossings in φ^* in which the edges of C participate. Another intuitive explanation is that, when $|S_2(C)|$ is low, there are fewer possible planar drawings of C, so we may think of all of them as being "close" to φ^*_C. Unfortunately, even this weaker property is challenging to achieve, since we need to ensure that it holds simultaneously with the Well-Linkedness and the Bridge properties, for all connected components of $G \setminus E''$. In order to overcome this obstacle, we allow ourselves to add a small set A of "fake" edges to the graph G, whose addition ensures that each component of $(G \setminus E'') \cup A$ is a 2-connected graph with few 2-separators, for which the Well-Linkedness and the Bridge properties hold. Intuitively, we show that the fake edges of A can be embedded into the graph G, so that, in a sense, we can augment the optimal drawing φ^* of G by adding the images of the edges of A to it, without increasing the number of crossings (though we note that this is an over-simplification that is only intended to provide intuition). The proof of Theorem 1 can then be thought of as consisting of two parts. In the first part, we present an efficient algorithm that computes the set E'' of edges of G with $E' \subseteq E''$, and the collection A of fake edges, such that, for every connected component C of graph $(G \setminus E'') \cup A$, one of the following holds: either $|\Gamma(C)| \leq \mathrm{poly}(\Delta \log n)$, or C has the Well-Linkedness and the Bridge properties, together with the weakened 3-Connectivity property. We also compute an embedding of the fake edges in A into G in this part. In the second part, we show that there exists a near-optimal drawing of G in which the edges of $G \setminus E''$ do not participate in crossings. The latter part formalizes and greatly generalizes ideas presented in [Chu11].

Organization: We start with Preliminaries in Section II, and then provide a high-level overview of the main technical result of our paper – the proof of Theorem 1 – in Section III. The detailed proof of the theorem, as well as proofs of Theorem 3, Theorem 4 and Corollary 6 are deferred to the full version of the paper, due to lack of space.

II. Preliminaries

By default, all logarithms are to the base of 2. All graphs are finite, simple and undirected. Graphs with parallel edges are explicitly referred to as multi-graphs.

We follow standard graph-theoretic notation. Assume that we are given a graph $G = (V, E)$. For a vertex $v \in V$, we denote by $\delta_G(v)$ the set of all edges of G that are incident to v. For two disjoint subsets A, B of vertices of G, we denote by $E_G(A, B)$ the set of all edges with one endpoint in A and another in B. For a subset $S \subseteq V$ of vertices, we denote by $E_G(S)$ the set of all edges with both endpoints in S, and we denote by $\mathsf{out}_G(S)$ the subset of edges of E with exactly one endpoint in S, namely $\mathsf{out}_G(S) = E_G(S, V \setminus S)$. We denote by $G[S]$ the subgraph of G induced by S. We sometimes omit the subscript G if it is clear from the context. We say that a graph G is ℓ-connected for some integer $\ell > 0$, if there are ℓ vertex-disjoint paths between every pair of vertices in G.

Given a graph $G = (V, E)$, a *drawing* φ of G is an embedding of the graph into the plane, that maps every vertex to a point and every edge to a continuous curve that connects the images of its endpoints. We require that the interiors of the curves representing the edges do not contain images of any of the vertices. We say that two edges *cross* at a point p, if the images of both edges contain p, and p is not the image of a shared endpoint of these edges; we call such a point p a *crossing*. We require that no three edges cross at the same point in a drawing of φ. We say that φ is a *planar drawing* of G iff no pair of edges of G cross in φ. For a vertex $v \in V(G)$, we denote by $\varphi(v)$ the image of v, and for an edge $e \in E(G)$, we denote by $\varphi(e)$ the image of e in φ. For any subgraph C of G, we denote by $\varphi(C)$ the union of images of all vertices and edges of C in φ. For a path $P \subseteq G$, we sometimes refer to $\varphi(P)$ as the *image of P in φ*. Note that a drawing of G in the plane naturally defines a drawing of G on the sphere and vice versa; we use both types of drawings. Given a graph G and a drawing φ of G in the plane, we use $\mathsf{cr}(\varphi)$ to denote the number of crossings in φ. Let φ' be the drawing of G that is a mirror image of φ. We say that φ and φ' are *identical* drawings of G, and that their *orientations* are different. We sometime say that φ' is obtained by *flipping* the drawing φ. If γ is a simple closed curve in φ that intersects G at vertices only, and S is the set of vertices of G whose images lie on γ, with $|S| \geq 3$, then we say that the circular orderings of the vertices of S along γ in φ and φ' are identical, but the orientations of the two orderings are different, or opposite.

Whitney [Whi92] proved that every 3-connected planar graph has a unique planar drawing. Throughout, for a 3-connected planar graph G, we denote by ρ_G the unique planar drawing of G.

Problem Definitions. The goal of the $\mathsf{Minimum\ Crossing\ Number}$ problem is to compute a drawing of the input graph G in the plane with smallest number of crossings. The value of the optimal solution, also called the *crossing number* of G, is denoted by $\mathsf{OPT}_{\mathsf{cr}}(G)$.

We also consider a closely related problem called Crossing Number with Rotation System (MCNwRS). In this problem, we are given a multi-graph G, and, for every vertex $v \in V(G)$, a circular ordering $\mathcal{O}_v$ of its incident edges. We denote $\Sigma = \{\mathcal{O}_v\}_{v \in V(G)}$, and we refer to Σ as a *rotation system for* G. We say that a drawing φ of G *respects* the rotation system Σ if the following holds. For every vertex $v \in V(G)$, let $\eta(v)$ be an arbitrarily small disc around v in φ. Then the images of the edges of $\delta_G(v)$ in φ must intersect the boundary of $\eta(v)$ in a circular order that is identical to $\mathcal{O}_v$ (but we can choose the orientation of this ordering, that may be either clock-wise or counter-clock-wise). In the MCNwRS problem, the input is a **multi-graph** G with a rotation system Σ, and the goal is to compute a drawing of G in the plane that respects Σ and minimizes the number of crossings.

Faces and Face Boundaries. Suppose we are given a planar graph G and a drawing φ of G in the plane. The set of faces of φ is the set of all connected regions of $\mathbb{R}^2 \setminus \varphi(G)$. We designate a single face of φ as the "outer", or the "infinite" face. The *boundary* $\delta(F)$ of a face F is a sub-graph of G consisting of all vertices and edges of G whose image is incident to F. Notice that, if graph G is not connected, then boundary of a face may also be not connected.

Bridges. Let G be a graph, and let $C \subseteq G$ be a sub-graph of G. A *bridge* for C in graph G is either (i) an edge $e = (u, v) \in E(G)$ with $u, v \in V(C)$ and $e \notin E(C)$; or (ii) a connected component of $G \setminus V(C)$. We denote by $\mathcal{R}_G(C)$ the set of all bridges for C in graph G. For each bridge $R \in \mathcal{R}_G(C)$, we define the set of vertices $L(R) \subseteq V(C)$, called the *legs of* R, as follows. If R consists of a single edge e, then $L(R)$ contains the endpoints of e. Otherwise, $L(R)$ contains all vertices $v \in V(C)$, such that v has a neighbor that belongs to R.

Sparsest Cut and Well-Linkedness. Suppose we are given a graph $G = (V, E)$, and a subset $\Gamma \subseteq V$ of its vertices. We say that a cut (X, Y) in G is a valid Γ-cut iff $X \cap \Gamma, Y \cap \Gamma \neq \emptyset$. The *sparsity* of a valid Γ-cut (X, Y) is $\frac{|E(X,Y)|}{\min\{|X \cap \Gamma|, |Y \cap \Gamma|\}}$. In the Sparsest Cut problem, given a graph G and a subset Γ of its vertices, the goal is to compute a valid Γ-cut of minimum sparsity. Arora, Rao and Vazirani [ARV09] have shown an $O(\sqrt{\log n})$-approximation algorithm for the sparsest cut problem. We denote this algorithm by $\mathcal{A}_{\text{ARV}}$, and its approximation factor by $\beta_{\text{ARV}}(n) = O(\sqrt{\log n})$.

We say that a set Γ of vertices of G is *α-well-linked* in G, iff the value of the sparsest cut in G with respect to Γ is at least α.

III. High-Level Overview

In this subsection we provide a high-level overview of the main technical result of our paper – the proof of Theorem 1. The detailed proof of the theorem, as well as proofs of Theorem 3, Theorem 4 and Corollary 6 appear in the full version of the paper. As in previous work, we start by considering a special case of the Minimum Crossing Number problem, where the input graph G is 3-connected. This special case seems to capture the main technical challenges of the whole problem, and the extension to non-3-connected graphs is relatively easy and follows the same framework as in previous work [CMS11]. We start by defining several central notions that our proof uses.

A. Acceptable Clusters and Decomposition into Acceptable Clusters

In this subsection we define acceptable clusters and decomposition into acceptable clusters. These definitions are central to all our results. Let G be an input graph on n vertices of maximum degree at most Δ; we assume that G is 3-connected. Let $\hat{E}$ be any planarizing set of edges for G, and let $H = G \setminus \hat{E}$. Let $\Gamma \subseteq V(G)$ be the set of all vertices that serve as endpoints of edges in $\hat{E}$; we call the vertices of Γ *terminals*. We will define a set A of *fake edges*; for every fake edge $e \in A$, both endpoints of e must lie in Γ. We emphasize that the edges of A do not necessarily lie in H or in G; in fact we will use these edges in order to augment the graph H.

We denote by $\mathcal{C}$ the set of all connected components of graph $H \cup A$, and we call elements of $\mathcal{C}$ *clusters*. For every cluster $C \in \mathcal{C}$, we denote by $\Gamma(C) = \Gamma \cap V(C)$ the set of all terminals that lie in C. We also denote by $A_C = A \cap C$ the set of all fake edges that lie in C.

Definition. *We say that a cluster $C \in \mathcal{C}$ is a* type-1 acceptable cluster *iff:*

- $A_C = \emptyset$*; and*
- $|\Gamma(C)| \leq \mu$ *for* $\mu = 512\Delta\beta_{\text{ARV}}(n)\log_{3/2} n = O(\Delta \log^{1.5} n)$.

Consider now some cluster $C \in \mathcal{C}$, and assume that it is 2-connected. For a pair (u, v) of vertices of C, we say that (u, v) is a 2-separator for C iff the graph $C \setminus \{u, v\}$ is not connected. We denote by $S_2(C)$ the set of all vertices of C that participate in 2-separators, that is, a vertex $v \in C$ belongs to $S_2(C)$ iff there is another vertex $u \in C$ such that (v, u) is a 2-separator for C. Next, we define type-2 acceptable clusters.

Definition. *We say that a cluster $C \in \mathcal{C}$ is a* type-2 acceptable cluster *with respect to its drawing ψ'_C on the sphere if the following conditions hold:*

- **(Connectivity):** *C is a simple 2-connected graph, and $|S_2(C)| \leq O(\Delta|\Gamma(C)|)$. Additionally, graph $C \setminus A_C$ is a 2-connected graph.*
- **(Planarity):** *C is a planar graph, and the drawing ψ'_C is planar. We denote by $\psi_{C\setminus A_C}$ the drawing of $C \setminus A_C$ is induced by ψ'_C.*
- **(Bridge Consistency Property):** *for every bridge $R \in \mathcal{R}_G(C \setminus A_C)$, there is a face F in the drawing $\psi_{C\setminus A_C}$ of $C \setminus A_C$, such that all vertices of $L(R)$ lie on the boundary of F; and*
- **(Well-Linkedness of Terminals):** *the set $\Gamma(C)$ of terminals is α-well-linked in $C \setminus A_C$, for $\alpha = \frac{1}{128\Delta\beta_{\text{ARV}}(n)\log_{3/2} n} = \Theta\left(\frac{1}{\Delta \log^{1.5} n}\right)$.*

Let $\mathcal{C}_1 \subseteq \mathcal{C}$ denote the set of all type-1 acceptable clusters. For a fake edge $e = (x, y) \in A$, an *embedding* of e is a path $P(e) \subseteq G$ connecting x to y. We will ensure that there exists an embedding of all fake edges in A that has additional useful properties summarized below.

Definition. *A* legal embedding *of the set A of fake edges is a collection $\mathcal{P}(A) = \{P(e) \mid e \in A\}$ of paths in G, such that the following hold:*

- *for every edge $e = (x, y) \in A$, path $P(e)$ has endpoints x and y, and moreover, there is a type-1 acceptable cluster $C(e) \in \mathcal{C}_1$ such that $P(e) \setminus \{x, y\}$ is contained in $C(e)$; and*
- *for any pair $e, e' \in A$ of distinct edges, $C(e) \neq C(e')$.*

Note that from the definition of the legal embedding, all paths in $\mathcal{P}(A)$ must be mutually internally disjoint. Finally, we define a decomposition of a graph G into acceptable clusters; this definition is central for the proof of our main result.

Definition. *A* decomposition of a graph G into acceptable clusters *consists of:*

- *a planarizing set $\hat{E} \subseteq E(G)$ of edges of G;*
- *a set A of fake edges (where the endpoints of each fake edge are terminals with respect to $\hat{E}$);*
- *a partition $(\mathcal{C}_1, \mathcal{C}_2)$ of all connected components (called clusters) of the resulting graph $(G \setminus \hat{E}) \cup A$ into two subsets, such that every cluster $C \in \mathcal{C}_1$ is a type-1 acceptable cluster;*
- *for every cluster $C \in \mathcal{C}_2$, a planar drawing ψ'_C of C on the sphere, such that C is a type-2 acceptable cluster with respect to ψ'_C; and*
- *a legal embedding $\mathcal{P}(A)$ of all fake edges.*

We denote such a decomposition by $\mathcal{K} =$

$\left(\hat{E}, A, \mathcal{C}_1, \mathcal{C}_2, \{\psi'_C\}_{C\in\mathcal{C}_2}, \mathcal{P}(A)\right)$.

Our first result is the following theorem, whose proof is deferred to the full version of the paper, that allows us to compute a decomposition of the input graph G into acceptable clusters. This result is one of the main technical contributions of our work.

Theorem 7 *There is an efficient algorithm, that, given a 3-connected n-vertex graph G with maximum vertex degree at most Δ and a planarizing set E' of edges for G, computes a decomposition $\mathcal{K} = \left(E'', A, \mathcal{C}_1, \mathcal{C}_2, \{\psi'_C\}_{C\in\mathcal{C}_2}, \mathcal{P}(A)\right)$ of G into acceptable clusters, such that $E' \subseteq E''$ and $|E''| \leq O((|E'| + \mathsf{OPT}_{\mathsf{cr}}(G)) \cdot \mathrm{poly}(\Delta \log n))$.*

B. Canonical Drawings

In this subsection, we assume that we are given a 3-connected n-vertex graph G with maximum vertex degree at most Δ, and a decomposition $\mathcal{K} = \left(E'', A, \mathcal{C}_1, \mathcal{C}_2, \{\psi'_C\}_{C\in\mathcal{C}_2}, \mathcal{P}(A)\right)$ of G into acceptable clusters. Next, we define drawings of G that are "canonical" with respect to the clusters in the decomposition. For brevity of notation, we refer to type-1 and type-2 acceptable clusters as type-1 and type-2 clusters, respectively.

Intuitively, in each such canonical drawing, we require that, for every type-2 cluster $C \in \mathcal{C}_2$, the edges of $C \setminus A_C$ do not participate in any crossings, and for every type-1 acceptable cluster $C \in \mathcal{C}_1$, the edges of C only participate in a small number of crossings (more specifically, we will define a subset $E^*(C)$ of edges for each cluster $C \in \mathcal{C}_1$ that are allowed to participate in crossings). We then show that any drawing of G can be transformed into a drawing that is canonical with respect to all clusters in $\mathcal{C}_1 \cup \mathcal{C}_2$, while only slightly increasing the number of crossings. This is sufficient in order to complete the proof of Theorem 1, by adding to E'' the set $\bigcup_{C\in\mathcal{C}_1} E^*(C)$ of edges. However, in order to be able to reduce the problem to the MCNwRS problem, as required in Theorem 3, we need stronger properties. We will define, for every type-1 cluster $C \in \mathcal{C}_1$, a fixed drawing ψ_C, and we will require that, in the final drawing of G, the induced drawing of each such cluster C is precisely ψ_C. For every type-2 cluster $C \in \mathcal{C}_2$, we have already defined a drawing $\psi_{C\setminus A_C}$ of $C \setminus A_C$ – the drawing of $C \setminus A_C$ that is induced by the drawing ψ'_C of C. We will require that the drawing of $C \setminus A_C$ that is induced by the final drawing of G is precisely $\psi_{C\setminus A_C}$. Additionally, for each cluster $C \in \mathcal{C}_1$, and for each bridge $R \in \mathcal{R}_G(C)$, we will define a disc $D(R)$ in the drawing ψ_C of C, and we will require that all vertices and edges of R are drawn inside $D(R)$ in the final drawing of G. Similarly, for each type-2 acceptable cluster $C \in \mathcal{C}_2$, for every bridge $R \in \mathcal{R}_G(C \setminus A_C)$, we define a disc $D(R)$ in the drawing $\psi_{C\setminus A_C}$ of $C \setminus A_C$, and we will require that all vertices and edges of R are drawn inside $D(R)$ in the final drawing of G. This will allow us to fix the locations of the components of $\mathcal{C}_1 \cup \mathcal{C}_2$ with respect to each other (that is, for each pair $C, C' \in \mathcal{C}_1 \cup \mathcal{C}_2$ of clusters, we will identify a face F in the drawing $\psi_{C\setminus A_C}$ of $C \setminus A_C$, and a face F' in the drawing $\psi_{C'\setminus A_{C'}}$ of $C' \setminus A_{C'}$, such that, in the final drawing φ of the graph G, graph $C' \setminus A_{C'}$ is drawn inside the face F (of the drawing of $C \setminus A$ induced by φ, which is identical to $\psi_{C\setminus A_C}$)), and similarly graph $C \setminus A_C$ is drawn inside the face F').

Before we continue, it would be convenient for us to ensure that, for every type-1 cluster $C \in \mathcal{C}_1$, the vertices of $\Gamma(C)$ have degree 1 in C, and degree 2 in G; it would also be convenient for us to ensure that no edge of E'' connects two vertices that lie in the same cluster. In order to ensure these properties, we subdivide some edges of G. Specifically, if $e = (u, v) \in E''$ is an edge with $u, v \in C$, for some cluster $C \in \mathcal{C}_1 \cup \mathcal{C}_2$, then we subdivide the edge (u, v) with two vertices, replacing it with a path (u, u', v', v). The edges (u, u') and (v', v) are then added to set E'' instead of the edge (u, v), and we add a new cluster to $\mathcal{C}_1$, that consists of the vertices u', v', and the edge (u', v'). This transformation ensures that no edge of E'' connects two vertices that lie in the same cluster. Consider now any type-1 cluster $C \in \mathcal{C}_1$. For every edge $e = (u, v) \in E''$ with $u \in V(C)$ and $v \notin V(C)$, we subdivide the edge with a new vertex u', thereby replacing the edge with the path (u, u', v). Vertex u' and edge (u, u') are added to the cluster C, while edge (u', v) replaces the edge (u, v) in set E''. Note that u' now becomes a terminal, and, once all edges of E'' that are incident to the vertices of C are processed, u will no longer be a terminal. Abusing the notation, the final cluster that is obtained after processing all edges of E'' incident to $V(C)$ is still denoted by C. Notice that now the number of terminals that lie in C may have grown by at most a factor Δ, and so $|\Gamma(C)| \leq \mu\Delta$ must hold. Abusing the notation, we will still refer to C as a type-1 acceptable cluster, and we will continue to denote by $\mathcal{C}_1$ the set of all such clusters in the decomposition. Observe that this transformation ensures that every vertex of $\Gamma(C)$ has degree 1 in C and degree 2 in G. Once every cluster $C \in \mathcal{C}_1$ is processed in this manner, we obtain the final graph G'. Observe that $|E''|$ may have increased by at most a constant factor. Notice also that any drawing of G' on the sphere immediately gives a drawing of G with the same number of crossings. Therefore, to simplify the notation, we will denote the graph G' by G, and we will assume that the decomposition $\mathcal{K}$ of G into acceptable clusters has the following two additional properties:

P1) For every edge $e \in E''$, the endpoints of e lie in different clusters of $\mathcal{C}_1 \cup \mathcal{C}_2$; and
P2) For every type-1 cluster $C \in \mathcal{C}_1$, for every terminal $t \in \Gamma(C)$, the degree of t in C is 1, and its degree in G is 2.

We now proceed to define canonical drawings of the graph G with respect to the clusters of $\mathcal{C}_1 \cup \mathcal{C}_2$.

1) Canonical Drawings for Type-2 Acceptable Clusters: Consider any type-2 cluster $C \in \mathcal{C}_2$. Recall that the decomposition $\mathcal{K}$ into acceptable clusters defines a planar drawing ψ'_C of C on the sphere, that induces a planar drawing $\psi_{C\setminus A_C}$ of $C \setminus A_C$ on the sphere. Recall that the Bridge Consistency Property of type-2 acceptable clusters ensures that, for every bridge $R \in \mathcal{R}_G(C \setminus A_C)$, there is a face F of the drawing $\psi_{C\setminus A_C}$, such that the vertices of $L(R)$ lie on the boundary of F (we note that face F is not uniquely defined; we break ties arbitrarily). Since graph $C\setminus A_C$ is 2-connected, the boundary of face F is a simple cycle, whose image is a simple closed curve. We denote by $D(R)$ the disc corresponding to the face F, so the boundary of $D(R)$ is the simple closed curve that serves as the boundary of the face F. Notice that the resulting set $\{D(R)\}_{R\in\mathcal{R}_G(C\setminus A_C)}$ of discs has the following properties:

D1) If $R \neq R'$ are two distinct bridges in $\mathcal{R}_G(C \setminus A_C)$, then either $D(R) = D(R')$, or $D(R) \cap D(R')$ only contains points on the boundaries of the two discs; and

D2) For every bridge $R \in \mathcal{R}_G(C \setminus A_C)$, the vertices of $L(R)$ lie on the boundary of $D(R)$ in the drawing $\psi_{C\setminus A_C}$.

We are now ready to define canonical drawings with respect to type-2 clusters.

Definition. *Let φ be any drawing of the graph G on the sphere. We say that the drawing φ is* canonical *with respect to a type-2 cluster $C \in \mathcal{C}_2$ iff:*

- *the drawing of $C \setminus A_C$ induced by φ is identical to $\psi_{C\setminus A_C}$ (but its orientation may be different);*
- *the edges of $C \setminus A_C$ do not participate in any crossings in φ; and*
- *for every bridge $R \in \mathcal{R}_G(C \setminus A_C)$, all vertices and edges of R are drawn in the interior of the disc $D(R)$ (that is defined with respect to the drawing $\psi_{C\setminus A_C}$ of $C \setminus A_C$).*

2) Canonical Drawings for Type-1 Acceptable Clusters: For convenience, we denote $\mathcal{C}_1 = \{C_1, \ldots, C_q\}$. We fix an arbitrary optimal drawing φ^* of the graph G. For each $1 \le i \le q$, we denote by χ_i the set of all crossings (e, e') such that either e or e' (or both) are edges of $E(C_i)$. The following observation is immediate.

Observation 8 $\sum_{i=1}^{q} |\chi_i| \le 2 \cdot \mathsf{cr}(\varphi^*) = 2 \cdot \mathsf{OPT}_{\mathsf{cr}}(G)$.

We use the following theorem in order to fix a drawing of each type-1 acceptable cluster C_i; the proof is deferred to the full version of the paper.

Theorem 9 *There is an efficient algorithm that, given a type-1 cluster $C_i \in \mathcal{C}_1$, computes a drawing ψ_{C_i} of C_i on the sphere with $O((|\chi_i| + |\Gamma(C_i)|) \cdot \mathrm{poly}(\Delta \log n))$ crossings. Additionally, the algorithm computes, for every bridge $R \in \mathcal{R}_G(C_i)$, a closed disc $D(R)$, such that:*

- *the vertices of $L(R)$ are drawn on the boundary of $D(R)$ in ψ_{C_i}, and all other vertices of C_i are drawn outside of $D(R)$;*
- *the image of every edge of C_i is disjoint from $D(R)$ in ψ_{C_i}, except possibly for an endpoint that belongs to $L(R)$ that is drawn on the boundary of $D(R)$; and*
- *the discs in $\{D(R)\}_{R\in\mathcal{R}_G(C_i)}$ are mutually disjoint.*

Note that in particular, Properties (D1) and (D2) also hold for the discs in $\{D(R)\}_{R\in\mathcal{R}_G(C)}$.

For each type-1 cluster $C_i \in \mathcal{C}_1$, let $E^*(C_i) \subseteq E(C_i)$ be the set of all edges of C_i that participate in crossings in ψ_{C_i}. Clearly, $|E^*(C_i)| \le O(\mathsf{cr}(\psi_{C_i})) \le O((|\chi_i| + |\Gamma(C_i)|)\,\mathrm{poly}(\Delta \log n))$. Let $E^* = \bigcup_{C_i\in\mathcal{C}_1} E^*(C_i)$. Then, from Observation 8 and Theorem 7:

$$
\begin{aligned}
|E^*| &\le \sum_{C_i\in\mathcal{C}_1} O((|\chi_i| + |\Gamma(C_i)|) \cdot \mathrm{poly}(\Delta \log n)) \\
&\le O\left((\mathsf{OPT}_{\mathsf{cr}}(G) + |E''|)\,\mathrm{poly}(\Delta \log n)\right) \\
&\le O\left((\mathsf{OPT}_{\mathsf{cr}}(G) + |E'|)\,\mathrm{poly}(\Delta \log n)\right).
\end{aligned}
$$

We now define canonical drawings with respect to type-1 clusters.

Definition. *Let φ be any drawing of the graph G on the sphere, and let $C_i \in \mathcal{C}_1$ be a type-1 cluster. We say that φ is a* canonical drawing *with respect to C_i, iff:*

- *the drawing of C_i induced by φ is identical to ψ_{C_i} (but orientation of the two drawings may be different). In particular, the only edges of C_i that participate in crossings of φ are the edges of $E^*(C_i)$; and*
- *for every bridge $R \in \mathcal{R}_G(C_i)$, all vertices and edges of R are drawn in the interior of the disc $D(R)$ (that is defined with respect to the drawing ψ_{C_i} of C_i).*

3) Obtaining a Canonical Drawing: Our next result shows that there exists a near-optimal drawing of the graph G that is canonical with respect to all clusters. The proof of the following theorem is deferred to the full version of the paper.

Theorem 10 *There is an efficient algorithm, that, given an n-vertex graph G of maximum vertex degree at most Δ, an arbitrary drawing φ of G, and a decomposition $\mathcal{K} = \left(E'', A, \mathcal{C}_1, \mathcal{C}_2, \{\psi_C\}_{C\in\mathcal{C}_2}, \mathcal{P}(A)\right)$ of G into acceptable clusters for which Properties (P1) and (P2) hold, together with a drawing ψ_{C_i} of each cluster $C_i \in \mathcal{C}_1$ as defined above, and, for each cluster $C \in \mathcal{C}_1 \cup \mathcal{C}_2$, a collection $\{D(R)\}_{R\in\mathcal{R}_G(C)}$ of discs on the sphere with Properties (D1) and (D2), computes a drawing φ' of G on the sphere with $O\left((|E''| + \mathsf{cr}(\varphi)) \cdot \mathrm{poly}(\Delta \log n)\right)$ crossings, such that φ' is canonical with respect to every cluster $C \in \mathcal{C}_1 \cup \mathcal{C}_2$.*

We note that for our purposes, an existential variant of the above theorem, that shows that a drawing φ' with the required properties exists, is sufficient. We provide the proof of the stronger constructive result in case it may be useful for future work on the problem.

C. Completing the Proof of Theorem 1 for 3-Connected Graphs

Notice that Theorem 10 concludes the proof of Theorem 1 for the special case where G is a 3-connected graph. Indeed, given a 3-connected graph G, we use Theorem 7 to compute a decomposition $\mathcal{K} = (E'', A, \mathcal{C}_1, \mathcal{C}_2, \{\psi_C\}_{C\in\mathcal{C}_2}, \mathcal{P}(A))$ of G into acceptable clusters, such that $E' \subseteq E''$ and $|E''| \leq O((|E'| + \mathsf{OPT}_{\mathsf{cr}}(G)) \cdot \mathrm{poly}(\Delta \log n))$. Next, we apply Theorem 9 to each type-1 cluster $C_i \in \mathcal{C}_1$, to obtain the set $E^*(C_i)$ of edges, the drawing ψ_{C_i} of C_i, and the discs $D(R)$ for all bridges $R \in \mathcal{R}_G(C_i)$. Let $E^* = \bigcup_{C_i\in\mathcal{C}_1} E^*(C_i)$, so $|E^*| \leq O\left((\mathsf{OPT}_{\mathsf{cr}}(G) + |E'|)\,\mathrm{poly}(\Delta \log n)\right)$, as observed above. The final output of the algorithm is the set $E'' \cup E^*$ of edges. Observe that $|E'' \cup E^*| \leq O((|E'| + \mathsf{OPT}_{\mathsf{cr}}(G)) \cdot \mathrm{poly}(\Delta \log n))$, as required. Moreover, by using Theorem 10 with the optimal drawing φ^* of G, we conclude that there exists a drawing φ' of G with $O\left((|E'| + \mathsf{OPT}_{\mathsf{cr}}(G)) \cdot \mathrm{poly}(\Delta \log n)\right)$ crossings, that is canonical with respect to all clusters in $\mathcal{C}_1 \cup \mathcal{C}_2$. In particular, the only edges that may participate in crossings in φ' are edges of $E'' \cup E^*$.

D. Extension to General Graphs

The extension of the proof of Theorem 1 to general graphs is deferred to the full version of the paper. The extension builds on techniques that were introduced in [CMS11]. Additionally, we prove the following theorem, that provides a black-box reduction from the problem of approximating Minimum Crossing Number in general graphs using paradigm Π', to the problem of approximating Minimum Crossing Number in 3-connected graphs, using the same paradigm. The proof is also deferred to the full version of the paper.

Theorem 11 *Suppose there is an efficient (possibly randomized) algorithm, that, given a 3-connected n-vertex graph G with maximum vertex degree Δ, and a planarizing set E' of its edges, computes a drawing of G with at most $f(n, \Delta) \cdot (\mathsf{OPT}_{\mathsf{cr}}(G) + |E'|)$ crossings, for any function f that is monotonously increasing in both n and Δ. Then there exists an efficient (possibly randomized) algorithm that, given a (not necessarily 3-connected) graph $\hat{G}$ on n vertices with maximum vertex degree Δ, and a planarizing set $\hat{E}'$ of its edges, computes a drawing of $\hat{G}$ with the number of crossings bounded by $O\left(f(n, \Delta) \cdot (\mathsf{OPT}_{\mathsf{cr}}(\hat{G}) + |\hat{E}'|) \cdot \mathrm{poly}(\Delta \log n)\right)$.*

Acknowledgment

Julia Chuzhoy was supported in part by NSF grant CCF-1616584. Zihan Tan was supported in part by NSF grant CCF-1616584.

References

[ACNS82] M. Ajtai, V. Chvátal, M. Newborn, and E. Szemerédi. Crossing-free subgraphs. *Theory and Practice of Combinatorics*, pages 9–12, 1982.

[AMS07] C. Ambuhl, M. Mastrolilli, and O. Svensson. Inapproximability results for sparsest cut, optimal linear arrangement, and precedence constrained scheduling. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 329–337, 2007.

[ARV09] Sanjeev Arora, Satish Rao, and Umesh V. Vazirani. Expander flows, geometric embeddings and graph partitioning. *J. ACM*, 56(2), 2009.

[BCG+13] Christoph Buchheim, Markus Chimani, Carsten Gutwenger, Michael Jünger, and Petra Mutzel. Crossings and planarization. In *Handbook on Graph Drawing and Visualization*, pages 43–85. 2013.

[BE14] Michael J Bannister and David Eppstein. Crossing minimization for 1-page and 2-page drawings of graphs with bounded treewidth. In *International Symposium on Graph Drawing*, pages 210–221. Springer, 2014.

[BL84] Sandeep N Bhatt and Frank Thomson Leighton. A framework for solving vlsi graph layout problems. *Journal of Computer and System Sciences*, 28(2):300–343, 1984.

[Cab13] Sergio Cabello. Hardness of approximation for crossing number. *Discrete & Computational Geometry*, 49(2):348–358, 2013.

[CH11] Markus Chimani and Petr Hliněný. A tighter insertion-based approximation of the crossing number. In *International Colloquium on Automata, Languages, and Programming*, pages 122–134. Springer, 2011.

[Chu11] Julia Chuzhoy. An algorithm for the graph crossing number problem. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 303–312. ACM, 2011.

[CMM16] Julia Chuzhoy, Vivek Madan, and Sepideh Mahabadi. In *Private Communication*, 2016.

[CMS11] Julia Chuzhoy, Yury Makarychev, and Anastasios Sidiropoulos. On graph crossing number and edge planarization. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete algorithms*, pages 1050–1069. SIAM, 2011.

[CS13] Chandra Chekuri and Anastasios Sidiropoulos. Approximation algorithms for euler genus and related problems. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 167–176. IEEE, 2013.

[DLM19] Walter Didimo, Giuseppe Liotta, and Fabrizio Montecchiani. A survey on graph drawing beyond planarity. *ACM Computing Surveys (CSUR)*, 52(1):1–37, 2019.

[EGS02] Guy Even, Sudipto Guha, and Baruch Schieber. Improved approximations of crossings in graph drawings and vlsi layout areas. *SIAM Journal on Computing*, 32(1):231–252, 2002.

[GJ83] M. R. Garey and D. S. Johnson. Crossing number is NP-complete. *SIAM J. Algebraic Discrete Methods*, 4(3):312–316, 1983.

[Gro04] M. Grohe. Computing crossing numbers in quadratic time. *J. Comput. Syst. Sci.*, 68(2):285–302, 2004.

[HD15] Petr Hliněnỳ and Marek Derňár. Crossing number is hard for kernelization. *arXiv preprint arXiv:1512.02379*, 2015.

[Hli06] P. Hlinený. Crossing number is hard for cubic graphs. *J. Comb. Theory, Ser. B*, 96(4):455–471, 2006.

[KR07] K. Kawarabayashi and B. A. Reed. Computing crossing number in linear time. In *STOC*, pages 382–390, 2007.

[KS17] Ken-ichi Kawarabayashi and Anastasios Sidiropoulos. Polylogarithmic approximation for minimum planarization (almost). In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 779–788, 2017.

[KS19] Ken-ichi Kawarabayashi and Anastasios Sidiropoulos. Polylogarithmic approximation for euler genus on bounded degree graphs. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 164–175. ACM, 2019.

[Lei83] F. T. Leighton. *Complexity issues in VLSI: optimal layouts for the shuffle-exchange graph and other networks*. MIT Press, 1983.

[LR99] Tom Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM (JACM)*, 46(6):787–832, 1999.

[LT79] Richard J. Lipton and Robert Endre Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979.

[Mat02] J. Matoušek. *Lectures on discrete geometry*. Springer-Verlag, 2002.

[PSŠ07] Michael J Pelsmajer, Marcus Schaefer, and Daniel Štefankovič. Crossing numbers and parameterized complexity. In *International Symposium on Graph Drawing*, pages 31–36. Springer, 2007.

[PSŠ11] Michael J Pelsmajer, Marcus Schaefer, and Daniel Štefankovič. Crossing numbers of graphs with rotation systems. *Algorithmica*, 60(3):679–702, 2011.

[PT00] J. Pach and G. Tóth. Thirteen problems on crossing numbers. *Geombinatorics*, 9(4):194–207, 2000.

[RS09] R. B. Richter and G. Salazar. Crossing numbers. In L. W. Beineke and R. J. Wilson, editors, *Topics in Topological Graph Theory*, chapter 7, pages 133–150. Cambridge University Press, 2009.

[Tur77] P. Turán. A note of welcome. *J. Graph Theory*, 1:1–5, 1977.

[Vrt] Imrich Vrto. Crossing numbers of graphs: A bibliography. http://www.ifi.savba.sk/~imrich.

[Whi92] Hassler Whitney. Congruent graphs and the connectivity of graphs. In *Hassler Whitney Collected Papers*, pages 61–79. Springer, 1992.