

Crossing Minimization in Perturbed Drawings

Radoslav Fulek · Csaba D. Tóth

Received: date / Accepted: date

Abstract Due to data compression or low resolution, nearby vertices and edges of a graph drawn in the plane may be bundled to a common node or arc. We model such a “compromised” drawing by a piecewise linear map $\varphi : G \rightarrow \mathbb{R}^2$. We wish to perturb φ by an arbitrarily small $\varepsilon > 0$ into a proper drawing (in which the vertices are distinct points, any two edges intersect in finitely many points, and no three edges have a common interior point) that minimizes the number of crossings. An ε -perturbation, for every $\varepsilon > 0$, is given by a piecewise linear map $\psi_\varepsilon : G \rightarrow \mathbb{R}^2$ with $\|\varphi - \psi_\varepsilon\| < \varepsilon$, where $\|\cdot\|$ is the uniform norm (i.e., sup norm).

We present a polynomial-time solution for this optimization problem when G is a cycle and the map φ has no **spurs** (i.e., no two adjacent edges are mapped to overlapping arcs). We also show that the problem becomes NP-complete (i) when G is an arbitrary graph and φ has no spurs, and (ii) when φ may have spurs and G is a cycle or a union of disjoint paths.

Keywords Map Approximation · C-Planarity · Crossing Number · NP-hardness

Mathematics Subject Classification (2010) 05C10 · 05C38 · 68R10

Research supported in part by the NSF awards DMS-1800734, CCF-1422311 and CCF-1423615, and the Austrian Science Fund (FWF): M2281-N35. A Preliminary version of this paper has been presented at the *26th International Symposium on Graph Drawing and Network Visualization (Barcelona, 2018)*.

R. Fulek
University of Arizona
Tucson, Arizona
E-mail: radoslav.fulek@gmail.com

C. D. Tóth
California State University Northridge
Los Angeles, CA, USA
E-mail: csaba.toth@csun.edu

1 Introduction

A graph $G = (V, E)$ is a 1-dimensional simplicial complex. A continuous piecewise linear map $\varphi : G \rightarrow \mathbb{R}^2$ maps the vertices in V into points in the plane, and the edges in E to piecewise linear arcs between the corresponding vertices. However, several vertices may be mapped to the same point, and two edges may be mapped to overlapping arcs. This scenario arises in applications in cartography, clustering, and visualization, due to data compression, graph semantics, or low resolution. Previous research focused on determining whether such a map φ can be “perturbed” into an embedding. Specifically, a continuous piecewise linear map $\varphi : G \rightarrow \mathbb{R}^2$ is a **weak embedding** if, for every $\varepsilon > 0$, there is an embedding $\psi_\varepsilon : G \rightarrow \mathbb{R}^2$ with $\|\varphi - \psi_\varepsilon\| < \varepsilon$, where $\|\cdot\|$ is the uniform norm (i.e., sup norm). Recently, Fulek and Kynčl [14] gave a polynomial-time algorithm for recognizing weak embeddings, and the running time was subsequently improved to $O(n \log n)$ for simplicial maps by Akitaya et al. [2]. Note, however, that only planar graphs admit embeddings and weak embeddings. The results in [2, 14] extend to weak embeddings $\varphi : G \rightarrow M$ of a graph G into any 2-dimensional manifold M endowed with a metric, but the machinery developed so far was not able to handle crossings. In this paper, we extend the concept of ε -perturbations to all graphs, and seek a perturbation with the minimum number of crossings.

A continuous map $\varphi : G \rightarrow M$ of a graph G to a 2-manifold M (e.g., $M = \mathbb{R}^2$) is a **drawing** if (i) the vertices in V are mapped to distinct points in M , (ii) each edge is mapped to a continuous arc between two vertices without passing through any other vertex, and (iii) any two edges intersect in finitely many points. A **crossing** between a pair of edges, $e_1, e_2 \in E$, is defined as an intersection point between the relative interiors of the arcs $\varphi(e_1)$ and $\varphi(e_2)$. For a piecewise linear map $\varphi : G \rightarrow \mathbb{R}^2$, let $\text{cr}(\varphi)$ be the minimum nonnegative integer k such that for every $\varepsilon > 0$, there exists a drawing $\psi_\varepsilon : G \rightarrow \mathbb{R}^2$ with $\|\varphi - \psi_\varepsilon\| < \varepsilon$ and k crossings, see Fig. 1 for an illustration.

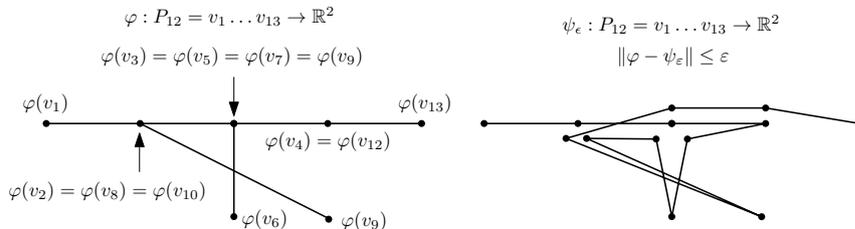


Fig. 1: An example for a map $\varphi : G \rightarrow \mathbb{R}^2$, where $G = P_{12}$, i.e., a path of length 12, with $\text{cr}(\varphi) = 5$ (left); and a perturbation ψ_ε witnessing that $\text{cr}(\varphi) \leq 5$ (right).

It is clear that φ is a weak embedding if and only if $\text{cr}(\varphi) = 0$. Note also that if $e_1, e_2 \in E$ and the arcs $\varphi(e_1)$ and $\varphi(e_2)$ cross transversely at some point

$p \in \mathbb{R}^2$, then $\psi_\varepsilon(e_1)$ and $\psi_\varepsilon(e_2)$ also cross in the ε -neighborhood of p for any sufficiently small $\varepsilon > 0$. An ε -perturbation may, however, remove tangencies and partial overlaps between edges.

The problem of determining $\text{cr}(\varphi)$ for a given map $\varphi : G \rightarrow \mathbb{R}^2$ is NP-complete: In the special case that $\varphi(G)$ is a single point, $\text{cr}(\varphi)$ equals the crossing number of G , and it is NP-complete to find the crossing number of a given graph [15] (even if G is a planar graph plus one edge [6]).

Chang et al. [7] identified two features of a piecewise linear map $\varphi : G \rightarrow \mathbb{R}^2$ that are difficult to handle: A **spur** is a vertex whose incident edges are mapped to the same arc or overlapping arcs, and a **fork** is a vertex mapped to the relative interior of the image of some nonincident edge (a vertex may be both a fork and a spur). Our results (Theorems 1 and 2 below) show that spurs are critical for the algorithmic complexity of determining $\text{cr}(\varphi)$. Forks, however, can easily be eliminated by a suitable subdivision of the edges, which increases the number of vertices by a polynomial factor and only impacts the running time of our algorithms. Similarly, by a suitable subdivision of the edges, we may also assume that $\varphi : G \rightarrow \mathbb{R}^2$ maps every edge of G to a straight-line segment in the plane; we call such a map a **straight-line** map. In the remainder of the paper, we assume that $\varphi : G \rightarrow \mathbb{R}^2$ is a straight-line map.

We prove the following results.

Theorem 1 *Given a cycle $C_n = (V, E)$ with n vertices and a straight-line map $\varphi : C_n \rightarrow \mathbb{R}^2$, then $\text{cr}(\varphi)$ can be computed*

1. in $O(n \log n)$ time if φ has neither spurs nor forks,
2. in $O(n^2 \log n)$ time if φ has no spurs.

As noted above, the problem of determining $\text{cr}(\varphi)$ is NP-complete when G is an arbitrary graph (even if φ is a constant map). We show that the problem remains NP-complete if G is a cycle and we drop the condition that φ has no spurs.

Theorem 2 *Given a positive integer k and a straight-line map $\varphi : G \rightarrow \mathbb{R}^2$, it is NP-complete to decide whether $\text{cr}(\varphi) \leq k$ and*

1. G is a cycle, or
2. G is a union of disjoint paths.

Related previous work. A series of recent results show that weak embeddings of a graph G with n vertices can be recognized in $O(n \log n)$ time. Specifically, Akitaya et al. [1] gave an $O(n \log n)$ -time algorithm for the special case that G is a cycle, improving on earlier work [7, 10]. When G is an arbitrary graph Akitaya et al. [2] gave an $O(n^2 \log n)$ -time algorithm in general, and an $O(n \log n)$ -time algorithm when the map $\varphi : G \rightarrow \mathbb{R}^2$ has no forks.

Finding efficient algorithms for the recognition of weak embeddings $\varphi : G \rightarrow M$, where G is an arbitrary graph and M is a 2-dimensional manifold, was posed as an open problem in [1, 7, 10]. The first polynomial-time solution for the general version follows from a recent variant [14] of the Hanani-Tutte

theorem [17,24], which was conjectured by M. Skopenkov [23] in 2003 and in a slightly weaker form already by Repovš and A. Skopenkov [22] in 1998. Weak embeddings of graphs also generalize various graph visualization models such as **strip planarity** [4] and **level planarity** [20]; and can be seen as a special case [3] of the notoriously difficult **cluster-planarity** (for short, **c-planarity**) [12,13], whose tractability has been a longstanding open problem.

Organization. We start in Sec. 2 with preliminary observations that show that determining $\text{cr}(\varphi)$ is equivalent to a purely combinatorial problem, which can be formulated without metric constraints. We describe and analyse a recognition algorithm, proving Theorem 1, in Sec. 3. We prove that the problem is NP-hard by a reduction from 3SAT in Sec. 4, and describe an algorithm that tests whether $\text{cr}(\varphi) \leq c$ in $n^{O(c)}$ time for any fixed $c > 0$; and conclude in Sec. 6.

2 Preliminaries

We rely on techniques introduced in [1,7,9,14], and complement them with additional tools to keep track of edge crossings. Let $\varphi : G \rightarrow \mathbb{R}^2$ be a piecewise linear function. We may assume, by subdividing the edges of G if necessary, that φ is a straight-line map (i.e., every edge is mapped to a line segment), and it has no forks (no vertex is mapped to the interior of an edge).

We define the **image graph** H by a graph homomorphism $\lambda : G \rightarrow H$ that identifies vertices in $V(G)$ that are mapped to the same point by φ , that is, we have $\lambda(u) = \lambda(v)$ for $u, v \in V(G)$ if and only if $\varphi(u) = \varphi(v)$. Since φ does not have forks, the map $\lambda : G \rightarrow H$ is **simplicial** (that is, it maps vertices to vertices and edges to edges). To distinguish the graphs G and H in our terminology (following Cortese et al. [10]), G has **vertices** $V(G)$ and **edges** $E(G)$, and H has **clusters** $V(H)$ and **pipes** $E(H)$.

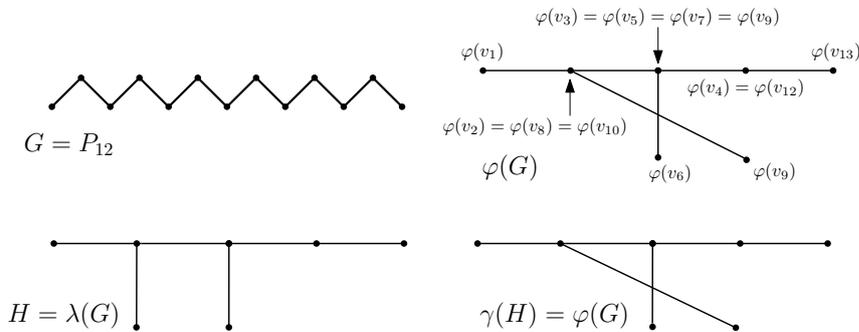


Fig. 2: Graph G (top left) and its straight-line map $\varphi : G \rightarrow \mathbb{R}^2$ (top right). The graph $H = \lambda(G)$ (bottom left) and its straight-line drawing $\gamma(H) = \varphi(G)$ (bottom right).

We can express $\varphi : G \rightarrow \mathbb{R}^2$ as a composition $\varphi = \gamma \circ \lambda$, where $\lambda : G \rightarrow H$ is a simplicial map from G to H (a continuous map between 1-dimensional simplicial complexes) and $\gamma : H \rightarrow \mathbb{R}^2$ is a drawing of H . Refer to Fig. 2. Since φ is a straight-line map, $\gamma : H \rightarrow \mathbb{R}^2$ is a straight-line drawing of H , where each edge in $E(H)$ is mapped to a line segment.

A perturbation ψ_ε of φ lies in the ε -neighborhood of $\varphi(G)$. We define suitable neighborhoods for the graph H and its drawing $\gamma(H) \subset \mathbb{R}^2$ as follows. For the graph H , a straight-line drawing $\gamma : H \rightarrow \mathbb{R}^2$, and an $\varepsilon \in (0, 1)$, we define the ε -**neighborhood** $\mathcal{N}(\varepsilon) \subset \mathbb{R}^2$ as the union of regions N_u and N_{uv} for every $u \in V(H)$ and $uv \in E(H)$, respectively, as follows. See Fig. 3 for an illustration. For every $u \in V(H)$, let N_u be the closed disk of radius ε centered at $\gamma(u)$. For every pipe $uv \in E(H)$, let N_{uv} be the set of points at distance at most ε^2 from $\gamma(uv)$ that lie in the interior of neither N_u nor N_v .

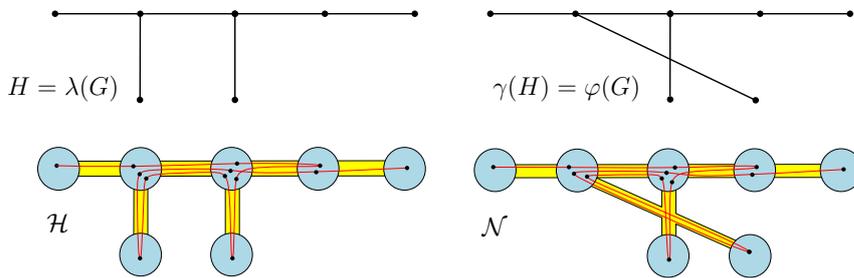


Fig. 3: The graph H (top left) and its straight-line drawing $\gamma(H)$ (top right). The thickening \mathcal{H} of H (i.e., a 2-dimensional manifold with boundary), and a map $\Lambda : G \rightarrow \mathcal{H}$ with one crossing (bottom left). The neighborhood \mathcal{N} of the drawing $\gamma(H)$ in the plane, and a map $\Gamma \circ \Lambda(G)$ with five crossings (bottom right).

Let $\varepsilon_0 \in (0, 1)$ be a sufficiently small constant such that for $\varepsilon = \varepsilon_0$, and for any triple of distinct clusters $\{u, v, w\} \subset V(H)$, the centers of the disks N_u and N_v are at distance at least 4ε apart (in particular, N_u and N_v are disjoint), N_u is disjoint from N_{vw} if $vw \in E(H)$, and the regions N_{uv} and $N_{u'v'}$ are disjoint from each other if $uv, u'v' \in E(H)$ (however, regions N_{uv} and $N_{u'v'}$ may intersect if the line segments $\gamma(uv)$ and $\gamma(u'v')$ cross). Such an $\varepsilon_0 > 0$ exists since φ is a straight-line map without forks. Clearly, these properties hold for every $\varepsilon \in (0, \varepsilon_0]$.

For the graph H and its drawing $\gamma : H \rightarrow \mathbb{R}^2$, we also define the **thickening** \mathcal{H} , $H \subset \mathcal{H}$, as a 2-dimensional manifold with boundary as follows. See Fig. 3 for an illustration. For every $u \in V(H)$, create a topological disk D_u , and for every pipe $uv \in E(H)$, create a rectangle R_{uv} . For every D_u and R_{uv} , fix an arbitrary orientation of ∂D_u and ∂R_{uv} , respectively. Choose $\deg(u)$ pairwise disjoint closed arcs on the boundary ∂D_u of D_u , and label them by $A_{u,v}$, for all $uv \in E(H)$, in the cyclic order around ∂D_u determined by the

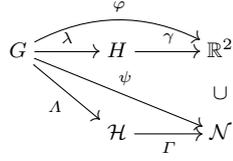


Fig. 4: The diagram of maps expressing the map φ and its perturbation ψ .

rotation of u in the drawing $\gamma(G)$. The manifold \mathcal{H} is obtained by identifying two opposite sides of every rectangle R_{uv} with $A_{u,v}$ and $A_{v,u}$ via an orientation preserving homeomorphism. Note that there is a natural map $\Gamma : \mathcal{H} \rightarrow \mathcal{N}$ such that $\Gamma|_H = \gamma$, furthermore $\Gamma|_{D_u}$ is a homeomorphism between D_u and N_u for every $u \in V(H)$, and $\Gamma|_{R_{uv}}$ is a homeomorphism between R_{uv} and N_{uv} for every $uv \in E(H)$.

Crossing Minimization Rephrased. Refer to Fig. 4. We reformulate a problem instance $\varphi : G \rightarrow \mathbb{R}^2$ as two functions $\lambda : G \rightarrow H$ and $\gamma : H \rightarrow \mathbb{R}^2$, where G and H are abstract graphs, λ is a simplicial map and γ is a straight-line drawing of H . We define a **perturbation** of the map $\varphi = \gamma \circ \lambda$ as a drawing $\psi = \Gamma \circ \Lambda$, where $\Lambda : G \rightarrow \mathcal{H}$ is a drawing of G on \mathcal{H} with the following properties:

- (P1) for every vertex $a \in V(G)$, $\Lambda(a) \in D_{\lambda(a)}$,
- (P2) for every edge $ab \in E(G)$, $\Lambda(ab) \subset D_{\lambda(a)} \cup R_{\lambda(a)\lambda(b)} \cup D_{\lambda(b)}$ such that it crosses the boundary of the disks $D_{\lambda(a)}$ and $D_{\lambda(b)}$ precisely once, and
- (P3) all crossing between arcs $\Lambda(e)$, $e \in E(G)$, lie in the disks D_u , $u \in V(H)$;

and $\Gamma : \mathcal{H} \rightarrow \mathbb{R}^2$ is a continuous map such that $\Gamma|_{D_u}$ is a homeomorphism between D_u and N_u for every $u \in V(H)$, and $\Gamma|_{R_{uv}}$ is a homeomorphism between R_{uv} and N_{uv} for every $uv \in E(H)$. Note, however, that Γ may map the rectangles R_{uv} and $R_{u'v'}$ to overlapping regions N_{uv} and $N_{u'v'}$ for two independent edges $uv, u'v' \in E(H)$.

Given an straight-line map $\varphi = \gamma \circ \lambda$, we seek a perturbation $\psi = \Gamma \circ \Lambda$ that minimizes the number of crossings. In the next few paragraphs, we show that finding a perturbation in this form is a purely combinatorial problem; and we show (cf. Lemma 1) that this problem is equivalent to finding $\text{cr}(\varphi)$.

Combinatorial Representation. Properties (P1)–(P3) allow for a combinatorial representation of the drawing $\Lambda : G \rightarrow \mathcal{H}$: For every pipe $uv \in E(H)$, let π_{uv} be a total order of the edges in $\lambda^{-1}[uv] \subseteq E(G)$ in $R_{\lambda(a)\lambda(b)}$; and let $\pi_\Lambda = \{\pi_{uv} : uv \in E(H)\}$ the collection of these total orders. In fact, we can assume that $\Lambda(G)$ consists of straight-line segments in every rectangle R_{uv} , and every disk D_u . The number of crossings in each disk D_u is determined by the cyclic order of the segment endpoints along ∂D_u . Thus, the number of crossings in all disks D_u , $u \in V(H)$, is determined by π_Λ .

Two Types of Crossings. By restricting the perturbations of a straight-line map $\varphi : G \rightarrow \mathbb{R}^2$ to the form $\psi = \Gamma \circ \Lambda$ defined above, we can distinguish

between two types of crossings: edge-crossings in the neighborhoods N_u , $u \in V(H)$, and edge-crossings between edges mapped to two pipes that cross each other.

The number of crossings between the edges of G inside a disk N_u , $u \in V(H)$, is the same as the number of crossings in D_u , since Γ is injective on D_u . We denote the total number of such crossings by

$$\text{cr}_1(\gamma, \lambda) = \min_A \left(\sum_{u \in V(H)} \text{CR}_A(u) \right),$$

where $\text{CR}_A(u)$ is the number of crossings of the drawing $A(G)$ in the disk D_u .

Let the **weight** of a pipe $e \in E(H)$ be the number of edges of G mapped to e , that is, $w(e) := |\lambda^{-1}[e]|$. If the arcs $\gamma(e_1)$ and $\gamma(e_2)$ cross in the plane, for some $e_1, e_2 \in E(H)$, then every edge in $\lambda^{-1}[e_1]$ crosses all edges in $\lambda^{-1}[e_2]$. The total number of crossings between the edges of G attributed to the crossings between pipes is

$$\text{cr}_2(\gamma, \lambda) = \sum_{\{e_1, e_2\} \in C} w(e_1)w(e_2),$$

where C is the set of pipe pairs $\{e_1, e_2\}$ such that $\gamma(e_1)$ and $\gamma(e_2)$ cross. It is now easy to show that $\text{cr}(\gamma \circ \lambda)$ equals to the sum of the two types of crossings.

Lemma 1 *Let $\varphi : G \rightarrow \mathbb{R}^2$ be a straight-line map without forks, where $\varphi = \gamma \circ \lambda$ for functions γ and λ defined above. Then*

$$\text{cr}(\varphi) = \text{cr}_1(\gamma, \lambda) + \text{cr}_2(\gamma, \lambda). \quad (1)$$

Proof We first show $\text{cr}(\varphi) \leq \text{cr}_1(\gamma, \lambda) + \text{cr}_2(\gamma, \lambda)$. Suppose that $\Lambda : G \rightarrow \mathcal{H}$ attains $\text{cr}_1(\gamma, \lambda)$; and recall that $\text{cr}_2(\gamma, \lambda)$ is determined by φ . We need to show that for every $\varepsilon > 0$, G admits a drawing $\psi_\varepsilon : G \rightarrow \mathbb{R}^2$ with at most $\text{cr}_1(\gamma, \lambda) + \text{cr}_2(\gamma, \lambda)$ crossings such that $\|\varphi - \psi_\varepsilon\| < \varepsilon$. Let $\varepsilon > 0$ be given. Let $\varepsilon_0 > 0$ be as defined above. Put $\varepsilon_1 = \min\{\varepsilon, \varepsilon_0\}$, and let $\mathcal{N}(\varepsilon_1)$ be the ε_1 -neighborhood of $\gamma(H)$. Let $\Gamma : \mathcal{H} \rightarrow \mathcal{N}(\varepsilon_1)$ be as described above. Then $\Gamma \circ \Lambda(G)$ is a drawing of G in $\mathcal{N}(\varepsilon_1) \subset \mathcal{N}(\varepsilon)$ with $\text{cr}_1(\gamma, \lambda) + \text{cr}_2(\gamma, \lambda)$ crossings; specifically, $\text{cr}_1(\gamma, \lambda)$ crossings within the disks N_u over all $u \in V(H)$, and $\text{cr}_2(\gamma, \lambda)$ crossings in the intersections of N_{uv} and $N_{u'v'}$ over all pairs of pipes $uv, u'v' \in E(H)$.

In the other direction, we need to prove that $\text{cr}_1(\gamma, \lambda) + \text{cr}_2(\gamma, \lambda) \leq \text{cr}(\varphi)$. Let $\varepsilon = \varepsilon_0^2$, and let $\psi_\varepsilon : G \rightarrow \mathbb{R}^2$ be a drawing of G with $\text{cr}(\varphi)$ crossings such that $\|\varphi - \psi_\varepsilon\| < \varepsilon$ and such that the cardinality of the crossings between $\psi_\varepsilon(G)$ and $\bigcup_{u \in V(H)} \partial N_u$, denoted by $\Xi(\psi_\varepsilon)$, is minimized. Clearly, $|\Xi(\psi_\varepsilon(G))|$ is finite, since we can assume that $\Xi(\psi_\varepsilon)$ consists of proper crossings between the edges $\psi_\varepsilon(e)$, $e \in E(G)$, and the curves ∂N_u , $u \in V(H)$. Then for every vertex $v \in V(G)$, we have $\|\varphi(v) - \psi_\varepsilon(v)\| < \varepsilon$, hence $\psi_\varepsilon(v) \in N_{\lambda(v)} = N_u$. Furthermore, for every edge $ab \in E(G)$, we have $\psi_\varepsilon(ab) \subset N_{\lambda(a)} \cup N_{\lambda(a)\lambda(b)} \cup N_{\lambda(b)}$. We prove that ψ_ε can be chosen so that $\psi_\varepsilon = \Gamma \circ \Lambda$, where $\Lambda : G \rightarrow \mathcal{H}$

satisfies (P1)–(P3). Then by the choice of ε , and the definition of $\text{cr}_1(\gamma, \lambda)$ and $\text{cr}_2(\gamma, \lambda)$, this implies $\text{cr}_1(\gamma, \lambda) + \text{cr}(\gamma, \lambda) \leq \text{cr}(\varphi)$.

Specifically, we can define $\Lambda : G \rightarrow \mathcal{H}$ for all vertices and edges in $\lambda^{-1}[u]$ as $(\Gamma|_{D_u})^{-1} \circ \psi_\varepsilon$, and for all edges in $\lambda^{-1}[uv]$ as $(\Gamma|_{D_u \cup D_v})^{-1} \circ \psi_\varepsilon$. Since $\psi_\varepsilon(v) \in N_{\lambda(v)}$ for every $v \in V(G)$, Λ satisfies (P1).

In order to establish (P2), we need to show that every edge $ab \in E(G)$ crosses $\partial N_{\lambda(a)}$ and $\partial N_{\lambda(b)}$ at most once. Suppose for the sake of contradiction that there is an edge $ab \in E(G)$ such that in $\psi_\varepsilon(ab)$ crosses $\partial N_{\lambda(a)}$ or $\partial N_{\lambda(b)}$ at least twice. By the definition of ε_0 , if $\lambda(a) \neq \lambda(b)$, then the centers of $N_{\lambda(a)}$ and $N_{\lambda(b)}$ are at distance at least 4ε . It follows that there exists exactly one connected component of $N_{\lambda(a)\lambda(b)} \cap \psi_\varepsilon(ab)$ joining a point on $\partial N_{\lambda(a)}$ with a point on $\partial N_{\lambda(b)}$. Hence, there exists a connected component α_{ab} of $N_{\lambda(a)\lambda(b)} \cap \psi_\varepsilon(ab)$ joining a pair of points p_1 and p_2 in $\partial N_{\lambda(a)}$ or a pair of such points in $\partial N_{\lambda(b)}$. Without loss of generality, suppose that both p_1 and p_2 , which belong to $\Xi(\psi_\varepsilon)$, are in $\partial N_{\lambda(a)}$.

In what follows we show that we can modify the drawing $\psi_\varepsilon(G)$ so that $|\Xi(\psi_\varepsilon)|$ decreases, which contradicts the choice of ψ_ε . Choose the edge $ab \in E(G)$ and the arc α_{ab} so that they minimize $\|p_1 p_2\|$. We modify $\psi_\varepsilon(G)$ by redrawing the arc α_{ab} as follows. We cut the edge ab in ψ_ε at p_1 and p_2 and reconnect the severed ends by an arc β_{ab} closely following $\partial N_{\lambda(a)}$ in the interior of $N_{\lambda(a)}$. As a result, the cardinality of $|\Xi(\psi_\varepsilon)|$ decreases by 2, and the number of crossings between edges does not increase. Indeed, since β_{ab} closely follows $\partial N_{\lambda(a)}$ and by the minimality of $\|p_1 p_2\|$, every edge $\psi_\varepsilon(e)$, $e \in E(G)$, crosses β_{ab} at most once. Since $\alpha_{ab} \cup \beta_{ab}$ forms a closed Jordan curve, if an edge $\psi_\varepsilon(e)$, $e \in E(G)$, crosses β_{ab} , it must also cross α_{ab} . It follows that β_{ab} has at most as many edge crossings as α_{ab} . This completes the proof of (P2).

Finally, in order to establish (P3), we modify the drawing ψ_ε as follows. We cut the edge $\psi_\varepsilon(ab)$, $ab \in E(G)$, at their intersection with $\partial N_{\lambda(b)}$ and squeeze the severed parts in a small neighborhood of $\psi_\varepsilon(G) \cap N_{\lambda(a)\lambda(b)}$ so that the severed ends belong to $\partial N_{\lambda(a)}$, but the rest is drawn in the interior of $N_{\lambda(a)}$ without introducing new edge crossings. The squeezing is performed so that the clockwise order of the severed ends of edges along $\partial N_{\lambda(a)}$ and $\partial N_{\lambda(b)}$ are reverse of each other. Finally, we reconnect the severed ends by pairwise noncrossing straight-line segments in $N_{\lambda(a)\lambda(b)}$. This is possible by the way how we performed the squeezing. \square

When the Image Graph is a Cycle. In Section 3, we successively modify an instance $\varphi = \gamma \circ \lambda$, while $\text{cr}(\varphi)$ remains invariant, until the image graph H becomes a cycle. We show that in this case it is easy to determine $\text{cr}_2(\gamma, \lambda)$, which is a consequence of the following folklore lemma.

Lemma 2 [18, Lemma 1.12] *If $G = C_n$, $H = C_k$, and $\lambda : G \rightarrow H$ is a simplicial map without spurs, where the cycle G winds around the cycle H precisely n/k times, then $\text{cr}_1(\gamma, \lambda) = \frac{n}{k} - 1$.*

3 Cycles without Spurs

Let $G = C_n$ be a cycle with n vertices, and H an arbitrary abstract graph, $\lambda : G \rightarrow H$ a simplicial map, and $\gamma : H \rightarrow \mathbb{R}^2$ a straight-line drawing such that $\varphi = \gamma \circ \lambda$ has no spurs (or, equivalently, λ does not map any two consecutive edges of C_n to the same edge in H). In this section, we prove that $\text{cr}(\gamma \circ \lambda)$ is invariant under the so-called **ClusterExpansion** and **PipeExpansion** operations. (Similar operations for weak embeddings have been introduced in [1, 7, 9, 14].) We show that a sequence of $O(n)$ operations produces an instance in which H is a cycle, where we can easily determine both $\text{cr}_1(\gamma, \lambda)$ and $\text{cr}_2(\gamma, \lambda)$, hence $\text{cr}(\gamma \circ \lambda)$.

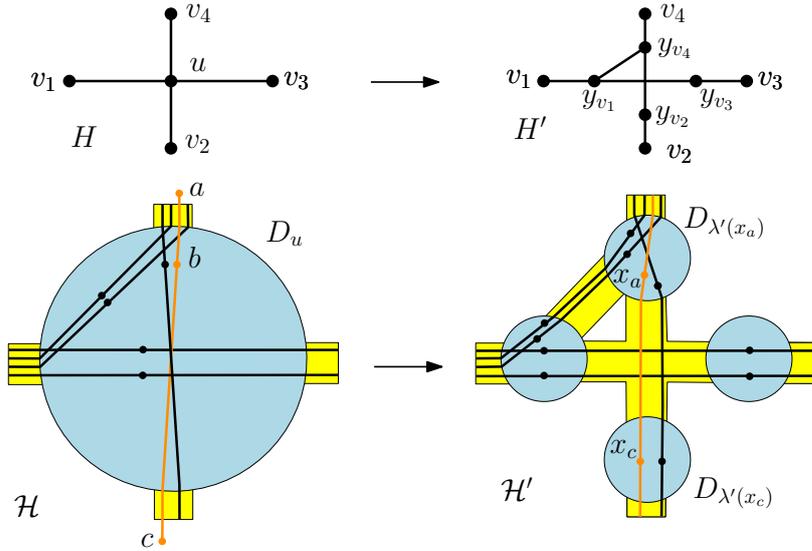


Fig. 5: **ClusterExpansion**(u). Changes in the graph H (top row), and changes in G (bottom row).

The first operation, **ClusterExpansion**(u), modifies $\varphi = \gamma \circ \lambda$ in a small neighborhood N_u of a cluster $u \in V(H)$. Intuitively, every maximal connected subcurve of $\varphi(G)$ in the disk N_u is replaced by a straight-line segment between the two endpoints of the subcurve; essentially applying shortcuts within the disk N_u . The formal definition (below) describes the changes incurred in both γ and λ . See Fig. 5 for an illustration.

ClusterExpansion(u). Input: a straight-line map $\varphi = (\gamma \circ \lambda) : G \rightarrow \mathbb{R}^2$ and a cluster $u \in V(H)$. (1) Let N_u be a sufficiently small disk centered at $\gamma(u)$ in \mathbb{R}^2 that intersects only the images of pipes incident to u . (2) Subdivide every pipe $uv \in E(H)$ incident to u with a new cluster y_v , let $\gamma(y_v) := \partial N_u \cap \gamma(uv)$. (3) Subdivide every edge $ab \in E(G)$ such that

$\lambda(b) = u$ with a new vertex x_a such that $\lambda(x_a) = y_{\lambda(a)}$. (4) For every vertex $b \in \lambda^{-1}[u]$, and its two neighbors x_a and x_c , insert an edge $x_a x_c$ in G , insert a pipe $\lambda(x_a)\lambda(x_c)$ in H if it is not already present, and draw this pipe in the plane as a straight-line segment between $\gamma(\lambda(x_a))$ and $\gamma(\lambda(x_c))$. (5) Delete cluster u from H , and delete all vertices in $\lambda^{-1}[u]$ from G . (6) Return the resulting instance $\varphi' = \gamma' \circ \lambda'$, where $\lambda' : G' \rightarrow H'$ and $\gamma' : H' \rightarrow \mathbb{R}^2$.

Lemma 3 *For every instance $\varphi : G \rightarrow \mathbb{R}^2$ without spurs, where G is a cycle and $u \in V(H)$, $\text{ClusterExpansion}(u)$ produces an instance $\varphi' : G' \rightarrow \mathbb{R}^2$ without spurs, where G' is a cycle, and $\text{cr}(\gamma' \circ \lambda') = \text{cr}(\gamma \circ \lambda)$.*

Proof Since G is a cycle, then every vertex $b \in \lambda^{-1}[u]$ has precisely two neighbors, say a and c . Step 3 subdivides the edges ab and bc with new vertices x_a and x_c ; Step 4 inserts an edge $x_a x_c$, and Step 5 deletes b . Consequently, the path (a, b, c) is replaced by a path (a, x_a, x_c, c) . Since G is a cycle, and the operation replaces edge-disjoint paths by new paths between the same pair of endpoints, $\text{ClusterExpansion}(u)$ returns a cycle.

Since $\gamma \circ \lambda$ has no spur, then for every vertex $b \in \lambda^{-1}[u]$, the neighbors a and c are in distinct clusters, that is, $\lambda(a) \neq \lambda(c)$. Consequently, $y_{\lambda(a)} \neq y_{\lambda(c)}$ and so $\lambda'(x_a) \neq \lambda'(x_c)$. Therefore the operation does not create spurs.

Let $\Lambda : G \rightarrow \mathcal{H}$ be a drawing that attains $\text{cr}_1(\gamma, \lambda)$. We may assume that every connected component of $\Lambda(G) \cap D_u$ is a line segment for every cluster $u \in V(H)$; and similarly every connected component of $\Lambda(G) \cap R_{uv}$ is a line segment for every pipe $uv \in E(H)$.

Let $u \in V(H)$ be a cluster. Assume that (a, b, c) and $(\hat{a}, \hat{b}, \hat{c})$ are paths in G such that $\lambda(b) = \lambda(\hat{b}) = u$. Then $\Lambda(G)$ has two possible types of crossings in D_u between paths (a, b, c) and $(\hat{a}, \hat{b}, \hat{c})$. In the first type, $\lambda(a)$ and $\lambda(c)$ interleave with $\lambda(\hat{a})$ and $\lambda(\hat{c})$ in the rotation at u . In the second type, we have $\{\lambda(a), \lambda(c)\} \cap \{\lambda(\hat{a}), \lambda(\hat{c})\} \neq \emptyset$ (hence $\{\lambda'(x_a), \lambda'(x_c)\} \cap \{\lambda'(x_{\hat{a}}), \lambda'(x_{\hat{c}})\} \neq \emptyset$).

For every cluster $u \in V(H)$, let $\text{CR}_\Lambda^\times(u)$ denote the number of crossings of the first type; and let $\text{CR}_\Lambda^\leq(u)$ denote the number of crossings of the second type. In the following we construct $\Lambda' : G' \rightarrow \mathcal{H}'$ witnessing $\text{cr}(\gamma' \circ \lambda') \leq \text{cr}(\gamma \circ \lambda)$ such that

$$\text{cr}_1(\gamma', \lambda') = \left(\sum_{v \in V(H), v \neq u} \text{CR}_{\Lambda'}(v) \right) + \text{CR}_\Lambda^\leq(u) = \text{cr}_1(\gamma, \lambda) - \text{CR}_\Lambda^\times(u)$$

and

$$\text{cr}_2(\gamma', \lambda') = \text{cr}_2(\gamma, \lambda) + \text{CR}_\Lambda^\times(u).$$

Note that the second condition does not depend on Λ' and follows by the construction of γ' .

Let h denote the natural homeomorphism between $\mathcal{H} \setminus \text{int}(D_u)$ and the connected components of $\mathcal{H}' \setminus \bigcup \{\text{int}(D_{y_v}) : v \in V(H), uv \in E(H)\}$ containing

the disks D_w for all surviving clusters $w \in V(H) \setminus \{u\}$. We put $\Lambda'(st) = h(\Lambda(st))$ for all $st \in E(G)$ where $u \notin \{\lambda(s), \lambda(t)\}$. We define Λ' on every path (a, x_a, x_c, c) in G' that replaced a path (a, b, c) in G , where $\lambda(b) = u$, as follows.

Let $p_{ab} = [\mathcal{H} \setminus \text{int}(D_u)] \cap \Lambda(ab)$ and $p_{bc} = [\mathcal{H} \setminus \text{int}(D_u)] \cap \Lambda(bc)$. We define $\Lambda'(a, x_a)$ as the concatenation of the arc from $h(\Lambda(a))$ to $h(p_{ab})$ contained in $\Lambda(ab)$ and a very short crossing-free line segment contained in $D_{\lambda'(x_a)}$. In the same manner we construct $\Lambda'(x_c, c)$.

Let $(\hat{a}, x_{\hat{a}}, x_{\hat{c}}, \hat{c})$ denote another such path, that is, $(\hat{a}, x_{\hat{a}}, x_{\hat{c}}, \hat{c})$ replaced $(\hat{a}, \hat{b}, \hat{c})$ in G such that $\lambda(\hat{b}) = u$.

We construct $\Lambda'(x_a, x_c)$ as an arc between $\Lambda'(x_a)$ and $\Lambda'(x_c)$ contained in $D_{y_{\lambda(a)}} \cup D_{y_{\lambda(c)}} \cup R_{y_{\lambda(a)}y_{\lambda(c)}} = D_{\lambda'(x_a)} \cup D_{\lambda'(x_c)} \cup R_{\lambda'(x_a)\lambda'(x_c)}$ such that $\Lambda'(x_a, x_c)$ and $\Lambda'(x_{\hat{a}}, x_{\hat{c}})$ cross if and only if p_{ab} and p_{bc} interleave with $p_{\hat{a}\hat{b}}$ and $p_{\hat{b}\hat{c}}$ along ∂D_u . In the case when $\Lambda'(x_a, x_c)$ and $\Lambda'(x_{\hat{a}}, x_{\hat{c}})$ cross, we also ensure that they cross exactly once. This establishes that $\Lambda(a, b, c)$ and $\Lambda(\hat{a}, \hat{b}, \hat{c})$ cross (i.e., they contribute 1 crossing to $\text{CR}_\Lambda^\times(u)$) if and only if $\Lambda'(x_a, x_c)$ and $\Lambda'(x_{\hat{a}}, x_{\hat{c}})$ cross, as desired.

Specifically, we draw $\Lambda'(x_a, x_c)$ as a polygonal line that consists of a line segment in each of $D_{\lambda'(x_a)}$, $D_{\lambda'(x_c)}$, and $R_{\lambda'(x_a)\lambda'(x_c)}$; and denote by $p_{ac}^a \in A_{\lambda'(x_a), \lambda'(x_c)}$ and $p_{ac}^c \in A_{\lambda'(x_c), \lambda'(x_a)}$ the intersection of $\Lambda'(x_a, x_c)$ with $\partial D_{\lambda'(x_a)}$ and $\partial D_{\lambda'(x_c)}$, respectively¹. In order to complete the construction of Λ' , it is enough to specify the order of these intersection points on each arc A_{y_v, y_w} , for all ordered pairs (y_v, y_w) .

Let us fix an arbitrary total order $<$ on the clusters in $V(H') \setminus V(H)$. Let $y_v, y_w \in V(H') \setminus V(H)$ be a pair of clusters, and assume without loss of generality that $y_v < y_w$. Let us arrange the intersection points along the arcs A_{y_v, y_w} and A_{y_w, y_v} such that the paths $\Lambda'(x_a, x_c)$, where $\lambda'(a) = y_v$ and $\lambda'(c) = y_w$, are pairwise noncrossing in both D_{y_v} and R_{y_v, y_w} . These intersection points on A_{y_v, y_w} , in turn, determine all crossings between these arcs in D_{y_w} . In particular a pair of such paths, $\Lambda'(x_a, x_c)$ and $\Lambda'(x_{\hat{a}}, x_{\hat{c}})$, cross in D_{y_w} if and only if p_{ab} and p_{bc} interleave with $p_{\hat{a}\hat{b}}$ and $p_{\hat{b}\hat{c}}$ along ∂D_u . Since both $\Lambda'(x_a, x_c) \cap D_{y_w}$ and $\Lambda'(x_{\hat{a}}, x_{\hat{c}}) \cap D_{y_w}$ are line segments, they cross at most once.

For any other pair of arcs, (a, b, c) and $(\hat{a}, \hat{b}, \hat{c})$, such that $\lambda(b) = \lambda(\hat{b})$ but $\{\lambda(a), \lambda(c)\} \neq \{\lambda(\hat{a}), \lambda(\hat{c})\}$, it is straightforward to check that $\Lambda(a, b, c)$ and $\Lambda(\hat{a}, \hat{b}, \hat{c})$ cross, if and only if $\Lambda'(x_a, x_c)$ and $\Lambda'(x_{\hat{a}}, x_{\hat{c}})$ cross.

To establish the other direction, we can start with a drawing $\Lambda' : G' \rightarrow \mathcal{H}'$ witnessing $\text{cr}(\gamma' \circ \lambda')$ and apply the inverse of h to construct Λ in $\mathcal{H} \setminus D_u$. Finally, it is enough to observe that the order of intersection points p_{ab} along ∂D_u specifies λ for which

$$\sum_{v \in V(H)} \text{CR}_\Lambda(v) \leq \text{cr}_1(\gamma', \lambda') - \text{CR}_\Lambda^\times(u)$$

and

$$\text{cr}_2(\gamma', \lambda') = \text{cr}_2(\lambda) + \text{CR}_\Lambda^\times(u),$$

¹ Recall the definition of $A_{u,v}$ from Section 2.

and that concludes the proof. \square

We remark that $\text{cr}(\gamma \circ \lambda)$ is invariant under the $\text{ClusterExpansion}(u)$ operation even in the presence of spurs, however the proof is somewhat simpler in the absence spurs, and Lemma 3 also establishes that $\text{ClusterExpansion}(u)$ does not create new spurs.

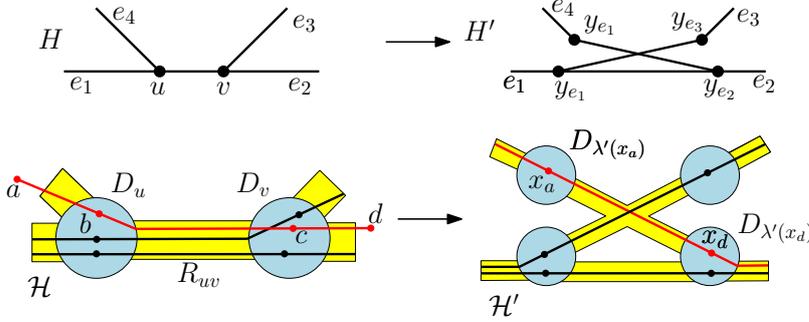


Fig. 6: $\text{PipeExpansion}(uv)$ for a safe pipe uv . Changes in the graph H (top row), and changes in G (bottom row).

Pipe Expansion. A cluster $u \in V(H)$ is a **base** of an incident pipe uv if every vertex in $\lambda^{-1}[u]$ is incident to an edge in $\lambda^{-1}[uv]$. A pipe $uv \in E(H)$ is **safe** if both u and v are bases of uv . The following operation is defined on safe pipes. (We note that our algorithm would be correct even if $\text{PipeExpansion}(uv)$ were defined on all pipes, unlike the result in [2], since λ does not contain spurs. We restrict this operation to safe pipe to simplify the runtime analysis.)

The second operation, $\text{PipeExpansion}(uv)$, is very similar to $\text{ClusterExpansion}(u)$. Instead of creating shortcuts within a disk N_u , it uses an ellipse L_{uv} that encloses the drawing of the pipe uv . The operation replaces every maximal connected subcurve of $\varphi(G)$ in the ellipse L_{uv} with a straight-line segment inside the ellipse. The formal definition below specifies the changes in both γ and λ . See Fig. 5 for an illustration.

$\text{PipeExpansion}(uv)$. Input: a straight-line map $\varphi = (\gamma \circ \lambda) : G \rightarrow \mathbb{R}^2$, and a safe pipe $uv \in E(H)$. (1) Let L_{uv} be a sufficiently narrow ellipse with foci at $\gamma(u)$ and $\gamma(v)$ that intersects only the images of pipes incident to u and v . (2) Subdivide every pipe $e \in E(H)$ incident to u or v with a new cluster y_e , let $\gamma(y_e) := \partial L_{uv} \cap \gamma(e)$. (3) Subdivide every edge $ab \in E(G)$ such that $\lambda(a) \notin \{u, v\}$ and $\lambda(b) \in \{u, v\}$ with a new vertex x_a such that $\lambda(x_a) = y_{\lambda(ab)}$. (4) For every edge $bc \in \lambda^{-1}[uv]$, and the two neighbors x_a and x_d of b and c , respectively, insert an edge $x_a x_d$ in G , insert a pipe $\lambda(x_a)\lambda(x_d)$ in H if it is not already present, and draw this pipe in the plane as a straight-line segment between $\gamma(\lambda(x_a))$ and $\gamma(\lambda(x_d))$. (5) Delete clusters u and v from H , and delete all vertices in

$\lambda^{-1}[uv]$ from G . (6) Return the resulting instance $\varphi' = \gamma' \circ \lambda$, where $\lambda' : G' \rightarrow H'$ and $\gamma' : H' \rightarrow \mathbb{R}^2$.

Lemma 4 *If G is a cycle, $\lambda : G \rightarrow H$ has no spur, and $uv \in E(H)$ is a safe pipe, then $\text{PipeExpansion}(uv)$ produces an instance where G' is a cycle, $\lambda' : G' \rightarrow H'$ has no spur, and $\text{cr}(\gamma' \circ \lambda) = \text{cr}(\gamma' \circ \lambda')$.*

The proof is analogous to the proof of Lemma 3. We only point out the differences.

Proof Since uv is a safe pipe, then every vertex in $\lambda^{-1}[u] \cup \lambda^{-1}[v]$ is incident to one edge in $\lambda^{-1}[uv]$. Since G is a cycle and φ has no spurs, every vertex in $\lambda^{-1}[u] \cup \lambda^{-1}[v]$ is incident to two edges, precisely one which is in $\lambda^{-1}[u] \cup \lambda^{-1}[v]$. That is, every vertex in $\lambda^{-1}[u] \cup \lambda^{-1}[v]$ is in some path (a, b, c, d) in which $\lambda(b) = u$ and $\lambda(c) = v$; and for every path (a, b, c, d) in G such that $\lambda(b) = u$ and $\lambda(c) = v$, we have $\lambda(ab) \neq uv$ and $\lambda(cd) \neq uv$. Step 3 subdivides the edges ab and cd with new vertices x_a and x_d ; Step 4 inserts an edge $x_a x_d$, and Step 5 deletes b and c . Consequently, the path (a, b, c, d) is replaced by a path (a, x_a, x_d, d) . Since G is a cycle, and the operation replaces edge-disjoint paths are replaced by new paths between the same pair of endpoints, $\text{PipeExpansion}(uv)$ returns a cycle.

Since $\gamma \circ \lambda$ has no spur, then for every vertex $b \in \lambda^{-1}[u] \cup \lambda^{-1}[v]$, the neighbors a and c are in distinct clusters, that is, $\lambda(a) \neq \lambda(c)$. Consequently, $y_{\lambda(a)} \neq y_{\lambda(c)}$ and so $\lambda'(x_a) \neq \lambda'(x_c)$. Therefore the operation does not create spurs.

The proof that $\text{cr}(\gamma' \circ \lambda') = \text{cr}(\gamma' \circ \lambda)$ is analogous to the proof of Lemma 3. The only difference is that $\text{CR}_\lambda^\times(u)$ and $\text{CR}_\lambda^<(u)$ (and the two types of crossings) are defined in terms of the cyclic order along ∂L_{uv} , rather than the cyclic order along ∂N_u (i.e., the rotation of u). We omit the details. \square

Main Algorithm. Given an instance $\lambda : G \rightarrow H$ and $\gamma : H \rightarrow \mathbb{R}^2$, we apply the two operations defined above as follows.

Algorithm 1. Input: (G, H, λ, γ)
 $U_0 \leftarrow V(H)$
for every $u \in U_0$ **do**
 | $\text{ClusterExpansion}(u)$
while there is a safe pipe $uv \in E(H)$ such that $\deg_H(u) \geq 3$ or $\deg_H(v) \geq 3$ **do**
 | $\text{PipeExpansion}(uv)$
 $uv \leftarrow$ an arbitrary edge in $E(H)$.
return $\text{cr}_2(\gamma, \lambda) + |\lambda^{-1}[uv]| - 1$.

Lemma 5 *Algorithm 1 terminates.*

Proof By Lemmas 3 and 4, $\lambda : G \rightarrow H$ has no spurs in any step of the algorithm. It is enough to show that the while loop of Algorithm 1 terminates.

We define the potential function $\Phi(G, H) = |E(G)| - |E(H)|$, and show that $\Phi(G, H) \geq 0$ and it decreases in every invocation of $\text{PipeExpansion}(uv)$. Since G is a cycle and λ has no spur, every edge in $\lambda^{-1}[uv]$ is adjacent to one edge in some other pipe incident to u and one edge in some other pipe incident to v . Each of these edges contributes to one edge in $E(G')$ inside the ellipse D_{uv} . Since uv is safe, G' has no other new edges. Consequently, $|E(G')| = |E(G)|$. Since $\deg_H(u) \geq 3$ or $\deg_H(v) \geq 3$, $\text{PipeExpansion}(uv)$ replaces the clusters u and v with at least 3 clusters, each of which is incident to at least one pipe in the ellipse D_{uv} . Consequently, $|E(H')| > |E(H)|$, and so $\Phi(G, H) > \Phi(G', H')$, as claimed. \square

Lemma 6 *At the end of the while loop of Algorithm 1, H is a cycle.*

Proof It is enough to show that if H is not a cycle in the while loop of Algorithm 1, then there is a safe pipe $uv \in E(H)$ such that $\deg_H(u) \geq 3$ or $\deg_H(v) \geq 3$. Note that in the entire course of the while loop, every cluster in $V(H)$ has been created by a previous $\text{ClusterExpansion}(u)$ or $\text{PipeExpansion}(uv)$ operation. Observe that every cluster created by $\text{ClusterExpansion}(u)$ (resp., $\text{PipeExpansion}(uv)$) is a base for the unique incident pipe in the exterior of the disk D_u (resp., ellipse D_{uv}). Let $s : V(H) \rightarrow E(H)$ be a function that maps every cluster to that incident pipe. Recall that the input does not have spurs, and no spurs are created in the algorithm by Lemmas 3 and 4. In the absence of spurs, the minimum degree in H is at least 2, and if $u \in V(H)$ and $\deg_H(u) = 2$, then u is a base for both incident pipes.

Assume that in some step of the while loop, H is not a cycle. Let $v_1 \in V(H)$ be an arbitrary cluster such that $\deg_H(v_1) \geq 3$. Construct a maximal simple path $(v_1, v_2, \dots, v_\ell)$ incrementally such that $s(v_i) = v_i v_{i+1}$ for $i = 1, 2, \dots, \ell$. If for some $j \geq 2$, $s(v_j) = s(v_{j-1}) = v_j v_{j-1}$, then the pipe $v_{j-1} v_j$ is safe, and we are done. Similarly, if $\deg_H(v_j) = 2$, then $v_j v_{j-1}$ is safe, and we are also done.

Otherwise, the path ends with a repeated cluster: $s(v_\ell) = v_\ell v_i$, for some $1 \leq i < \ell - 1$, and so we obtain a cycle $(v_i, v_{i+1}, \dots, v_\ell)$ of at least 3 vertices. Let v_j , $i \leq j \leq \ell$, be a cluster created by the algorithm in the most recent invocation of a $\text{ClusterExpansion}(u)$ or $\text{PipeExpansion}(uv)$ operation among v_i, \dots, v_ℓ . Then $s(v_j)$ is a pipe in the exterior of a disk D_u or an ellipse D_{uv} . Hence, the pipe $v_{j-1} v_j$ or $v_\ell v_i$ if $j > i$ or $j = i$, respectively, is in the interior of D_u or D_{uv} . Without loss of generality suppose the former. It follows that v_j and v_{j-1} were created by the same invocation of $\text{ClusterExpansion}(u)$ or $\text{PipeExpansion}(uv)$. However, this implies that $s(v_{j-1}) \neq v_{j-1} v_j$, contradicting the assumption that $(v_i, v_{i+1}, \dots, v_\ell)$ is a cycle.

We conclude that the path v_1, \dots, v_ℓ contains a safe pipe. \square

Lemma 7 *Algorithm 1 returns $\text{cr}(\gamma \circ \lambda)$.*

Proof By Lemma 1, $\text{cr}(\gamma \circ \lambda) = \text{cr}_1(\gamma, \lambda) + \text{cr}_2(\gamma, \lambda)$. Here $\text{cr}_2(\gamma, \lambda)$ can be computed by a line sweep of the drawing $\gamma(H)$. By Lemmas 2 and 6, at the end of the algorithm, $\text{cr}_1(\gamma, \lambda) = |\lambda^{-1}[uv]| - 1$ for an arbitrary edge $uv \in E(H)$. By Lemmas 3 and 4, $\text{cr}(\gamma \circ \lambda)$ is invariant under the operations, so the algorithm reports $\text{cr}(\gamma \circ \lambda)$ for the input instance. \square

Running Time. The efficient implementation of our algorithm relies on the following data structures. For every cluster $u \in V(H)$ we maintain the set of vertices of $V(G)$ in $\lambda^{-1}[u]$. For every pipe $uv \in E(H)$, we maintain $\lambda^{-1}[uv] \subset E(G)$, the weight $w(uv) = |\lambda^{-1}[uv]|$, and the sum of weights of all pipes that cross uv , that we denote by $W(uv)$. Then we have

$$\text{cr}_2(\gamma, \lambda) = \frac{1}{2} \sum_{uv \in E(H)} w(uv)W(uv).$$

We also maintain the current value of $\text{cr}_2(\gamma, \lambda)$. We further maintain indicator variables that support checking the conditions of the while loop in Algorithm 1: (i) whether the cluster is a base for the pipe, (ii) whether a cluster has degree 2, and (iii) whether a pipe is safe.

Lemma 8 *With the above data structures, Algorithm 1 runs in $O((M+R) \log M)$ time, where $M = |E(H)| + |E(G)|$ and $R = \text{cr}(\gamma \circ \lambda) < M^2$.*

Proof At preprocessing, we can compute $\lambda^{-1}[u]$, $\lambda^{-1}[uv]$, and $w(uv)$ by a simple traversal of G in $O(|E(G)|)$ time. Since every crossing in the drawing $\gamma(H)$ corresponds to at least one crossing in any perturbation, $\gamma(H)$ has at most R crossings. Hence the complexity of the arrangement of all edges in $\gamma(H)$ is $O(M+R)$. A standard line sweep algorithm can find all crossings of $\gamma(H)$ in $O((M+R) \log(M+R)) = O((M+R) \log M)$ time. The same algorithm can also compute $W(uv)$ for all $uv \in E(H)$, and $\text{cr}_2(\gamma, \lambda)$.

Algorithm 1 starts with a for-loop over all $u \in U_0$. We can update $\lambda^{-1}[u]$, $\lambda^{-1}[uv]$, and $w(uv)$ in $O(\deg_H(u) + |\lambda^{-1}(u)|)$ time per $\text{ClusterExpansion}(u)$. This sums to $O(|E(H)| + |E(G)|)$ time for all $u \in U_0$. All new crossings in $\gamma(H)$ occur between the pipes created in the interior of the disks D_u , for all $u \in U_0$. These crossings can be found in $O((M+R) \log M)$ total time.

Note also that $\text{ClusterExpansion}(u)$, for all $u \in U_0$ doubles the number of edges in G . However, $|E(G)|$ is invariant under PipeExpansion operations. In fact, $\text{PipeExpansion}(uv)$ partitions the set $\lambda^{-1}[uv] \subset E(G)$ into two or more subsets, which are mapped to pipes in the ellipse D_{uv} , and the $\lambda^{-1}(e)$ for every other pipe $e \in E(H)$ remains unchanged. We maintain $\lambda^{-1}[u]$, $\lambda^{-1}[v]$, $\lambda^{-1}[uv]$, and $w(uv)$ in the while loop of Algorithm 1 using a heavy-path decomposition. Suppose $\text{PipeExpansion}(uv)$ replaces uv with pipes u_1v_1, \dots, u_kv_k , which correspond to pairs of clusters in the neighborhood of u and v , respectively. The naive implementation would take $O(w(uv))$ time, but we can reduce it to $O(w(uv) - \max_i w(u_iv_i))$: Put $S = \lambda^{-1}[uv]$ and compute the sets $\lambda^{-1}[u_iv_i]$ incrementally in parallel by deleting edges from S ; when all but maximal set has been computed, then all remaining elements of S can be added to this maximal set in $O(1)$ time. The time $O(w(uv) - \max_i w(u_iv_i))$ can then be charged to the edges that move from $\lambda^{-1}[uv]$ to a set $\lambda^{-1}[u_iv_i]$ with $w(u_iv_i) \leq w(uv)/2$. Over all operations of the while loop of Algorithm 1, edges that are initially mapped to a pipe of weight w receive a charge of at most $O(\sum_{i=0}^{\infty} 2^i \lfloor w/2^i \rfloor) = O(w \log w)$. Summation over all edges of $E(H)$ yields

$$O\left(\sum_{uv \in E(H)} w(uv) \log w(uv)\right) \leq O(|E(G)| \log |E(G)|) = O(M \log M).$$

When $\text{PipeExpansion}(uv)$ replaces a pipe uv with new pipes u_1v_1, \dots, u_kv_k , then every pipe that crossed uv will cross u_1v_1, \dots, u_kv_k . So $W(u_iv_i)$, $i = 1, \dots, k$, can be computed by adding the number of *new* crossings to $W(uv)$. All new crossings created by $\text{PipeExpansion}(uv)$ are between new pipes in the ellipse D_{uv} . Since pipe crossings are never removed, the total number of such pipe crossings is at most R , and they can be computed in $O((M + R) \log M)$ time over all operations of the while loop of Algorithm 1.

At the end of the algorithm, both $\text{cr}_1(\gamma, \lambda) = w(uv) - 1$ for an arbitrary pipe $uv \in E(H)$, and $\text{cr}_2(\gamma, \lambda) = \frac{1}{2} \sum_{uv \in E(H)} w(uv)W(uv)$ can be calculated in $O(M)$ time. \square

4 NP-Completeness in the Presence of Spurs

In this section, we prove Theorem 2. In a problem instance, we are given a simplicial map $\lambda : G \rightarrow H$, a straight-line drawing $\gamma : H \rightarrow \mathbb{R}^2$, and a nonnegative integer K , and ask whether $\text{cr}(\gamma \circ \lambda) \leq K$.

Lemma 9 *The above problem is in NP.*

Proof A feasible drawing $\Gamma \circ \Lambda : G \rightarrow \mathbb{R}^2$ with $\text{cr}(\Gamma \circ \Lambda) \leq K$ can be witnessed by a combinatorial representation of Λ . Specifically, we can determine $\text{cr}_2(\gamma, \lambda)$ by computing the weight of each pipe $uv \in E(H)$ in $O(|E(G)| + |E(H)|)$ time, and finding all edge-crossings in the drawing $\gamma(H)$ in $O(|E(H)| \log |E(H)|)$ time. Given a combinatorial representation of a drawing $\Lambda : G \rightarrow \mathcal{H}$, we can determine the number of crossings at all nodes $u \in V(H)$ in $O(\sum_{u \in V(H)} |\lambda^{-1}[u]|) = O(|E(G)|)$ time. \square

We prove NP-hardness by a reduction from 3SAT. Let Φ be a Boolean formula in 3CNF with a set $\mathcal{X} = \{x_1, \dots, x_n\}$ of variables and a set $\mathcal{C} = \{c_1, \dots, c_m\}$ of clauses. We construct graphs G and H , a simplicial map $\lambda : G \rightarrow H$, a straight-line drawing $\gamma : H \rightarrow \mathbb{R}^2$, and an integer $K \in \mathbb{N}$ such that $\text{cr}(\gamma \circ \lambda) \leq K$ if and only if Φ is satisfiable.

We present the construction of G , H , λ , and γ for Theorem 2(2), where G is a disjoint union of paths, in Section 4.1. In Section 4.2, we describe a slight modification of the construction for Theorem 2(1), where G is a cycle.

4.1 First Construction: Disjoint Union of Paths

We start with an overview of the construction that highlights the key ideas and then continue with the details. Let P be a path. A map $\varphi : P \rightarrow \mathbb{R}^2$ is

a **crimp** if P is a concatenation of 3 paths P_1 , P_2 , and P_3 of equal lengths in the given order such that $\varphi(P_1) = \varphi(P_2) = \varphi(P_3)$, and $\varphi(P_i)$ is injective for $i = 1, 2, 3$. An crossing-free perturbation of φ can be loosely regarded to have the shape of the letter Z or its mirror image. We encode the truth value of a Boolean variable by the two possible embeddings, or equivalently by the above-below relationship between P_1 and P_3 ; see Fig. 7(left).

Let $\varphi : G \rightarrow \mathbb{R}^2$ and let P_0 be a path in G with internal vertices of degree 2 such that $\varphi(P_0)$ is injective. We define the operation on φ and G , called the **crumpling** of P_0 , that results in $\varphi' : G' \rightarrow \mathbb{R}^2$ such that G' is obtained from G by tripling the length of P_0 thereby turning P_0 into a path P , and φ' is obtained from φ as follows. We set $\varphi'(v) = \varphi(v)$ for all $V \in V(G)$ outside the interior of P , and $\varphi'|_P$ is a crimp such that $\varphi'(P) = \varphi(P_0)$.

It is not hard to see that given a straight-line map $\varphi : G \rightarrow \mathbb{R}^2$ and a path P_0 in G with internal vertices of degree 2 such that $\varphi(P_0)$ is injective, then the crumpling of P_0 does not require crossings as long as φ is a weak embedding, that is, $\text{cr}(\varphi) = 0$ implies $\text{cr}(\varphi') = 0$. Our construction is based on the fact that if $\text{cr}(\varphi) > 0$, then the crumpling of P_0 may increase the number of crossings.

Roughly speaking, in the reduction we model each Boolean variable x of a 3SAT formula by a path in G obtained as follows. First, we introduce a crimp in a path thereby obtaining a path P^x as above consisting of three subpaths P_1^x , P_2^x , and P_3^x of equal length (Fig. 7, left). Second, for every occurrence of the variable x in a clause c we introduce a crimp in a short (crimp-free) subpath of both P_1^x and P_3^x ; offset by one unit. Then each clause c is modeled by a short path Q on four vertices that crosses the crimps on P_1^x and P_3^x corresponding to the occurrences of variables x in c (Fig. 7, right). We show that the path Q creates either 3 or 5 crossings with the path P^x in every perturbation of φ , and 3 crossings are possible if and only if the perturbation corresponds to a truth assignment in which at least one literal in c has positive truth value.

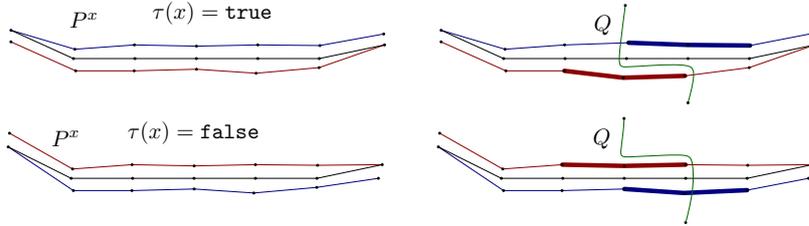


Fig. 7: Left: two crossing-free perturbations of a crimp P^x that encodes the truth value of x . Right: Two additional crimps on P_1^x and P_3^x (thick subpaths); the path Q has either 3 or 5 crossings with P^x , depending on the truth value of x .

The formal argument follows. We first define H and its straight-line drawing $\gamma : H \rightarrow \mathbb{R}^2$; and then define G and the simplicial map $\lambda : G \rightarrow H$.

Construction of H and $\gamma : H \rightarrow \mathbb{R}^2$. For every variable $x \in \mathcal{X}$, create a path of $5m + 3$ vertices, denoted $H_x = (u_3^x, u_4^x, \dots, u_{5m+5}^x)$; see Fig. 8.

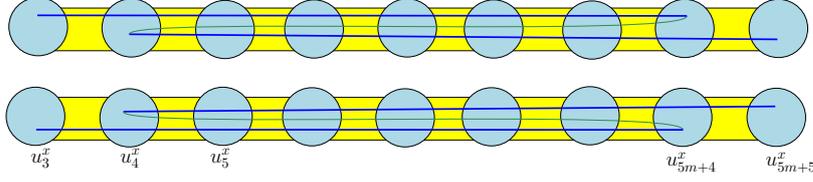


Fig. 8: Two embeddings of G_x for a Boolean variable $x \in \mathcal{X}$. Top: P_1^x is above P_3^x . Bottom: P_1^x is below P_3^x .

For $i = 1, \dots, m$, clause $c_i \in \mathcal{C}$ is associated to at most three (negated or nonnegated) variables. If c_i is associated with variables $x, y, z \in \mathcal{X}$, then we identify the clusters $u_{5i+\ell}^x, u_{5i+\ell}^y, u_{5i+\ell}^z$ and denote the resulting cluster by $u_{5i+\ell}$ for $\ell = 0, 1, 2, 3$. Add two new clusters v_i and w_i , and two new pipes $v_i u_{5i+1}^x$ and $w_i u_{5i+2}^x$. This completes the description of H .

We now describe a straight-line drawing $\gamma : H \rightarrow \mathbb{R}^2$; see Fig. 9. For $i = 1, \dots, m$, map the clusters u_{5i}, \dots, u_{5i+3} (associated with clauses) to integer points $(5i, 0), \dots, (5i + 3, 0)$ on the x -axis, respectively. For two additional clusters, v_i and w_i , are mapped to points $\gamma(v_i) = (5i + 1, 1)$ and $\gamma(w_i) = (5i + 2, -1)$, above and below the x -axis. The remaining clusters of each path H_{x_j} , $x_j \in \mathcal{X}$, are mapped to integer points in the horizontal line $y = j + 1$ for $j = 1, \dots, n$. Specifically, $\gamma(u_i^{x_j}) = (i, j + 1)$, for $3 \leq i \leq 5m + 5$, except for clusters $u_i^{x_j}$ that have been associated with clauses.

Observation 1 For every $x \in \mathcal{X}$, $\gamma(H_x)$ is an x -monotone polygonal path in the plane. Consequently, if $c_i \in \mathcal{C}$ contains variables x, y , and z , then the pipes of H_x, H_y , and H_z that enter u_{5i} and exit u_{5i+3} appear in reverse ccw order in the rotation of u_{5i} and u_{5i+3} , respectively.

Construction of G and $\lambda : G \rightarrow H$. For each clause $c_i \in \mathcal{C}$, create a path G_i of 4 vertices, and map it to the path $(v_i, u_{5i+1}, u_{5i+2}, w_i)$ in H . For each variable $x \in \mathcal{X}$, create a path G_x as follows. First, create a crimp composed of a path of $15m + 5$ vertices as a concatenation of three paths: P_1^x, P_2^x , and P_3^x , which are mapped by λ to $(u_3^x, \dots, u_{5m+4}^x), (u_{5m+4}^x, \dots, u_4^x)$, and $(u_4^x, \dots, u_{5m+5}^x)$, respectively. We shall further crumple subpaths of P_1^x and P_3^x . Regardless of these local crumplings, in every embedding of G_x , the path P_2^x lies between P_1^x and P_3^x . The truth value of variable x is encoded by the above-below relationship between P_1^x and P_3^x (Fig. 8).

Each pair $(x, c_i) \in \mathcal{X} \times \mathcal{C}$, where a literal x or \bar{x} appears in c_i , corresponds to the subpath $(u_{5i}, \dots, u_{5i+3})$ of H_x . If c_i contains the nonnegated x , then crumple the subpath of P_1^x mapped by λ to $(u_{5i+1}, u_{5i+2}, u_{5i+3})$ and the subpath of P_3^x mapped to $(u_{5i}, u_{5i+1}, u_{5i+2})$. Otherwise (if c_i contains the negated \bar{x}), then crumple the subpath of P_3^x mapped by λ to $(u_{5i+1}, u_{5i+2}, u_{5i+3})$ and

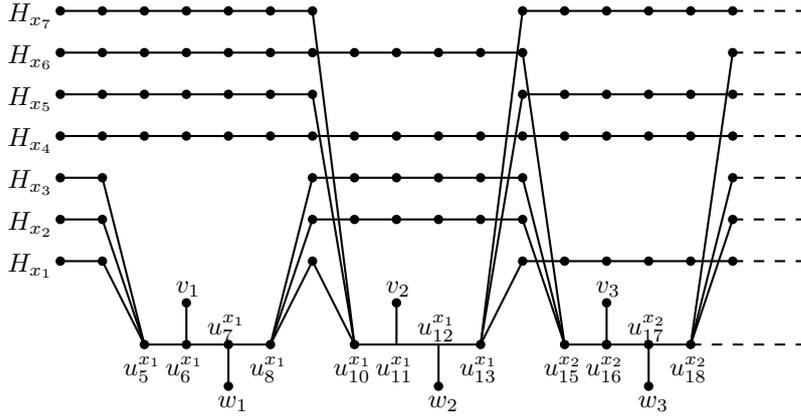


Fig. 9: A bottom left part of a straight-line drawing γ of H of the NP-hardness reduction corresponding to variables x_1, \dots, x_7 and clauses c_1, c_2 and c_3 . The truth value of the clause c_1 depends on the variables x_1, x_2, x_3 , the truth value of c_2 on x_1, x_5, x_7 , and the truth value of c_3 on x_2, x_3, x_6 .

the subpath of P_x^1 mapped to $(u_{5i}, u_{5i+1}, u_{5i+2})$. This completes the definition of G . See Fig. 10 for an illustration.

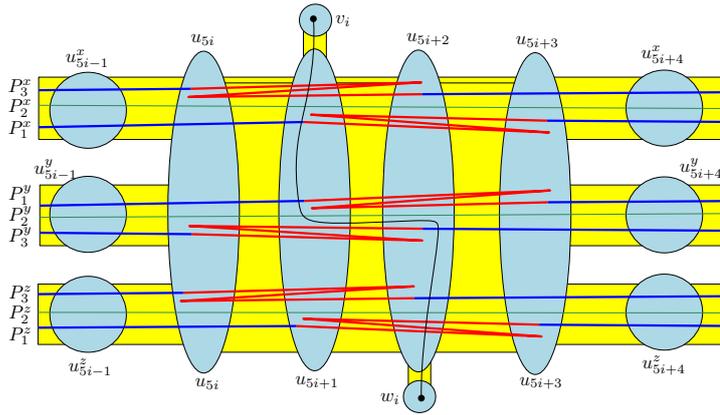


Fig. 10: A clause gadget for $c_i = (x \vee y \vee z)$, where $\tau(x) = \tau(z) = \text{false}$ and $\tau(y) = \text{true}$. The neighborhood of the four middle “vertically prolonged” clusters and pipes between them forms \mathcal{N}_i .

The maps $\lambda : G \rightarrow H$ and $\gamma : H \rightarrow \mathbb{R}^2$ determine $\text{cr}_2(\gamma, \lambda)$. Finally, let $K = \text{cr}_2(\gamma, \lambda) + 13m$. This completes the construction for an instance corresponding to the instance Φ of 3SAT.

Note that G and H have $O(mn)$ vertices and edges, and the drawing γ maps the clusters in $V(H)$ to integer points in an $O(m) \times O(n)$ grid. In particular, the size of the instance is polynomial in m and n .

Equivalence. First, we show that if the Boolean formula Φ is satisfiable, then $\text{cr}(\gamma, \lambda) \leq K$. Assume that Φ is satisfiable, and let $\tau : \mathcal{X} \rightarrow \{\text{true}, \text{false}\}$ be a satisfying truth assignment. Fix $\varepsilon \in (0, \varepsilon_0)$. For every $x \in \mathcal{X}$, denote by \mathcal{N}_x the union of disks N_u and regions N_{uv} for all clusters $v \in V(H_x)$ and pipes $uv \in E(H_x)$; and similarly for every $i = 1, \dots, m$ let \mathcal{N}_i be the union of such regions for the path $H_i = (u_{5i+1}, \dots, u_{5i+3})$ associated with clause c_i in H . For every $x \in \mathcal{X}$, incrementally, embed the path G_x in \mathcal{N}_x as follows: each edge is an x -monotone Jordan arc; if $\tau(x) = \text{true}$, then P_1^x lies above P_3^x ; otherwise P_3^x lies above P_1^x . If a clause c_i contains variables $x, y, z \in \mathcal{X}$, we also ensure that the embeddings of $G_x, G_y,$ and G_z are pairwise disjoint within \mathcal{N}_i . This is possible by Observation 1. Finally, for $i = 1, \dots, m$, embed the path G_i as follows. Assume that c_i contains the variables $x, y, z \in \mathcal{X}$, where x corresponds to a true literal in c_i . Then $\Gamma(G_i)$ starts from $\gamma(v_i)$ along the vertical line $x = 5i + 1$ until it crosses the arc $\Gamma(P_2^x)$, then follows $\Gamma(P_2^x)$ to the vertical line $x = 5i + 2$, and continues to $\gamma(w_i)$ along that line. Note that $\Gamma(P_2^x)$ crosses only 3 edges in $\Gamma(G_x)$, and 5 edges in $\Gamma(G_y)$ and $\Gamma(G_z)$. So there are 13 crossings in \mathcal{N}_i for $i = 1, \dots, m$; and the total number of crossings is $\text{cr}_2(\gamma, \lambda) + 13m$, as required.

Second, we show that if $\text{cr}(\gamma, \lambda) \leq K$, then Φ is satisfiable by constructing a satisfying truth assignment. Consider functions $\Lambda : G \rightarrow \mathcal{H}$ and $\Gamma : \mathcal{H} \rightarrow \mathbb{R}^2$ such that $\Gamma \circ \Lambda : G \rightarrow \mathbb{R}^2$ is a drawing in which $\text{cr}(\Gamma \circ \Lambda) \leq K$. Note that $\text{cr}_2(\gamma, \lambda)$ crossings are unavoidable due to edge-crossings in the drawing $\gamma(H)$. Hence, by the definition of K , there are at most $13m$ crossings in the neighborhoods of clusters. We show that (1) there must be precisely 13 crossings in the neighborhood \mathcal{N}_i of each clause c_i , (2) $\Gamma \circ \Lambda(G_x)$ is an embedding for every variable $x \in \mathcal{X}$, and (3) the embeddings of G_x , for all $x \in \mathcal{X}$, jointly encode a satisfying truth assignment for Φ . Properties (1) and (2) are established by the following lemma.

Lemma 10 *Let $i \in \{1, \dots, m\}$ and let $x, y, z \in \mathcal{X}$ be the three variables in c_i . In $\Gamma \circ \Lambda$, there are at least 13 crossings in the neighborhood \mathcal{N}_i , and equality is possible only if none of the drawings $\Gamma \circ \Lambda(G_x)$, $x \in \mathcal{X}$, has self-crossings in \mathcal{N}_i , and at least one of $G_x, G_y,$ and G_z is crossed exactly 3 times by G_i .*

Proof Let $i \in \{1, \dots, m\}$, and assume that clause c_i contains the variables $x, y, z \in \mathcal{X}$. Recall that c_i is associated with the path $H_i = (u_{5i}, \dots, u_{5i+3})$ in H . We may assume without loss of generality that the ccw order of the neighbors of u_{5i} in $\gamma(H)$ is $(u_{5i-1}^x, u_{5i-1}^y, u_{5i-1}^z, u_{5i+1})$.

Each of the graphs $G_x, G_y,$ and G_z have 3 vertex disjoint connected subgraphs in $\lambda^{-1}[H_i]$. Due to the rotation of cluster u_{5i+1} and u_{5i+2} , the path G_i has to cross each of them, which yields at least 3 crossings in \mathcal{N}_i with each of $G_x, G_y,$ and G_z . Furthermore, $G_x, G_y,$ and G_z each has 6 components

(each of which is formed by a single vertex) in $\lambda^{-1}[u_{5i+1}]$ (resp., $\lambda^{-1}[u_{5i+2}]$), and 5 of these components are adjacent to vertices in both adjacent clusters of H_i ; we call these **through** components. For each of G_x , G_y , and G_z , there exist altogether exactly 7 edges incident to these subgraphs (vertices) in $\lambda^{-1}[u_{5i+1}u_{5i+2}]$. Note that G_i has only one edge in $\lambda^{-1}[u_{5i+1}u_{5i+2}]$, which we denote by e_i .

Without loss of generality we assume that all the edge crossings of G_i with G_x , G_y , and G_z in the drawing $\Gamma \circ \Lambda$ occur along e_i , and outside of $N_{u_{5i+1}u_{5i+2}}$ by (P3). Consequently, the drawing $\Gamma \circ \Lambda$ defines a total “top to bottom” order of the $7 \cdot 3 + 1 = 22$ edges in $\lambda^{-1}[u_{5i+1}u_{5i+2}]$; given by the order of the intersection points of edges of G along $N_{u_{5i+1}u_{5i+2}} \cap N_{u_{5i+2}}$. Let I_x , I_y , and I_z be the minimum intervals in this order spanned by the edges of $\lambda^{-1}[u_{5i+1}u_{5i+2}]$ in G_x , G_y , and G_z , respectively. If the edge e_i is above (resp., below) all the 7 edges of G_x in $\lambda^{-1}[u_{5i+1}u_{5i+2}]$, then it creates at least 5 crossings with the edges incident to the 5 through components in $N_{u_{5i+1}}$ (resp., $N_{u_{5i+2}}$). Analogous statements hold for G_y and G_z , as well. That is, if e_i is not in I_x (resp., I_y and I_z), then G_i crosses G_x (resp., G_y and G_z) at least 5 times in \mathcal{N}_i .

We distinguish several cases based on the relative positions of the intervals I_x , I_y , and I_z . If I_x , I_y , and I_z are pairwise disjoint, then e_i lies in at most one of these intervals, and G_i crosses G_x , G_y , and G_z altogether at least $3 + 5 + 5 = 13$ times. If e_i lies in exactly two of these intervals, say I_x and I_y , then there are at least 2 crossings between G_x and G_y in \mathcal{N}_i , and G_i crosses G_x , G_y , and G_z at least $3 + 3 + 5 = 11$ times. Hence, altogether there exist at least $11 + 2 = 13$ crossings in this case. Finally, if e_i lies in all three intervals, then there must be at least 6 crossings between G_x , G_y , and G_z in \mathcal{N}_i , that is, 2 between each pair. Furthermore, G_i crosses G_x , G_y , and G_z altogether at least $3 + 3 + 3 = 9$ times. Hence, altogether there exist at least $6 + 9 = 15$ crossings in this case. In all cases, the number of crossings among G_i , G_x , G_y , and G_z in \mathcal{N}_i is at least 13, as required. Equality is possibly only if none of G_x , G_y , and G_z has self-crossings, and at least one of G_x , G_y , and G_z is crossed by G_i exactly 3 times. \square

By Lemma 10, $\text{cr}_1(\gamma, \lambda) \leq 13m$ implies that $\Gamma \circ \Lambda$ defines an embedding of G_x , for all $x \in \mathcal{X}$, in each region \mathcal{N}_i , $i = 1, \dots, m$. Consequently, $\Gamma \circ \Lambda$ defines an embedding of G_x in \mathbb{R}^2 for all $x \in \mathcal{X}$. In every embedding $\Gamma \circ \Lambda(G_x)$, for $x \in \mathcal{X}$, either P_1^x lies above P_2^x , or vice versa. We can now define a truth assignment $\tau : \mathcal{X} \rightarrow \{\text{true}, \text{false}\}$ such that for every $x \in \mathcal{X}$, $\tau(x) = \text{true}$ if and only if P_1^x lies above P_2^x in $\Gamma \circ \Lambda(G_x)$.

Lemma 11 *Assume that $\Gamma \circ \Lambda(G_x)$ is an embedding for every $x \in \mathcal{X}$, which determines the truth assignment $\tau : \mathcal{X} \rightarrow \{\text{true}, \text{false}\}$ described above. For every $i = 1, \dots, m$, if variable x appears in clause c_i , and G_i crosses G_x at most 3 times in \mathcal{N}_i , then x appears as a true literal in c_i .*

Proof Consider the highest and lowest path P_h and P_ℓ , respectively, among P_1^x , P_2^x , and P_3^x in $\mathcal{N}_i \cap \Gamma \circ \Lambda(G_x)$, none of which can be P_2^x since $\Gamma \circ \Lambda(G_x)$

is an embedding. By the construction of λ , either there exists exactly one through component of P_h in $\lambda^{-1}[u_{5i+1}]$ and exactly one through component of P_ℓ in $\lambda^{-1}[u_{5i+2}]$, or vice versa.

By the construction of λ , G_i crosses each of P_1^x , P_2^x , and P_3^x at least once in \mathcal{N}_i . By the hypothesis of the lemma, it crosses each exactly once. Then P_h has only one through component in $\lambda^{-1}[u_{5i+1}]$, and P_ℓ has only one through component in $\lambda^{-1}[u_{5i+2}]$. By the construction of λ , if x appears as a nonnegated variable in c_i , then $P_h = P_1^x$ lies above P_2^x and therefore $\tau(x) = \text{true}$. Similarly, if x appears as a negated variable in c_i , then $P_3^x = P_h$ lies above P_2^x and therefore $\tau(x) = \text{false}$. Consequently, x appears as a true literal in c_i and that concludes the proof. \square

Since $\text{cr}_1(\gamma, \lambda) \leq 13m$, for every $i = 1, \dots, m$, there are exactly 13 crossings in \mathcal{N}_i by Lemma 10. Moreover, by Lemma 10 the drawing $\Gamma \circ \Lambda(G_x)$ is an embedding for every $x \in \mathcal{X}$, and in every c_i for one its variables x the drawing of G_x is crossed by G_i exactly 3 times. By Lemma 11, the assignment τ makes at least one literal in each clause c_i of Φ true. We conclude that Φ is satisfiable, as required. This completes the proof of NP-hardness.

4.2 Second Construction: Cycle.

In our first construction in Section 4.1, G was a disjoint union of paths, and for every path endpoint $a \in V(G)$, a is the only vertex mapped to the cluster $\lambda(a) \in V(H)$. This property allows us to expand the construction as follows. We augment G into a cycle \bar{G} by adding 2-edge paths connecting the path endpoints, and we augment H with 2-pipe paths between the corresponding clusters, $u = \lambda(a)$ and $v = \lambda(b)$, by drawing a straight-line path $\gamma(uv)$ between $\gamma(u)$ and $\gamma(v)$ that does not pass through the image of any other cluster (but may cross images of other pipes). The augmentation does not change $\text{cr}_1(\gamma, \lambda)$, and we can easily compute the increase in $\text{cr}_2(\gamma, \lambda)$ due to new crossings. Consequently, finding $\text{cr}(\gamma \circ \lambda)$ remains NP-hard.

5 An Efficient Algorithm for a Constant Number of Crossings

Similarly to the standard crossing number, we can decide in polynomial time whether $\text{cr}(\varphi)$ is below a constant $k \in \mathbb{N}$. For the standard crossing number, the analogous problem is reduced to $O(n^{O(k)})$ calls of a planarity test. Here, instead of planarity testing, we use an algorithm for recognizing weak embeddings, which has been extended to testing embeddability into a given manifold [2]. Roughly speaking, we first guess how the edges of G cross under a map $\Lambda : G \rightarrow \mathcal{H}$ witnessing $\text{cr}(\varphi) \leq k$. Then we turn the crossings in Λ into vertices, which yields a graph G' . Finally, we construct an equivalent instance of the weak embeddability problem, which can be solved in polynomial time [2].

Theorem 3 *For every $k \in \mathbb{N}$, there is an algorithm that decides in $O(n^{O(k)})$ time whether $\text{cr}(\varphi) \leq k$ for a given map $\varphi : G \rightarrow \mathbb{R}^2$, where G has n vertices.*

Proof We are given a straight-line map $\varphi : G \rightarrow \mathbb{R}^2$ and a positive integer $k \in \mathbb{N}$, and we need to decide whether $\text{cr}(\varphi) \leq k$. Recall that φ is given as a composition of $\lambda : G \rightarrow H$ and $\gamma : H \rightarrow \mathbb{R}^2$, where H is a graph. Recall also that $\text{cr}(\gamma \circ \lambda) = \text{cr}_1(\gamma, \lambda) + \text{cr}_2(\gamma, \lambda)$ by Lemma 1, where $\text{cr}_2(\gamma, \lambda)$ is the number of edge pairs in G mapped to crossing pipes in H , and it can be computed in polynomial time. If $\text{cr}_2(\gamma, \lambda) > k$, then we can report that $\text{cr}(\varphi) > k$. In the remainder of the proof, we assume $\text{cr}_2(\gamma, \lambda) \leq k$, and let $k_1 = k - \text{cr}_2(\gamma, \lambda)$. We need to decide whether $\text{cr}_1(\gamma, \lambda) \leq k_1$, that is, whether there exists a drawing $\Lambda : G \rightarrow \mathcal{H}$ with properties (P1)–(P3) and at most k_1 crossings.

If such a drawing Λ exists, and we insert new vertices at the crossings, we obtain a graph G' with at most $n + k_1$ vertices that admits an embedding into \mathcal{H} with properties (P1)–(P3). We can guess the number of crossings x , $0 \leq x \leq k_1$. For each crossing, we guess the pair of edges that cross and the disk D_u in which they cross. If an edge crosses several other edges within a disk D_u , we also guess the order of these crossings. Based on these guesses, we can construct a graph G' in which crossings are turned into vertices. Specifically G' is obtained from G by first subdividing edges by vertices representing crossings, and then identifying the pairs of vertices that correspond to the same crossing. Then we extend the map $\lambda : G \rightarrow H$ to a map $\lambda' : G' \rightarrow H$ as follows. If a vertex $c \in V(G') \setminus V(G)$ represents a crossing in a disk D_u , then we put $\lambda'(c) := u$. There are $O(n^3)$ possibilities for each crossing, and so there are $O(k! n^{3k})$ possible guesses for $k_1 \leq k$ crossings.

Conversely, if $\lambda' : G' \rightarrow H$ admits a crossing-free embedding $\Lambda' : G' \rightarrow \mathcal{H}$ satisfying properties (P1)–(P3), then it yields a drawing $\Lambda : G \rightarrow \mathcal{H}$ with properties (P1)–(P3), and at most $x \leq k_1$ crossings, which witnesses that $\text{cr}_2(\gamma, \lambda) \leq k_1$. (In the drawing $\Lambda(G)$, each vertex in $V(G') \setminus V(G)$ is a common interior point of two edges, which either properly cross or have a point of tangency at that point.) Testing whether such an embedding $\Lambda' : G' \rightarrow \mathcal{H}$ exists is precisely an instance of the weak embeddability problem, and can be solved in $O(n^2 \log n)$ time [2]. The overall running time of the algorithm is $O(k! n^{3k+2} \log n)$. \square

It remains an open problem whether computing $\text{cr}(\varphi)$ is FPT when parameterized by the solution value. In other words, we are asking whether the running time $O(n^{O(k)})$ in Theorem 3 can be improved to $f(k)n^{O(1)}$, similar to the (standard) crossing number by a result of Grohe [16] (see also [21]). Note, however, that deciding whether the crossing number of a graph is less than or equal to a given threshold k does not admit a polynomial kernel [19].

6 Conclusions

Motivated by recent efficient algorithms that can decide whether a piecewise linear map $\varphi : G \rightarrow \mathbb{R}^2$ can be perturbed into an embedding, we investigate the problem of computing the minimum number of crossings in a perturbation. We have described an efficient algorithm when G is a cycle and φ has no spurs

(Theorem 1); and the problem becomes NP-hard if G is an arbitrary graph, or if G is a cycle but φ may have spurs (Theorem 2). However, perhaps one can minimize the number of crossings efficiently under milder assumptions. We formulate one promising scenario as follows: Is there a polynomial-time algorithm that finds $\text{cr}(\gamma \circ \lambda)$ when $\lambda^{-1}[u]$ is a planar graph (resp., an edgeless graph) for every cluster $u \in V(H)$ and λ has no spurs?

Another interesting research direction, raised in Section 5, is whether there is an FPT algorithm for computing $\text{cr}(\varphi)$ when parameterized by the solution value, that is, whether it can be computed in time $f(k)n^{O(1)}$ for a given straight-line map $\varphi : G \rightarrow \mathbb{R}^2$, where G has n vertices.

Finally, we are also interested in approximating $\text{cr}(\varphi)$ in polynomial time when G has maximum degree 2. Every such graph is planar, hence its crossing number is 0. In contrast, it is APX-hard to approximate the crossing number of an arbitrary graph, or even a 3-regular graph [5]. For graphs of bounded maximum degree, the first sublinear approximation algorithm by a factor of $\tilde{O}(n^{0.9})$ was achieved by Chuzhoy in [8]. For bounded degree graphs, there also exists a polynomial time algorithm [11] that approximates the quantity $n + \text{cr}(G)$ within a factor of $O(\log^2 n)$ as explained in [8].

References

1. Hugo A. Akitaya, Greg Aloupis, Jeff Erickson, and Csaba D. Tóth. Recognizing weakly simple polygons. *Discrete Comput. Geom.*, 58(4):785–821, 2017. doi:10.1007/s00454-017-9918-3.
2. Hugo A. Akitaya, Radoslav Fulek, and Csaba D. Tóth. Recognizing weak embeddings of graphs. In *Proc. 29th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 274–292. SIAM, 2018. doi:10.1137/1.9781611975031.20.
3. Patrizio Angelini and Giordano Da Lozzo. Clustered planarity with pipes. *Algorithmica*, 81(6):2484–2526, 2019. doi:10.1007/s00453-018-00541-w.
4. Patrizio Angelini, Giordano Da Lozzo, Giuseppe Di Battista, and Fabrizio Frati. Strip planarity testing for embedded planar graphs. *Algorithmica*, 77(4):1022–1059, 2017. doi:10.1007/s00453-016-0128-9.
5. Sergio Cabello. Hardness of approximation for crossing number. *Discrete & Computational Geometry*, 49(2):348–358, 2013. doi:10.1007/s00454-012-9440-6.
6. Sergio Cabello and Bojan Mohar. Adding one edge to planar graphs makes crossing number and 1-planarity hard. *SIAM Journal on Computing*, 42(5):1803–1829, 2013. doi:10.1137/120872310.
7. Hsien-Chih Chang, Jeff Erickson, and Chao Xu. Detecting weakly simple polygons. In *Proc. 26th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1655–1670, 2015. doi:10.1137/1.9781611973730.110.
8. Julia Chuzhoy. An algorithm for the graph crossing number problem. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC)*, pages 303–312. ACM, 2011. Preprint, arXiv:1012.0255 isbn = 978-1-4503-0691-1. doi:10.1145/1993636.1993678.
9. Pier Francesco Cortese, Giuseppe Di Battista, Maurizio Patrignani, and Maurizio Pizzonia. Clustering cycles into cycles of clusters. *J. Graph Alg. Appl.*, 9(3):391–413, 2005. doi:10.7155/jgaa.00115.
10. Pier Francesco Cortese, Giuseppe Di Battista, Maurizio Patrignani, and Maurizio Pizzonia. On embedding a cycle in a plane graph. *Discrete Math.*, 309(7):1856–1869, 2009. doi:10.1016/j.disc.2007.12.090.
11. Guy Even, Sudipto Guha, and Baruch Schieber. Improved approximations of crossings in graph drawings and VLSI layout areas. *SIAM J. Comput.*, 32(1):231–252, 2002. URL: <https://doi.org/10.1137/S0097539700373520>, doi:10.1137/S0097539700373520.

12. Qing-Wen Feng, Robert F. Cohen, and Peter Eades. How to draw a planar clustered graph. In Ding-Zhu Du and Ming Li, editors, *Proc. 1st Conference on Computing and combinatorics (COCOON)*, volume 959 of *LNCS*, pages 21–30. Springer, Berlin, 1995. doi:10.1007/BFb0030816.
13. Qing-Wen Feng, Robert F. Cohen, and Peter Eades. Planarity for clustered graphs. In Paul Spirakis, editor, *Proc. 3rd European Symposium on Algorithms (ESA)*, volume 979 of *LNCS*, pages 213–226, Berlin, 1995. Springer. doi:10.1007/3-540-60313-1_145.
14. Radoslav Fulek and Jan Kynčl. Hanani-Tutte for approximating maps of graphs. In *Proc. 34th Symposium on Computational Geometry (SoCG)*, volume 99 of *LIPICs*, pages 39:1–39:15, Dagstuhl, Germany, 2018. doi:10.4230/LIPICs.SocG.2018.39.
15. M. R. Garey and D. S. Johnson. Crossing number is NP-complete. *SIAM. J. on Algebraic and Discrete Methods*, 4(3):312–316, 1982. doi:10.1137/0604033.
16. Martin Grohe. Computing crossing numbers in quadratic time. *Journal of Computer and System Sciences*, 68(2):285–302, 2004. doi:10.1016/j.jcss.2003.07.008.
17. Haim Hanani. Über wesentlich unplättbare Kurven im drei-dimensionalen Raume. *Fundamenta Mathematicae*, 23:135–142, 1934. doi:10.4064/fm-23-1-135-142.
18. Joel Hass and Peter Scott. Intersections of curves on surfaces. *Israel Journal of Mathematics*, 51(1):90–120, 1985. doi:10.1007/BF02772960.
19. Petr Hliněný and Marek Dernár. Crossing number is hard for kernelization. In *32nd International Symposium on Computational Geometry (SoCG)*, volume 51 of *LIPICs*, pages 42:1–42:10. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016. doi:10.4230/LIPICs.SocG.2016.42.
20. Michael Jünger, Sebastian Leipert, and Petra Mutzel. Level planarity testing in linear time. In Sue H. Whitesides, editor, *Proc. 6th Symposium on Graph Drawing (GD)*, volume 1547 of *LNCS*, pages 224–237. Springer, Berlin, 1998. doi:10.1007/3-540-37623-2_17.
21. Ken-ichi Kawarabayashi and Bruce A. Reed. Computing crossing number in linear time. In *Proc. 39th ACM Symposium on Theory of Computing (STOC)*, pages 382–390. ACM, 2007. doi:10.1145/1250790.1250848.
22. Dušan Repovš and Arkadij B. Skopenkov. A deleted product criterion for approximability of maps by embeddings. *Topology Appl.*, 87(1):1–19, 1998. doi:10.1016/S0166-8641(97)00121-1.
23. Mikhail Skopenkov. On approximability by embeddings of cycles in the plane. *Topology Appl.*, 134(1):1–22, 2003. doi:10.1016/S0166-8641(03)00069-5.
24. William T. Tutte. Toward a theory of crossing numbers. *J. Combin. Theory*, 8:45–53, 1970. doi:10.1016/S0021-9800(70)80007-2.