Extending Orthogonal Planar Graph Drawings is Fixed-Parameter Tractable

Sujoy Bhore[1], Robert Ganian[2], Liana Khazaliya[2], Fabrizio Montecchiani[3], and Martin Nöllenburg[2]

1 **Indian Institute of Technology Bombay, India**
`sujoy.bhore@gmail.com`

2 **Technische Universität Wien, Austria**
`{rganian,lkhazaliya,noellenburg}@ac.tuwien.ac.at`

3 **University of Perugia**
`fabrizio.montecchiani@unipg.it`

Abstract

We consider the extension problem for bend-minimal orthogonal drawings of planar graphs, which is among the most fundamental geometric graph drawing representations. While the problem was known to be NP-hard, it is natural to consider the case where the drawn part is connected and only a small part of the graph is still to be drawn. Here, we prove the problem is in FPT when parameterized by the size of the missing subgraph.

Related Version `http://arxiv.org/abs/2302.10046`

1 Introduction

Drawing extension problems are motivated, for instance, by visualizing networks, in which certain subgraphs represent important motifs that require a specific drawing, or by visualizing dynamic networks, in which new edges and vertices must be integrated in an existing, stable drawing. Generally speaking, we are given a graph G and a (typically connected) subgraph H of G with a drawing $\Gamma(H)$, which is called a *partial* drawing of G. The drawing $\Gamma(H)$ satisfies certain topological or geometric properties, e.g., planarity, upward planarity, or 1-planarity, and the goal of the corresponding extension problem is to extend $\Gamma(H)$ to a drawing $\Gamma(G)$ of the whole graph G (if possible) by inserting the missing vertices and edges into $\Gamma(H)$ while maintaining the required drawing properties.

In this paper, we study the geometric drawing extension problem arising in the context of one of the most fundamental graph drawing styles: *orthogonal drawings* [3, 4, 6, 10]. In a planar orthogonal drawing, edges are represented as polylines comprised of (one or more) horizontal and vertical segments, ideally with as few overall bends as possible, where edges are not allowed to intersect except at common endpoints. Orthogonal drawings find applications in various domains from VLSI and printed circuit board (PCB) design, to schematic network visualizations, e.g., UML diagrams in software engineering, argument maps, or flow charts.

Given the above, a key optimization goal in orthogonal drawings is bend minimization. This task is known to be NP-hard [8] when optimizing over all possible combinatorial embeddings of a given graph, but can be solved in polynomial time for a fixed combinatorial embedding using the network flow model of Tamassia [11].

Despite the general popularity of planar orthogonal graph drawings, the corresponding extension problem has only been considered recently by Angelini et al. [1]. While they showed that the existence of a planar orthogonal extension can be decided in linear time,

39th European Workshop on Computational Geometry, Barcelona, Spain, March 29–31, 2023.
This is an extended abstract of a presentation given at EuroCG'23. It has been made public for the benefit of the community and should be considered a preprint rather than a formally reviewed paper. Thus, this work is expected to appear eventually in more final form at a conference with formal proceedings and/or in a journal.

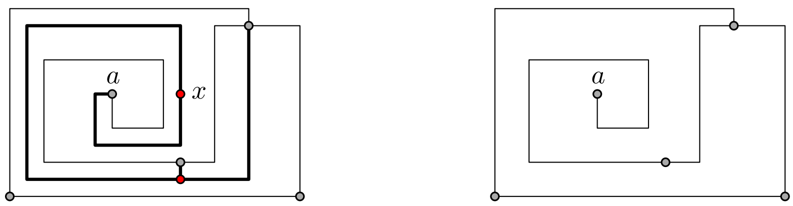

(a) $\Gamma(G)$ **(b)** $\Gamma(H)$

Figure 1 An orthogonal drawing of (a) a graph G and (b) a subgraph H of G.

the orthogonal bend-minimal drawing extension problem in general is easily seen to be NP-hard as it generalizes the case in which the pre-drawn part of the graph is empty [8]. Our paper addresses the parameterized complexity of the bend-minimal extension problem for planar orthogonal graph drawings under the most natural parameterization of the problem, which is the size of the subgraph that is still missing from the drawing.

Problem Statement. Let G be a planar graph and H be a connected subgraph of G. We call the complement $X = V(G) \setminus V(H)$ the *missing vertex set* of G, and $E_X = E(G) \setminus E(H)$ the *missing edge set*. Let $\Gamma(H)$ be a planar orthogonal drawing of H. A planar orthogonal drawing $\Gamma(G)$ *extends* $\Gamma(H)$ if its restriction to the vertices and edges of H coincides with $\Gamma(H)$. Moreover, $\Gamma(G)$ is a β-extension of $\Gamma(H)$ if it extends $\Gamma(H)$ and the total number of bends along the edges of E_X is at most β, for some $\beta \in \mathbb{N}$. For example, Figure 1a shows a 7-extension $\Gamma(G)$ of the drawing $\Gamma(H)$ in Figure 1b, with the missing vertices drawn in red.

BEND-MINIMAL ORTHOGONAL EXTENSION (BMOE)
Input: $(G, H, \Gamma(H))$, integer β
Problem: Is there a β-extension $\Gamma(G)$ of $\Gamma(H)$?

Our parameter of interest, denoted by κ, is the number of vertices and edges missing from H, i.e., $\kappa = |V(G) \setminus V(H)| + |E(G) \setminus E(H)|$.

Contributions and overview. We establish the fixed-parameter tractability of BMOE when parameterized by κ. While there have been numerous recent advances in the parameterized study of drawing extension problems [5,7,9], the specific drawing styles considered in those papers were primarily topological in nature, while for bend minimization the geometry of the instance is crucial. In order to overcome this difficulty, we develop a new set of tools summarized below. We first apply an initial branching step to simplify the problem (Section 2). This step allows us to reduce our target problem to BEND-MINIMAL ORTHOGONAL EXTENSION ON A FACE (F-BMOE), where the missing edges and vertices are drawn only in a marked face f and we have some additional information about how the edges are geometrically connected. Next, we focus on solving an instance of F-BMOE (Section 3). We show that certain parts of the marked face f are irrelevant and can be pruned away, and also use an involved argument to reduce the case of f being the outer face to the case of f being an inner face. Once that is done, we enter the centerpiece of our approach (Section 4), where the aim is to obtain a suitable discretization of our instance. To this end, we split the face f into so-called *sectors*, which group together points that have the same "bend distances" to all of the connecting points on the boundary of f. Furthermore, we construct a *sector-grid*—a point-set such that each sector contains a bounded number of points from this set, and every bend-minimal extension can be modified to only use points from this set for all vertices and bends. While this latter result would make it easy to handle each individual sector by brute force, the issue is that the number of sectors can be very large, hindering tractability. To deal with this obstacle, we capture the connections between sectors via a *sector graph* whose vertices are the sectors and whose edges represent geometric adjacencies between sectors. Crucially, we show that the sector graph has treewidth bounded by a function of κ. Having obtained this bound on the treewidth, the last step simply combines the already constructed sector grid with dynamic programming to solve F-BMOE (and hence also BMOE).

Many technicalities and proofs have been omitted; see [2] for the full paper.

2 Initial Branching

Let $\langle(G, H, \Gamma(H)), \beta\rangle$ be an instance of BMOE. A vertex $w \in V(H)$ is called an *anchor* if it is incident to an edge in E_X. For a missing edge $vw \in E_X$ incident to a vertex $v \in V(H)$, we will use "ports" to specify a direction that vw could potentially use to reach v in an extension of $\Gamma(H)$; we denote these directions as d, which is an element from $\{\downarrow$ (north), $\uparrow$ (south), $\leftarrow$ (east), $\rightarrow$ (west)$\}$. Formally, a *port candidate* for $vw \in E_X$ and $v \in V(H)$ is a pair (v, d). A *port-function* is an ordered set of port candidates which contains precisely one port candidate for each $vw \in E_X, v \in V(H)$, ordered lexicographically by v and then by w.

Bend-Minimal Orthogonal Extension on a Face (F-BMOE)
Input: Planar graph G_f; induced subgraph H_f of G_f with $k = |X_f|$, where $X_f = V(G_f) \setminus V(H_f)$; drawing $\Gamma(H_f)$ of H_f consisting of a single inner face f; port-function $\mathcal{P}$. **Task:** Compute the minimum β for which a β-extension of $\Gamma(H_f)$ exists s.t. (1) all missing edges and vertices are drawn in face f, (2) each edge $xa \in E_X$ where $a \in V(H)$ connects to a via its port candidate defined by $\mathcal{P}$, or determine that no such extension exists.

▶ **Lemma 2.1.** *There is an algorithm that solves an instance $\mathcal{I}$ of* BMOE *in time $3^{\mathcal{O}(\kappa)} \cdot T(|\mathcal{I}|, k)$, where $T(a, b)$ is the time required to solve an instance of* F-BMOE *with instance size a and parameter value b.*

The algorithm in [1] can be used to test whether an instance of F-BMOE admits some β-extension. Hence, we will assume to be dealing with instances where such an extension exists. We will call a β-extension minimizing the value of β a *solution*.

3 Preprocessing

The first two steps that will allow us to solve F-BMOE include pruning out certain parts of the face which are provably irrelevant, and reducing the case of f being the outer face to the case of f being an inner face.

Let $\Gamma(G)$ be an orthogonal drawing of a graph G and let f be a face of $\Gamma(G)$. A *feature point* of $\Gamma(G)$ is a point representing either a vertex or a bend of an edge. A *reflex corner* p of f is a feature point that makes an angle larger than π inside f. Also, if p is an anchor, then it is called an *essential* reflex corner. A projection ℓ of a reflex corner p is a horizontal or vertical line-segment in the interior of f that starts at p and ends at its first intersection with the boundary of f. Figure 2 (left) shows two projections ℓ_1 and ℓ_2 of a reflex corner p.

Observe that each projection ℓ of a reflex corner p divides the face f into two connected regions. If p is not essential and one of the two regions contains no reflex corners of its own and no anchors, we call the region *redundant*. Our aim will be to show that such regions can be safely removed from the instance. Namely, we can prove the following, where a *clean*

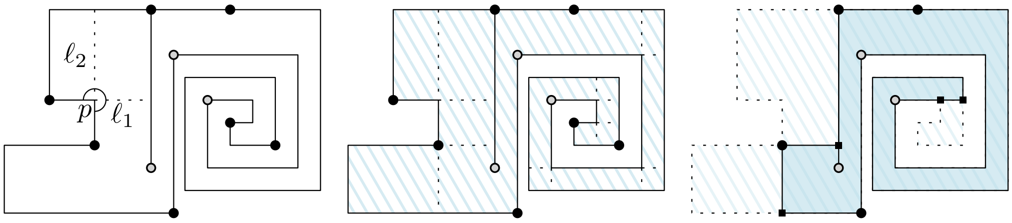

Figure 2 Left: A reflex corner p and its projections ℓ_1 and ℓ_2. Middle: A face (striped) with all its non-essential reflex corners and projections (anchor vertices have a gray filling while non-anchors are solid). Right: The corresponding clean instance (dummy vertices are drawn as small squares).

instance is such that each projection of each non-essential reflex corner in f splits f into two faces, each of which has at least one port on its boundary; see Figure 2 (right).

▶ **Lemma 3.1.** *There is a polynomial-time algorithm that takes as input an arbitrary instance of* F-BMOE *and outputs an equivalent instance which is clean.*

Given Lemma 3.1, we will hereinafter assume that our instances of F-BMOE are clean. Next, consider an instance of F-BMOE where the marked face is the outer face of $\Gamma(H_f)$, and let us begin by constructing a rectangle that bounds $\Gamma(H_f)$ and will serve as a "frame" for any solution. More formally, given an instance $\mathcal{I}$ of F-BMOE and a rectangle R that contains $\Gamma(H_f)$ in its interior, one easily sees that $\mathcal{I}$ admits a solution that lies in the interior of R. Based on this fact, we shall assume that any instance $\mathcal{I}$ is modified such that the outer face of $\Gamma(H_f)$ is a rectangle R containing no anchors (e.g., with four dummy vertices at its corners connected in a cycle). Notice that, while this ensures that f is no longer the outer face, f now contains a hole (that is, H_f is not connected anymore). The goal is now to remove this hole by connecting it to the boundary of R. To do so, let us consider an arbitrary horizontal or vertical line-segment ζ that connects the boundary of R with an edge-segment in the drawing $\Gamma(H_f)$ and intersects no other edge-segment of $\Gamma(H_f)$. Observe that, w.l.o.g., we can assume that each edge-segment in a solution $\Gamma(G_f)$ only intersects ζ in single points (and not in a line-segment); otherwise, one may shift ζ by a sufficiently small ϵ to avoid such intersections. Roughly speaking, we can show that the instance $\mathcal{I}$ can be "cut open" along ζ to construct an equivalent instance where the boundary of the polygon includes R, and to branch in order to determine how the edges in a hypothetical solution cross through ζ. However, to do so we need to ensure that there is a solution, in which the number of such crossings through ζ is bounded. To summarize, we can prove the following.

▶ **Lemma 3.2.** *There is an algorithm that takes as input an instance $\mathcal{I}$ of* F-BMOE *where f is the outer face and solves it in time $2^{\mathcal{O}(k^2 \log k)} \cdot Q(|\mathcal{I}|, k)$, where $Q(a, b)$ is the time to solve an instance of* F-BMOE *of size a and parameter value b such that f is the inner face.*

4 The Sector Graph

For a point $p \in f$, the *bend distance* $\mathrm{bd}(p, (a, d))$ to a port candidate (a, d) is the minimum integer q such that there exists an orthogonal polyline with q bends connecting p and a in the interior of f which arrives to a from direction d.

▶ **Definition 4.1.** Let $\mathcal{P} = ((a_1, d_1), \ldots, (a_q, d_q))$ be an ordered set of port candidates. For each point $p \in f$, we define its bend-vector as the tuple $\mathrm{vect}(p) = (\mathrm{bd}(p, (a_1, d_1)), \ldots, \mathrm{bd}(p, (a_q, d_q)))$.

▶ **Definition 4.2.** Given an ordered set of port candidates $\mathcal{P}$, a *sector* F is a maximal connected set of points with the same bend-vector w.r.t. $\mathcal{P}$.

When $\mathcal{P}$ is not specified explicitly, we will assume it to be the set of port candidates provided by the considered instance of F-BMOE. The face f is now partitioned into a set $\mathcal{F}$ of sectors. It is worth noting that sectors are connected regions in the face f, they do not overlap, and they cover the whole interior of f. We further notice that a sector can be degenerate, it may be a single point or a line-segment, and that pairs of (non-adjacent) sectors may have the same bend-vectors. At this point, we can define a graph representation capturing the adjacencies between the sectors in our instance; see Figure 3.

▶ **Definition 4.3.** Sectors A and B are *adjacent* if there exists a point p in A and a direction $d \in \{\uparrow, \downarrow, \leftarrow, \rightarrow\}$ such that the first point outside of A hit by the ray starting from p in direction d is in B.

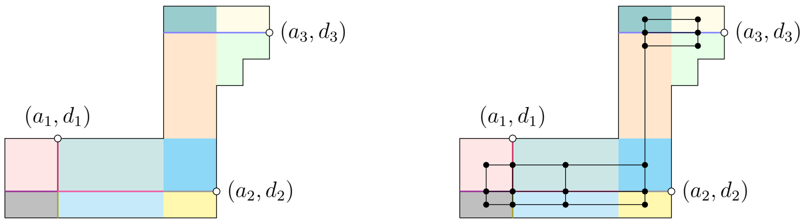

Figure 3 Left: partioning a face f into a set $\mathcal{F}$ of sectors, with three anchors marked using white circles. Right: the graph representation of $\mathcal{F}$.

The *sector graph* $\mathcal{G}$ is the graph whose vertex set is the set of sectors $\mathcal{F}$, and adjacencies of vertices are defined via the adjacency of sectors. It is not difficult to observe that the sector graph is a connected planar graph. Concerning its size, we observe that each sector contains at least one intersection point between two projections and that any such intersection point can be shared by at most nine sectors (four non-degenerate sectors plus five degenerate sectors). Hence the number of vertices in $\mathcal{G}$ is upper-bounded by $9x^2$, where x is the number of feature points in $\Gamma(H_f)$.

We now construct a "universal" point-set with the property that there exists a solution which places feature points only on these points, and where the intersection of the point-set with each sector is upper-bounded by a function of the parameter. Namely, let $\texttt{gridsize}(k) = c \cdot k^8$ (for some constant $c \approx 10^6$). Then we can prove the following:

▶ **Lemma 4.4.** *Given an instance $\mathcal{I}$ of* F-BMOE *we can construct a point-set (called a* sector grid*) in time $\mathcal{O}(|\mathcal{I}|)$ with the following properties: (1) $\mathcal{I}$ admits a solution whose feature points all lie on the sector grid, and (2) each sector contains at most $\texttt{gridsize}(k)$ points of the sector grid.*

To complete the proof of our fixed-parameter tractability result we proceed by first showing that the sector graphs in fact have treewidth bounded by a function of the parameter k, and then by using this fact to design a dynamic programming algorithm solving F-BMOE.

▶ **Theorem 4.5.** *Let $\mathcal{G}$ be a sector graph of a face f of the drawing $\Gamma(G)$. Then $\text{tw}(\mathcal{G}) \leq (4{+}4k)^{4k}$. Based on this, there is an algorithm that solves* F-BMOE *in time $2^{k^{\mathcal{O}(1)}} \cdot |V(G_f)|$.*

By combining Theorem 4.5 with Lemma 2.1, we conclude:

▶ **Corollary 4.6.** BMOE *can be solved in time $2^{\kappa^{\mathcal{O}(1)}} \cdot n$, where n is the number of feature points of $\Gamma(H)$.*

References

1. Patrizio Angelini, Ignaz Rutter, and T. P. Sandhya. Extending partial orthogonal drawings. *J. Graph Algorithms Appl.*, 25(1):581–602, 2021. `doi:10.7155/jgaa.00573`.
2. Sujoy Bhore, Robert Ganian, Liana Khazaliya, Fabrizio Montecchiani, and Martin Nöllenburg. Extending orthogonal planar graph drawings is fixed-parameter tractable, 2023. URL: `https://arxiv.org/abs/2302.10046`, `doi:10.48550/ARXIV.2302.10046`.
3. Giuseppe Di Battista, Peter Eades, Roberto Tamassia, and Ioannis G. Tollis. *Graph Drawing: Algorithms for the Visualization of Graphs.* Prentice-Hall, 1999.
4. Christian A. Duncan and Michael T. Goodrich. Planar orthogonal and polyline drawing algorithms. In Roberto Tamassia, editor, *Handbook of Graph Drawing and Visualization*, chapter 7, pages 223–246. CRC Press, 2013.
5. Eduard Eiben, Robert Ganian, Thekla Hamm, Fabian Klute, and Martin Nöllenburg. Extending partial 1-planar drawings. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *Automata, Languages, and Programming (ICALP'20)*, volume 168 of *LIPIcs*, pages 43:1–43:19. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020. `doi:10.4230/LIPIcs.ICALP.2020.43`.
6. Markus Eiglsperger, Sándor P. Fekete, and Gunnar W. Klau. Orthogonal graph drawing. In Michael Kaufmann and Dorothea Wagner, editors, *Drawing Graphs: Methods and Models*, volume 2025 of *LNCS*, chapter 6, pages 121–171. Springer-Verlag, 2001. `doi:10.1007/3-540-44969-8_6`.
7. Robert Ganian, Thekla Hamm, Fabian Klute, Irene Parada, and Birgit Vogtenhuber. Crossing-optimal extension of simple drawings. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021*, volume 198 of *LIPIcs*, pages 72:1–72:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. `doi:10.4230/LIPIcs.ICALP.2021.72`.
8. Ashim Garg and Roberto Tamassia. On the computational complexity of upward and rectilinear planarity testing. *SIAM J. Comput.*, 31(2):601–625, 2001. `doi:10.1137/S0097539794277123`.
9. Thekla Hamm and Petr Hlinený. Parameterised partially-predrawn crossing number. In Xavier Goaoc and Michael Kerber, editors, *38th International Symposium on Computational Geometry, SoCG 2022*, volume 224 of *LIPIcs*, pages 46:1–46:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. `doi:10.4230/LIPIcs.SoCG.2022.46`.
10. Takao Nishizeki and Md. Saidur Rahman. *Planar Graph Drawing*, volume 12 of *Lecture Notes Series on Computing.* World Scientific, 2004. `doi:10.1142/5648`.
11. Roberto Tamassia. On embedding a graph in the grid with the minimum number of bends. *SIAM J. Comput.*, 16(3):421–444, 1987. `doi:10.1137/0216030`.