



This is a repository copy of *A tight lower bound for Steiner Orientation*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/200965/>

Version: Accepted Version

Proceedings Paper:

Chitnis, R. and Feldmann, A.E. orcid.org/0000-0001-6229-5332 (2018) A tight lower bound for Steiner Orientation. In: Fomin, F.V. and Podolskii, V.V., (eds.) Computer Science – Theory and Applications: 13th International Computer Science Symposium in Russia, CSR 2018, Moscow, Russia, June 6–10, 2018, Proceedings. 13th International Computer Science Symposium in Russia, CSR 2018, 06-10 Jun 2018, Moscow, Russia. Lecture Notes in Computer Science, LNTCS 10846 . Springer International Publishing , pp. 65-77. ISBN 9783319905297

https://doi.org/10.1007/978-3-319-90530-3_7

This version of the contribution has been accepted for publication, after peer review (when applicable) but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: http://dx.doi.org/10.1007/978-3-319-90530-3_7. Use of this Accepted Version is subject to the publisher's Accepted Manuscript terms of use <https://www.springernature.com/gp/open-research/policies/accepted-manuscript-terms>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

A Tight Lower Bound for Planar Steiner Orientation*

Rajesh Chitnis^{1 **}, Andreas Emil Feldmann^{2 ***}, and Ondřej Suchý^{3 †}

¹ Department of Computer Science, University of Warwick, UK.

Email: rajeshchitnis@gmail.com

² Charles University in Prague, Czechia. Email: feldmann.a.e@gmail.com

³ Faculty of Information Technology, Czech Technical University in Prague, Czechia.

Email: ondrej.suchy@fit.cvut.cz

Abstract. In the STEINER ORIENTATION problem, the input is a mixed graph G (it has both directed and undirected edges) and a set of k terminal pairs \mathcal{T} . The question is whether we can orient the undirected edges in a way such that there is a directed $s \rightsquigarrow t$ path for each terminal pair $(s, t) \in \mathcal{T}$. Arkin and Hassin [DAM '02] showed that the STEINER ORIENTATION problem is NP-complete. They also gave a polynomial time algorithm for the special case when $k = 2$.

From the viewpoint of exact algorithms, Cygan, Kortsarz and Nutov [ESA '12, SIDMA '13] designed an XP algorithm running in $n^{O(k)}$ time for all $k \geq 1$. Pilipczuk and Wahlström [SODA '16, TOCT '18] showed that the STEINER ORIENTATION problem is W[1]-hard parameterized by k . As a byproduct of their reduction, they were able to show that under the Exponential Time Hypothesis (ETH) of Impagliazzo, Paturi and Zane [JCSS '01] the STEINER ORIENTATION problem does not admit an $f(k) \cdot n^{o(k/\log k)}$ algorithm for any computable function f .

In this paper, we give a short and easy proof that the $n^{O(k)}$ algorithm of Cygan et al. is asymptotically optimal, even if the input graph is planar. Formally, we show that the PLANAR STEINER ORIENTATION problem is W[1]-hard parameterized by the number k of terminal pairs, and, under ETH, cannot be solved in $f(k) \cdot n^{o(k)}$ time for any computable function f .

Moreover, under a stronger hypothesis called Gap-ETH of Dinur [ECCC '16] and Manurangsi and Raghavendra [ICALP '17], we are able to show that there is no constant $\vartheta > 0$ such that PLANAR STEINER ORIENTATION admits an $(\frac{19}{20} + \vartheta)$ -approximation in FPT time, i.e., no $f(k) \cdot n^{o(k)}$ time algorithm can distinguish between the case when all k pairs are satisfiable versus the case when less than $k \cdot (\frac{19}{20} + \vartheta)$ pairs are satisfiable. To the best of our knowledge, this is the first FPT inapproximability result on planar graphs.

* A preliminary version of this article [6] showing the W[1]-hardness and the $f(k) \cdot n^{o(k)}$ running time lower bound under ETH for graphs of genus 1 appeared in CSR '18. This version extends the results to graphs of genus 0 (planar graphs) and also contains an FPT inapproximability result.

** Supported by ERC grant CoG 647557 “Small Summaries for Big Data”. Part of this work was done when the author was at the Weizmann Institute of Science (and supported by Israel Science Foundation grant #897/13), and visiting Charles University in Prague, Czechia.

*** Supported by project CE-ITI (GAČR no. P202/12/G061) of the Czech Science Foundation, and by the Center for Foundations of Modern Computer Science (Charles Univ. project UNCE/SCI/004).

† Supported by grant 17-20065S of the Czech Science Foundation.

1 Introduction

In the STEINER ORIENTATION problem, the input is a mixed graph $G = (V, E)$ (it has both directed and undirected edges) and a set of terminal pairs $\mathcal{T} \subseteq V \times V$. The question is whether we can orient the undirected edges in a way such that there is a directed $s \rightsquigarrow t$ path for each terminal pair $(s, t) \in \mathcal{T}$. This problem has applications in traffic control (cf. [25] proving Robbins' theorem on graphs with strong orientations) and protein-protein interaction (PPI) networks [22].

STEINER ORIENTATION

Input: A mixed graph $G = (V, E)$, and a set $\mathcal{T} \subseteq V \times V$ of k terminal pairs.

Question: Is there an orientation of the undirected edges of G such that the resulting graph has an $s \rightsquigarrow t$ path for each $(s, t) \in \mathcal{T}$?

Parameter: k

Hassin and Megiddo [13] showed that STEINER ORIENTATION is polynomial time solvable if the input graph G is completely undirected, i.e., has no directed edges. If the input graph G is actually mixed, then Arkin and Hassin [1] showed that STEINER ORIENTATION is NP-complete. They also gave a polynomial time algorithm for the special case when $k = 2$. Cygan, Kortsarz and Nutov [10] generalized this by giving an $n^{O(k)}$ -time algorithm for all $k \geq 1$, i.e., STEINER ORIENTATION is in XP parameterized by k . Although the algorithm of Cygan et al. [10] is polynomial time for fixed k , the degree of the polynomial changes as k changes. This left open the question of whether one could design a much more efficient FPT algorithm for STEINER ORIENTATION parameterized by k , i.e., an algorithm which runs in time $f(k) \cdot n^{O(1)}$ for some computable function f independent of n .

Pilipczuk and Wahlström [23] answered this question negatively by showing that STEINER ORIENTATION is W[1]-hard parameterized by k . As a byproduct of their reduction, they were able to show that under the Exponential Time Hypothesis (ETH) of Impagliazzo, Paturi and Zane [14, 15] the STEINER ORIENTATION problem does not admit an $f(k) \cdot n^{o(k/\log k)}$ -time algorithm for any computable function f [24]. That is, the $n^{O(k)}$ -time algorithm of Cygan et al. is asymptotically almost optimal. This left open the following two questions:

- Can we close the gap between the $n^{O(k)}$ algorithm and the $f(k) \cdot n^{o(k/\log k)}$ hardness for STEINER ORIENTATION on general graphs?
- Is PLANAR STEINER ORIENTATION (see below) FPT parameterized by k , or can we at least obtain an improved runtime such as $f(k) \cdot n^{O(\sqrt{k})}$?

PLANAR STEINER ORIENTATION, the main problem of interest in this paper, is the restriction of STEINER ORIENTATION to planar mixed graphs. Naturally, these are mixed graphs for which the union of the graph formed by the undirected edges and the underlying undirected graph of the graph formed by the directed edges is planar.

In a preliminary version of this paper [6] the first question was answered by showing a tight lower bound of $f(k) \cdot n^{o(k)}$ for STEINER ORIENTATION under ETH. However, the constructed graph in [6] had genus 1 which left open the complexity for PLANAR

STEINER ORIENTATION since planar graphs have genus 0. In this full version, we completely answer both the above questions. Formally, we show that:

Theorem 1. *The PLANAR STEINER ORIENTATION problem is W[1]-hard parameterized by the number k of terminal pairs. Moreover, under ETH, PLANAR STEINER ORIENTATION cannot be solved in $f(k) \cdot n^{o(k)}$ time for any computable function f .*

Previously, the only result about the complexity of PLANAR STEINER ORIENTATION was an NP-completeness proof due to Beck et al. [2]. Since our reduction can be performed in polynomial time, it also shows NP-completeness of PLANAR STEINER ORIENTATION.

Furthermore, using the same reduction with a stronger hypothesis we are able to show that it is even hard to approximate the number of satisfiable pairs very precisely.

Theorem 2. *Assuming Gap-ETH, for each $0 < \vartheta \leq \frac{1}{20}$ there exists a constant $\zeta > 0$ such that for any computable function f there is no algorithm for PLANAR STEINER ORIENTATION, which given an instance (G, \mathcal{T}) with G planar, can distinguish between the following cases in $f(k) \cdot n^{\zeta k}$ time, where $n = |V(G)|$ and $k = |\mathcal{T}|$:*

- *there is an orientation of G that satisfies all pairs of \mathcal{T} and*
- *any orientation of G satisfies less than $(\frac{19}{20} + \vartheta)$ -fraction of pairs in \mathcal{T} .*

To the best of our knowledge, the only known results on approximability of STEINER ORIENTATION are for the case when the input graph is undirected. In this case, there are several known polynomial time approximation algorithms with ratios $O(\frac{\log |\mathcal{T}|}{\log \log |\mathcal{T}|})$ [10] and $O(\frac{\log n}{\log \log n})$ [12, 22]. We are not aware of any non-trivial approximation algorithms for STEINER ORIENTATION when the input graph is mixed.

Our reduction uses some ideas given by Pilipeczuk and Wahlström [23], who obtained a lower bound of $f(k) \cdot n^{o(\sqrt{k})}$ via a rather involved reduction from MULTICOLORED CLIQUE. This was later [24] improved by the same authors to $f(k) \cdot n^{o(k/\log k)}$ via the standard trick of reducing from the PARTITIONED SUBGRAPH ISOMORPHISM problem [18] instead. In this paper, we give a parameterized reduction from k -CLIQUE to STEINER ORIENTATION on planar graphs. Typically, W[1]-hardness results for problems on planar graphs are shown via reductions from the GRID TILING problem introduced by Marx [17]⁴. Indeed, the GRID TILING problem has turned out to be quite useful in showing W[1]-hardness for several problems on planar graphs [5, 7, 8, 19–21]. However, in this paper we chose to give a parameterized reduction from k -CLIQUE instead so that we can use the same reduction for the FPT inapproximability result of Theorem 2 as well. Note that there is a simple reduction [9, Thm 14.28] from k -CLIQUE to $k \times k$ GRID TILING which is also used implicitly in our reduction.

2 The reduction

We begin by describing the reduction from an instance (G', k) of k -CLIQUE to an instance of PLANAR STEINER ORIENTATION with $O(k)$ terminal pairs. We prove the exact properties of the reduction in subsequent subsections.

⁴ In [17] the GRID TILING problem (under the name of MATRIX TILING) was actually used to rule out PTASs for several problems under ETH.

2.1 Construction




Fig. 1. The instance of STEINER ORIENTATION with $O(k)$ terminal pairs created from an instance of k -CLIQUE (before the splitting operation). At this point, the only undirected edges are the green path edges. A gadget $G_{i,j}$ is highlighted by a dotted rectangle.

Consider an instance (G', k) of k -CLIQUE. We now build an instance (G, \mathcal{F}) of PLANAR STEINER ORIENTATION as follows (refer to Figure 1). For simplicity we assume that $|V(G')| = n - 1$ and $V(G') = \{u_1, \dots, u_{n-1}\}$ (this allows us to make less case distinctions in some further definitions and descriptions, for example Definition 2 and Definition 3).

- **Black Grid Edges:** For each $1 \leq i, j \leq 2k$ we introduce the $n \times n$ grid $G_{i,j}$. The grids $G_{2i,2j}$ will later be modified to represent adjacencies from the graph G' .

The other grids serve to transfer information. In [Figure 1](#) we highlight a gadget $G_{i,j}$ by a dotted rectangle.

We place the grids $G_{i,j}$ into a big grid of grids left to right according to growing i and from top to bottom according to growing j . That is $G_{1,1}$ is at the top left corner of the construction, the $G_{2k,1}$ at the top right corner, etc.

- In $G_{i,j}$ the column numbers grow from left to right, from 1 to n , and the row numbers grow from top to bottom, from 1 to n .
- For each $1 \leq h, \ell \leq n$ the unique vertex which is the intersection of the h^{th} column and ℓ^{th} row of $G_{i,j}$ is denoted by $v_{i,j}^{h,\ell}$.

Orient each horizontal edge of the grid $G_{i,j}$ to the right, and each vertical edge to the bottom.

– We now define four special sets of vertices for the gadget $G_{i,j}$ given by

- $\text{Left}(G_{i,j}) = \{v_{i,j}^{1,\ell} \mid \ell \in [n]\}$
- $\text{Right}(G_{i,j}) = \{v_{i,j}^{n,\ell} \mid \ell \in [n]\}$
- $\text{Top}(G_{i,j}) = \{v_{i,j}^{\ell,1} \mid \ell \in [n]\}$
- $\text{Bottom}(G_{i,j}) = \{v_{i,j}^{\ell,n} \mid \ell \in [n]\}$

– **Horizontal orange inter-grid edges:**

For each $i \in [2k-1], j \in [2k]$

- Add the directed perfect matching from vertices of $\text{Right}(G_{i,j})$ to $\text{Left}(G_{i+1,j})$ given by the set of edges $\{v_{i,j}^{n,\ell} \rightarrow v_{i+1,j}^{1,\ell} \mid \ell \in [n]\}$.

– **Vertical orange inter-grid edges:**

For each $i \in [2k], j \in [2k-1]$

- Add the directed perfect matching from vertices of $\text{Bottom}(G_{i,j})$ to $\text{Top}(G_{i,j+1})$ given by the set of edges $\{v_{i,j}^{h,n} \rightarrow v_{i,j+1}^{h,1} \mid h \in [n]\}$.

Note that after adding all inter-grid edges, the graph forms one large grid of size $2kn \times 2kn$ with all edges oriented to the right or to the bottom.

– We introduce $12k \cdot n$ **red** vertices given by

- $\hat{A} := \{\hat{a}_j^\ell \mid j \in [k], \ell \in [n]\}$
- $\hat{B} := \{\hat{b}_j^\ell \mid j \in [k], \ell \in [n]\}$
- $\hat{C} := \{\hat{c}_i^h \mid i \in [k], h \in [n]\}$
- $\hat{D} := \{\hat{d}_i^h \mid i \in [k], h \in [n]\}$
- $A := \{a_j^\ell \mid j \in [2k], \ell \in [n]\}$
- $B := \{b_j^\ell \mid j \in [2k], \ell \in [n]\}$
- $C := \{c_i^h \mid i \in [2k], h \in [n]\}$
- $D := \{d_i^h \mid i \in [2k], h \in [n]\}$

– **Blue connector edges:**

- For each $j \in [2k], \ell \in [n]$
 - * add the directed edge $a_j^\ell \rightarrow v_{1,j}^{1,\ell}$,
 - * add the directed edge $v_{2k,j}^{n,\ell} \rightarrow b_j^\ell$.
- For each $j \in [k], \ell \in [n]$
 - * add the directed edge $a_{2j-1}^\ell \rightarrow \hat{a}_j^\ell$,
 - * add the directed edge $a_{2j}^\ell \rightarrow \hat{a}_j^{n+1-\ell}$,

- * add the directed edge $\widehat{b}_j^\ell \rightarrow b_{2j-1}^\ell$,
- * add the directed edge $\widehat{b}_j^\ell \rightarrow b_{2j}^{n+1-\ell}$.
- For each $i \in [2k], h \in [n]$
 - * add the directed edge $c_i^h \rightarrow v_{i,1}^{h,1}$,
 - * add the directed edge $v_{i,2k}^{h,n} \rightarrow d_i^h$.
- For each $i \in [k], h \in [n]$
 - * add the directed edge $c_{2i-1}^h \rightarrow \widehat{c}_i^h$,
 - * add the directed edge $c_{2i}^h \rightarrow \widehat{c}_i^{n+1-h}$,
 - * add the directed edge $\widehat{d}_i^h \rightarrow d_{2i-1}^h$,
 - * add the directed edge $\widehat{d}_i^h \rightarrow d_{2i}^{n+1-h}$.

– **Green path edges:**

For each $j \in [k]$

- add the undirected path $\widehat{a}_j^1 - \widehat{a}_j^2 - \widehat{a}_j^3 - \dots - \widehat{a}_j^{n-1} - \widehat{a}_j^n$, and denote this path⁵ by \widehat{A}_j ,
- add the undirected path $\widehat{b}_j^1 - \widehat{b}_j^2 - \widehat{b}_j^3 - \dots - \widehat{b}_j^{n-1} - \widehat{b}_j^n$, and denote this path by \widehat{B}_j .

For each $j \in [2k]$

- add the undirected path $a_j^1 - a_j^2 - a_j^3 - \dots - a_j^{n-1} - a_j^n$, and denote this path by A_j ,
- add the undirected path $b_j^1 - b_j^2 - b_j^3 - \dots - b_j^{n-1} - b_j^n$, and denote this path by B_j .

For each $i \in [k]$

- add the undirected path $\widehat{c}_i^1 - \widehat{c}_i^2 - \widehat{c}_i^3 - \dots - \widehat{c}_i^{n-1} - \widehat{c}_i^n$, and denote this path by \widehat{C}_i ,
- add the undirected path $\widehat{d}_i^1 - \widehat{d}_i^2 - \widehat{d}_i^3 - \dots - \widehat{d}_i^{n-1} - \widehat{d}_i^n$, and denote this path by \widehat{D}_i .

For each $i \in [2k]$

- add the undirected path $c_i^1 - c_i^2 - c_i^3 - \dots - c_i^{n-1} - c_i^n$, and denote this path by C_i ,
- add the undirected path $d_i^1 - d_i^2 - d_i^3 - \dots - d_i^{n-1} - d_i^n$, and denote this path by D_i .

– **Splitting operation:**

We first define the sets $S_{i,j} \subseteq [n] \times [n]$ for each $i, j \in [k]$:

- For each $i \in [k]$ we let $S_{i,i} = \{(x, x) \mid x \in [n-1]\}$.
- For each $i, j \in [k]$ with $i \neq j$ we let $S_{i,j} = \{(x, y), (y, x) \mid \{u_x, u_y\} \in E(G')\}$. That is, $S_{i,j}$ represents the ones in the adjacency matrix of G' (with one extra row and one extra column of zeroes added).

Then for each $i, j \in [k]$ and each $x, y \in [n]$ we perform the following operation on the vertex $v_{2i,2j}^{x,y}$:

- If $(x, y) \in S_{i,j}$ then we keep the vertex $v_{2i,2j}^{x,y}$ as is.

⁵ Sometimes we also abuse notation slightly and use \widehat{A}_j to denote this set of vertices




Fig. 2. The splitting operation for vertex $v_{2i,2j}^{h,\ell}$ when $(h, \ell) \notin S_{i,j}$. The idea behind this splitting is that no matter which way we orient the undirected dotted edge we cannot go **both** from left to right and from top to bottom. However, if we just want to go from left to right (top to bottom) then it is possible by orienting the dotted edge to the right (left), respectively.




Fig. 3. Illustration of the terminal pairs in two consecutive rows. Each pair is represented by a magenta arrow between the appropriate vertices. The black labels represent the sets the pair belongs to, i.e., L stands for \mathcal{T}_j^L , R stands for \mathcal{T}_j^R , and H stands for \mathcal{T}_j^H .

- Otherwise we **split** the vertex $v_{2i,2j}^{x,y}$ into two vertices $v_{2i,2j,LB}^{x,y}$ and $v_{2i,2j,TR}^{x,y}$. Note that $v_{2i,2j}^{x,y}$ had 4 incident edges: two incoming (one each from the left and the top) and two outgoing (one each to the right and the bottom). We change the edges as follows (see Figure 2):
 - * Make the left incoming edge and bottom outgoing edge incident on $v_{2i,2j,LB}^{x,y}$ (denoted by red color in Figure 2).
 - * Make the top incoming edge and right outgoing edge incident on $v_{2i,2j,TR}^{x,y}$ (denoted by blue color in Figure 2).
 - * Add an **undirected** edge between $v_{2i,2j,LB}^{x,y}$ and $v_{2i,2j,TR}^{x,y}$ (denoted by the dotted edge in Figure 2).

By slightly abusing the notation we still refer to the grid as $G_{2i,2j}$, even though some of its vertices were split.

- The set \mathcal{T} of terminal pairs is given by $\bigcup_{i \in [k]} \mathcal{T}_i^L \cup \mathcal{T}_i^R \cup \mathcal{T}_i^T \cup \mathcal{T}_i^B \cup \mathcal{T}_i^H \cup \mathcal{T}_i^V$, where (see also Figure 3)
 - $\mathcal{T}_j^L = \{(a_{2j-1}^n, \hat{a}_j^1), (a_{2j-1}^1, \hat{a}_j^n), (a_{2j}^1, \hat{a}_j^1), (a_{2j}^n, \hat{a}_j^n)\}$ for each $j \in [k]$,
 - $\mathcal{T}_j^R = \{(\hat{b}_j^1, b_{2j-1}^1), (\hat{b}_j^n, b_{2j-1}^n), (\hat{b}_j^1, b_{2j}^1), (\hat{b}_j^n, b_{2j}^n)\}$ for each $j \in [k]$,

- $\mathcal{F}_i^T = \{(c_{2i-1}^n, \widehat{c}_i^1), (c_{2i-1}^1, \widehat{c}_i^n), (c_{2i}^1, \widehat{c}_i^1), (c_{2i}^n, \widehat{c}_i^n)\}$ for each $i \in [k]$,
- $\mathcal{F}_i^B = \{(\widehat{d}_i^1, d_{2i-1}^n), (\widehat{d}_i^n, d_{2i-1}^1), (\widehat{d}_i^1, d_{2i}^1), (\widehat{d}_i^n, d_{2i}^n)\}$ for each $i \in [k]$,
- $\mathcal{F}_j^H = \{(a_{2j-1}^n, b_{2j-1}^1), (a_{2j}^n, b_{2j}^1)\}$ for each $j \in [k]$, and
- $\mathcal{F}_i^V = \{(c_{2i-1}^n, d_{2i-1}^1), (c_{2i}^n, d_{2i}^1)\}$ for each $i \in [k]$

Note that the total number of terminal pairs is $(4 \cdot 4 + 2 \cdot 2)k = 20k$.

Remark 1. Note that the constructed graph G is planar: [Figure 1](#) gives a planar embedding of the graph before the splitting operation, and the splitting operation (see [Figure 2](#)) clearly preserves planarity.

Remark 2. Note that the index n does not appear in either first or second coordinate of any pair in the sets $S_{i,j}$ (for any $i, j \in [k]$) as there is no vertex u_n in G' . This technical condition is not necessary for our reduction, but we added it to help simplify⁶ some of the definitions and proofs which appear later in the paper. For example, one can now observe easily that all vertices in $\text{Right}(G_{2i,2j})$ and $\text{Bottom}(G_{2i,2j})$ are split (for every $i, j \in [k]$).

Before giving an intuition about and proving the properties of the reduction, we first introduce some notation concerning orientations of the green path edges (i.e. potential solutions) of the instance.

Definition 1. For any $i \in [n]$, a path on n vertices $a_1 - a_2 - \dots - a_n$ is said to be oriented towards (away from) i if every edge $a_{j-1} - a_j$ is oriented towards (away from) a_j for every $j \leq i$ and every edge $a_j - a_{j+1}$ is oriented towards (away from) a_j for every $j \geq i$, respectively.

We now give some intuition on how the individual terminal pairs enforce the desired structure on the solution (see [Figure 3](#) and [Figure 4](#)). To satisfy pairs in \mathcal{F}_j^L the path \widehat{A}_j must be oriented away from some index, path A_{2j-1} towards the same index and path A_{2j} towards the opposite index (see [Lemma 4](#)). Similarly, to satisfy pairs in \mathcal{F}_j^R the path \widehat{B}_j must be oriented towards some index, the path B_{2j-1} must be oriented away from the same index and the path must be oriented B_{2j} away from the opposite index. Then, the pairs in \mathcal{F}_j^H enforce one inequality on the indexes selected on left and right side each, and to satisfy them, there must be a horizontal path from left to right in each row (see [Lemma 5](#)). Pairs in \mathcal{F}_i^T , \mathcal{F}_i^B , and \mathcal{F}_i^V enforce similar properties on paths $\widehat{C}_i, C_{2i-1}, C_{2i}, \widehat{D}_i, D_{2i-1}, D_{2i}$ and the existence of a vertical path in each column. The splitting in grids on the diagonal enforces that the indexes selected in rows are the same as those selected in the corresponding columns and thus actually correspond to selecting a vertex of G' (see [Lemma 6](#)). The splitting in grids off the diagonal ensures that the selected vertices are adjacent (see [Lemma 7](#)).

2.2 Completeness of the reduction

In this subsection we prove the following lemma.

⁶ In [Definitions 2](#) and [3](#), instead of having three conditions (either j is odd or even, and for the case when j is even we would have two further cases depending on whether or not that vertex is split) for the last black vertex of a horizontal canonical path, we now only have two conditions (either j is odd and the vertex is not split, or j is even and the vertex is definitely split).

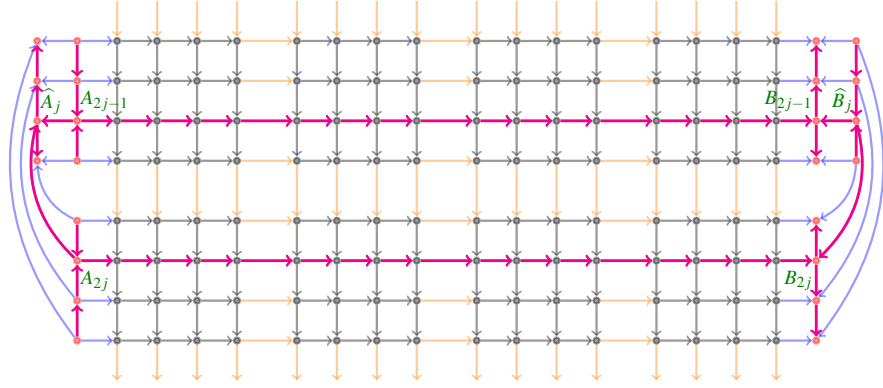


Fig. 4. Illustration of the oriented paths in an intended solution in two consecutive rows of the construction. The arcs used by the paths are drawn in magenta, vertical paths are not drawn. The path \widehat{A}_j is oriented away from $n+1-x_j$ in Lemma 1 and away from $n+1-\lambda_j$ in Lemma 4 and Lemma 6. The path A_{2j-1} is oriented towards $n+1-x_j$ in Lemma 1, towards $n+1-\lambda_j$ in Lemma 4 and Lemma 6, and towards α_{2j-1} in Lemma 5. The path A_{2j} is oriented towards x_j in Lemma 1, towards λ_j in Lemma 4 and Lemma 6, and towards α_{2j} in Lemma 5. The path B_{2j-1} is oriented away from $n+1-x_j$ in Lemma 1, away from $n+1-\mu_j$ in Lemma 4 and Lemma 6, and away from β_{2j-1} in Lemma 5. The path B_{2j} is oriented away from x_j in Lemma 1, away from μ_j in Lemma 4 and Lemma 6, and away from β_{2j} in Lemma 5. The path \widehat{B}_j is oriented towards $n+1-x_j$ in Lemma 1 and towards $n+1-\mu_j$ in Lemma 4 and Lemma 6.

Lemma 1. *On an instance (G', k) of k -CLIQUE such that G' contains a clique of size k , the reduction produces an instance (G, \mathcal{T}) of PLANAR STEINER ORIENTATION which has a solution satisfying all terminal pairs.*

Suppose that G' has a clique X of size k . Let the clique be formed by the vertices $u_{x_1}, u_{x_2}, \dots, u_{x_k}$. We know that $x_i \in [n]$ for each $i \in [k]$. We now show that the constructed instance (G, \mathcal{T}) of PLANAR STEINER ORIENTATION has a solution that satisfies all the terminal pairs. Orient the undirected green path edges as follows (see Figure 4 for an illustration). For each $i, j \in [k]$

- orient \widehat{A}_j and B_{2j-1} away from $n+1-x_j$,
- orient A_{2j-1} and \widehat{B}_j towards $n+1-x_j$,
- orient A_{2j} towards x_j ,
- orient B_{2j} away from x_j ,
- orient \widehat{C}_i and D_{2i-1} away from $n+1-x_i$,
- orient C_{2i-1} and \widehat{D}_i towards $n+1-x_i$,
- orient C_{2i} towards x_i ,
- orient D_{2i} away from x_i ,

Note that so far we have only oriented the undirected green path edges, and that the dotted undirected edges arising from the splitting operation (Figure 2) are not yet oriented at this step. It is easy to see that the above orientations ensure that all terminal

pairs of $\bigcup_{i \in [k]} \mathcal{T}_i^L \cup \mathcal{T}_i^R \cup \mathcal{T}_i^T \cup \mathcal{T}_i^B$ are satisfied (see Figure 4). We now show that terminal pairs of $\bigcup_{i \in [k]} \mathcal{T}_i^H \cup \mathcal{T}_i^V$ are also satisfied. First we need some definitions:

Definition 2. (horizontal canonical paths) For $j \in [2k]$ and $\ell \in [n]$, we denote by Q_j^ℓ the unique (horizontal) directed $a_j^\ell \rightarrow b_j^\ell$ path whose second vertex is $v_{1,j}^{1,\ell}$ and second-last vertex is $v_{2k,j}^{n,\ell}$ if j is odd and $v_{2k,j,TR}^{n,\ell}$ if j is even. This path starts with the blue connector edge from a_j^ℓ to $v_{1,j}^{1,\ell}$ and ends with the blue connector edge from $v_{2k,j}^{n,\ell}$ or $v_{2k,j,TR}^{n,\ell}$ to b_j^ℓ . The intermediate edges are obtained by selecting the paths of black edges given by the ℓ^{th} rows of each gadget $G_{i,j}$ for $i \in [2k]$, and connecting these small paths by horizontal orange edges.

However, we need to address what to do when we encounter a split vertex on this path. Consider the vertex $v_{i,j}^{h,\ell}$ for some $i \in [2k]$ and $h \in [n]$. If $v_{i,j}^{h,\ell}$ is not split, then we do not have to do anything. Otherwise, if $v_{i,j}^{h,\ell}$ is split, then we add the edge $v_{i,j,LB}^{h,\ell} \rightarrow v_{i,j,TR}^{h,\ell}$ to Q_j^ℓ .

Note that the orientation of G which orients all dotted edges rightwards, i.e., $LB \rightarrow TR$, contains each of the horizontal canonical paths defined above.

Definition 3. (vertical canonical paths) For $i \in [2k]$ and $h \in [n]$, we denote by P_i^h the unique (vertical) $c_i^h \rightarrow d_i^h$ path whose second vertex is $v_{i,1}^{h,1}$ and second-last vertex is $v_{i,2k}^{h,n}$ if i is odd and $v_{i,2k,LB}^{h,n}$ if i is even. This path starts with the blue connector edge $(c_i^h, v_{i,1}^{h,1})$ and ends with the blue connector edge from $v_{i,2k}^{h,n}$ or $v_{i,2k,LB}^{h,n}$ to d_i^h . The intermediate edges are obtained by selecting the paths of black edges given by the h^{th} columns of each gadget $G_{i,j}$ for $j \in [2k]$, and connecting these small paths by vertical orange edges.

However, we need to address what to do when we encounter a split vertex on this path. Consider the vertex $v_{i,j}^{h,\ell}$ for some $j \in [2k]$ and $\ell \in [n]$. If $v_{i,j}^{h,\ell}$ is not split, then we do not have to do anything. Otherwise, if $v_{i,j}^{h,\ell}$ is split, then we add the edge $v_{i,j,LB}^{h,\ell} \leftarrow v_{i,j,TR}^{h,\ell}$ to P_i^h .

Note that the orientation of G which orients all dotted edges leftwards, i.e., $LB \leftarrow TR$, contains each of the vertical canonical paths defined above. Observe that both the horizontal canonical paths and vertical canonical paths assign orientations to the dotted edges arising from splitting vertices. Hence, one needs to be careful because the splitting operation (see Figure 2) is designed to ensure that the existence of a horizontal canonical path implies that some vertical canonical path cannot exist (recall that we are allowed to orient each undirected edge in exactly one direction).

Definition 4. (realizable set of paths) A set of directed paths \mathcal{P} in a mixed graph G is realizable if there is an orientation G^* of G such that each path $P \in \mathcal{P}$ appears in G^* .

Lemma 2. The set of vertical canonical paths $\{P_{2i-1}^{n+1-x_i}, P_{2i}^{x_i} \mid i \in [k]\}$ together with the set of horizontal canonical paths $\{Q_{2j-1}^{n+1-x_j}, Q_{2j}^{x_j} \mid j \in [k]\}$ are realizable in G .

Proof. Orient each undirected dotted edge which appears on some horizontal path from $\{Q_{2j-1}^{n+1-x_j}, Q_{2j}^{x_j} \mid j \in [k]\}$ in LB \rightarrow TR direction and each undirected dotted edge which appears on some vertical path from $\{P_{2i-1}^{n+1-x_i}, P_{2i}^{x_i} \mid i \in [k]\}$ in TR \rightarrow LB direction. Since, if $i = j$, then we have $(x_i, x_i) \in S_{i,i}$ by definition and if $i \neq j$, then u_{x_i} and u_{x_j} are adjacent in G' as X is a clique in G' , implying that we also have $(x_i, x_j) \in S_{i,j}$ by definition, no black vertex $v_{2i,2j}^{x_i, x_j}$ was split. Thus there is no undirected dotted edge which would appear on both a horizontal and a vertical path from the set. To prove the lemma it remains to arbitrarily orient the edges which do not have a prescribed orientation yet. \square

Observe that for each $j \in [k]$, the horizontal path $Q_{2j-1}^{n+1-x_j}$ together with the oriented paths $a_{2j-1}^n \rightarrow a_{2j-1}^{n+1-x_j}$ in A_{2j-1} and $b_{2j-1}^{n+1-x_j} \rightarrow b_{2j-1}^1$ in B_{2j-1} satisfies the terminal pair (a_{2j-1}^n, b_{2j-1}^1) of \mathcal{T}_j^H for each $j \in [k]$ (see Figure 4). The same holds for the horizontal path $Q_{2j}^{x_j}$, paths in A_{2j} and B_{2j} and the terminal pair (a_{2j}^n, b_{2j}^1) of \mathcal{T}_j^H for each $j \in [k]$.

Similarly, for each $i \in [k]$, the path $P_{2i-1}^{n+1-x_i}$ together with the oriented paths $c_{2i-1}^n \rightarrow c_{2i-1}^{n+1-x_i}$ in C_{2i-1} and $d_{2i-1}^{n+1-x_i} \rightarrow d_{2i-1}^1$ in D_{2i-1} satisfies the terminal pair (c_{2i-1}^n, d_{2i-1}^1) of \mathcal{T}_i^V . Finally, the path $P_{2i}^{x_i}$ together with paths in C_{2i} and D_{2i} satisfies pair (c_{2i}^n, d_{2i}^1) of \mathcal{T}_i^V for each $i \in [k]$.

Lemma 2 guarantees that these families of canonical vertical and horizontal paths can be realized by some orientation of G (note that the canonical paths only orient black edges, and not green path edges whose orientation was already fixed at the start of this subsection). This implies that (G, \mathcal{T}) has a solution which satisfies all terminal pairs, concluding the proof Lemma 1.

2.3 Soundness of the reduction

In this subsection we prove the following lemma.

Lemma 3. *Let (G', k) be an input instance of k -CLIQUE, (G, \mathcal{T}) be instance of PLANAR STEINER ORIENTATION produced by the reduction, and $0 < \vartheta \leq \frac{1}{20}$. If there is an orientation G^* of G which satisfies at least $(\frac{19}{20} + \vartheta)$ -fraction of the terminal pairs, then there is a clique of size at least $20\vartheta \cdot k$ in G' . In particular, if G^* satisfies all terminal pairs, then G' has a clique of size k .*

Let G^* be the orientation which satisfies at least $(\frac{19}{20} + \vartheta)$ -fraction of terminal pairs from \mathcal{T} . Note that the set of vertices $A \cup \widehat{B} \cup C \cup \widehat{D}$ has no incoming edges. Similarly, the set of vertices $\widehat{A} \cup B \cup \widehat{C} \cup D$ has no outgoing edges.

The next lemma restricts the orientations of the undirected paths of green path edges.

Lemma 4. *In the orientation G^* , for each $i, j \in [k]$ we have that*

- if G^* satisfies all pairs of \mathcal{T}_j^L , then there exists an integer $\lambda_j \in [n]$ such that the path \widehat{A}_j is oriented away from $n+1 - \lambda_j$, the path A_{2j-1} is oriented towards $n+1 - \lambda_j$ and the path A_{2j} is oriented towards λ_j .

- if G^* satisfies all pairs of \mathcal{T}_j^R , then there exists an integer $\mu_j \in [n]$ such that the path \widehat{B}_j is oriented away from $n+1-\mu_j$, the path B_{2j-1} is oriented away from $n+1-\mu_j$ and the path B_{2j} is oriented towards μ_j .
- if G^* satisfies all pairs of \mathcal{T}_i^T , then there exists an integer $\delta_i \in [n]$ such that the path \widehat{C}_i is oriented away from $n+1-\delta_i$, the path C_{2i-1} is oriented towards $n+1-\delta_i$ and the path C_{2i} is oriented towards δ_i .
- if G^* satisfies all pairs of \mathcal{T}_i^B , then there exists an integer $\varepsilon_i \in [n]$ such that the path \widehat{D}_i is oriented away from $n+1-\varepsilon_i$, the path D_{2i-1} is oriented away from $n+1-\varepsilon_i$ and the path D_{2i} is oriented towards ε_i .

Proof. Fix $j \in [k]$. We just prove the lemma for the pairs of \mathcal{T}_j^L and the paths \widehat{A}_j, A_{2j-1} , and A_{2j} since the proof for other cases is similar. See again Figure 4 for an illustration. Since the only edges incoming to \widehat{A} are blue connector edges from A (and A has no incoming edges), it follows that the terminal pairs $(a_{2j-1}^n, \widehat{a}_j^1)$ and $(a_{2j-1}^1, \widehat{a}_j^n)$ of \mathcal{T}_j^L are satisfied by edges from the graph $G^*[\widehat{A}_j \cup A_{2j-1}]$. The path satisfying the terminal pair $(a_{2j-1}^n, \widehat{a}_j^1)$ has to travel upwards along A_{2j-1} , use a blue connector edge and then finally travel upwards along \widehat{A}_j . Similarly, the path satisfying the terminal pair $(a_{2j-1}^1, \widehat{a}_j^n)$ has to travel downwards along A_{2j-1} , use a blue connector edge and then finally travel downwards along \widehat{A}_j . Since we can only orient each green path edge in exactly one direction, it follows that both these paths must use the same blue connector edge, i.e., there exists an integer $\lambda'_i \in [n]$ such that the paths \widehat{A}_j, A_{2j-1} are oriented away from, and towards λ'_i , respectively. A similar argument shows that both the paths satisfying the terminal pairs $(a_{2j}^1, \widehat{a}_j^1)$ and $(a_{2j}^n, \widehat{a}_j^n)$ have to use the blue connector edge $a_{2j}^{n+1-\lambda'_i} \rightarrow \widehat{a}_j^{\lambda'_i}$. Hence, the path A_{2j} is oriented towards $n+1-\lambda'_i$. Setting $\lambda_i = n+1-\lambda'_i$, we obtain the desired result. \square

Lemma 5. *Let G^* be an orientation of G . For each $j \in [2k]$ if A_j is oriented towards $\alpha_j \in [n]$, path B_j is oriented away from $\beta_j \in [n]$, and the pair (a_j^n, b_j^1) is satisfied in G^* , then $\alpha_j \leq \beta_j$. If moreover $\alpha_j = \beta_j$, then G^* contains the horizontal canonical path $Q_j^{\alpha_j}$. For each $i \in [2k]$ if C_i is oriented towards $\gamma_i \in [n]$, path D_i is oriented away from $\zeta_i \in [n]$, and the pair (c_i^n, b_i^1) is satisfied in G^* , then $\gamma_i \leq \zeta_i$. If moreover $\gamma_i = \zeta_i$, then G^* contains the vertical canonical path $P_i^{\gamma_i}$.*

Proof. We only prove the first two claims, the proof for the other two is completely analogous. See again Figure 4 for an illustration. Fix any $j \in [2k]$ and consider the terminal pair (a_j^n, b_j^1) . Since G^* satisfies the pair (a_j^n, b_j^1) , it contains a path P from a_j^n to b_j^1 . Since \widehat{A} has no outgoing edges, P does not contain any vertices of \widehat{A} . Similarly, P does not contain any vertices of \widehat{B} . Hence, the path P consist of the following five parts: the vertical upwards path $a_j^n \rightarrow a_j^{n-1} \rightarrow \dots \rightarrow a_j^\tau$ followed by the blue connector edge $a_j^\tau \rightarrow v_{1,j}^{1,\tau}$, then a path in the graph $G^* \left[\bigcup_{i=1}^{2k} \bigcup_{j'=1}^{2k} V(G_{i,j'}) \right]$ followed by a blue connector edge $v_{2k,j}^{n,v} \rightarrow b_j^v$ and a vertical upwards path $b_j^v \rightarrow b_j^{v-1} \dots \rightarrow b_j^1$. If j is even, then the second blue connector edge is actually $v_{2k,j,TR}^{n,v} \rightarrow b_j^v$.

Since A_j is oriented towards α_j , for the first part of P to exist we have $\alpha_j \leq \tau$. Similarly, since B_j is oriented away from β_j , we have $\nu \leq \beta_j$. Furthermore, all edges in $G^* \left[\bigcup_{i=1}^{2k} \bigcup_{j'=1}^{2k} V(G_{i,j'}) \right]$, except for those formed by splitting vertices, are oriented either downward or to the right. Moreover, the splitting does not really change the row/column level. Therefore, the central part of P only goes down or right and $\tau \leq \nu$. Together we obtain $\alpha_j \leq \tau \leq \nu \leq \beta_j$, and the first claim follows.

Moreover, if $\alpha_j = \beta_j$, then $\tau = \nu = \alpha_j$ and the central part of P only goes right and from LB to TR within the split vertices. It follows that it is the horizontal canonical path $Q_j^{\alpha_j}$. \square

Lemma 6. *For each $i \in [k]$ and $j \in [k]$ and integers $\lambda_j, \mu_j, \delta_i, \varepsilon_i$ as given by Lemma 4 we have the following.*

- If G^* satisfies all pairs in $\mathcal{T}_j^L \cup \mathcal{T}_j^R \cup \mathcal{T}_j^H$, then $\lambda_j = \mu_j$. Moreover, G^* contains the horizontal canonical path $Q_{2j}^{\lambda_j}$.
- If G^* satisfies all pairs in $\mathcal{T}_i^T \cup \mathcal{T}_i^B \cup \mathcal{T}_i^V$, then $\delta_i = \varepsilon_i$. Moreover, G^* contains the vertical canonical path $P_{2i}^{\delta_i}$.
- If $i = j$ and both previous conditions hold, then moreover $\lambda_i = \delta_i$.

Proof. Fix $j \in [k]$. See again Figure 4 for an illustration. We first prove that $\lambda_j = \mu_j$, if G^* satisfies the pairs in $\mathcal{T}_j^L \cup \mathcal{T}_j^R \cup \mathcal{T}_j^H$. In this case, by Lemma 4, A_{2j-1} is oriented towards $n+1-\lambda_j$ and B_{2j-1} is oriented away from $n+1-\mu_j$ and from Lemma 5 it follows that $n+1-\lambda_j \leq n+1-\mu_j$. That is, $\lambda_j \geq \mu_j$. Since A_{2j} is oriented towards λ_j and B_{2j} is oriented away from μ_j , from Lemma 5 it follows that $\lambda_j \leq \mu_j$. Thus $\lambda_j = \mu_j$ and applying Lemma 5 again to paths A_{2j} and B_{2j} we get that G^* contains the horizontal canonical path $Q_{2j}^{\lambda_j}$.

The proof for the second part is completely analogous to the proof of the first part. Thus we move to the third part. From the previous two parts we know that G^* contains both the horizontal canonical path $Q_{2i}^{\lambda_i}$ and the vertical canonical path $P_{2i}^{\delta_i}$. It follows that the vertex $v_{2i,2i}^{\delta_i, \lambda_i}$ was not split, that is $(\delta_i, \lambda_i) \in S_{i,i}$. This implies $\delta_i = \lambda_i$ by the definition of $S_{i,i}$, concluding the proof. \square

Definition 5. *An index $i \in [k]$ is called good (for G^*) if all terminal pairs from $\mathcal{T}_i^L \cup \mathcal{T}_i^R \cup \mathcal{T}_i^T \cup \mathcal{T}_i^B \cup \mathcal{T}_i^H \cup \mathcal{T}_i^V$ are satisfied in G^* .*

Lemma 7. *If both $i, j \in [k]$, $i \neq j$ are good in G^* , then $\{u_{\delta_i}, u_{\delta_j}\} \in E(G')$. In particular $\delta_i \neq \delta_j$.*

Proof. By Lemma 6, if i, j are good, then G^* contains the horizontal canonical path $Q_{2j}^{\delta_j}$ and the vertical canonical path $P_{2i}^{\delta_i}$. It follows that the vertex $v_{2i,2j}^{\delta_i, \delta_j}$ was not split, that is $(\delta_i, \delta_j) \in S_{i,j}$. This implies $\{u_{\delta_i}, u_{\delta_j}\} \in E(G')$ by the definition of $S_{i,j}$. The second part follows as we do not allow loops. \square

Lemma 8. *Let $Y \subseteq [k]$ be the set of good indices. For each $0 \leq \vartheta \leq \frac{1}{20}$, if G^* satisfies at least $(\frac{19}{20} + \vartheta)$ -fraction of the terminal pairs, then $|Y| \geq 20\vartheta \cdot k$.*

Proof. If i is good then all 20 pairs from $\mathcal{T}_i^L \cup \mathcal{T}_i^R \cup \mathcal{T}_i^T \cup \mathcal{T}_i^B \cup \mathcal{T}_i^H \cup \mathcal{T}_i^V$ are satisfied in G^* . Otherwise, at most 19 pairs from this set are satisfied. Hence, the total number of satisfied pairs is at most $20|Y| + 19(k - |Y|) = 19k + |Y|$. However, we know that G^* satisfies at least $(\frac{19}{20} + \vartheta)|\mathcal{T}| = (\frac{19}{20} + \vartheta)20k = 19k + 20\vartheta \cdot k$ pairs. Hence, we have $|Y| \geq 20\vartheta \cdot k$. \square

Proof (of Lemma 3). If G^* satisfies at least $(\frac{19}{20} + \vartheta)$ -fraction of the terminal pairs, then by Lemma 8 there is a set Y of good indices of size at least $20\vartheta \cdot k$. By Lemma 7 the set $X = \{u_{\delta_i} \mid i \in Y\}$ forms a clique in G' of size at least $20\vartheta \cdot k$. \square

2.4 Proving the lower bounds

In this subsection we finish the proof of both our theorems.

Proof (of Theorem 1). It is easy to see that the graph G has $O(n^2k^2)$ vertices and can be constructed in $\text{poly}(n+k)$ time. Recall that (Remark 1) the graph G constructed in the STEINER ORIENTATION instance is planar. Combining this with the two directions from Lemma 1 and Lemma 3, we get a parameterized reduction from k -CLIQUE to PLANAR STEINER ORIENTATION. Hence, the $W[1]$ -hardness of PLANAR STEINER ORIENTATION follows from the $W[1]$ -hardness of k -CLIQUE. Chen et al. [4] showed that, for any function f , the existence of an $f(k) \cdot n^{o(k)}$ algorithm for k -CLIQUE violates ETH. Our reduction transforms an instance of k -CLIQUE into an equivalent instance of PLANAR STEINER ORIENTATION with $O(k)$ demand pairs. We obtain that under ETH there is no $f(k) \cdot n^{o(k)}$ time algorithm for PLANAR STEINER ORIENTATION. This concludes the proof. \square

For the inapproximability proof we need some more definitions. We will use the recent result on parameterized inapproximability of k -CLIQUE from [3]. To state the result precisely, let us first state the underlying assumption, the Gap Exponential Time Hypothesis (Gap-ETH).

Hypothesis 1 ((Randomized) Gap-ETH [11, 16]) *There exists a constant $\delta > 0$ such that, given a 3CNF formula Φ on n variables, no (possibly randomized) $2^{o(n)}$ -time algorithm can distinguish between the following two cases correctly with probability at least $2/3$:*

- Φ is satisfiable.
- Every assignment to the variables violates at least a δ -fraction of the clauses of Φ .

Here we do not attempt to reason why Gap-ETH is a plausible assumption; for more detailed discussions on the topic, please refer to [11] or [3]. For now, let us move on to state the inapproximability result from [3] that we need. Chalermsook et al. [3] showed that the only way to distinguish graphs with large cliques from those with small cliques is essentially to enumerate all vertex subsets of certain size, as stated more formally below.

Theorem 3 ([3, Theorem 18]). *Assuming Gap-ETH, there exist constants $\delta, r_0 > 0$ such that, for any function g and for any positive integers $q \geq r \geq r_0$, there is no algorithm that, given a graph G' , can distinguish between the following cases in $g(q, r) \cdot n^{\delta r}$ time, where $n = |V(G')|$:*

- $\text{CLIQUE}(G') \geq q$ and
- $\text{CLIQUE}(G') < r$,

where $\text{CLIQUE}(G')$ denotes the maximum size of a clique in G' .

While the construction presented in [3] is randomized, the authors claim that it can be derandomized so that the assumption can be limited to the (deterministic) Gap-ETH.

Proof (of Theorem 2). Let δ and r_0 be the constants from Theorem 3, ϑ be a constant such that $0 < \vartheta \leq \frac{1}{20}$ and let $\zeta = \frac{\vartheta\delta}{4}$. Let $r = \max\{3, r_0\}$ and $q = \lceil \frac{r}{20\vartheta} \rceil$. Note that since $r > 2$, we have $q > 2$ and, thus $q \leq 2\frac{r}{20\vartheta}$.

Assume for contradiction that there is a function f and an algorithm \mathbb{A} that can distinguish between the two cases:

- there is an orientation of G that satisfies all pairs of \mathcal{T} and
- any orientation of G satisfies less than $\frac{19}{20} + \vartheta$ fraction of pairs in \mathcal{T}

in $f(|\mathcal{T}|)(|V(G)|^{\zeta|\mathcal{T}|})$ time.

We design an algorithm \mathbb{B} that will contradict Theorem 3 for q and r as follows. Given a graph G' we apply our reduction from Subsection 2.1 to the instance (G', q) of CLIQUE. Then we apply algorithm \mathbb{A} to the resulting instance (G, \mathcal{T}) of STEINER ORIENTATION. If G' contains a clique of size q , then, by Lemma 1, there is an orientation of G satisfying all pairs from \mathcal{T} and the algorithm \mathbb{A} recognizes that. If each clique in G' is of size less than $r \leq 20\vartheta q$, then each orientation of G satisfies less than a $\frac{19}{20} + \vartheta$ fraction of pairs in \mathcal{T} . Indeed, by Lemma 3, if any orientation satisfied a $\frac{19}{20} + \vartheta$ fraction of pairs in \mathcal{T} , then there would be a clique of size $20\vartheta q \geq r$ in G' . Again, the algorithm \mathbb{A} recognizes that.

Now since $|V(G)| = O(q^2|V(G')|^2)$, the reduction can be performed in linear time in the size of the output, and \mathbb{A} runs in time $f(|\mathcal{T}|)(|V(G)|^{\zeta|\mathcal{T}|})$, we obtain that \mathbb{B} runs in time

$$\begin{aligned} f(|\mathcal{T}|)(|V(G)|^{\zeta|\mathcal{T}|}) &= f(20q)((O(q^2|V(G')|^2))^{\zeta \cdot 20q}) = \\ &= f(20q)((O(q^2))^{\zeta \cdot 20q})(|V(G')|^2)^{\zeta \cdot 20q} = g(q)(|V(G')|^2)^{\zeta \cdot 20q} \leq g(q)(|V(G')|^{\delta \cdot r}), \end{aligned}$$

for suitable g . The last inequality holds since $2\zeta \cdot 20q = 2\zeta \cdot 20 \cdot \lceil \frac{r}{20\vartheta} \rceil \leq 2\zeta \cdot 20 \cdot 2\frac{r}{20\vartheta} = \frac{4\zeta}{\vartheta} r = \frac{4\vartheta\delta}{4\vartheta} r = \delta r$ by the definition of ζ .

Thus the existence of \mathbb{A} would indeed contradict Theorem 3. \square

3 Open Problems

This work completely closes the gap between upper and lower bounds for the STEINER ORIENTATION problem parameterized by the number of terminal pairs k , even on planar graphs. However, there are still some interesting open questions to pursue. We list some of them below:

- Does STEINER ORIENTATION belong to the class $W[1]$?

- Can we show $W[2]$ -hardness for STEINER ORIENTATION parameterized by k ?
- Can we obtain $O(1)$ -FPT approximation for PLANAR STEINER ORIENTATION parameterized by k ?

Acknowledgements: We would like to thank the reviewers of both CSR '18 and of this version for their detailed and insightful comments which have greatly helped the presentation and readability.

References

1. E. M. Arkin and R. Hassin. A note on orientations of mixed graphs. *Discrete Applied Mathematics*, 116(3):271–278, 2002.
2. M. Beck, J. Blum, M. Kryven, A. Löfller, and J. Zink. Planar Steiner Orientation is NP-complete. *CoRR*, abs/1804.07496, 2018.
3. P. Chalermsook, M. Cygan, G. Kortsarz, B. Laekhanukit, P. Manurangsi, D. Nanongkai, and L. Trevisan. From Gap-ETH to FPT-Inapproximability: Clique, Dominating Set, and More. In *FOCS 2017*, pages 743–754, 2017.
4. J. Chen, X. Huang, I. A. Kanj, and G. Xia. Strong Computational Lower Bounds via Parameterized Complexity. *J. Comput. Syst. Sci.*, 72(8):1346–1367, 2006.
5. R. Chitnis, H. Esfandiari, M. T. Hajiaghayi, R. Khandekar, G. Kortsarz, and S. Seddighin. A Tight Algorithm for Strongly Connected Steiner Subgraph on Two Terminals with Demands. *Algorithmica*, 77(4):1216–1239, 2017.
6. R. Chitnis and A. E. Feldmann. A Tight Lower Bound for Steiner Orientation. In *CSR 2018*, pages 65–77, 2018.
7. R. Chitnis, A. E. Feldmann, and P. Manurangsi. Parameterized approximation algorithms for bidirected steiner network problems. In *26th Annual European Symposium on Algorithms, ESA 2018, August 20-22, 2018, Helsinki, Finland*, pages 20:1–20:16, 2018.
8. R. H. Chitnis, M. Hajiaghayi, and D. Marx. Tight bounds for planar strongly connected Steiner subgraph with fixed number of terminals (and extensions). In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 1782–1801, 2014.
9. M. Cygan, F. V. Fomin, L. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized Algorithms*. Springer, 2015.
10. M. Cygan, G. Kortsarz, and Z. Nutov. Steiner forest orientation problems. *SIAM J. Discrete Math.*, 27(3):1503–1513, 2013.
11. I. Dinur. Mildly exponential reduction from gap 3SAT to polynomial-gap label-cover. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:128, 2016.
12. I. Gamzu, D. Segev, and R. Sharan. Improved orientations of physical networks. In *International Workshop on Algorithms in Bioinformatics (WABI)*, pages 215–225, 2010.
13. R. Hassin and N. Megiddo. On orientations and shortest paths. *Linear Algebra and its Applications*, 114:589–602, 1989. Special Issue Dedicated to Alan J. Hoffman.
14. R. Impagliazzo and R. Paturi. On the Complexity of k -SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.
15. R. Impagliazzo, R. Paturi, and F. Zane. Which Problems Have Strongly Exponential Complexity? *J. Comput. System Sci.*, 63(4):512–530, 2001.
16. P. Manurangsi and P. Raghavendra. A Birthday Repetition Theorem and Complexity of Approximating Dense CSPs. In *ICALP*, pages 78:1–78:15, 2017.
17. D. Marx. On the optimality of planar and geometric approximation schemes. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007), October 20-23, 2007, Providence, RI, USA, Proceedings*, pages 338–348, 2007.

18. D. Marx. Can you beat treewidth? *Theory of Computing*, 6(1):85–112, 2010.
19. D. Marx. A Tight Lower Bound for Planar Multiway Cut with Fixed Number of Terminals. In *ICALP (1)*, pages 677–688, 2012.
20. D. Marx and M. Pilipczuk. Everything you always wanted to know about the parameterized complexity of subgraph isomorphism (but were afraid to ask). In *31st International Symposium on Theoretical Aspects of Computer Science (STACS 2014), STACS 2014, March 5-8, 2014, Lyon, France*, pages 542–553, 2014.
21. D. Marx and M. Pilipczuk. Optimal parameterized algorithms for planar facility location problems using voronoi diagrams. In *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, pages 865–877, 2015.
22. A. Medvedovsky, V. Bafna, U. Zwick, and R. Sharan. An algorithm for orienting graphs based on cause-effect pairs and its applications to orienting protein networks. In *International Workshop on Algorithms in Bioinformatics (WABI)*, pages 222–232, 2008.
23. M. Pilipczuk and M. Wahlström. Directed multicut is $W[1]$ -hard, even for four terminal pairs. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1167–1178, 2016.
24. M. Pilipczuk and M. Wahlström. Directed multicut is $W[1]$ -hard, even for four terminal pairs. *TOCT*, 10(3):13:1–13:18, 2018.
25. H. E. Robbins. A theorem on graphs, with an application to a problem of traffic control. *The American Mathematical Monthly*, 46(5):281–283, 1939.