

New Hardness Results for Routing on Disjoint Paths

Julia Chuzhoy*

David H. K. Kim[†]

Rachit Nimavat[‡]

November 16, 2016

Abstract

In the classical Node-Disjoint Paths (NDP) problem, the input consists of an undirected n -vertex graph G , and a collection $\mathcal{M} = \{(s_1, t_1), \dots, (s_k, t_k)\}$ of pairs of its vertices, called source-destination, or demand, pairs. The goal is to route the largest possible number of the demand pairs via node-disjoint paths. The best current approximation for the problem is achieved by a simple greedy algorithm, whose approximation factor is $O(\sqrt{n})$, while the best current negative result is an $\Omega(\log^{1/2-\delta} n)$ -hardness of approximation for any constant δ , under standard complexity assumptions. Even seemingly simple special cases of the problem are still poorly understood: when the input graph is a grid, the best current algorithm achieves an $\tilde{O}(n^{1/4})$ -approximation, and when it is a general planar graph, the best current approximation ratio of an efficient algorithm is $\tilde{O}(n^{9/19})$. The best currently known lower bound on the approximability of both these versions of the problem is APX-hardness.

In this paper we prove that NDP is $2^{\Omega(\sqrt{\log n})}$ -hard to approximate, unless all problems in NP have algorithms with running time $n^{O(\log n)}$. Our result holds even when the underlying graph is a planar graph with maximum vertex degree 3, and all source vertices lie on the boundary of a single face (but the destination vertices may lie anywhere in the graph). We extend this result to the closely related Edge-Disjoint Paths problem, showing the same hardness of approximation ratio even for sub-cubic planar graphs with all sources lying on the boundary of a single face.

*Toyota Technological Institute at Chicago. Email: cjulia@ttic.edu. Supported in part by NSF grants CCF-1318242 and CCF-1616584.

[†]Computer Science Department, University of Chicago. Email: hongk@cs.uchicago.edu. Supported in part by NSF grant CCF-1318242.

[‡]Toyota Technological Institute at Chicago. Email: nimavat@ttic.edu. Supported in part by NSF grant CCF-1318242.

1 Introduction

The main focus of this paper is the Node-Disjoint Paths (NDP) problem: given an undirected n -vertex graph G , together with a collection $\mathcal{M} = \{(s_1, t_1), \dots, (s_k, t_k)\}$ of pairs of its vertices, called *source-destination*, or *demand* pairs, route the largest possible number of the demand pairs via node-disjoint paths. In other words, a solution to the problem is a collection \mathcal{P} of node-disjoint paths, with each path connecting a distinct source-destination pair, and the goal is to maximize $|\mathcal{P}|$. The vertices participating in the demand pairs of \mathcal{M} are called *terminals*. NDP is a classical routing problem, that has been extensively studied in both Graph Theory and Theoretical Computer Science communities. One of the key parts of Robertson and Seymour's Graph Minors series is an efficient algorithm for the special case of the problem, where the number k of the demand pairs is bounded by a constant [RS90, RS95]; the running time of their algorithm is $f(k) \cdot \text{poly}(n)$ for some large function f . However, when k is a part of input, the problem becomes NP-hard [Kar72, EIS76], even on planar graphs [Lyn75], and even on grid graphs [KvL84]. The following simple greedy algorithm provides an $O(\sqrt{n})$ -approximation for NDP [KS04]: Start with $\mathcal{P} = \emptyset$. While G contains a path connecting any demand pair, select the shortest such path P , add it to \mathcal{P} , and delete all vertices of P from G . Surprisingly, despite the extensive amount of work on the problem and its variations, this elementary algorithm remains the best currently known approximation algorithm for the problem, and until recently, this was true even for the special cases where G is a planar graph, or a grid graph. The latter two special cases have slightly better algorithms now: a recent result of Chuzhoy and Kim [CK15] gives a $\tilde{O}(n^{1/4})$ -approximation for NDP on grid graphs, and Chuzhoy, Kim and Li [CKL16] provide a $\tilde{O}(n^{9/19})$ -approximation algorithm for the problem on planar graphs. The best current negative result shows that NDP has no $O(\log^{1/2-\delta} n)$ -approximation algorithms for any constant δ , unless $\text{NP} \subseteq \text{ZPTIME}(n^{\text{poly} \log n})$ [AZ05, ACG⁺10]. For the special case of grids and planar graphs only APX-hardness is currently known on the negative side [CK15].

The main result of this paper is that NDP is $2^{\Omega(\sqrt{\log n})}$ -hard to approximate unless $\text{NP} \subseteq \text{DTIME}(n^{O(\log n)})$, even if the underlying graph is a planar graph with maximum vertex degree at most 3, and all source vertices $\{s_1, \dots, s_k\}$ lie on the boundary of a single face. We note that NDP can be solved efficiently on graphs whose maximum vertex degree is 2^1 .

A problem closely related to NDP is Edge-Disjoint Paths (EDP). The input to this problem is the same as to NDP, and the goal is again to route the largest possible number of the demand pairs. However, the routing paths are now allowed to share vertices, as long as they remain disjoint in their edges. The two problems are closely related: it is easy to see that EDP is a special case of NDP, by using the line graph of the EDP instance to obtain an equivalent NDP instance (but this transformation may inflate the number of vertices, and so approximation factors depending on n may not be preserved). This relationship is not known for planar graphs, as the line graph of a planar graph is not necessarily planar. The current approximability status of EDP is similar to that of NDP: the best current approximation algorithm achieves an $O(\sqrt{n})$ -approximation factor [CKS06], and the best current negative result is an $\Omega(\log^{1/2-\delta} n)$ -hardness of approximation for any constant δ , unless $\text{NP} \subseteq \text{ZPTIME}(n^{\text{poly} \log n})$ [AZ05, ACG⁺10]. An analogue of the special case of NDP on grid graphs for the EDP problem is when the input graph is a wall, and the work of [CK15] gives an $\tilde{O}(n^{1/4})$ -approximation algorithm for EDP on wall graphs. However, for planar graphs, no better than $O(\sqrt{n})$ -approximation is currently known for EDP, and it is not clear

¹A graph whose maximum vertex degree is 2 is a collection of disjoint paths and cycles. It is enough to solve the problem on each such path and cycle separately. The problem is then equivalent to computing a maximum independent set of intervals on a line or on a circle, and can be solved efficiently by standard methods.

whether the algorithm of [CKL16] can be adapted to this setting in order to break the $O(\sqrt{n})$ -approximation barrier for EDP in planar graphs. Our hardness result extends to EDP on planar sub-cubic graphs, where all source vertices lie on the boundary of a single face.

We say that two instances (G, \mathcal{M}) and (G', \mathcal{M}') of the NDP problem are *equivalent* iff the sets of terminals in both instances are the same, $\mathcal{M} = \mathcal{M}'$, and for every subset $\tilde{\mathcal{M}} \subseteq \mathcal{M}$ of demand pairs, the pairs in $\tilde{\mathcal{M}}$ are routable via node-disjoint paths in G iff they are routable in G' . Equivalence of EDP problem instances is defined similarly, with respect to edge-disjoint routing. We show in Section C of the Appendix that for every integer $d > 3$, there is an instance (G, \mathcal{M}) of NDP, where G is a planar graph, such that for every instance (G', \mathcal{M}') of NDP that is equivalent to (G, \mathcal{M}) , some vertex of G' has degree at least d . Therefore, the class of all planar graphs is strictly more general than the class of all planar graphs with maximum vertex degree at most 3 for NDP. In contrast, it is well known that for any instance (G, \mathcal{M}) of EDP, there is an equivalent instance (G', \mathcal{M}') of the problem, with maximum vertex degree at most 4^2 . Moreover, if G is planar then G' can also be made planar. Informally, graph G' is obtained from G by replacing every large-degree vertex with a grid: if the degree of a vertex v is $d_v > 4$, then we replace v with the $(d_v \times d_v)$ -grid, and connect the edges incident to v to the vertices on the first row of the grid. This transformation inflates the number of vertices, but can be performed in a way that preserves planarity. However, as we show in Section D, there is an instance (G, \mathcal{M}) of EDP, where G is a planar graph, such that for every instance (G', \mathcal{M}') of EDP that is equivalent to (G, \mathcal{M}) , such that G' is planar, the maximum vertex degree in G' is at least 4. This shows that we cannot reduce the degree all the way to 3.

Interestingly, better algorithms are known for several special cases of EDP on planar graphs. Kleinberg [Kle05], building on the work of Chekuri, Khanna and Shepherd [CKS05, CKS04], has shown an $O(\log^2 n)$ -approximation algorithm for even-degree planar graphs. Aumann and Rabani [AR95] showed an $O(\log^2 n)$ -approximation algorithm for EDP on grids, and Kleinberg and Tardos [KT98, KT95] showed $O(\log n)$ -approximation algorithms for broader classes of nearly-Eulerian uniformly high-diameter planar graphs, and nearly-Eulerian densely embedded graphs. Recently, Kawarabayashi and Kobayashi [KK13] gave an $O(\log n)$ -approximation algorithm for EDP on 4-edge-connected planar graphs and on Eulerian planar graphs. It seems that the restriction of the graph G to be planar Eulerian, or nearly-Eulerian, makes the EDP problem significantly more tractable. In contrast, the graphs we construct in our hardness of approximation proof are sub-cubic, and far from being Eulerian.

A variation of the NDP and EDP problems, where small congestion is allowed, has been a subject of extensive study. In the NDP with congestion (NDPwC) problem, the input is the same as in the NDP problem, and we are additionally given an integer $c \geq 1$. The goal is to route as many of the demand pairs as possible with congestion at most c : that is, every vertex may participate in at most c paths in the solution. EDP with Congestion (EDPwC) is defined similarly, except that now the congestion bound is imposed on the graph edges and not vertices. The classical randomized rounding technique of Raghavan and Thompson [RT87] gives a constant-factor approximation for both problems, if the congestion c is allowed to be as high as $\Theta(\log n / \log \log n)$. A long line of work [CKS05, R ac02, And10, RZ10, Chu12, CL12, CE13, CC] has led to an $O(\text{poly log } k)$ -approximation for both NDPwC and EDPwC problems, with congestion bound $c = 2$. For planar graphs, a constant-factor approximation with congestion 2 is known for EDP [SCS11]. Our new

²We use a slightly different definition of equivalence for EDP: namely, we only require that for every subset $\tilde{\mathcal{M}}$ of demand pairs, such that every terminal participates in at most one demand pair in $\tilde{\mathcal{M}}$, set $\tilde{\mathcal{M}}$ is routable in G iff it is routable in G' . This definition is more appropriate for the EDP problem, as we discuss in Section D.

hardness results demonstrate that there is a dramatic difference in the approximability of routing problems with congestion 1 and 2.

Our Results and Techniques. Our main result is the proof of the following two theorems.

Theorem 1.1 *There is a constant c , such that no efficient algorithm achieves a factor $2^{c\sqrt{\log n}}$ -approximation for NDP, unless $\text{NP} \subseteq \text{DTIME}(n^{O(\log n)})$. This result holds even for planar graphs with maximum vertex degree 3, where all source vertices lie on the boundary of a single face.*

Theorem 1.2 *There is a constant c , such that no efficient algorithm achieves a factor $2^{c\sqrt{\log n}}$ -approximation for EDP, unless $\text{NP} \subseteq \text{DTIME}(n^{O(\log n)})$. This result holds even for planar graphs with maximum vertex degree 3, where all source vertices lie on the boundary of a single face.*

We now provide an informal high-level overview of the proof of Theorem 1.1. It is somewhat easier to describe the proof of the theorem for the case where the maximum vertex degree is allowed to be 4 instead of 3. This proof can then be easily modified to ensure that the maximum vertex degree in the instances we obtain does not exceed 3, and also extended to the EDP problem. We perform a reduction from the 3SAT(5) problem. In this problem, we are given a SAT formula φ defined over a set of n Boolean variables. The formula consists of m clauses, each of which is an OR of three literals, where every literal is either a variable or its negation. Every variable of φ participates in exactly 5 distinct clauses, and the literals of every clause correspond to three distinct variables. We say that φ is a YES-INSTANCE, if there is an assignment to its variables that satisfies all clauses, and we say that it is a NO-INSTANCE, if no assignment satisfies more than a $(1 - \epsilon)$ -fraction of the clauses, for some fixed constant $0 < \epsilon < \frac{1}{2}$. The famous PCP theorem [AS98, ALM⁺98] shows that, unless $\text{P} = \text{NP}$, no efficient algorithm can distinguish between the YES- and the NO-INSTANCES of 3SAT(5).

We perform $\Theta(\log n)$ iterations, where in iteration i we construct what we call a *level- i* instance. We use two parameters, N_i and N'_i , and ensure that, if the reduction is performed from a formula φ which is a YES-INSTANCE, then there is a solution to the level- i instance of NDP that routes N_i demand pairs, while if φ is a NO-INSTANCE, then no solution routes more than N'_i demand pairs. We let $g_i = N_i/N'_i$ be the gap achieved by the level- i instance. Our construction ensures that the gap grows by a small constant factor in every iteration, so $g_i = 2^{\Theta(i)}$, while the instance size grows by roughly factor- $\Theta(n \cdot g_{i-1})$ in iteration i . Therefore, after $\Theta(\log n)$ iterations, the gap becomes $2^{\Omega(\log n)}$, while the instance size becomes $n' = 2^{O(\log^2 n)}$, giving us the desired $2^{\Omega(\sqrt{\log n'})}$ -hardness of approximation, unless $\text{NP} \subseteq \text{DTIME}(n^{O(\log n)})$.

In all our instances of NDP, the underlying graph is a subgraph of a grid, with all sources lying on the top boundary of the grid; all vertices participating in the demand pairs are distinct. In the first iteration, a level-1 instance is constructed by a simple reduction from 3SAT(5), achieving a small constant gap g_1 . Intuitively, once we construct a level- i instance, in order to construct a level- $(i+1)$ instance, we replace every demand pair from a level-1 instance with a collection of level- i instances. In order to be able to do so, we need the instances to be “flexible”, so that, for example, we have some freedom in choosing the locations of the source and the destination vertices of a given level- i instance in the grid.

We achieve this flexibility by defining, for each level i , a family of level- i instances. The graph associated with a level- i instance \mathcal{I} is a subgraph of a large enough grid G_i . The construction of the instance consists of two parts. First, we construct a path $Z(\mathcal{I})$, and place all source vertices on

this path. Second, we construct a vertex-induced subgraph $B(\mathcal{I})$ of a relatively small grid G'_i , and we call $B(\mathcal{I})$ a *box*. We place all the destination vertices inside the box $B(\mathcal{I})$. Graphs $Z(\mathcal{I})$ and G'_i are completely disjoint from the grid G_i and from each other. In order to construct a specific level- i instance, we select a placement of the path $Z(\mathcal{I})$ on the first row of the grid G_i , and a placement of the box $B(\mathcal{I})$ in G_i , far enough from its boundaries (see Figure 1). In other words, we choose a sub-path P of the first row of the grid, of the same length as $Z(\mathcal{I})$, and map the vertices of $Z(\mathcal{I})$ to P in a natural way. We also choose a sub-grid G''_i of G_i , of the same dimensions as G'_i , and map the vertices of G'_i to the vertices of G''_i in a natural way. Since $B(\mathcal{I}) \subseteq G'_i$, this also defines a mapping of the vertices of $B(\mathcal{I})$ to the vertices of G_i . Once these placements are selected, the mapping of the vertices of $Z(\mathcal{I})$ to the vertices of P determines the identities of the source vertices, and the mapping of the vertices of $B(\mathcal{I})$ to the vertices of G''_i determines the identities of the destination vertices. We delete from G_i all vertices to which the vertices of $G'_i \setminus B(\mathcal{I})$ are mapped. In other words, all vertices that were removed from G'_i to construct $B(\mathcal{I})$, are also removed from G''_i , and hence from G_i . We note that box $B(\mathcal{I})$ may be constructed recursively, for example, by placing several boxes $B(\mathcal{I}')$ corresponding to lower-level instances \mathcal{I}' inside it. The mapping of the vertices of $B(\mathcal{I}')$ to the vertices of $B(\mathcal{I})$, the placement of the destination vertices, and the removal of the vertices of $B(\mathcal{I})$ corresponding to the grid vertices removed from $B(\mathcal{I}')$ is done similarly. In order to reduce the maximum vertex degree to 3, we can use wall graphs instead of grid graphs and employ a similar proof. Alternatively, a simple modification of the final instance we obtain can directly reduce its maximum vertex degree to 3.

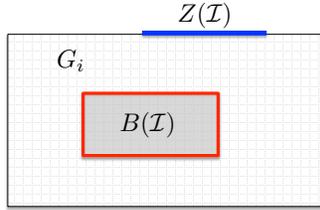


Figure 1: A schematic view of a level- i instance \mathcal{I} . All source vertices lie on $Z(\mathcal{I})$, whose location can be chosen arbitrarily on the first row of G_i . All destination vertices belong to $B(\mathcal{I})$, that can be located anywhere in G_i , far enough from its boundaries.

The most natural intuitive way to think about our construction is the one described above. An equivalent, and somewhat easier way to define our construction is slightly different: we let a level-0 instance be an instance consisting of a single demand pair (s, t) , with s lying on the first row of the grid and t lying far from the grid boundary. We then show, for each $i > 0$, a procedure that constructs a level- i instance by combining a number of level- $(i - 1)$ instances. The latter definition is somewhat more convenient, because it saves us the need to provide a separate correctness proof for level-1 instances, which is essentially identical to the proof for higher-level instances. However, we still feel that defining level-1 instances explicitly is useful for the sake of intuition. Therefore, we start with preliminaries in Section 2 and describe our construction of level-1 instances in Section 3, together with an intuition for constructing higher-level instances. We only provide a sketch of the correctness proof, as the complete correctness proof appears in the following sections. In Section 4, we define our construction in two steps: by first defining level-0 instances, and then showing how to construct level- i instances from level- $(i - 1)$ instances. The resulting level-1 instances will be similar to those defined in Section 3. We provide a complete proof of correctness in Sections 4–6. This provides a proof of Theorem 1.1 for the case where the maximum vertex degree is allowed to

be 4. Section B of the Appendix extends our result to EDP in planar graphs, and show how to reduce the degree of the hard NDP instances to 3, completing the proofs of Theorems 1.1 and 1.2. Finally, sections C and D discuss degree reduction in NDP and EDP instances, respectively.

2 Preliminaries

For a pair $\ell, h > 0$ of integers, we let $G^{\ell, h}$ denote a grid of length ℓ and height h . The set of vertices of $G^{\ell, h}$ is $V(G^{\ell, h}) = \{v(i, j) \mid 1 \leq i \leq \ell, 1 \leq j \leq h\}$, and the set of its edges is the union of two subsets: the set of horizontal edges $E^H = \{(v_{i,j}, v_{i,j+1}) \mid 1 \leq i \leq \ell, 1 \leq j < \ell\}$ and the set of vertical edges $E^V = \{(v_{i,j}, v_{i+1,j}) \mid 1 \leq i < \ell, 1 \leq j \leq h\}$. The subgraph $G^{\ell, h}$ induced by the edges of E^H consists of h paths, that we call the *rows* of the grid; for $1 \leq i \leq h$, the i th row R_i is the row containing the vertex $v(i, 1)$. Similarly, the subgraph induced by the edges of E^V consists of ℓ paths that we call the *columns* of the grid, and for $1 \leq j \leq \ell$, the j th column W_j is the column containing $v(1, j)$. We think of the rows as ordered from top to bottom and the columns as ordered from left to right. Row $R_{\lceil h/2 \rceil}$ is called the *middle row* of the grid $G^{\ell, h}$. Given two vertices $u = v(i, j), u' = v(i', j')$ of the grid, the distance between them is $d(u, u') = |i - i'| + |j - j'|$. Given two vertex subsets $X, Y \subseteq V(G^{\ell, h})$, the distance between them is $d(X, Y) = \min_{u \in X, u' \in Y} \{d(u, u')\}$. Given a vertex $v = v(i, j)$ of the grid, we denote by $\text{row}(v)$ and $\text{col}(v)$ the row and the column, respectively, that contain v .

Given a set \mathcal{R} of consecutive rows of a grid $G = G^{\ell, h}$ and a set \mathcal{W} of consecutive columns of G , we let $\Upsilon(\mathcal{R}, \mathcal{W})$ be the subgraph of G induced by the set $\{v(j, j') \mid R_j \in \mathcal{R}, W_{j'} \in \mathcal{W}\}$ of vertices. We say that $\Upsilon = \Upsilon(\mathcal{R}, \mathcal{W})$ is the *sub-grid of G spanned by the set \mathcal{R} of rows and the set \mathcal{W} of columns*.

Assume now that we are given a grid G , a sequence $\mathcal{S} = (G_1, \dots, G_r)$ of disjoint sub-grids of G , and an integer N . We say that the grids of \mathcal{S} are *aligned and N -separated* iff the middle row of each grid G_i is a sub-path of the middle row of G ; the grids in $\{G_1, \dots, G_r\}$ appear in this left-to-right order inside G ; every pair of consecutive grids G_i is separated by at least N columns of G ; and every grid in \mathcal{S} is separated by at least N columns from the right and the left boundaries of G .

Let R', R'' be two distinct rows of some grid G , and let \mathcal{P} be a set of node-disjoint paths, such that every path in \mathcal{P} has one endpoint (called a source) on row R' and another (called a destination) on row R'' . We say that the set \mathcal{P} of paths is *order-preserving* iff the source vertices of the paths in \mathcal{P} appear on row R' in exactly the same left-to-right order as their destination vertices on R'' .

Throughout our construction, we use the notion of a box. A box B of length ℓ and height h is a vertex-induced subgraph of $G^{\ell, h}$. We denote $U(B) = V(G^{\ell, h}) \setminus V(B)$, and we sometimes think of set $U(B)$ as the “set of vertices deleted from B ”. We say that B is a *cut-out box* iff $U(B)$ contains all vertices lying on the left, right, and bottom boundaries of $G^{\ell, h}$; note that $U(B)$ may contain additional vertices of $G^{\ell, h}$. The remaining vertices of the top boundary of $G^{\ell, h}$ that belong to $V(B)$ are called *the opening of B* . We sometimes say that the vertices of B that belong to row $R_{\lceil h/2 \rceil}$ of $G^{\ell, h}$ lie on the middle row of B .

Given any set \mathcal{M} of demand pairs, we let $S(\mathcal{M})$ denote the set of all source vertices participating in \mathcal{M} and $T(\mathcal{M})$ the set of all destination vertices. Given a path P , the length of the path is the number of vertices on it. We say that two paths P, P' are *internally disjoint* iff every vertex in $P \cap P'$ is an endpoint of both paths. Given a set \mathcal{P} of paths and some set U of vertices, we say that the paths in \mathcal{P} are internally disjoint from U iff for every path $P \in \mathcal{P}$, every vertex in $P \cap U$ is an endpoint of P . Given a subgraph G' of a graph G and a set \mathcal{P} of paths in G , we say that the

paths in \mathcal{P} are internally disjoint from G' iff they are internally disjoint from $V(G')$.

As already described in the introduction, for every level $0 \leq i \leq \Theta(\log n)$, we construct a level- i instance \mathcal{I} . In fact, it is a family of instances, but it is more convenient to think of it as one instance with different instantiations. A definition of a level- i instance \mathcal{I} consists of the following ingredients:

- integral parameters L_i, L'_i and an even integer H_i ;
- a path $Z(\mathcal{I})$ of length L_i ;
- a grid G'_i of length L'_i and height H_i , together with a cut-out box $B(\mathcal{I}) \subseteq G'_i$; and
- a set \mathcal{M} of demand pairs, together with a mapping of the vertices of $S(\mathcal{M})$ to distinct vertices of $Z(\mathcal{I})$ and a mapping of the vertices of $T(\mathcal{M})$ to distinct vertices on the middle row of $B(\mathcal{I})$.

In order to instantiate a level- i instance \mathcal{I} , we select a grid G_i of length at least $2L_i + 2L'_i + 4H_i$ and height at least $3H_i$, a sub-path P of the first row of G_i of length L_i , and a sub-grid G''_i of G_i of height H_i and length L'_i , so that the distance from the vertices of G''_i to the vertices lying on the boundary of G_i is at least H_i . We map every vertex of $Z(\mathcal{I})$ to the corresponding vertex of P in a natural way, and this determines the identities of the source vertices in the instance we construct. We also map every vertex of G'_i to the corresponding vertex of G''_i , and this determines the identities of the destination vertices. Finally, for every vertex $u \in U(B(\mathcal{I}))$, we delete the vertex of G''_i to which u is mapped from G_i . This defines an instance of NDP on a subgraph of G_i , where all the sources lie on the top boundary of G_i and all source and destination vertices are distinct.

Assume now that we are given an instantiation of a level- i instance \mathcal{I} and a set \mathcal{P} of node-disjoint paths routing a subset \mathcal{M}' of the demand pairs in that instance. Assume for convenience that $\mathcal{M}' = \{(s_1, t_1), \dots, (s_r, t_r)\}$, that the vertices s_1, \dots, s_r appear in this left-to-right order on $Z(\mathcal{I})$, and that $\mathcal{P} = \{P_1, \dots, P_r\}$, where path P_j connects s_j to t_j . Let A be the set of all vertices of the top row of the grid G''_i that were not deleted (that is, these are the vertices lying on the opening of $B(\mathcal{I})$). We say that the set \mathcal{P} of paths *respects* the box $B(\mathcal{I})$ iff for all $1 \leq j \leq r$, $P_j \cap A$ is a single vertex, that we denote by u_j , and u_j is the j th vertex of A from the left. Intuitively, the paths in \mathcal{P} connect the sources to a set of consecutive vertices on the opening of $B(\mathcal{I})$ in a straightforward manner, and the actual routing occurs inside the box $B(\mathcal{I})$.

We perform a reduction from the 3SAT(5) problem. In this problem, we are given a SAT formula φ on a set $\{x_1, \dots, x_n\}$ of n Boolean variables and a set $\mathcal{C} = \{C_1, \dots, C_m\}$ of $m = 5n/3$ clauses. Each clause contains 3 literals, each of which is either a variable or its negation. The literals of each clause correspond to 3 distinct variables, and each variable participates in exactly 5 clauses. We denote the literals of the clause C_q by ℓ_{q_1}, ℓ_{q_2} and ℓ_{q_3} . A clause is satisfied by an assignment to the variables iff at least one of its literals evaluates to TRUE. We say that φ is a YES-INSTANCE if there is an assignment to its variables satisfying all its clauses. We say that it is a NO-INSTANCE with respect to some parameter ϵ , if no assignment satisfies more than $(1 - \epsilon)m$ clauses. The following well-known theorem follows from the PCP theorem [AS98, ALM⁺98].

Theorem 2.1 *There is a constant $\epsilon : 0 < \epsilon < \frac{1}{2}$, such that it is NP-hard to distinguish between the YES-INSTANCES and the NO-INSTANCES (defined with respect to ϵ) of the 3SAT(5) problem.*

Given an input formula φ , we will construct an instance (G, \mathcal{M}) of the NDP problem with $|V(G)| = n' = n^{O(\log n)}$, that has the following properties: if φ is a YES-INSTANCE, then there is a solution

to the NDP instance routing N demand pairs, for some parameter N ; if φ is a NO-INSTANCE, then at most N/g demand pairs can be routed, where $g = 2^{\Omega(\log n)} = 2^{\Omega(\sqrt{\log n})}$. This will prove that no efficient algorithm can achieve a better than factor $2^{O(\sqrt{\log n})}$ -approximation for NDP, unless $\text{NP} \subseteq \text{DTIME}(n^{O(\log n)})$. The instance we construct is a subgraph of a grid with all source vertices lying on its top boundary, so the hardness result holds for planar graphs with maximum vertex degree 4, with all sources lying on the boundary of a single face. In Section B of the Appendix, we modify this instance to reduce its maximum vertex degree to 3.

3 The Level-1 Instance

In this section we define our level-1 instance \mathcal{I} and provide intuition for generalizing it to higher-level instances. Since Sections 4–6 contain all formal definitions and proofs, including those for the level-1 instance, the description here is informal, and we only provide proof sketches.

We assume that we are given a 3SAT(5) formula φ defined over a set $\{x_1, \dots, x_n\}$ of variables and a set $\mathcal{C} = \{C_1, \dots, C_m\}$ of clauses, so $m = 5n/3$. For every variable x_j of φ , we will define a set $\mathcal{M}(x_j)$ of demand pairs that represent that variable, and similarly, for every clause $C_q \in \mathcal{C}$ we will define a set $\mathcal{M}(C_q)$ of demand pairs representing it. We call the demand pairs in set $\mathcal{M}^V = \bigcup_{j=1}^n \mathcal{M}(x_j)$ *variable-pairs* and the demand pairs in set $\mathcal{M}^C = \bigcup_{C_q \in \mathcal{C}} \mathcal{M}(C_q)$ *clause-pairs*.

Let $h = 1000/\epsilon$ and $\delta = 8\epsilon^2/10^{12}$, where ϵ is the parameter from Theorem 2.1. We set $N_1 = (200h/3 + 1)n$ and $N'_1 = (1 - \delta)N_1$. Our construction will ensure that, if the input formula φ is a YES-INSTANCE, then for every instantiation of \mathcal{I} , there is a collection \mathcal{P} of node-disjoint paths that respects the box $B(\mathcal{I})$ and routes N_1 demand pairs. On the other hand, if φ is a NO-INSTANCE, then no solution can route more than N'_1 demand pairs in any instantiation of \mathcal{I} . This gives a gap of $1/(1 - \delta)$ between the YES- and NO-INSTANCE solution costs. In the following levels we gradually amplify this gap.

We set $L_1 = (80h + 2)n$, $L'_1 = 20N_1^3$ and $H_1 = 20N_1$. In order to construct a level-1 instance \mathcal{I} , we start with a path $Z(\mathcal{I})$ of length L_1 and a grid G'_1 of length L'_1 and height H_1 . We delete all vertices lying on the bottom, left and right boundaries of G'_1 to obtain the initial cut-out box $B(\mathcal{I})$; we will later delete some additional vertices from $B(\mathcal{I})$.

We define two sub-grids of G'_1 : grid B^V , that will contain all vertices of $T(\mathcal{M}^V)$ (the destination vertices of the demand pairs in \mathcal{M}^V), and grid B^C , that will contain all vertices of $T(\mathcal{M}^C)$. Both grids have sufficiently large length and height: length $9N_1^3$ and height $16N_1$ for each grid are sufficient. We place both grids inside G'_1 , so that the middle row of each grid is contained in the middle row of G'_1 , there is a horizontal spacing of at least $2N_1$ between the two grids, and both grids are disjoint from the left and the right boundaries of G'_1 . It is easy to see that at least $2N_1$ rows of G'_1 lie above and below both grids (see Figure 2).

Next, we define smaller sub-grids of the grids B^V and B^C . For every variable x_j of φ , we select a sub-grid $B(x_j)$ of B^V of height $H^V = 4N_1 + 3$ and length $L^V = 4N_1 + (70h + 2)$. This is done so that the sequence $B(x_1), \dots, B(x_n)$ of grids is aligned and $(2N_1)$ -separated in B^V (see Figure 2). Recall that this means that the middle row of each grid coincides with the middle row of B^V , and the horizontal distance between every pair of these grids, and between each grid and the right and the left boundaries of B^V is at least $2N_1$. It is easy to verify that grid B^V is large enough to allow this. For each variable x_j of φ , the vertices of $T(\mathcal{M}(x_j))$ will lie in $B(x_j)$. Note that there are at least $2N_1$ rows of B^V above and below each box $B(x_1), \dots, B(x_n)$.

Similarly, for every clause $C_q \in \mathcal{C}$, we define a sub-grid $B(C_q)$ of B^C of length $L^C = 3h$ and height $H^C = 3$. We select the sub-grids $B(C_1), \dots, B(C_m)$ of B^C so that they are aligned and $(4N_1)$ -separated. As before, box B^C is sufficiently large to allow this, and there are at least $2N_1$ rows of B^C both above and below each such grid $B(C_q)$ (see Figure 2).

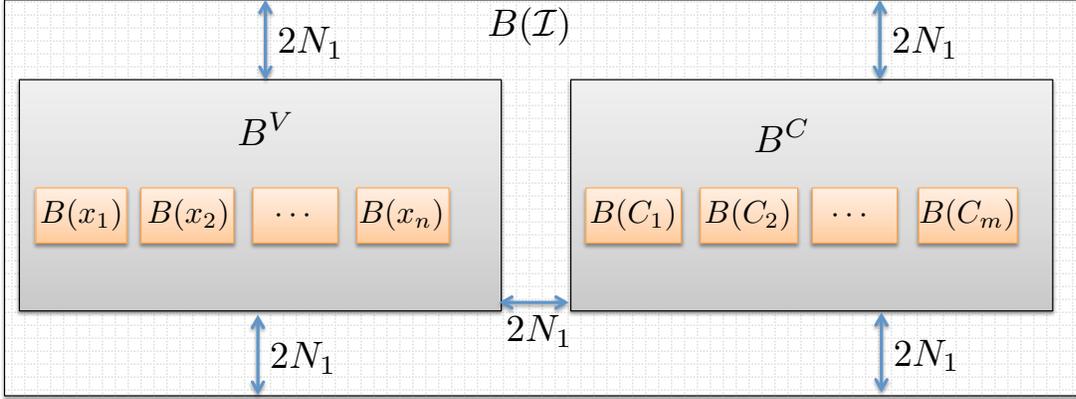


Figure 2: High-level view of the level-1 construction. Boxes $B(x_1), \dots, B(x_n)$ are at a distance at least $2N_1$ from each other and from the boundaries of B^V , and boxes $B(C_1), \dots, B(C_m)$ are at a distance at least $4N_1$ from each other and from the left and right boundaries of B^C .

Recall that the length of the path $Z(\mathcal{I})$ is $L_1 = (80h + 2)n$. We partition $Z(\mathcal{I})$ into n disjoint sub-paths $I(x_1), I(x_2), \dots, I(x_n)$ of length $80h + 2$ each, that we refer to as intervals. For each $1 \leq j \leq n$, vertices of $S(\mathcal{M}(x_j))$ will lie on $I(x_j)$. Additionally, for every clause C_q in which variable x_j participates, path $I(x_j)$ will contain some vertices of $S(\mathcal{M}(C_q))$. The remainder of the construction consists of two parts — variable gadgets and clause gadgets, that we define next, starting with the variable gadgets.

Variable gadgets. Consider some variable x of the formula φ and its corresponding interval $I(x)$ of $Z(\mathcal{I})$. We partition $I(x)$ as follows. Let $I^T(x), I^F(x) \subseteq I(x)$ denote the sub-intervals of $I(x)$ containing the first and the last $(10h + 1)$ consecutive vertices of $I(x)$, respectively, and let $I^X(x)$ be the interval containing the remaining $60h$ vertices. Consider the box $B(x)$, and recall that it has length $4N_1 + 70h + 2$ and height $4N_1 + 3$. Let $B'(x) \subseteq B(x)$ be a sub-grid of $B(x)$ of length $70h + 2$ and height 3, so that there are exactly $2N_1$ rows of $B(x)$ above and below $B'(x)$, and $2N_1$ columns of $B(x)$ to the left and to the right of $B'(x)$. Notice that the middle row of $B'(x)$, that we denote by $R'(x)$, is aligned with the middle row of $B(x)$ and hence of $B(\mathcal{I})$. We delete all vertices of $B'(x)$ that lie on its bottom row, and we will place all destination vertices of the demand pairs in $\mathcal{M}(x)$ on $R'(x)$. In order to do so, we further partition $R'(x)$ into three intervals: interval $\hat{I}^T(x)$ contains the first $5h + 1$ vertices of $R'(x)$; interval $\hat{I}^F(x)$ contains the following $5h + 1$ vertices of $R'(x)$; and interval $\hat{I}^X(x)$ contains the remaining $60h$ vertices of $R'(x)$ (see Figure 3(a)). The set $\mathcal{M}(x)$ of demand pairs consists of three subsets:

- **(Extra Pairs).** Let $\mathcal{M}^X(x) = \{(s_y^X(x), t_y^X(x))\}_{y=1}^{60h}$ be a set of $60h$ demand pairs that we call the EXTRA pairs for x . The vertices $s_1^X(x), \dots, s_{60h}^X(x)$ appear on $I^X(x)$ in this order, and the vertices $t_1^X(x), \dots, t_{60h}^X(x)$ appear on $\hat{I}^X(x)$ in this order.

of $B(C_q)$. Partition $R(C_q)$ into h intervals $\hat{I}_1(C_q), \dots, \hat{I}_h(C_q)$, each containing three consecutive vertices (see Figure 3(b)). Fix some $1 \leq j \leq h$. We identify the three vertices of $\hat{I}_j(C_q)$ with the destination vertices $t_j(C_q, \ell_{q_1}), t_j(C_q, \ell_{q_2}),$ and $t_j(C_q, \ell_{q_3})$ in this order. For each $1 \leq z \leq 3$, the corresponding source vertex $s_j(C_q, \ell_{q_z})$ has already been defined as part of the definition of the variable gadget corresponding to the literal ℓ_{q_z} . Let $\mathcal{M}(C_q) = \bigcup_{z=1}^3 \mathcal{M}(C_q, \ell_{q_z})$ be the set of all demand pairs representing C_q , so $|\mathcal{M}(C_q)| = 3h$. Let $\mathcal{M}^C = \bigcup_{q=1}^m \mathcal{M}(C_q)$ be the set of all clause-pairs, and let $\mathcal{M}^V = \bigcup_{j=1}^n \mathcal{M}(x_j)$ be the set of all variable-pairs. Our final set of demand pairs is $\mathcal{M} = \mathcal{M}^V \cup \mathcal{M}^C$. This concludes the definition of the level-1 instance. We now proceed to analyze it.

Yes-Instance Analysis. Assume that φ is a YES-INSTANCE. We show that for every instantiation of the level-1 instance \mathcal{I} , there is a set \mathcal{P} of node-disjoint paths routing $N_1 = (200h/3 + 1)n$ demand pairs, that respect the box $B(\mathcal{I})$. Assume that we are given some instantiation of \mathcal{I} . We first select the set $\hat{\mathcal{M}} \subseteq \mathcal{M}$ of demand pairs to route, and then define the routing. Fix some assignment \mathcal{A} to the variables of φ that satisfies all the clauses. For every variable x , if \mathcal{A} assigns the value TRUE to x , then we let $\hat{\mathcal{M}}(x) = \mathcal{M}^T(x) \cup \mathcal{M}^X(x)$, and otherwise, we let $\hat{\mathcal{M}}(x) = \mathcal{M}^F(x) \cup \mathcal{M}^X(x)$. Notice that in either case, $|\hat{\mathcal{M}}(x)| = 65h + 1$. For each clause $C_q = (\ell_{q_1} \vee \ell_{q_2} \vee \ell_{q_3}) \in \mathcal{C}$, let ℓ_{q_z} be a literal which evaluates to TRUE by \mathcal{A} (if there are several such literals, select one of them arbitrarily). We let $\hat{\mathcal{M}}(C_q) = \mathcal{M}(C_q, \ell_{q_z})$. Let $\hat{\mathcal{M}}^V = \bigcup_{j=1}^n \hat{\mathcal{M}}(x_j)$ and $\hat{\mathcal{M}}^C = \bigcup_{q=1}^m \hat{\mathcal{M}}(C_q)$. We then set $\hat{\mathcal{M}} = \hat{\mathcal{M}}^V \cup \hat{\mathcal{M}}^C$, so $|\hat{\mathcal{M}}| = (65h + 1)n + mh = (200h/3 + 1)n = N_1$, as $m = 5n/3$. We now show that all demand pairs in $\hat{\mathcal{M}}$ can be routed by a set \mathcal{P} of paths that respects the box $B(\mathcal{I})$. We only provide an intuitive description of the routing here; a formal proof appears in Section 5.

Consider some variable x , and assume that it is assigned the value TRUE. Let $\hat{\mathcal{M}}'(x) \subseteq \hat{\mathcal{M}}^C$ be the set of all clause-pairs, whose source vertices lie on the interval $I(x)$. Since the current assignment to x must satisfy their corresponding clauses, all vertices of $S(\hat{\mathcal{M}}'(x))$ lie on $I^F(x)$. Similarly, if x is assigned the value FALSE, then all vertices of $S(\hat{\mathcal{M}}'(x))$ lie on $I^T(x)$. Therefore, for every variable x , the sources of the demand pairs in $\hat{\mathcal{M}}(x)$ appear consecutively on $Z(\mathcal{I})$ (relatively to the vertices of $S(\hat{\mathcal{M}})$), and the same holds for the sources of the demand pairs in $\hat{\mathcal{M}}(C_q)$, for every clause C_q .

We build the paths in \mathcal{P} gradually, growing them from the source vertices. We start by selecting, for every variable x , a set $A(x)$ of $|\hat{\mathcal{M}}(x)|$ vertices on the top boundary of $B(x)$, and similarly, for every clause C_q , a set $A(C_q)$ of $|\hat{\mathcal{M}}(C_q)|$ vertices on the top boundary of $B(C_q)$. We discuss this selection later. In the first step, we route the paths from their source vertices to these newly selected vertices, so that for every variable x , all paths originating from the vertices of $S(\hat{\mathcal{M}}(x))$ terminate at the vertices of $A(x)$ and they are order-preserving, and similarly for every clause $C_q \in \mathcal{C}$, all paths originating from the vertices of $S(\hat{\mathcal{M}}(C_q))$ terminate at the vertices of $A(C_q)$ and they are order-preserving. In order to execute this step, we carefully select a set Γ of $|\hat{\mathcal{M}}|$ vertices on the top boundary of B^V ; a set Γ'' of $|\hat{\mathcal{M}}^C|$ vertices on the bottom boundary of B^V ; and the set Γ''' of $|\hat{\mathcal{M}}^C|$ left-most vertices on the top boundary of B^C . We first connect all vertices in $S(\hat{\mathcal{M}})$ to the $|\hat{\mathcal{M}}|$ leftmost vertices on the opening of $B(\mathcal{I})$ via a set of order-preserving node-disjoint paths, and then extend them to the vertices of Γ via node-disjoint order-preserving paths inside $B(\mathcal{I})$. This part of the routing is straightforward. For every variable x , the paths originating at the vertices of $S(\hat{\mathcal{M}}(x))$ then continue to their corresponding boxes $B(x)$, while for each clause C_q , the paths originating at the vertices of $S(\hat{\mathcal{M}}(C_q))$ continue to the bottom boundary of B^V and terminate at a consecutive set of vertices of Γ'' (see Section 5 and Figure 10 for more details). We select the vertices of Γ and Γ'' in a way that ensures that every path that we route inside the box B^V is a sub-path of a column of B^V . We then connect the vertices of Γ'' to the vertices of Γ''' via a set of

node-disjoint order-preserving paths, that are internally disjoint from the boxes B^V and B^C ; these paths exploit the spacing between the two boxes. Finally, we complete the routing inside the box B^C , ensuring that for every clause C_q , the paths originating at the vertices of $S(\hat{\mathcal{M}}(C_q))$ terminate at the vertices of $A(C_q)$. This is done via a standard snake-like routing. This routing critically uses the fact that the endpoints of the paths that we have constructed so far, which originate at the vertices of $S(\hat{\mathcal{M}}(C_q))$, appear consecutively on Γ''' , for every clause C_q .

By appropriately choosing, for every clause $C_q \in \mathcal{C}$, the set $A(C_q)$ of vertices on the top boundary of $B(C_q)$, it is easy to extend the paths originating at the vertices of $S(\hat{\mathcal{M}}(C_q))$, so that they terminate at the vertices in $T(\hat{\mathcal{M}}(C_q))$. Since the resulting paths are order-preserving, we will route all demand pairs in $\hat{\mathcal{M}}(C_q)$.

We now consider some variable x and show how to complete the routing of the demand pairs in $\hat{\mathcal{M}}(x)$ inside the box $B(x)$. Assume first that x is assigned the value TRUE. Then the paths routing the demand pairs in $\hat{\mathcal{M}}(x)$ arrive at the top boundary of $B(x)$ in the same order as the ordering of their source vertices on $Z(\mathcal{I})$. The order of their corresponding destination vertices on the second row of $B'(x)$ is identical, and so it is immediate to extend the paths inside $B(x)$ to complete the routing.

Assume now that x is assigned the value FALSE. Let J and J' be the intervals of the top boundary of $B(x)$ where the paths routing the pairs in $\hat{\mathcal{M}}^F(x)$ and $\hat{\mathcal{M}}^X(x)$ arrive, respectively. Unfortunately, J lies to the right of J' , while interval $\hat{I}^F(x)$ lies to the left of the interval $\hat{I}^X(x)$ on the second row of $B'(x)$. Therefore, we need to “switch” the ordering of these two sets of paths before we can complete the routing. It is easy to do so by exploiting the ample spacing between the box $B'(x)$ and the boundaries of the box $B(x)$ (see Figure 11).

No-Instance Analysis. Assume now that φ is a NO-INSTANCE, and that we are given some instantiation of the level-1 instance \mathcal{I} and a set $\tilde{\mathcal{P}}$ of node-disjoint paths routing some subset $\tilde{\mathcal{M}} \subseteq \mathcal{M}$ of demand pairs. Our goal is to prove that $|\tilde{\mathcal{M}}| \leq N'_1 = (1 - \delta)N_1 = (1 - \delta)(200h/3 + 1)n$. Assume for contradiction that $|\tilde{\mathcal{M}}| > N'_1$. In order to analyze the NO-INSTANCE, it is convenient to view the construction slightly differently. Let \mathcal{C}' be the set of clauses obtained by adding, for each clause $C_q \in \mathcal{C}$, h copies C_q^1, \dots, C_q^h of C_q to \mathcal{C}' . We will refer to the clauses in \mathcal{C} as the *original clauses*, and the clauses in \mathcal{C}' as the *new clauses*. Notice that $|\mathcal{C}'| = mh$, and it is easy to verify that no assignment to the variables of φ can satisfy more than $(1 - \epsilon)hm$ clauses of \mathcal{C}' . We will reach a contradiction by defining an assignment to the variables of φ that satisfies more than $(1 - \epsilon)hm$ clauses of \mathcal{C}' . For each new clause $C_q^j \in \mathcal{C}'$, we let $\mathcal{M}(C_q^j) \subseteq \mathcal{M}$ be the set of all demand pair whose destinations lie on interval $\hat{I}_j(C_q)$; we view these demand pairs as representing the clause C_q^j .

For every variable x of φ , let $\tilde{\mathcal{M}}(x) = \tilde{\mathcal{M}} \cap \mathcal{M}(x)$, and let $\tilde{\mathcal{M}}^T(x), \tilde{\mathcal{M}}^F(x), \tilde{\mathcal{M}}^X(x)$ denote $\mathcal{M}^T(x) \cap \tilde{\mathcal{M}}, \mathcal{M}^F(x) \cap \tilde{\mathcal{M}}$ and $\mathcal{M}^X(x) \cap \tilde{\mathcal{M}}$, respectively. For every new clause $C_q^j \in \mathcal{C}'$, let $\tilde{\mathcal{M}}(C_q^j) = \tilde{\mathcal{M}} \cap \mathcal{M}(C_q^j)$. We also denote by $\tilde{\mathcal{M}}^V = \bigcup_{j=1}^n \tilde{\mathcal{M}}(x_j)$ and by $\tilde{\mathcal{M}}^C = \bigcup_{C_q^j \in \mathcal{C}'} \tilde{\mathcal{M}}(C_q^j)$, the sets of the variable-pairs and the clause-pairs, respectively, that are routed by $\tilde{\mathcal{P}}$. We use the following claim.

Claim 3.1 *For each variable x of φ , at least one of the sets $\tilde{\mathcal{M}}^T(x), \tilde{\mathcal{M}}^F(x), \tilde{\mathcal{M}}^X(x)$ is empty.*

The proof is omitted here; we prove a somewhat stronger claim in Section 6 (see also Observation 6.6 from which the claim follows immediately). Notice that from the above claim, $|\tilde{\mathcal{M}}^V| \leq (65h + 1)n$.

Consider now some variable x . Assume first that $\tilde{\mathcal{M}}^F(x) = \emptyset$. We then assign x the value TRUE. We say that an index $1 \leq j \leq 5h + 1$ is *bad* for x , iff the pair $(s_j^T(x), t_j^T(x)) \notin \tilde{\mathcal{M}}(x)$. Otherwise, if $\tilde{\mathcal{M}}^F(x) \neq \emptyset$, then we assign x the value FALSE. In this case, we say that an index $1 \leq j \leq 5h + 1$

is bad for x , iff $(s_j^F(x), t_j^F(x)) \notin \tilde{\mathcal{M}}(x)$. We later show that the total number of pairs (x, j) , where j is a bad index for variable x , is small.

Consider some new clause C_q^j . We say that it is an *interesting* clause if $|\tilde{\mathcal{M}}(C_q^j)| \geq 1$ (in other words, at least one pair in the set $\{(s_j(C_q, \ell_{q_z}), t_j(C_q, \ell_{q_z}))\}_{z=1}^3$ is in $\tilde{\mathcal{M}}$), and we say that it is uninteresting otherwise. We say that clause C_q^j is *troublesome* iff $|\tilde{\mathcal{M}}(C_q^j)| > 1$. The proof of the following simple observation is omitted here; we prove a more general statement in Section 6.

Observation 3.2 *For each clause C_q , at most three of its copies are troublesome.*

We conclude that $|\tilde{\mathcal{M}}^C| \leq m(h+6) = 5n(h+6)/3$. Let $\mathcal{C}'_1 \subseteq \mathcal{C}'$ be the set of all interesting new clauses. A simple accounting shows that, if $|\tilde{\mathcal{M}}| \geq (1-\delta)(200h/3+1)n$, then $|\mathcal{C}'_1| \geq (1-\epsilon/10)hm$ must hold. Notice that for each new clause $C_q^j \in \mathcal{C}'_1$, at least one demand pair from the set $\{(s_j(C_q, \ell_{q_z}), t_j(C_q, \ell_{q_z}))\}_{z=1}^3$ is in $\tilde{\mathcal{M}}$. We select any literal $\ell \in \{\ell_{q_1}, \ell_{q_2}, \ell_{q_3}\}$ such that $(s_j(C_q, \ell), t_j(C_q, \ell)) \in \tilde{\mathcal{M}}$, and we say that clause C_q^j *chooses* the literal ℓ . Let x be the variable corresponding to the literal ℓ . We say that C_q^j is a *cheating* clause iff the assignment that we chose for x is not consistent with the literal ℓ : that is, if $\ell = x$, then $\mathcal{A}(x) = \text{FALSE}$, and if $\ell = \neg x$, then $\mathcal{A}(x) = \text{TRUE}$. Notice that, if C_q^j is an interesting and a non-cheating clause, then the current assignment satisfies C_q^j . Therefore, in order to complete the analysis, it is enough to prove the following claim.

Claim 3.3 *The number of cheating clauses in \mathcal{C}'_1 is bounded by $\epsilon mh/2$.*

We prove a stronger claim in Section 6 (see Lemma 6.11), and provide a proof sketch here. Let $C_q^j \in \mathcal{C}'$ be a cheating clause, and suppose it has chosen the literal ℓ , whose corresponding variable is x . We say that C_q^j is a *bad* cheating clause, iff at least one of the indices $j, j+1$ is a bad index for variable x (recall that j is a bad index for x if $\mathcal{A}(x) = \text{TRUE}$ and $(s_j^T(x), t_j^T(x)) \notin \tilde{\mathcal{M}}$, or $\mathcal{A}(x) = \text{FALSE}$ and $(s_j^F(x), t_j^F(x)) \notin \tilde{\mathcal{M}}$). Otherwise, we say that C_q^j is a good cheating clause. A simple accounting shows that the number of pairs (x, j) , where j is a bad index for x is bounded by $\epsilon mh/16$. Each such pair (x, j) may contribute to at most two bad cheating clauses, and so there are at most $\epsilon mh/8$ bad cheating clauses.

Our final step is to show that the number of good cheating clauses is bounded by $\epsilon mh/4$. We show that for each original clause C_q , at most 3 copies of C_q are good cheating clauses. It then follows that the total number of good cheating clauses is at most $3m < \epsilon mh/4$, since $h = 1000/\epsilon$. Consider some original clause C_q . It is enough to show that for each literal ℓ of C_q , the number of copies of C_q that choose ℓ and are good cheating clauses is at most 1. Assume for contradiction that there are two such copies C_q^j and $C_q^{j'}$. Assume w.l.o.g. that the variable x that corresponds to ℓ is assigned the value TRUE , so $\ell = \neg x$. Then vertex $s_j(C_q, \ell)$ lies on interval $I^T(x)$, between $s_j^T(x)$ and $s_{j+1}^T(x)$, while vertex $s_{j'}(C_q, \ell)$ lies on interval $I^T(x)$, between $s_{j'}^T(x)$ and $s_{j'+1}^T(x)$. Assume w.l.o.g. that $j < j'$. Consider the plane with only the top boundary of the grid G_1 , the bottom boundary of the box $B^T(x)$, and the images of the paths of $\tilde{\mathcal{P}}$ routing the pairs $(s_j^T(x), t_j^T(x)), (s_{j+1}^T(x), t_{j+1}^T(x))$, and $(s_{j'+1}^T(x), t_{j'+1}^T(x))$ present. Let f, f' be the two faces of this drawing that differ from the outer face, such that f has $s_j(C_q, \ell)$ on its boundary and f' has $s_{j'}(C_q, \ell)$ on its boundary. Then $f \neq f'$, and the bottom boundary of $B(C_q)$ must belong to a single face of the resulting drawing. Assume w.l.o.g. that this face is $f^* \neq f$. Then $t_j(C_q, \ell)$ lies on f^* , and so it is impossible that a path of $\tilde{\mathcal{P}}$ connects $s_j(C_q, \ell)$ to $t_j(C_q, \ell)$. We conclude that the current assignment satisfies at least $(1-\epsilon/10)hm - \epsilon mh/2 > (1-\epsilon)hm$ clauses of \mathcal{C}' , a contradiction.

Generalization to Higher Levels and the Hardness Gap. Assume now that we are given a construction of a level- i instance, and we would like to construct a level- $(i+1)$ instance. Intuitively, we would like to start with the level-1 instance described above, and then replace each source-destination pair (s, t) with a distinct copy of a level- i instance \mathcal{I}' . So we would replace the vertex s with the path $Z(\mathcal{I}')$, and the vertex t with the cut-out box $B(\mathcal{I}')$. We say that a level- i instance \mathcal{I}' is routed by a solution \mathcal{P} to this resulting instance, iff \mathcal{P} routes a significant number of the demand pairs in \mathcal{I}' . The idea is that, due to the level-1 instance analysis, the number of such level- i instances that can be routed in the YES- and the NO-INSTANCE cases differ by a constant factor, while within each such instance we already have some gap g_i between the YES- and the NO-INSTANCE solutions, and so the gap grows by a constant factor in every level. Unfortunately this idea does not quite work. If we consider, for example, a level-1 instance \mathcal{I}' , then its destination vertices appear quite far – at distance $\Theta(N_1)$ – from the bottom boundary of the box $B(\mathcal{I}')$. In general, in a level- i instance, this distance needs to be roughly $\Theta(N_i)$, to allow the routing in the YES-INSTANCE case (recall that N_i is the number of the demand pairs that can be routed in the YES-INSTANCE case). Therefore, if we replace each level-1 demand pair by a level- i instance, some of the paths in the new level- $(i+1)$ instance may “cheat” by passing through the boxes $B(\mathcal{I}')$ of level- i instances \mathcal{I}' , and exploiting the spacing between the destination vertices and the bottom boundary of each such box. For example, it is now possible that in a variable gadget, we will be able to route many demand pairs from each set $\mathcal{M}^X(x), \mathcal{M}^T(x)$ and $\mathcal{M}^F(x)$ simultaneously. A simple way to get around this problem is to create more level- i instances, namely: we replace each source-destination pair from a level-1 instance by a collection of c_{i+1} level- i instances. The idea is that, if the number of the demand pairs we try to route in many such c_{i+1} -tuples of level- i instances is large enough, then on average only a small fraction of the routing paths may cheat by exploiting the spacing between the destination vertices and their corresponding box boundaries, and this will not affect the overall accounting by too much. However, if the formula φ is a NO-INSTANCE, then we will only attempt to route N'_i demand pairs from each level- i instance, and therefore we need to ensure that $c_{i+1}N'_i \gg N_i$ in order for the gap to grow in the current level. In other words, the number of copies of the level- i instances that we use in the level- $(i+1)$ instance construction should be proportional to the gap between the YES- and the NO-INSTANCE cost at level i (times n). A simple calculation shows that, if we follow this approach, we will obtain a gap of $2^{\Omega(i)}$ in level- i instances, with construction size roughly $n^{\Theta(i)} \cdot 2^{\Theta(i^2)}$. Therefore, after $i^* = \Theta(\log n)$ iterations, we obtain a gap of $2^{\Omega(\sqrt{\log n})}$, where n' is the size of the level- i^* instance. This rapid growth in the instance size is the main obstacle to obtaining a stronger hardness of approximation factor using this approach.

4 The Full Construction

In this section we provide a full description of our construction. The resulting graphs will have maximum vertex degree 4. We show in Section B of the Appendix how to modify the resulting instances in order to obtain the proof of Theorem 1.1 for sub-cubic graphs. We start with setting the parameters.

Parameters. The two main parameters that we use are $h = 1000/\epsilon$ and $\delta = 8\epsilon^2/10^{12}$, where ϵ is the constant from Theorem 2.1. We define the remaining parameters in terms of these two parameters.

For every level $i \geq 0$ of our construction, we use two parameters, N_i and N'_i . We will ensure

that for every instantiation of the level- i instance \mathcal{I} , if the initial 3SAT(5) formula φ is a YES-INSTANCE, then there is a solution to \mathcal{I} routing N_i demand pairs, that respects the box $B(\mathcal{I})$, and if φ is a NO-INSTANCE, then no solution to \mathcal{I} can route more than N'_i demand pairs. We define the parameters N_i, N'_i inductively, starting with $N_0 = N'_0 = 1$. Assume now that for some $i \geq 0$, we are given the values of N_i and N'_i . Let $g_i = N_i/N'_i$ be the gap between the YES- and the NO-INSTANCE solution values at level i , and let $c_{i+1} = 10^8 h^2 g_i = O(g_i)$. Parameter c_{i+1} will be used in our construction of the level- $(i+1)$ instance. We then set $N_{i+1} = n c_{i+1} (200h/3 + 1) N_i$ and $N'_{i+1} = (1 - \delta) n c_{i+1} (200h/3 + 1) N'_i$. It is immediate to verify that $g_{i+1} = \frac{g_i}{1-\delta}$, and so for all $i \geq 0$, $g_i = \left(\frac{1}{1-\delta}\right)^i$, and $N_i = O(n \cdot g_{i-1}) \cdot N_{i-1} = (\rho n)^i \cdot 2^{O(i^2)}$, for some absolute constant ρ . We set the parameters L_i, L'_i and H_i below, but we will ensure that each of these parameters is bounded by $20N_i^3$. Our construction has $i^* = \log n$ levels, giving us a gap of $2^{\Omega(\log n)}$ between the YES- and the NO-INSTANCE solution values. For our final level- i^* instance, we can choose the grid G_{i^*} to be of size $(Q \times Q)$, where $Q = 2L_{i^*} + 2L'_{i^*} + 4H_{i^*} = O(N_{i^*}^3)$, and so the instance size is bounded by $n' = O(N_{i^*}^6) = n^{O(\log n)} \cdot 2^{O(\log^2 n)} = 2^{O(\log^2 n)}$. Overall, we obtain a factor $2^{\Omega(\sqrt{\log n'})}$ -hardness of approximation, unless all problems in NP have deterministic algorithms running in time $n^{O(\log n)}$.

For $i \geq 0$, we set the parameter $H_i = 20N_i$. The following bound on H_i follows immediately from the definitions of our parameters, and we use it several times in our analysis:

$$H_i = 20N_i = 20g_i N'_i = \frac{20c_{i+1} N'_i}{10^8 h^2} = \frac{2c_{i+1} \epsilon^2 N'_i}{10^{13}} \quad (1)$$

For all $i \geq 0$, we set $L'_i = 20N_i^3$. Parameter L_i is defined as follows: $L_0 = 1$, and for $i > 0$, $L_i = (80h + 2)c_i L_{i-1} n \leq (80h + 2)c_i \cdot 20N_{i-1}^3 n \leq 20N_i^3$.

Level-0 Instance. A level-0 instance \mathcal{I} consists of a single demand pair (s, t) . In order to be consistent with our definitions, we let $Z(\mathcal{I})$ be a path of length $L_0 = 1$, with s mapped to the unique vertex of $Z(\mathcal{I})$. Recall that $N_0 = N'_0 = 1$, $H_0 = 20N_0 = 20$, and $L'_0 = 20N_0^3 = 20$. Let G'_0 be a grid of length $L'_0 = 20$ and height $H_0 = 20$. We obtain the box $B(\mathcal{I})$ from G'_0 by deleting all vertices lying on the left, bottom, and right boundaries of G'_0 . Let R' be the middle row of G'_0 . We map t to any vertex of $B(\mathcal{I})$ that belongs to R' . It is immediate to verify that for every instantiation of this level-0 instance, there is a solution that routes one demand pair and respects $B(\mathcal{I})$, regardless of whether we are in the YES or the NO-INSTANCE.

From now on we focus on constructing instances of levels $i > 0$. As already mentioned, the construction we obtain for level $i = 1$ is similar to that described in Section 3.

4.1 Level- $(i+1)$ Construction

A level- $(i+1)$ instance is obtained by combining a number of level- i instances. We start with a quick overview of the level- i construction.

Level- i Construction Overview. Recall that a definition of a level- i instance \mathcal{I} consists of a path $Z(\mathcal{I})$ of length L_i , a grid G'_i of height H_i and length L'_i , together with a cut-out box $B(\mathcal{I}) \subseteq G'_i$, and a collection \mathcal{M} of demand pairs, such that all vertices of $S(\mathcal{M})$ are mapped to vertices of $Z(\mathcal{I})$, while all vertices of $T(\mathcal{M})$ are mapped to distinct vertices of $B(\mathcal{I}) \cap R'$, where R' is the middle row

of G'_i . Path $Z(\mathcal{I})$ will eventually become a sub-path of the first row of the larger grid G_i , and box $B(\mathcal{I})$ will be placed inside G_i , within distance at least H_i from its boundaries.

For every destination vertex t , we draw a straight line Q_t from t to the bottom of $B(\mathcal{I})$. This line contains at most $H_i/2$ vertices of the graph. We will use these lines in the analysis of the NO-INSTANCE case of the level- $(i+1)$ construction.

It will sometimes be useful to place several level- i instances side-by-side. For an integer $c > 0$, a c -wide level- i instance \mathcal{I} is constructed as follows. Intuitively, we construct c disjoint level- i instances $\mathcal{I}_1, \dots, \mathcal{I}_c$, placing their intervals $Z(\mathcal{I}_j)$ side-by-side on $Z(\mathcal{I})$ and placing their boxes $B(\mathcal{I}_j)$ side-by-side inside $B(\mathcal{I})$. Formally, for each $1 \leq j \leq c$, let \mathcal{M}_j be the set of the demand pairs of the level- i instance \mathcal{I}_j , and let G^j be the corresponding grid G'_i for that instance. The set of the demand pairs of instance \mathcal{I} is $\mathcal{M} = \bigcup_{j=1}^c \mathcal{M}_j$. We let $Z(\mathcal{I})$ be a path of length $c \cdot L_i$, partitioned into c equal-length intervals A_1, \dots, A_c . We let G' be a grid of length cL'_i and height H_i , that we partition into c sub-grids of length L'_i and height H_i each. For $1 \leq j \leq c$, we map the vertices of $Z(\mathcal{I}_j)$ to the vertices of A_j in a natural way. This defines the mapping of the vertices of $S(\mathcal{M})$ to the vertices of $Z(\mathcal{I})$. For each $1 \leq j \leq c$, we map the vertices of G^j to the j th sub-grid of G' , and delete from G' all vertices to which the vertices of $G^j \setminus B(\mathcal{I}_j)$ are mapped. The resulting subgraph of G' becomes the box $B(\mathcal{I})$, and the above mapping defines the mapping of the destination vertices in $T(\mathcal{M})$ to the vertices of $B(\mathcal{I})$. Note that if R' denotes the middle row of G' , then all vertices of $T(\mathcal{M})$ lie on R' . In order to instantiate this instance, we need to select a grid G of length at least $c(2L_i + 2L'_i + 4H_i)$ and height at least $3cH_i$, a sub-path P of the first row of G of length cL_i , to which $Z(\mathcal{I})$ will be mapped, and a sub-grid G'' of G of the same dimensions as G' , to which the vertices of G' will be mapped. The vertices of G'' must be at a distance at least cH_i from the boundaries of G . Clearly, for any instantiation of this instance, in the YES-INSTANCE case, there is a solution \mathcal{P} routing cN_i demand pairs, such that, if we denote, for each $1 \leq j \leq c$, by $\mathcal{P}_j \subseteq \mathcal{P}$ the set of paths routing demand pairs in \mathcal{M}_j , then $|\mathcal{P}_j| = N_i$ and \mathcal{P}_j respects the box $B(\mathcal{I}_j)$. On the other hand, in the NO-INSTANCE case, no solution to \mathcal{I} can route more than cN'_i demand pairs.

We now assume that we are given a construction of a level- i instance, for $i \geq 0$, and describe a construction of a level- $(i+1)$ instance \mathcal{I} . For convenience, we denote c_{i+1} by c . We use parameters $L_{i+1}, H_{i+1}, L'_{i+1}$ described above, so $H_{i+1} = 20N_{i+1}$, $L'_{i+1} = 20N_{i+1}^3$, and $L_{i+1} = (80h + 2)cL_{i+1}$.

In order to construct the box $B(\mathcal{I})$, we start with a grid G'_{i+1} of length L'_{i+1} and height H_{i+1} . We define two sub-grids of $B(\mathcal{I})$, of length $9N_{i+1}^3$ and height $16N_{i+1}$ each: grid B^V that will contain all destination vertices of the demand pairs representing the variables of the formula φ , and grid B^C that will contain all destination vertices of the demand pairs representing the clauses of the formula φ . In order to construct both grids, let \mathcal{R} be the set of all rows of G'_{i+1} , excluding the top $2N_{i+1}$ and the bottom $2N_{i+1}$ rows, so that $|\mathcal{R}| = H_{i+1} - 4N_{i+1} = 16N_{i+1}$. Let \mathcal{W} be a consecutive set of $9N_{i+1}^3$ columns of G'_{i+1} , starting from the second column, and let \mathcal{W}' be a consecutive set of $9N_{i+1}^3$ columns of G'_{i+1} , terminating at the penultimate column. We then let B^V be the sub-grid of G'_{i+1} spanned by the rows in \mathcal{R} and the columns in \mathcal{W} , and we let B^C be the sub-grid of G'_{i+1} spanned by the rows in \mathcal{R} and the columns in \mathcal{W}' (see Figure 4). Notice that at least $2N_{i+1}$ columns of G'_{i+1} separate the two grids. We delete the bottom, left, and right boundaries of G'_{i+1} to turn it into a cut-out box, that we refer to as $B(\mathcal{I})$ from now on. We will later delete some additional vertices from $B(\mathcal{I})$.

Next, we define smaller sub-grids of the two grids B^V and B^C . For every variable x_j of φ , we define a sub-grid $B(x_j)$ of B^V , of length $L^V = 4N_{i+1} + (70h + 2)cL'_i$ and height $H^V = H_{i+1} + 2N_{i+1}$. This box will contain all destination vertices of the demand pairs that represent the variable x_j . We place

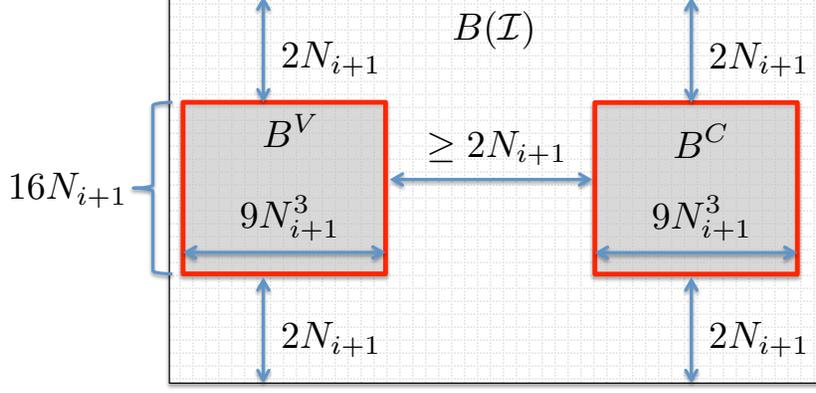


Figure 4: High-level view of box $B(\mathcal{I})$

the boxes $B(x_1), \dots, B(x_n)$ inside grid B^V , so that they are aligned and $2N_{i+1}$ -separated. In other words, the middle row of each box is contained in the middle row of B^V , and the horizontal distance between every pair of these boxes, and between each box and the left and right boundaries of B^V is at least $2N_{i+1}$. Since $n \cdot L^V + (n+1) \cdot 2N_{i+1} \leq 7nN_{i+1} + (70h+2)cL'_i n \leq 7nN_{i+1} + 1500hcN_i^3 n < 9N_{i+1}^3$, we can find such grids $B(x_1), \dots, B(x_n)$. Since $H^V = H_i + 2N_{i+1} = 20N_i + 2N_{i+1} < 3N_{i+1}$, there are at least $2N_{i+1}$ rows of B^V above and below these new grids (see Figure 5).

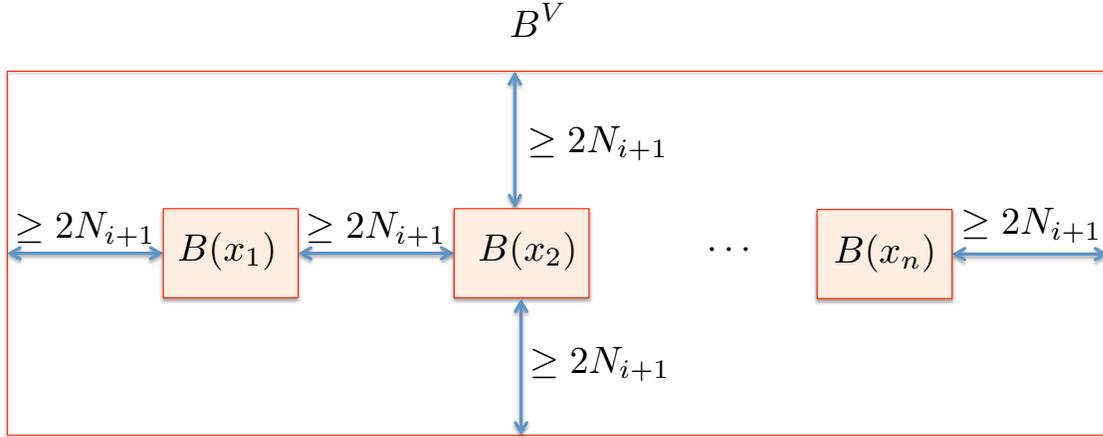


Figure 5: Box B^V . Each box $B(x_j)$ has length $L^V = 4N_{i+1} + (70h+2)cL'_i$ and height $H^V = H_i + 2N_{i+1}$.

We similarly define sub-grids $B(C_1), \dots, B(C_m)$ of B^C . Each such sub-grid has height $H^C = H_i$ and width $L^C = 3chL'_i$. For each clause $C_q \in \mathcal{C}$, box $B(C_q)$ will contain all destination vertices of the demand pairs that represent this clause. We let $B(C_1), \dots, B(C_m)$ be sub-grids of H^C that are aligned and $4N_{i+1}$ -separated. Since $m \cdot L^C + (m+1) \cdot 4N_{i+1} \leq 20nN_{i+1} + 15hcL'_i n \leq 20nN_{i+1} + 300hcN_i^3 n < 9N_{i+1}^3$, we can find such grids (see Figure 6). Since $H^C = H_i = 20N_i < N_{i+1}$, there are at least $2N_{i+1}$ rows of B^C above and below these new grids.

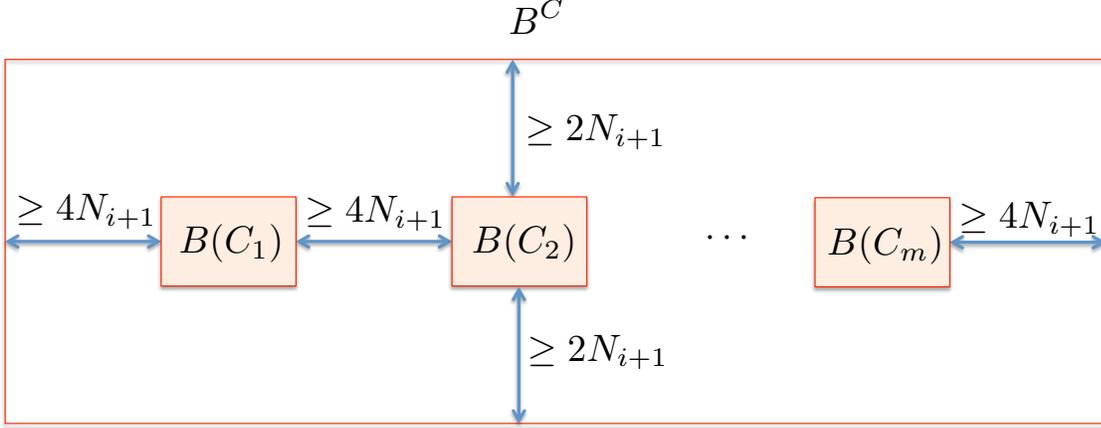


Figure 6: Box B^C . Each box $B(C_q)$ has length $L^C = 3chL'_i$ height $H^C = H_i$.

Our construction consists of two parts, called variable gadgets and clause gadgets. For each variable x_j , we construct a number of level- i instances \mathcal{I} , whose corresponding boxes $B(\mathcal{I})$ are placed inside $B(x_j)$. Whenever we do so, we delete the corresponding vertices of $B(x_j)$ as described in the preliminaries. We also construct clause gadgets similarly.

Variable Gadgets Let $Z(\mathcal{I})$ be a path of length L_{i+1} , and let Π be a partition of $Z(\mathcal{I})$ into disjoint contiguous sub-paths (that we sometimes refer to as *intervals*) of length cL_i each. For each $1 \leq j \leq n$, we let $I(x_j)$ be a sub-path of $Z(\mathcal{I})$, containing exactly $80h + 2$ consecutive intervals of Π , so that $I(x_1), I(x_2), \dots, I(x_n)$ appear on $Z(\mathcal{I})$ in this left-to-right order.

Consider some variable x of the 3SAT(5) formula φ and the corresponding interval $I(x)$ of $Z(\mathcal{I})$, containing $80h + 2$ consecutive intervals of Π . We further partition $I(x)$ as follows. Let $I^T(x), I^F(x) \subseteq I(x)$ denote the subpaths of $I(x)$ containing the first $(10h + 1)$ and the last $(10h + 1)$ consecutive intervals of Π , respectively. Let $I^X(x)$ denote the union of the remaining $60h$ consecutive intervals of Π (see Figure 7).

- **(Extra Pairs).** We use $60h$ copies of c -wide level- i instances, that we denote by $\mathcal{I}_j^X(x)$, for $1 \leq j \leq 60h$. For each $1 \leq j \leq 60h$, we let $Z(\mathcal{I}_j^X(x))$ be the j th interval of $I^X(x)$. We place the corresponding boxes $B(\mathcal{I}_1^X(x)), \dots, B(\mathcal{I}_{60h}^X(x))$ side-by-side, obtaining one box $B_X(x)$ of width $60hcL'_i$ and height H_i . We define the placement of this box inside $B(x)$ later. We denote by $\mathcal{M}^X(x)$ the resulting set of demand pairs, and we refer to them as *the EXTRA demand pairs of x* .
- **(True Pairs).** We denote the intervals of Π appearing on $I^T(x)$ by:

$$A_1^T, Y_1^T, A_2^T, Y_2^T, \dots, Y_{5h}^T, A_{5h+1}^T,$$

and we assume that they appear on $I^T(x)$ in this left-to-right order. We use $(5h + 1)$ copies of the c -wide level- i instance, that we denote by $\mathcal{I}_j^T(x)$, for $1 \leq j \leq 5h + 1$. For each $1 \leq j \leq 5h + 1$, we let $Z(\mathcal{I}_j^T(x))$ be the interval A_j^T . We place the corresponding boxes $B(\mathcal{I}_1^T(x)), \dots, B(\mathcal{I}_{5h+1}^T(x))$ side-by-side, obtaining one box $B_T(x)$ of width $(5h + 1)cL'_i$ and

height H_i . We denote by $\mathcal{M}^T(x)$ the resulting set of demand pairs, and we refer to them as *the TRUE demand pairs of x* .

- (**FALSE Pairs**). Similarly, we denote the intervals of Π appearing on $I^F(x)$ by:

$$A_1^F, Y_1^F, A_2^F, Y_2^F, \dots, Y_{5h}^F, A_{5h+1}^F,$$

and we assume that they appear on $I^F(x)$ in this left-to-right order. We use $(5h + 1)$ copies of the c -wide level- i instance, that we denote by $\mathcal{I}_j^F(x)$, for $1 \leq j \leq 5h + 1$. For each $1 \leq j \leq 5h + 1$, we let $Z(\mathcal{I}_j^F)$ be the interval A_j^F . We place the corresponding boxes $B(\mathcal{I}_1^F(x)), \dots, B(\mathcal{I}_{5h+1}^F(x))$ side-by-side, obtaining one box $B_F(x)$ of width $(5h + 1)cL'_i$ and height H_i . We denote by $\mathcal{M}^F(x)$ the resulting set of demand pairs, and we refer to them as *the FALSE demand pairs of x* .

We let $\mathcal{M}(x) = \mathcal{M}^X(x) \cup \mathcal{M}^T(x) \cup \mathcal{M}^F(x)$. We call the demand pairs in $\mathcal{M}(x)$ *variable-pairs representing x* .

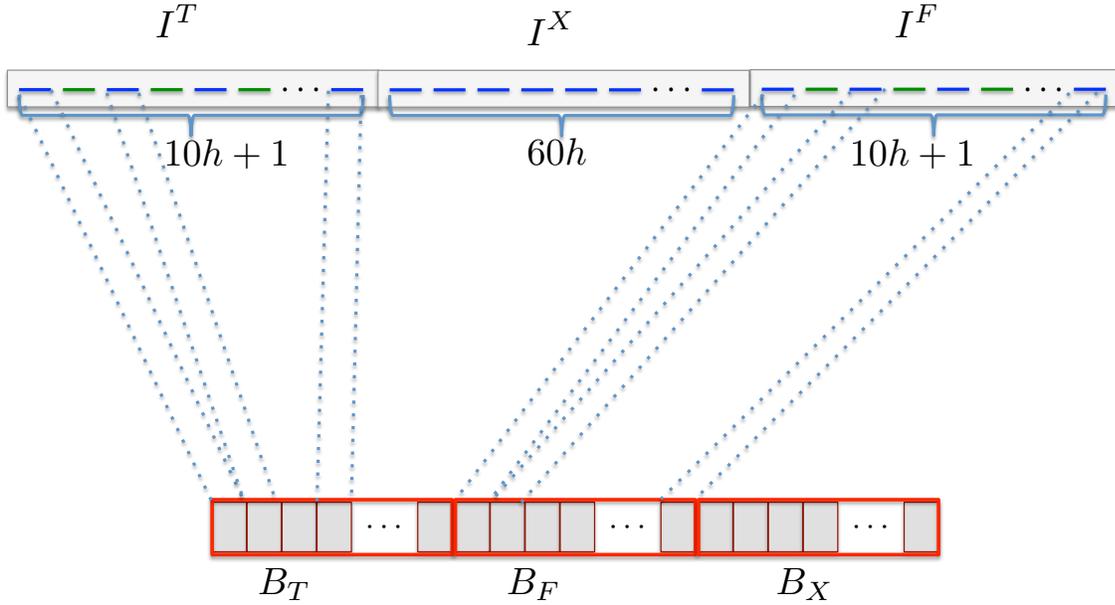


Figure 7: Variable gadget for level- $(i + 1)$ instance. Index x is omitted for convenience. Blue sub-intervals of I^F belong to instances \mathcal{I}_j^F whose destinations lie in B_F ; green sub-intervals belong to instances $\mathcal{I}_j(C, x)$ associated with clauses $C \in \mathcal{C}(x)$. The $(5h)$ green intervals are partitioned into 5 groups of h consecutive intervals each, and each group belongs to a distinct clause. Blue and green intervals of I^T are dealt with similarly.

Recall that the length of box $B_X(x)$ is $60hcL'_i$, while boxes $B_T(x), B_F(x)$ have length $(5h + 1)cL'_i$ each. The height of each box is H_i . Recall also that box $B(x)$ has length $L^V = 4N_{i+1} + (70h + 2)cL'_i$ and height $H^V = H_i + 2N_i$.

We place the boxes $B_T(x), B_F(x)$ and $B_X(x)$ side-by-side inside $B(x)$ in this order, so that the middle row of each box is contained in the middle row of $B(x)$, and there is a horizontal spacing of $2N_{i+1}$ between the left boundaries of $B_T(x)$ and $B(x)$, and between the right boundaries of $B_X(x)$

and $B(x)$ (see Figure 8). Notice that there is no horizontal spacing between $B_T(x)$, $B_F(x)$ and $B_X(x)$, and all destination vertices lying in $B(x)$ belong to the middle row of $B(x)$, and hence to the middle row of $B(\mathcal{I})$.

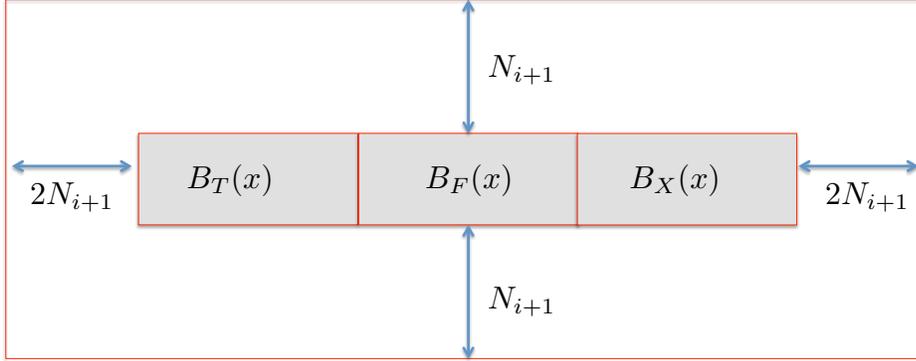


Figure 8: Box $B(x)$. Height: $H^V = H_i + 2N_{i+1}$, length: $L^V = 4N_{i+1} + (70h + 2)cL'_i$.

Consider the set $\mathcal{C}(x) \subseteq \mathcal{C}$ of clauses in which variable x appears without negation. Assume without loss of generality that $\mathcal{C}(x) = \{C_1, \dots, C_r\}$, where $r \leq 5$. For each $1 \leq r' \leq r$, we will create h level- i instances of width c , that represent the variable x of $C_{r'}$. We denote these instances by $\mathcal{I}_j(C_{r'}, x)$, for $1 \leq j \leq h$. Consider the interval $I^F(x)$. We will use the sub-intervals Y_j^F of $I^F(x)$ as intervals $Z(\mathcal{I}_j(C_{r'}, x))$, where, intuitively, intervals corresponding to the same clause-variable pair appear consecutively. Formally, for each $1 \leq r' \leq r$, for each $1 \leq j \leq h$, we use the interval $Y_{(r'-1)h+j}^F$ of $I^F(x)$ as $Z(\mathcal{I}_j(C_{r'}, x))$, and we say that it is the sub-interval of $I^F(x)$ that belongs to instance $\mathcal{I}_j(C_{r'}, x)$. Intuitively, if x is assigned the value FALSE, then we will route a large number of demand pairs in $\mathcal{M}^F(x)$. The paths routing these pairs will “block” the intervals Y_j^F of $I^F(x)$, thus preventing us from routing demand pairs that belong to instances $\mathcal{I}_j(C_{r'}, x)$, for $1 \leq j \leq h$ and $C_{r'} \in \mathcal{C}(x)$.

We treat the subset $\mathcal{C}'(x) \subseteq \mathcal{C}$ of clauses containing the negation of x similarly, except that we assign to each resulting instance an interval Y_j^T of $I^T(x)$.

Clause Gadgets. Consider some clause $C_q = (\ell_{q_1} \vee \ell_{q_2} \vee \ell_{q_3})$. For each one of the three literals of C_q , we construct h level- i width- c instances, with instances $\{\mathcal{I}_j(C_q, \ell_{q_z})\}_{j=1}^h$ representing the literal ℓ_{q_z} , for $1 \leq z \leq 3$. Recall that $B(C_q)$ is a grid of height $H^C = H_i$ and length $L^C = 3chL'_i$. We partition $B(C_q)$ into h sub-grids $B^1(C_q), \dots, B^h(C_q)$, each of which has height H_i and length $3cL'_i$. For each $1 \leq j \leq h$, we place the boxes $B(\mathcal{I}_j(C_q, \ell_{q_1})), B(\mathcal{I}_j(C_q, \ell_{q_2})), B(\mathcal{I}_j(C_q, \ell_{q_3}))$ inside $B^j(C_q)$ side-by-side in this order (see Figure 9).

The intervals $Z(\mathcal{I}_j(C_q, \ell_{q_z}))$ are the same as the ones defined in the constructions of the variable gadgets. We denote by $\mathcal{M}(C_q)$ the set of all demand pairs whose destinations lie in $B(C_q)$, and we call them *clause-pairs representing C_q* . For each $1 \leq z \leq 3$, we denote by $\mathcal{M}(C_q, \ell_{q_z})$ the set of all demand pairs that belong to instances $\mathcal{I}_j(C_q, \ell_{q_z})$, for $1 \leq j \leq h$, and we sometimes say that they represent literal ℓ_{q_z} of clause C_q . We then let $\mathcal{M}^C = \bigcup_{q=1}^m \mathcal{M}(C_q)$ be the set of all clause-pairs, and $\mathcal{M}^V = \bigcup_{j=1}^n \mathcal{M}(x_j)$ the set of all variable-pairs. Our final set of demand pairs is $\mathcal{M} = \mathcal{M}^V \cup \mathcal{M}^C$. This completes the definition of the level- $(i + 1)$ instance.

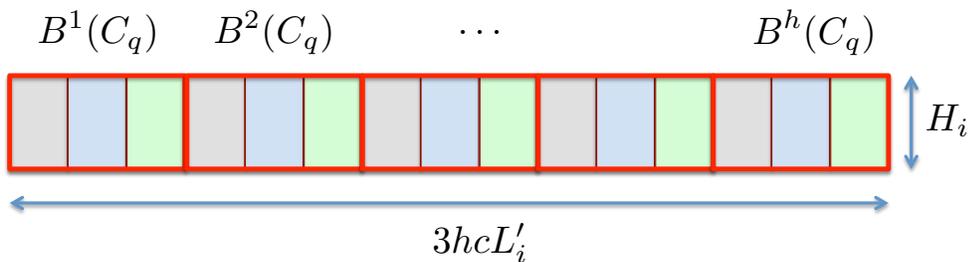


Figure 9: Box $B(C_q)$. Different colors show boxes that represent the three different literals.

5 Yes-Instance Analysis

The goal of this section is to prove the following theorem.

Theorem 5.1 *Assume that the input 3SAT(5) formula φ is a YES-INSTANCE. Then for all $i \geq 0$, for every instantiation of the level- i instance \mathcal{I} , there is a solution routing N_i demand pairs, that respects the box $B(\mathcal{I})$.*

The remainder of this section is devoted to proving this theorem. The proof proceeds by induction. For $i = 0$, $N_0 = 1$, and it is easy to see that for any instantiation of the level-0 instance \mathcal{I} , there is a solution that routes the unique demand pair of this instance and respects the box $B(\mathcal{I})$. We now assume that the theorem holds for some $i \geq 0$, and prove it for a level- $(i+1)$ instance, that we denote by \mathcal{I} . We assume that we are given an instantiation of instance \mathcal{I} , that consists of a grid G_{i+1} of length at least $2L_{i+1} + 2L'_{i+1} + 4H_{i+1}$ and height at least $3H_{i+1}$, the placement of the path $Z(\mathcal{I})$ on the top boundary of G_{i+1} , and the placement of the box $B(\mathcal{I})$ inside G_{i+1} , at distance at least H_{i+1} from its boundaries. We denote the resulting graph by G , and the resulting set of demand pairs by \mathcal{M} . Recall that our construction combines a number of level- i instances. From the induction hypothesis, for each such instance \mathcal{I}' , for every instantiation of instance \mathcal{I}' , there is a set $\mathcal{P}(\mathcal{I}')$ of disjoint paths, routing a set of N_i demand pairs of \mathcal{I}' , such that the paths in $\mathcal{P}(\mathcal{I}')$ respect the box $B(\mathcal{I}')$. It is easy to verify that, if a set $\mathcal{M}^*(\mathcal{I}')$ of demand pairs of \mathcal{I}' has a routing that respects $B(\mathcal{I}')$ in one instantiation of \mathcal{I}' , then it has such a routing in every instantiation of $B(\mathcal{I}')$. Therefore, for every level- i instance \mathcal{I}' , we can fix one such set $\mathcal{M}^*(\mathcal{I}')$ of demand pairs with $|\mathcal{M}^*(\mathcal{I}')| = N_i$, and assume that $\mathcal{M}^*(\mathcal{I}')$ has a routing that respects $B(\mathcal{I}')$ in every instantiation of \mathcal{I}' .

Recall that our construction combines level- i instances into c_{i+1} -wide level- i instances. Let \mathcal{I}'' be any such c_{i+1} -wide level- i instance, and assume that it consists of level- i instances $\mathcal{I}_1, \dots, \mathcal{I}_{c_{i+1}}$. We set $\mathcal{M}^*(\mathcal{I}'') = \bigcup_{j=1}^{c_{i+1}} \mathcal{M}^*(\mathcal{I}_j)$, so $|\mathcal{M}^*(\mathcal{I}'')| = c_{i+1}N_i$. It is easy to see that for any instantiation of \mathcal{I}'' , there is a routing of all demand pairs in $\mathcal{M}^*(\mathcal{I}'')$, such that for each $1 \leq j \leq c_{i+1}$, the demand pairs in $\mathcal{M}^*(\mathcal{I}_j)$ are routed via paths that respects the box $B(\mathcal{I}_j)$.

Consider now the given instantiation (G, \mathcal{M}) of the level- $(i+1)$ instance \mathcal{I} . We first select the set $\hat{\mathcal{M}} \subseteq \mathcal{M}$ of demand pairs that we route, and then compute their routing. We fix some assignment \mathcal{A} to the variables $\{x_1, \dots, x_n\}$ of φ , that satisfies all clauses.

Variable Pairs. Let x be some variable, and let $\hat{\mathcal{M}}^X(x) = \bigcup_{j=1}^{60h} \mathcal{M}^*(\mathcal{I}_j^X(x))$ — the set of all demand pairs that are routed by the solutions to the c_{i+1} -wide level- i instances $\mathcal{I}_1^X(x), \dots, \mathcal{I}_{60h}^X(x)$.

Notice that $|\hat{\mathcal{M}}^X(x)| = 60hc_{i+1}N_i$. If x is assigned the value TRUE, then we let $\hat{\mathcal{M}}^T(x) = \bigcup_{j=1}^{5h+1} \mathcal{M}^*(\mathcal{I}_j^T(x))$, and we set $\hat{\mathcal{M}}^F(x) = \emptyset$. Notice that $|\hat{\mathcal{M}}^T(x)| = (5h+1)c_{i+1}N_i$ in this case. Otherwise, we let $\hat{\mathcal{M}}^F(x) = \bigcup_{j=1}^{5h+1} \mathcal{M}^*(\mathcal{I}_j^F(x))$, so $|\hat{\mathcal{M}}^F(x)| = (5h+1)c_{i+1}N_i$, and we set $\hat{\mathcal{M}}^T(x) = \emptyset$. We denote $\hat{\mathcal{M}}(x) = \hat{\mathcal{M}}^X(x) \cup \hat{\mathcal{M}}^T(x) \cup \hat{\mathcal{M}}^F(x)$, and we let $\hat{\mathcal{M}}^V = \bigcup_x \hat{\mathcal{M}}(x)$, so $|\hat{\mathcal{M}}^V| = nc_{i+1}(65h+1)N_i$.

Clause Pairs. Let $C_q \in \mathcal{C}$ be a clause, and let ℓ_q be a literal of C_q whose value is TRUE (if there are several such literals, we select any one of them arbitrarily). We say that clause C_q *chooses* the literal ℓ_q . For simplicity, we denote $\hat{\mathcal{M}}_j(C_q) = \mathcal{M}^*(\mathcal{I}_j(C_q, \ell_q))$, for all $1 \leq j \leq h$, and we let $\hat{\mathcal{M}}(C_q) = \bigcup_{j=1}^h \hat{\mathcal{M}}_j(C_q)$. Let $\hat{\mathcal{M}}^C = \bigcup_{q=1}^m \hat{\mathcal{M}}(C_q)$. Clearly, for each $1 \leq q \leq m$, $|\hat{\mathcal{M}}(C_q)| = hc_{i+1}N_i$, and overall, $|\hat{\mathcal{M}}^C| = mhc_{i+1}N_i = 5nhc_{i+1}N_i/3$.

Finally, we let $\hat{\mathcal{M}} = \hat{\mathcal{M}}^V \cup \hat{\mathcal{M}}^C$, so $|\hat{\mathcal{M}}| = nc_{i+1} \cdot (65h+1)N_i + 5nc_{i+1}hN_i/3 = n \cdot N_i c_{i+1} (200h/3+1) = N_{i+1}$. It is now enough to prove the following lemma.

Lemma 5.2 *There is a set \mathcal{P} of node-disjoint paths in graph G , routing all demand pairs in $\hat{\mathcal{M}}$, such that \mathcal{P} respects box $B(\mathcal{I})$.*

For convenience, we denote the first row of the grid G_{i+1} by R . Let $S' \subseteq S(\hat{\mathcal{M}})$ be any subset of the source vertices of the demand pairs in $\hat{\mathcal{M}}$. We say that the sources of S' appear *consecutively* on R , iff there is a sub-path P of R that contains all the vertices of S' and does not contain any vertex of $S(\hat{\mathcal{M}}) \setminus S'$. We let \mathcal{O} be the left-to-right ordering of the vertices of $S(\hat{\mathcal{M}})$ on row R .

Consider some variable x , and denote by $\hat{\mathcal{M}}'(x)$ the set of all demand pairs in $\hat{\mathcal{M}}^C$ whose sources lie on the interval $I(x)$ — these are the demand pairs representing the clauses C_q that chose either x or $\neg x$ as their literal. Assume first that x is assigned the value TRUE. Then all sources of the demand pairs in $\hat{\mathcal{M}}(x)$ lie on the intervals $I^X(x)$ and $I^T(x)$, while all sources in set $S(\hat{\mathcal{M}}'(x))$ lie on $I^F(x)$: indeed, since x is assigned the value TRUE, the corresponding clauses must contain x without negation and so their sources lie on $I^F(x)$. Therefore, the sources of each set $\hat{\mathcal{M}}(x)$ and $\hat{\mathcal{M}}'(x)$ are consecutive on R . If x is assigned the value FALSE, then similarly all sources of the demand pairs in $\hat{\mathcal{M}}(x)$ are consecutive on R , while all sources of the demand pairs in $\hat{\mathcal{M}}'(x)$ appear on $I^T(x)$ and are therefore also consecutive on R . In either case, for every clause $C_q \in \mathcal{C}$, the vertices of $S(\hat{\mathcal{M}}(C_q))$ are consecutive on R .

In order to construct the routing, it is convenient to use special subgraphs of G that we call snakes, in which routing can be done easily. We start by defining a corridor, and then combine several corridors to define a snake.

Recall that our graph G is a subgraph of a grid G_{i+1} . Recall also that, given a set \mathcal{R} of consecutive rows of G_{i+1} and a set \mathcal{W} of consecutive columns of G_{i+1} , we denote by $\Upsilon(\mathcal{R}, \mathcal{W})$ the subgraph of G_{i+1} induced by the set $\{v(j, j') \mid R_j \in \mathcal{R}, W_{j'} \in \mathcal{W}\}$ of vertices. We say that $\Upsilon = \Upsilon(\mathcal{R}, \mathcal{W})$ is a *corridor* iff every vertex of Υ belongs to G . For convenience, we will say that Υ is a corridor spanned by the rows in \mathcal{R} and the columns of \mathcal{W} . Let R' and R'' be the first and the last row of \mathcal{R} respectively, and let W' and W'' be the first and the last column of \mathcal{W} respectively. Each of the four paths $\Upsilon \cap R'$, $\Upsilon \cap R''$, $\Upsilon \cap W'$ and $\Upsilon \cap W''$ is called a *boundary edge* of Υ , and their union is called the *boundary* of Υ . We say that two corridors Υ, Υ' are *internally disjoint*, iff every vertex $v \in \Upsilon \cap \Upsilon'$ belongs to the boundaries of both corridors. We say that two internally disjoint corridors Υ, Υ' are *neighbors* iff $\Upsilon \cap \Upsilon' \neq \emptyset$.

We are now ready to define snakes. A snake \mathcal{Y} of length ℓ is a sequence $\Upsilon_1, \Upsilon_2, \dots, \Upsilon_\ell$ of ℓ corridors

that are pairwise internally disjoint. Moreover, for all $1 \leq \ell', \ell'' < \ell$, $\Upsilon_{\ell'}$ is a neighbor of $\Upsilon_{\ell''}$ iff $|\ell' - \ell''| = 1$. We say that the width of the snake is w iff for each $1 \leq \ell' \leq \ell$, $\Upsilon_{\ell'}$ is spanned by a set of at least w rows and by a set of at least w columns, and for all $1 \leq \ell' < \ell$, $\Upsilon_{\ell'} \cap \Upsilon_{\ell'+1}$ contains at least w vertices. We use the following simple claim, whose proof is deferred to the Appendix, for routing in snakes.

Claim 5.3 *Let $\mathcal{Y} = (\Upsilon_1, \dots, \Upsilon_\ell)$ be a snake of width w , and let A, A' be two sets of vertices with $|A| = |A'| \leq w - 2$, such that the vertices of A lie on a single boundary edge of Υ_1 and the vertices of A' lie on a single boundary edge of Υ_ℓ . Then there is a set \mathcal{Q} of node-disjoint paths contained in $\bigcup_{\ell'=1}^{\ell} \Upsilon_{\ell'}$, that connect every vertex of A to a distinct vertex of A' .*

Our routing consists of two steps. In the first step, we connect each source vertex $s \in S(\hat{\mathcal{M}})$ to the top boundary of the unique box in $\mathcal{B} = \{B(x_1), \dots, B(x_n), B(C_1), \dots, B(C_m)\}$, that contains its corresponding destination vertex. We will later select specific vertices on the top boundary of each such box to which the sources are routed. The resulting paths will be internally disjoint from the boxes in \mathcal{B} . In the second step, we complete the routing inside each box. The first step is summarized in the following claim.

Claim 5.4 *Suppose we are given, for every variable x_j , a set $A(x_j)$ of $|\hat{\mathcal{M}}(x_j)|$ vertices on the top boundary of $B(x_j)$, and for every clause $C_q \in \mathcal{C}$, a set $A(C_q)$ of $|\hat{\mathcal{M}}(C_q)|$ vertices on the top boundary of $B(C_q)$. Then there is a collection \mathcal{P}' of node-disjoint paths in G with the following properties:*

- *the paths of \mathcal{P}' are internally disjoint from all boxes in $\mathcal{B} = \{B(x_1), \dots, B(x_n), B(C_1), \dots, B(C_m)\}$;*
- *for every variable x_j of φ , there is a subset $\mathcal{P}(x_j) \subseteq \mathcal{P}'$ of paths, that connect every vertex in $S(\hat{\mathcal{M}}(x_j))$ to a vertex of $A(x_j)$, so that the paths in $\mathcal{P}(x_j)$ are order-preserving; and*
- *for every clause $C_q \in \mathcal{C}$, there is a subset $\mathcal{P}(C_q) \subseteq \mathcal{P}'$ of paths, connecting every vertex in $S(\hat{\mathcal{M}}(C_q))$ to a vertex of $A(C_q)$, so that the paths in $\mathcal{P}(C_q)$ are order-preserving.*

Proof: We provide here a high-level sketch of the proof. Turning it into a formal proof is straightforward but somewhat tedious; we defer the formal proof to the Appendix. We construct 5 sets of node-disjoint paths, $\mathcal{P}_0, \dots, \mathcal{P}_4$. Let Z' be the set of $|\hat{\mathcal{M}}|$ leftmost vertices on the opening of $B(\mathcal{I})$. Set \mathcal{P}_0 of paths routes all vertices in $S(\hat{\mathcal{M}})$ to the vertices of Z' via node-disjoint paths that are order-preserving and internally disjoint from $B(\mathcal{I})$, in a straightforward manner. Next, we select a subset Γ of $|\hat{\mathcal{M}}|$ vertices on the top row of B^V . Set \mathcal{P}_1 of paths connects every vertex in Z' to a distinct vertex of Γ , so that the paths in \mathcal{P}_1 are node-disjoint, order-preserving, internally disjoint from $B^V \cup B^C$, and contained in $B(\mathcal{I})$. In order to route the paths in \mathcal{P}_1 , we construct an appropriate snake inside $B(\mathcal{I})$ (see Figure 10). If we combine the paths in \mathcal{P}_0 and \mathcal{P}_1 , we obtain a set of node-disjoint order-preserving paths, connecting every vertex of $S(\hat{\mathcal{M}})$ to a distinct vertex of Γ . Paths originating at the vertices corresponding to variable demand pairs then continue to their corresponding boxes $B(x_j)$, while paths originating at the vertices corresponding to clause demand pairs continue directly to the bottom boundary of B^V , exploiting the spacing between the boxes $B(x_j)$. The vertices of Γ are selected in such a way that each such path can be implemented as a sub-path of a column of B^V (see Figure 10). The resulting set of paths is denoted by \mathcal{P}_2 . We then connect the endpoints of the paths in \mathcal{P}_2 that appear on the bottom boundary of B^V to a subset Γ''' of vertices on the top boundary of B^C , by defining an appropriate snake that exploits

the spacing between B^V and B^C . This latter set of paths is denoted by \mathcal{P}_3 . Let \mathcal{P}'' be the set of paths obtained by combining the paths in $\mathcal{P}_0, \dots, \mathcal{P}_3$. For every variable x_j of φ , there is a subset $\mathcal{P}(x_j) \subseteq \mathcal{P}''$ of paths, connecting every vertex of $S(\hat{\mathcal{M}}(x_j))$ to a vertex of $A(x_j)$. For every clause $C_q \in \mathcal{C}$, there is a subset $\mathcal{P}'(C_q) \subseteq \mathcal{P}''$ of paths, connecting every vertex of $S(\mathcal{M}(C_q))$ to a distinct vertex of Γ''' . We denote by $\Gamma(C_q) \subseteq \Gamma'''$ the set of vertices that serve as endpoints of the paths in $\mathcal{P}'(C_q)$. Note that for each clause $C_q \in \mathcal{C}$, the vertices of $\Gamma(C_q)$ appear consecutively on the top boundary of B^C .

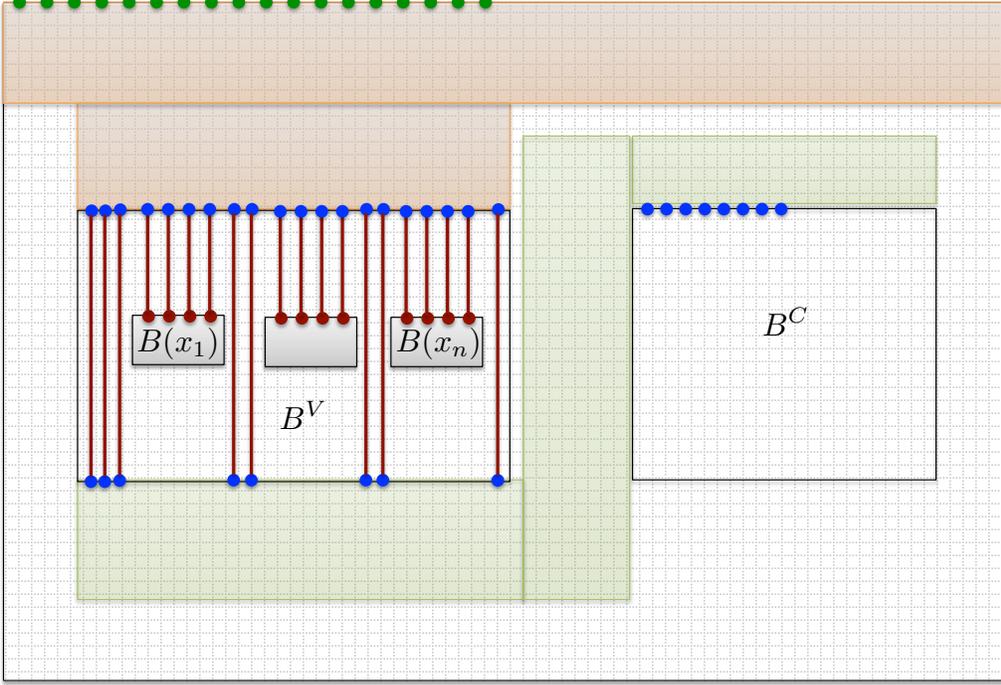


Figure 10: Routing the sets $\mathcal{P}_1, \mathcal{P}_2$ and \mathcal{P}_3 of paths inside $B(\mathcal{I})$. The paths in \mathcal{P}_2 are shown in red; the paths of \mathcal{P}_1 are routed inside the orange snake, and the paths of \mathcal{P}_3 are routed inside the green snake.

We can now define an ordering \mathcal{O}' of the clauses in \mathcal{C} as follows: clause C_q appears before clause $C_{q'}$ in this ordering iff the vertices of $\Gamma(C_q)$ appear to the left of the vertices of $\Gamma(C_{q'})$ on the top row of B^C . Our final step is to define a set \mathcal{P}_4 of node-disjoint paths, that are contained in B^C and are internally disjoint from all boxes $\{B(C_q) \mid C_q \in \mathcal{C}\}$, such that for each clause $C_q \in \mathcal{C}$, there is a subset $\mathcal{P}'(C_q) \subseteq \mathcal{P}_4$ of order-preserving paths, connecting the vertices of $\Gamma(C_q)$ to the vertices of $A(C_q)$. Set \mathcal{P}_4 is constructed by using a standard snake-like routing inside the box B^C . For each clause C_q , we extend the box $B(C_q)$ by N_{i+1} columns to the right and to the left, obtaining a larger box $B'(C_q)$. The idea is to route the paths, so that they visit the boxes $B'(C_1), \dots, B'(C_m)$ in this order. For all $1 \leq q \leq m$, only paths corresponding to the clauses C_q, C_{q+1}, \dots, C_m visit the box $B'(C_q)$. The paths corresponding to clause C_q enter the box $B(C_q)$ and terminate at its top boundary. For each clause $C \in \{C_{q+1}, \dots, C_m\}$, if C appears before C_q in the ordering \mathcal{O}' , then the paths corresponding to C traverse the box $B'(C_q)$ to the left of $B(C_q)$, and otherwise they traverse it to the right of $B(C_q)$. The routing itself is implemented by constructing an appropriate collection of snakes, where each snake is used to connect the vertices on the bottom boundary of $B'(C_q)$ to the vertices on the top boundary of $B'(C_{q+1})$ and exploits the spacing between the two

boxes (see Section A.2 of Appendix for more details). \square

We now define the sets $\{A(x_j) \mid 1 \leq j \leq n\} \cup \{A(C_q) \mid C_q \in \mathcal{C}\}$ of vertices and complete the routing.

Consider first some clause $C_q \in \mathcal{C}$, and assume that it has chosen the literal ℓ_q . Fix some index $1 \leq j \leq h$, and consider the width- c_{i+1} level- i instance $\mathcal{I}' = \mathcal{I}_j(C_q, \ell_q)$. Instance \mathcal{I}' consists of c_{i+1} level- i instances, that we denote by $\mathcal{I}'_1, \dots, \mathcal{I}'_{c_{i+1}}$. Then set $\mathcal{M}(C_q)$ contains $c_{i+1}N_i$ demand pairs that belong to instance \mathcal{I}' , with exactly N_i pairs from each instance \mathcal{I}'_r . For each such instance \mathcal{I}'_r , we select N_i leftmost vertices on the opening of the box $B(\mathcal{I}'_r)$. We then let $A(C_q)$ be the set of all such vertices we have selected, for all $1 \leq j \leq h$ and $1 \leq r \leq c_{i+1}$, so $|A(C_q)| = hc_{i+1}N_i = |\hat{\mathcal{M}}(C_q)|$. Notice that, if we are given any set $\mathcal{P}(C_q)$ of $|\hat{\mathcal{M}}(C_q)|$ node-disjoint order-preserving paths, that connect the vertices of $S(\hat{\mathcal{M}}(C_q))$ to the vertices of $A(C_q)$, such that the paths in $\mathcal{P}(C_q)$ are internally disjoint from the box $B(C_q)$, then we can extend the paths in $\mathcal{P}(C_q)$ inside $B(C_q)$, so that they route the set $\hat{\mathcal{M}}(C_q)$ of demand pairs. In order to do so, for each c_{i+1} -wide level- i instance $\mathcal{I}' = \mathcal{I}_j(C_q, \ell_q)$, for each level- i instance \mathcal{I}'_r that was used to construct \mathcal{I}' , we use the routing that respects the box $B(\mathcal{I}'_r)$, which is guaranteed by the induction hypothesis and by our choice of the demand pairs to route, in order to connect the vertices we have selected on the opening of $B(\mathcal{I}'_r)$ to the corresponding destination vertices.

Consider now some variable x_j , for $1 \leq j \leq n$. We now show how to select a set $A(x_j)$ of $|\hat{\mathcal{M}}(x_j)|$ vertices on the top boundary of box $B(x_j)$. We start with $A(x_j) = \emptyset$ and then add vertices to it.

Assume first that x_j is assigned the value TRUE. Let \mathcal{I}' be any of the $(65h+1)c_{i+1}$ level- i instances whose box $B(\mathcal{I}')$ is contained in $B_X(x) \cup B_T(x)$. Recall that $\hat{\mathcal{M}}(x_j)$ contains N_i demand pairs that belong to this instance. Let $A'(\mathcal{I}')$ be the set of N_i leftmost vertices on the opening of the box $B(\mathcal{I}')$. For each vertex $v \in A'(\mathcal{I}')$, we add to $A(x_j)$ the vertex v' that belongs to the top boundary of $B(x_j)$ and lies on the same column W as v . We let P_v be the sub-path of W between v and v' . Notice that $|A(x_j)| = |\hat{\mathcal{M}}(x_j)|$, and, given any set $\mathcal{P}(x_j)$ of node-disjoint order-preserving paths connecting the vertices of $S(\hat{\mathcal{M}}(x_j))$ to the vertices of $A(x_j)$, such that the paths in $\mathcal{P}(x_j)$ are internally disjoint from $B(x_j)$, we can extend these paths inside the box $B(x_j)$, so that they route the set $\hat{\mathcal{M}}(x_j)$ of demand pairs. This is done using the box-respecting routing of each level- i instance as before, together with the set $\{P_v \mid v' \in A(x_j)\}$ of paths.

Finally, assume that x_j is assigned the value FALSE. Notice that the sources of the EXTRA demand pairs for x_j appear to the left of the sources of the FALSE demand pairs of x_j , while the box $B_F(x_j)$ appears to the left of the box $B_X(x_j)$. Therefore, the straightforward routing as above cannot be employed here and we need to “switch” the two sets of paths. In order to do so, we exploit the spacing between the boxes $B_T(x_j), B_F(x_j), B_X(x_j)$ and the boundaries of $B(x_j)$ (see Figure 8).

For each of the $(5h+1)c_{i+1}$ level- i instances whose box $B(\mathcal{I}')$ is contained in $B_F(x)$, we define the set $A'(\mathcal{I}')$ of N_i vertices on the opening of the box $B(\mathcal{I}')$, and for each such vertex v , the corresponding path P_v and vertex v' that is added to $A(x_j)$ exactly as before. Let J denote the set of vertices added to $A(x_j)$ so far. We also add to $A(x_j)$ the set J' of $60hc_{i+1}N_i$ leftmost vertices on the top boundary of $B(x_j)$. This ensures that $|A(x_j)| = (65h+1)hc_{i+1}N_i = |\hat{\mathcal{M}}(x_j)|$. Assume now that we are given any set $\mathcal{P}(x_j)$ of node-disjoint order-preserving paths connecting the vertices of $S(\hat{\mathcal{M}}(x_j))$ to the vertices of $A(x_j)$, such that the paths in $\mathcal{P}(x_j)$ are internally disjoint from $B(x_j)$. We now show how to extend these paths inside the box $B(x_j)$, so that they route the set $\hat{\mathcal{M}}(x_j)$ of demand pairs. We partition the set $\mathcal{P}(x_j)$ into two subsets: set \mathcal{P}^X contains all paths that originate at the sources of the demand pairs in $\hat{\mathcal{M}}^X(x_j)$ — that is, the EXTRA demand pairs of x_j , and \mathcal{P}^F contains all remaining demand pairs, that must originate at the source vertices of the FALSE demand pairs for x_j . Since the set $\mathcal{P}(x_j)$ of paths is order-preserving, the paths in

\mathcal{P}^X terminate at the vertices of J' , while the paths in \mathcal{P}^F terminate at the vertices of J . We extend the paths in set \mathcal{P}^F exactly as before, using the paths in set $\{P_v \mid v' \in J\}$, and then using the box-preserving routing of each corresponding level- i instance to connect each source vertex of $S(\hat{\mathcal{M}}^F(x_j))$ to its destination.

In order to extend the paths in \mathcal{P}^X we do the following. For each level- i instance \mathcal{I}' whose box $B(\mathcal{I}')$ is contained in $B_X(x_j)$, we let $A'(\mathcal{I}')$ be the set of N_i left-most vertices on the opening of $B(\mathcal{I}')$. Let A' be the set of all such vertices, for all such instances \mathcal{I}' , so $|A'| = |J'|$. It is now enough to construct a set \mathcal{Q} of order-preserving node-disjoint paths contained in $B(x_j)$ that connect every vertex in J to a distinct vertex of A' , such that the paths in \mathcal{Q} are internally disjoint from $B_T(x_j) \cup B_F(x_j) \cup B_X(x_j)$ and are completely disjoint from $\{P_v \mid v' \in J\}$. We show the existence of the set \mathcal{Q} of paths by constructing a snake \mathcal{Y} , that consists of four corridors. The first corridor, Υ_1 , is the set of the first N_{i+1} columns of $B(x_j)$. The third corridor, Υ_3 , is the set of the last N_{i+1} columns of $B(x_j)$. Let W' be the last column of Υ_1 and let W'' be the first column of Υ_3 . The second corridor, Υ_2 , is spanned by the bottom N_{i+1} rows of $B(x_j)$ and the columns of $B(x_j)$ from W' to W'' , including these two columns. Let W''' be the leftmost column of G_{i+1} that intersects $B_X(x_j)$; let R' be the topmost row of G_{i+1} that intersects $B_X(x_j)$, and let \mathcal{R} be the set of N_{i+1} consecutive rows lying above R' , including R' . The last corridor, Υ_4 , is spanned by the set \mathcal{R} of rows, and the set of all columns from W''' to W'' (see Figure 11). Using Claim 5.3, it is immediate to verify that the desired set \mathcal{Q} of paths exists inside the resulting snake. We then extend the routing inside each box $B(\mathcal{I}')$ of each corresponding level- i instance \mathcal{I}' exactly as before.

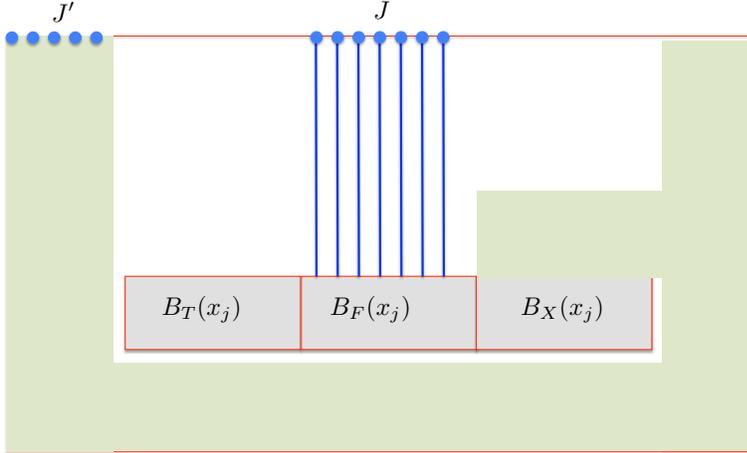


Figure 11: Routing inside box $B(x_j)$ if x_j is assigned value F .

Using Claim 5.4 with our definitions of the sets $\{A(x_j) \mid 1 \leq j \leq n\} \cup \{A(C_q) \mid 1 \leq q \leq m\}$ of vertices, and the above discussion, we can route all demand pairs in $\hat{\mathcal{M}}$ via a set \mathcal{P} of node-disjoint paths that respects the box $B(\mathcal{I})$.

6 No-Instance Analysis

In this section we analyze the NO-INSTANCE case, by proving the following theorem.

Theorem 6.1 *Assume that φ is a NO-INSTANCE. Then for every integer $i \geq 0$, for every instantiation of the level- i instance \mathcal{I} , and for every solution \mathcal{P} to this instance, $|\mathcal{P}| \leq N'_i$.*

The proof is again by induction. For the base case of $i = 0$, $N'_i = 1$, and the corresponding level-0 instance contains a single demand pair, so the theorem clearly holds. We now assume that the theorem holds for some value $i \geq 0$ and prove it for a level- $(i + 1)$ instance \mathcal{I} . We assume that we are given some instantiation of \mathcal{I} , and from now on our goal is to prove that no solution to this instance of NDP can route more than $N'_{i+1} = (1 - \delta)nc_{i+1} \cdot (200h/3 + 1)N'_i$ demand pairs, where $\delta = 8\epsilon^2/10^{12}$. We assume for contradiction that this is not the case, and we let $\tilde{\mathcal{P}}$ be a collection of more than N'_{i+1} node-disjoint paths, routing a set $\tilde{\mathcal{M}} \subseteq \mathcal{M}$ of demand pairs. For every demand pair $(s, t) \in \tilde{\mathcal{M}}$, we let $P(s, t) \in \tilde{\mathcal{P}}$ be the path routing this pair in the solution.

Recall that our construction of a level- $(i + 1)$ instance \mathcal{I} consists of a number of copies of c_{i+1} -wide level- i instances: For every variable x of φ , we have constructed $(70h + 2)$ such instances ($60h$ instances for the extra pairs, and $(5h + 1)$ instances each for TRUE and FALSE pairs); for every clause $C \in \mathcal{C}$, we have constructed $3h$ such instances. Therefore, overall we use $(70h + 2)n + 3hm = 75nh + 2n$ copies of c_{i+1} -wide level- i instances (we have used the fact that $m = 5n/3$). We assume (by induction) that at most $c_{i+1}N'_i$ pairs from each such instance are in $\tilde{\mathcal{M}}$. We say that a c_{i+1} -wide level- i instance \mathcal{I}' is *interesting* iff at least $25H_i$ demand pairs of $\mathcal{M}(\mathcal{I}')$ belong to $\tilde{\mathcal{M}}$; otherwise we say that it is uninteresting. We let $\hat{\mathcal{M}} \subseteq \tilde{\mathcal{M}}$ be the set of all demand pairs that belong to uninteresting instances, and we call them *excess pairs*. We need the following simple observation.

Observation 6.2 $|\hat{\mathcal{M}}| \leq \delta(200h/3 + 1)nc_{i+1}N'_i$.

Proof: As observed above, there are at most $75nh + 2n$ uninteresting instances, each of which contributes at most $25H_i \leq \frac{50c_{i+1}\epsilon^2N'_i}{10^{13}}$ excess demand pairs. Therefore, it is enough to show that:

$$(75nh + 2n) \cdot \frac{50c_{i+1}\epsilon^2N'_i}{10^{13}} \leq \delta \left(\frac{200h}{3} + 1 \right) nc_{i+1}N'_i,$$

which is immediate to verify, substituting $\delta = 8\epsilon^2/10^{12}$. \square

It would be convenient for us to assume that no excess pairs exist. In order to do so, we discard all excess pairs from $\tilde{\mathcal{M}}$. From Observation 6.2, $|\tilde{\mathcal{M}}| \geq (1 - 2\delta)nc_{i+1}(200h/3 + 1)N'_i$ still holds.

For every variable x , we let $\tilde{\mathcal{M}}(x) = \tilde{\mathcal{M}} \cap \mathcal{M}(x)$, and for every clause C_q , we let $\tilde{\mathcal{M}}(C_q) = \tilde{\mathcal{M}} \cap \mathcal{M}(C_q)$. We also denote by $\tilde{\mathcal{M}}^V = \bigcup_{j=1}^n \tilde{\mathcal{M}}(x_j)$ and by $\tilde{\mathcal{M}}^C = \bigcup_{q=1}^m \tilde{\mathcal{M}}(C_q)$.

For the sake of the NO-INSTANCE-analysis, it is convenient to view our construction slightly differently. Let φ be the input 3SAT(5) formula, and recall that $\mathcal{C} = \{C_1, \dots, C_m\}$ is the set of its clauses. For each clause $C_q \in \mathcal{C}$, we create h new clauses C_q^1, \dots, C_q^h , each of which is a copy of the original clause. We let $\mathcal{C}' = \{C_q^j \mid 1 \leq q \leq m, 1 \leq j \leq h\}$ be the resulting set of clauses, and φ' the corresponding 3SAT formula. In order to avoid confusion, we refer to the clauses in \mathcal{C} as the *original* clauses, to the clauses of \mathcal{C}' as the *new clauses*, and for each $1 \leq q \leq m$, $1 \leq j \leq h$, we call C_q^j the *jth copy of clause C_q* . Recall that the clause gadget for $C_q \in \mathcal{C}$ contains h boxes $B^1(C_q), \dots, B^h(C_q)$, where box $B^j(C_q)$ is the union of three boxes: $B(\mathcal{I}_j(C_q, \ell_{q_1}))$, $B(\mathcal{I}_j(C_q, \ell_{q_2}))$ and $B(\mathcal{I}_j(C_q, \ell_{q_3}))$ (see Figure 9). We think of the box $B^j(C_q)$ as representing the new clause $C_q^j \in \mathcal{C}'$. For convenience, we denote by $\tilde{\mathcal{M}}(C_q^j) \subseteq \tilde{\mathcal{M}}(C_q)$ the set of all demand pairs routed by our solution whose destinations lie in $B^j(C_q)$. This set is further partitioned into three subsets, $\tilde{\mathcal{M}}(C_q^j, \ell_{q_1})$, $\tilde{\mathcal{M}}(C_q^j, \ell_{q_2})$, $\tilde{\mathcal{M}}(C_q^j, \ell_{q_3})$, each of which contains demand pairs from the instances $\mathcal{I}_j(C_q, \ell_{q_1})$, $\mathcal{I}_j(C_q, \ell_{q_2})$, and $\mathcal{I}_j(C_q, \ell_{q_3})$ respectively. The following observation is immediate:

Observation 6.3 *If φ is a NO-INSTANCE, then for any assignment to its variables, at most $(1 - \epsilon)mh$ clauses of C' are satisfied.*

Encircling and its Resolution Let $(s, t) \in \tilde{\mathcal{M}}$ be any demand pair routed by the solution. Recall that (s, t) belongs to some level- i instance \mathcal{I}' , and we have defined a line Q_t containing at most $H_i/2$ vertices of the graph, that connects t to the bottom of the box $B(\mathcal{I}')$ (which is a cut-out box). We say that a demand pair $(s', t') \in \tilde{\mathcal{M}}$ *encircles* pair (s, t) iff path $P(s', t')$ contains a vertex lying on Q_t . Since Q_t contains at most $H_i/2$ vertices, at most $H_i/2$ demand pairs may encircle (s, t) . We repeatedly use the following simple lemma.

Lemma 6.4 *Let S_1, \dots, S_r be a collection of disjoint subsets of $\tilde{\mathcal{M}}$, such that for all $1 \leq j \leq r$, $|S_j| \geq r^2 H_i/2$. Then there is a collection $\mathcal{M}' = \{(s_1, t_1), \dots, (s_r, t_r)\}$ of demand pairs, such that for all $1 \leq j \leq r$, $(s_j, t_j) \in S_j$, and for all distinct $(s, t), (s', t') \in \mathcal{M}'$, pair (s', t') does not encircle pair (s, t) .*

Proof: We perform $(r - 1)$ iterations. At the beginning of iteration ℓ , we are given a set $\mathcal{M}' = \{(s_1, t_1), \dots, (s_{\ell-1}, t_{\ell-1})\}$ of demand pairs, where for all $1 \leq j \leq \ell - 1$, $(s_j, t_j) \in S_j$, and no two pairs in \mathcal{M}' encircle each other. Additionally, for each $\ell \leq j \leq r$, we are given a subset $S'_j \subseteq S_j$ of $(r^2 - r(\ell - 1))H_i/2$ demand pairs, such that no pair in S'_j encircles a pair in \mathcal{M}' , and no pair in S'_j is encircled by a pair in \mathcal{M}' . Therefore, at the end of iteration $(r - 1)$, set \mathcal{M}' contains one pair from each set S_1, \dots, S_{r-1} , and $S_r \neq \emptyset$. We add an arbitrary pair of S_r to \mathcal{M}' to obtain the desired output.

At the beginning of the algorithm, $\mathcal{M}' = \emptyset$, and for each $1 \leq j \leq r$, set S'_j contains any subset of $r^2 H_i/2$ demand pairs of S_j . We now describe an execution of some iteration ℓ . Let $\mathcal{M}'' = \bigcup_{j=\ell+1}^r S'_j$, and let U be the set of all vertices appearing on lines Q_t , where t is a destination vertex of a demand pair in \mathcal{M}'' . Then $|\mathcal{M}''| \leq (r - 1)(r^2 - r(\ell - 1))H_i/2$ and $|U| \leq (r - 1)(r^2 - r(\ell - 1))H_i^2/4$, while S_ℓ contains $(r^2 - r(\ell - 1))H_i/2$ demand pairs. Therefore, there is some demand pair $(s_\ell, t_\ell) \in S_\ell$ that contains at most $(r - 1)H_i/2$ vertices of U . We add (s_ℓ, t_ℓ) to \mathcal{M}' . Consider now some set S'_j , for some $\ell + 1 \leq j \leq r$. Then pair (s_ℓ, t_ℓ) may encircle at most $(r - 1)H_i/2$ pairs of S'_j , and at most $H_i/2$ pairs in S'_j may encircle (s_ℓ, t_ℓ) . We discard from S'_j all pairs that either encircle (s_ℓ, t_ℓ) or are encircled by it. At the end of this procedure, $|S'_j| \geq (r^2 - r(\ell - 1))H_i/2 - rH_i/2 \geq (r^2 - r\ell)H_i/2$. If $|S'_j| > (r^2 - r\ell)H_i/2$, then we discard arbitrary pairs from S'_j until the equality holds. \square

Variable Gadget Analysis We fix some variable x and consider its corresponding gadget. We start with the following lemma.

Lemma 6.5 *Let $\tilde{\mathcal{M}}_X = \tilde{\mathcal{M}} \cap \mathcal{M}^X(x)$, $\tilde{\mathcal{M}}_T = \tilde{\mathcal{M}} \cap \mathcal{M}^T(x)$, and $\tilde{\mathcal{M}}_F = \tilde{\mathcal{M}} \cap \mathcal{M}^F(x)$ be the subsets of $\mathcal{M}^X(x)$, $\mathcal{M}^T(x)$, and $\mathcal{M}^F(x)$, respectively, that are routed by our solution. Then at least one of the sets $\tilde{\mathcal{M}}_X, \tilde{\mathcal{M}}_T, \tilde{\mathcal{M}}_F$ is empty.*

Proof: Assume otherwise. Since we have discarded all excess pairs, each one of the three sets $\tilde{\mathcal{M}}_X, \tilde{\mathcal{M}}_T, \tilde{\mathcal{M}}_F$ contains at least $25H_i$ demand pairs. From Lemma 6.4, we can find three pairs, $(s_X, t_X) \in \tilde{\mathcal{M}}_X$, $(s_T, t_T) \in \tilde{\mathcal{M}}_T$ and $(s_F, t_F) \in \tilde{\mathcal{M}}_F$, with neither pair encircling the other. We now claim that it is impossible to route all three pairs simultaneously via node-disjoint paths. In order to do so, we use the following simple observation.

Observation 6.6 Let Σ be the sphere, and let D, D' be two disjoint closed simple discs on Σ , whose boundaries are denoted by Γ and Γ' , respectively. Let Σ' be the cylinder obtained from Σ by removing the open discs $D \setminus \Gamma$, $D' \setminus \Gamma'$ from it. Assume further that we have three distinct points a_1, a_2, a_3 appearing on Γ in this circular order, and three distinct points a'_1, a'_3, a'_2 appearing on Γ' in this circular order (see Figure 12). Let $\gamma_1, \gamma_2, \gamma_3$ be three simple curves on Σ' , where for each $1 \leq j \leq 3$, γ_j connects a_j to a'_j . Then the three curves cannot be pairwise disjoint.

Proof: Assume otherwise. We can twist the cylinder so that the curve γ_1 becomes a straight vertical line, and then cut the cylinder along this line, obtaining a square, with a_2 appearing to the left of a_3 on the top of the square, and a'_3 appearing to the left of a'_2 on the bottom of the square. It is now immediate to see that it is impossible to connect a_2 to a'_2 and a_3 to a'_3 via disjoint curves.

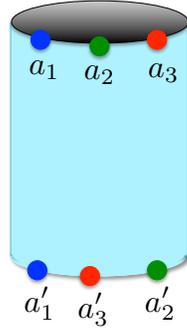


Figure 12: Illustration to Observation 6.6.

□

Assume for contradiction that there are three node-disjoint paths, routing demand pairs (s_X, t_X) , (s_T, t_T) , and (s_F, t_F) . We embed the grid G_{i+1} onto a sphere in a natural way, and then construct a curve γ_X , by concatenating the image of the path $P(s_X, t_X)$ with the line Q_{t_X} , denoting by b_X the endpoint of this curve that is distinct from s_X . If this curve is not simple, we delete all loops from it until it becomes simple. We define curves γ_T , γ_F , and their endpoints b_T and b_F for the pairs (s_T, t_T) , and (s_F, t_F) , respectively, in a similar way. Let Γ be the top boundary of the grid, and let Γ' be the union of the bottom boundaries of the boxes $B_T(x)$, $B_F(x)$ and $B_X(x)$. By slightly thickening Γ and Γ' , we can create two disjoint discs D and D' in the plane, with Γ being the boundary of D and Γ' the boundary of D' , so that the curves γ_X, γ_T and γ_F are internally disjoint from D and D' . Vertices s_T, s_X and s_F appear on Γ in this circular order, while vertices b_T, b_F and b_X appear on Γ' in this circular order. But then we obtain a collection of three disjoint curves on a cylinder that contradicts Observation 6.6. □

The following corollary is now immediate.

Corollary 6.7 $|\tilde{\mathcal{M}}^V| \leq (65h + 1) \cdot n \cdot c_{i+1} N'_i$.

Consider some variable x of φ . If $\tilde{\mathcal{M}} \cap \mathcal{M}^T(x) = \emptyset$, then we assign it the value FALSE, and otherwise we assign it the value TRUE.

Fix some variable $x \in X$ and some index $1 \leq j \leq 5h + 1$. We say that index j is bad for variable x if either (i) x is assigned the value TRUE, and instance $\mathcal{I}_j^T(x)$ is uninteresting; or (ii) x is assigned

the value FALSE, and instance $\mathcal{I}_j^F(x)$ is uninteresting. We will later show that the total number of pairs (x, j) where j is a bad index for x is bounded by $12n$.

Clause Gadget Analysis Consider a new clause $C_q^j \in \mathcal{C}'$ and its three literals $\ell_{q_1}, \ell_{q_2}, \ell_{q_3}$. We say that clause C_q^j is a *troublesome clause*, or a *troublesome copy* of C_q , iff there are at least two values $1 \leq z < z' \leq 3$, for which instances $\mathcal{I}_j(C_q, \ell_{q_z}), \mathcal{I}_j(C_q, \ell_{q_{z'}})$ are both interesting.

Lemma 6.8 *For every original clause $C_q \in \mathcal{C}$, at most three of its copies are troublesome.*

Proof: Assume otherwise, and let $C_q \in \mathcal{C}$ be a clause, such that at least four of its copies are troublesome. Then there are two indices $1 \leq z < z' \leq 3$, and two indices $1 \leq j < j' \leq h$, such that each one of the four instances $\mathcal{I}_j(C_q, \ell_{q_z}), \mathcal{I}_j(C_q, \ell_{q_{z'}}), \mathcal{I}_{j'}(C_q, \ell_{q_z}), \mathcal{I}_{j'}(C_q, \ell_{q_{z'}})$ is interesting, and so each of the four sets $\tilde{\mathcal{M}}(C_q^j, \ell_{q_z}), \tilde{\mathcal{M}}(C_q^j, \ell_{q_{z'}}), \tilde{\mathcal{M}}(C_q^{j'}, \ell_{q_z}), \tilde{\mathcal{M}}(C_q^{j'}, \ell_{q_{z'}})$ contains at least $25H_i$ demand pairs. From Lemma 6.4, we can find demand pairs $(s_1, t_1) \in \tilde{\mathcal{M}}(C_q^j, \ell_{q_z}), (s_2, t_2) \in \tilde{\mathcal{M}}(C_q^j, \ell_{q_{z'}}), (s'_1, t'_1) \in \tilde{\mathcal{M}}(C_q^{j'}, \ell_{q_z}),$ and $(s'_2, t'_2) \in \tilde{\mathcal{M}}(C_q^{j'}, \ell_{q_{z'}})$, such that none of the four resulting pairs encircles another. Notice that the sources of these pairs appear on the top boundary of the grid in one of the following two orders: (s_1, s'_1, s_2, s'_2) , if the variable corresponding to ℓ_{q_z} appears before the variable corresponding to $\ell_{q_{z'}}$ in the ordering (x_1, x_2, \dots, x_n) of the variables, or (s_2, s'_2, s_1, s'_1) otherwise.

We draw a curve γ_1 by concatenating the image of the path $P(s_1, t_1)$ and the line Q_{t_1} . We draw curves γ'_1, γ_2 and γ'_2 for the remaining pairs similarly. On the one hand, none of the the resulting curves may cross each other, while on the other hand the vertices lying on the bottom boundary of the box $B(C_q)$ have been deleted, which is impossible. \square

Corollary 6.9 *For each original clause $C_q \in \mathcal{C}$, $|\tilde{\mathcal{M}}(C_q)| \leq (6 + h)N'_i c_{i+1}$, and overall $|\tilde{\mathcal{M}}^C| \leq 5n(6 + h)N'_i c_{i+1}/3$.*

Proof: Consider some original clause $C_q \in \mathcal{C}$. Each of its non-troublesome copies C_q^j contributes at most $c_{i+1}N'_i$ demand pairs to $\tilde{\mathcal{M}}^C$, while a troublesome copy may contribute at most $3c_{i+1}N'_i$ demand pairs. Since at most three copies are troublesome, overall $|\tilde{\mathcal{M}}(C_q)| \leq (6 + h)N'_i c_{i+1}$ and $|\tilde{\mathcal{M}}^C| \leq m(6 + h)N'_i c_{i+1} = 5n(6 + h)N'_i c_{i+1}/3$, since $m = 5n/3$. \square

In the rest of our proof, we will reach a contradiction by proving that the current assignment to the variables of φ satisfies more than $(1 - \epsilon)hm$ clauses in \mathcal{C}' . In order to do so, we gradually discard clauses from \mathcal{C}' , until we obtain a large enough subset of clauses that is guaranteed to be satisfied by the current assignment.

Our first step is to define uninteresting clauses. Recall that for each new clause $C_q^j \in \mathcal{C}'$, there are three corresponding c_{i+1} -wide level- i instances, $\mathcal{I}_j(C_q, \ell_{q_1}), \mathcal{I}_j(C_q, \ell_{q_2}),$ and $\mathcal{I}_j(C_q, \ell_{q_3})$. We say that clause C_q^j is interesting iff at least one of these three instances is interesting, and we say that it is uninteresting otherwise. Let $\mathcal{C}'_0 \subseteq \mathcal{C}'$ be the set of all uninteresting clauses.

Claim 6.10 $|\mathcal{C}'_0| \leq 12n$.

Proof: Assume otherwise. From Corollary 6.7, $|\tilde{\mathcal{M}}^V| \leq (65h + 1)nc_{i+1}N'_i$, while from Corollary 6.9, it is easy to see that $|\tilde{\mathcal{M}}^C| \leq \frac{5n(6+h)N'_i c_{i+1}}{3} - |\mathcal{C}'_0|N'_i c_{i+1} \leq \frac{5n(6+h)N'_i c_{i+1}}{3} - 12nN'_i c_{i+1} = nN'_i c_{i+1}(\frac{5h}{3} - 2)$. But then:

$$|\tilde{\mathcal{M}}| \leq (65h+1)nc_{i+1}N'_i + \left(\frac{5h}{3} - 2\right) nN'_i c_{i+1} \leq \left(\frac{200h}{3} - 1\right) nN'_i c_{i+1} < (1-2\delta) \left(\frac{200h}{3} + 1\right) nN'_i c_{i+1},$$

since $h = 1000/\epsilon$ and $\delta = 8\epsilon^2/10^{12}$, a contradiction. \square

Consider some clause $C_q^j \in \mathcal{C}'$ that is interesting. Then there is an index $z \in \{1, 2, 3\}$, such that instance $\mathcal{I}_j(C_q, \ell_{q_z})$ is interesting. If there are several such indices z (if C_q^j is troublesome), then we choose one of them arbitrarily. We say that clause C_q^j *chooses* the literal ℓ_{q_z} . We say that C_q^j is a *cheating clause* iff the variable x_{q_z} corresponding to literal ℓ_{q_z} is assigned the opposite value: In other words, if $\ell_{q_z} = x_{q_z}$, then x is assigned the value FALSE, and otherwise, $\ell_{q_z} = \neg x_{q_z}$, and x_{q_z} is assigned the value TRUE. We further say that it is a *bad cheating clause* iff at least one of the indices $j, j+1$ is bad for the variable x_{q_z} , and we say that it is a *good cheating clause* otherwise. Let $\mathcal{C}'_1 \subseteq \mathcal{C}' \setminus \mathcal{C}'_0$ be the set of all the cheating clauses. In the following lemma we bound the number of the cheating clauses.

Lemma 6.11 *There are at most $24n$ bad cheating clauses, and at most $3m$ good cheating clauses.*

We provide the proof of the lemma below, after we complete the analysis of the NO-INSTANCE using it. Notice that if a clause C_q^j is an interesting non-cheating clause, then the current assignment must satisfy it. From Claim 6.10 and Lemma 6.11, there are at least $hm - 12n - 24n - 3m = hm - 123m/5 = (1 - 123\epsilon/5000)hm > (1 - \epsilon)hm$ such clauses, contradicting Observation 6.3. It now remains to complete the proof of Lemma 6.11.

Proof of Lemma 6.11. In order to bound the number of bad cheating clauses, we first bound the total number of bad indices for variables. Let β denote the total number of pairs (x, j) , where x is a variable of φ and $1 \leq j \leq 5h + 1$ is an index, such that j is a bad index for x . Recall that, assuming that x is assigned the value TRUE, this means that instance $\mathcal{I}_j^T(x)$ is uninteresting, and if x is assigned the value FALSE, then instance $\mathcal{I}_j^F(x)$ is uninteresting. For every variable x , let $\beta(x)$ be the total number of indices $1 \leq j \leq 5h + 1$ that are bad for x . Then $|\tilde{\mathcal{M}}(x)| \leq (65h + 1)N'_i c_{i+1} - \beta(x) \cdot N'_i c_{i+1}$. Therefore, overall,

$$\begin{aligned} |\tilde{\mathcal{M}}| &\leq n \cdot (65h + 1)N'_i c_{i+1} - \beta \cdot N'_i c_{i+1} + |\tilde{\mathcal{M}}^C| \\ &\leq n \cdot (65h + 1)N'_i c_{i+1} - \beta \cdot N'_i c_{i+1} + 5n(h + 6)N'_i c_{i+1}/3 \\ &= (200h/3 + 11)nN'_i c_{i+1} - \beta \cdot N'_i c_{i+1}. \end{aligned}$$

Assume now for contradiction that $\beta > 12n$. Then $|\tilde{\mathcal{M}}| < (200h/3 - 1)nN'_i c_{i+1} < (1 - 2\delta)(200h/3 + 1)nN'_i c_{i+1}$, a contradiction. It is immediate to verify that each pair (x, j) where j is a bad index for variable x may be responsible for at most two bad cheating clauses, and so the total number of bad cheating clauses is bounded by $24n$.

We now turn to bound the number of good cheating clauses. In order to do so, it is enough to prove the following claim:

Claim 6.12 *For each original clause $C_q \in \mathcal{C}$ and index $z \in \{1, 2, 3\}$, at most one copy of C_q that chooses the literal ℓ_{q_z} is a good cheating clause.*

Notice that from the above claim, for each original clause $C_q \in \mathcal{C}$, there are at most 3 copies of C_q that are good cheating clauses, and so overall there are at most $3m$ good cheating clauses in $\mathcal{C}' \setminus \mathcal{C}'_0$. It now remains to prove Claim 6.12.

Proof of Claim 6.12. Let $C_q \in \mathcal{C}$ be an original clause, and let $\ell \in \{\ell_{q_1}, \ell_{q_2}, \ell_{q_3}\}$ be one of its literals. Assume for contradiction that two distinct copies of C_q , that we denote by C_q^j and $C_q^{j'}$ are good cheating clauses and that they both select the literal ℓ of C_q . Let x be the variable corresponding to ℓ . We assume without loss of generality that x appears in C_q without negation (so $\ell = x$), but we assigned x the value FALSE; the other case is dealt with similarly. We also assume without loss of generality that $j < j'$.

Since both $C_q^j, C_q^{j'}$ are good cheating clauses, indices $j, j+1$ and $j'+1$ are good indices for x . Therefore, each of the instances $\mathcal{I}_j^F(x), \mathcal{I}_{j+1}^F(x)$ and $\mathcal{I}_{j'+1}^F(x)$ is an interesting instance, and $\tilde{\mathcal{M}}$ contains at least $25H_i$ demand pairs that belong to each of the three instances. Moreover, $\tilde{\mathcal{M}}$ contains at least $25H_i$ demand pairs that belong to each of the two instances $\mathcal{I}_j(C_q, \ell)$ and $\mathcal{I}_{j'}(C_q, \ell)$, from the choice of the literal ℓ by each corresponding copies of C_q .

We let \mathcal{M}_1 be the set of all demand pairs of $\mathcal{I}_j(C_q, \ell)$ that belong to $\tilde{\mathcal{M}}$, and we define \mathcal{M}_2 for instance $\mathcal{I}_{j'}(C_q, \ell)$ similarly. We also let $\mathcal{M}_3, \mathcal{M}_4$ and \mathcal{M}_5 be the sets of the demand pairs from instances $\mathcal{I}_j^F(x), \mathcal{I}_{j+1}^F(x)$, and $\mathcal{I}_{j'+1}^F(x)$ that belong to $\tilde{\mathcal{M}}$, respectively. From the above discussion, each of the five sets contain at least $25H_i$ demand pairs. From Lemma 6.4, we can find, for each $1 \leq y \leq 5$, a demand pair $(s_y, t_y) \in \mathcal{M}_y$, such that none of the five resulting demand pairs encircle each other.

Consider the plane into which the grid G_i is embedded with only the following curves. First, we add the boundary of the grid. We also add the bottom boundary of the box $B_F(x)$, and for every $3 \leq y \leq 5$, a curve γ_y , obtained by concatenating the image of $P(s_y, t_y)$, with the line Q_{t_y} . Notice that the three curves cannot intersect. Let f and f' be the faces of the resulting drawing, that are distinct from the outer face, and are incident to the intervals Y_j^F and $Y_{j'}^F$, respectively. Let Γ be the bottom boundary of the box $B(C_q)$. Then Γ must lie in the same face of the drawing as t_1 and t_2 , since none of the five pairs encircles the other. Assume without loss of generality that this face is $f^* \neq f$. Then the source of the pair (s_1, t_1) lies on the part Y_j^F of boundary of f that is also incident to the infinite face, while its destination lies strictly inside another face (and this other face is disjoint from $Y_{j'}^F$), which is impossible. \square \square

Appendix

A Proofs Omitted from Section 5

A.1 Proof of Claim 5.3

The proof is by induction on ℓ . For the base case, assume that $\ell = 1$, and let A, A' any pair of vertex sets (that are not necessarily disjoint), such that the vertices of each set lie on a single boundary edge of the corridor $\Upsilon = \Upsilon_1$, and $|A| = |A'| = w' \leq w - 2$. We show that there is a set \mathcal{Q} of node-disjoint paths in Υ , connecting every vertex of A to a vertex of A' , with a slightly stronger property: namely, the paths in \mathcal{Q} are internally disjoint from the boundary of Υ . Let Υ' be the graph obtained from Υ , by deleting all vertices lying on the boundary of Υ , except for the vertices of $A \cup A'$. It is enough to show that there is a set \mathcal{Q} of w' node-disjoint paths in Υ' , connecting vertices of A to vertices of A' . Assume for contradiction that such a set of paths does not exist. Then from Menger's theorem, there is a set J of $w' - 1$ vertices, such that $\Upsilon' \setminus J$ has no path connecting a vertex of A to the vertex of A' . We consider three cases.

The first case is when A and A' lie on the opposite boundary edges of Υ . Assume without loss of generality that the vertices of A lie on the top boundary edge, and the vertices of A' on the bottom boundary edge of Υ . Let \mathcal{R} and \mathcal{W} be the sets of rows and columns of G_{i+1} , respectively, that span Υ . Let $\mathcal{W}', \mathcal{W}'' \subseteq \mathcal{W}$ be the sets of columns of Υ where the vertices of A and A' lie, respectively, so $|\mathcal{W}'| = |\mathcal{W}''| = w'$. Since $|J| < w'$, there is a column $W' \in \mathcal{W}'$ and a column $W'' \in \mathcal{W}''$, with $W' \cap J, W'' \cap J = \emptyset$. There is also some row $R' \in \mathcal{R}$ of Υ' , that is not its top or bottom row, such that $R' \cap J = \emptyset$. But $W' \cup W'' \cup R'$ is a connected subgraph of $\Upsilon' \setminus J$, that contains a vertex of A and a vertex of A' , a contradiction.

The other two cases, when A and A' lie on adjacent boundary edges of Υ , and when A and A' lie on the same boundary edge of Υ are analyzed similarly.

Assume now that the claim holds for some value $\ell \geq 0$. We now prove it for $\ell + 1$. Denote $|A| = |A'| = w'$, and let U be any set of w' vertices in $\Upsilon_\ell \cap \Upsilon_{\ell+1}$, so the vertices of U lie on the boundaries of both corridors. Notice that the vertices of U must belong to a single boundary edge of Υ_ℓ , and a single boundary edge of $\Upsilon_{\ell+1}$. Using the induction hypothesis, there is a set \mathcal{P}_1 of w' node-disjoint paths in $\Upsilon_1 \cup \dots \cup \Upsilon_\ell$, connecting every vertex of A to a distinct vertex of U . From our analysis of the base case, there is a set \mathcal{P}_2 of w' node-disjoint paths in $\Upsilon_{\ell+1}$, connecting every vertex of U to a distinct vertex of A' . Moreover, the paths in \mathcal{P}_2 are internally disjoint from the boundary of $\Upsilon_{\ell+1}$. We obtain the desired set of paths by concatenating the paths in \mathcal{P}_1 and the paths in \mathcal{P}_2 .

A.2 Proof of Claim 5.4

Consider two consecutive variables x_j, x_{j+1} , for $1 \leq j < n$. Recall that the vertices $S(\hat{\mathcal{M}}(x_j))$ appear consecutively on R , and so do the vertices of $S(\hat{\mathcal{M}}(x_{j+1}))$. Let $V_j \subseteq S(\hat{\mathcal{M}}^C)$ be the set of all source vertices that correspond to clause-pairs, and lie between the vertices of $S(\hat{\mathcal{M}}(x_j))$ and the vertices of $S(\hat{\mathcal{M}}(x_{j+1}))$ on R . Notice that the vertices of V_j may only correspond to clauses in which x_j or x_{j+1} participate, so $|V_j| \leq 10hN_i c_{i+1} < N_{i+1}$. Similarly, we let $V_0, V_n \subseteq S(\hat{\mathcal{M}}^C)$ be the sets of all source vertices that correspond to clause-pairs and lie before the vertices of $S(\hat{\mathcal{M}}(x_1))$ and after the vertices of $S(\hat{\mathcal{M}}(x_n))$ on R , respectively. We still have $|V_0|, |V_n| \leq N_{i+1}$. For all

$0 \leq j \leq n$, let $M'_j = |V_j|$. For all $1 \leq j \leq n$, let $M_j = |\hat{\mathcal{M}}(x_j)| = c_{i+1}N_i(65h+1) < N_{i+1}$.

We now select a set Γ of vertices on the top row of B^V , as follows. First, for every variable x_j of φ , for every vertex $a \in A(x_j)$, we select a vertex a' on the top row of B^V , lying in the same column as a , and we let P_a be the sub-path of the column $\text{col}(a)$ between a and a' . Let Γ_j denote the resulting set of vertices that we have selected for x_j .

For $1 \leq j < n$, let \mathcal{W}_j be the set of columns of B^V that lie between the boxes $B(x_j)$ and $B(x_{j+1})$. We also let \mathcal{W}_0 be the set of all columns of B^V that lie before $B(x_1)$, and \mathcal{W}_n the set of all columns lying after $B(x_n)$. For all $0 \leq j \leq n$, we select an arbitrary set Γ'_j of M'_j distinct vertices on the top row of B^V that belong to the columns of \mathcal{W}_j . Since $|\mathcal{W}_j| \geq N_{i+1}$, while $M'_j \leq N_{i+1}$, such a set of vertices exists. For each selected vertex $v \in \Gamma'_j$, we let P_v be the column of B^V in which v lies, and we let v' be the other endpoint of the column. Let $\Gamma''_j = \{v' \mid v \in \Gamma'_j\}$. We denote $\Gamma = \left(\bigcup_{j=1}^n \Gamma_j\right) \cup \left(\bigcup_{j=0}^n \Gamma'_j\right)$, and $\Gamma'' = \bigcup_{j=0}^n \Gamma''_j$. Finally, we let Γ''' be the set of $|\hat{\mathcal{M}}^C|$ consecutive left-most vertices on the top boundary of B^C .

We construct five sets of node-disjoint paths, $\mathcal{P}_0, \dots, \mathcal{P}_4$. The final set \mathcal{P}' of paths will be obtained by combining these sets of paths.

Set \mathcal{P}_0 : Let Z' be the set of $|\hat{\mathcal{M}}|$ leftmost vertices on the opening of $B(\mathcal{I})$. Set \mathcal{P}_0 consists of $|\hat{\mathcal{M}}|$ node-disjoint paths, connecting every vertex of $S(\hat{\mathcal{M}})$ to a vertex of Z' , so that the paths in \mathcal{P}_0 are order-preserving and internally disjoint from $B(\mathcal{I})$. Since the distance between $B(\mathcal{I})$ and the boundaries of G_{i+1} is at least $H_{i+1} \geq N_{i+1}$, while $|\hat{\mathcal{M}}| = N_{i+1}$, it is easy to verify that such a set of paths exists.

Set \mathcal{P}_1 : This set of paths connects every vertex of Z' to a distinct vertex of Γ in a node-disjoint and order-preserving manner. In order to construct it, we use a snake \mathcal{Y}^1 , consisting of two corridors, Υ_1^1 and Υ_2^1 . The first corridor Υ_1^1 is simply the set of the top $N_{i+1} + 2$ rows of $B(\mathcal{I})$. In order to construct the second corridor, Υ_2^1 , we denote by \mathcal{W}^V the set of all columns of G_{i+1} that intersect the box B^V , and we denote by \hat{R} the row of G_{i+1} that contains the topmost row of B^V . Let $\hat{\mathcal{R}}$ be a consecutive set of rows of $B(\mathcal{I})$, starting from row $R_{N_{i+1}+2}$ and ending at row \hat{R} . Let Υ_2^1 be the corridor spanned by the set \mathcal{W}^V of columns and the set $\hat{\mathcal{R}}$ of rows (see Figure 13). By combining the two corridors, we obtain a snake \mathcal{Y}^1 of width at least $N_{i+1} + 2$. From Claim 5.3, there is a set \mathcal{P}_1 of node-disjoint paths, connecting every vertex of Z' to a distinct vertex of Γ inside the snake. It is immediate to verify that the set \mathcal{P}_1 of paths must be order-preserving.

Set \mathcal{P}_2 : Set \mathcal{P}_2 contains, for every vertex $v \in \Gamma$, the path P_v that we have defined above. Note that for each $1 \leq j \leq n$, paths in \mathcal{P}_2 connect the vertices of Γ_j to the vertices of $A(x_j)$, lying on the top row of $B(x_j)$, and for each $0 \leq j \leq n$, paths in \mathcal{P}_2 connect the vertices of Γ'_j to the vertices of Γ''_j , lying on the bottom row of B^V .

Set \mathcal{P}_3 : This set contains $|\hat{\mathcal{M}}^C|$ node-disjoint paths, connecting every vertex of Γ'' to a distinct vertex of Γ''' . The paths in \mathcal{P}_3 will be internally disjoint from B^V and B^C , and order-preserving. In order to construct the set \mathcal{P}_3 of paths, we construct a snake \mathcal{Y}^2 . Let \mathcal{W}^V and \mathcal{W}^C be the sets of columns of G_{i+1} that intersect B^V and B^C , respectively, and let \mathcal{W}^M be the set of columns lying

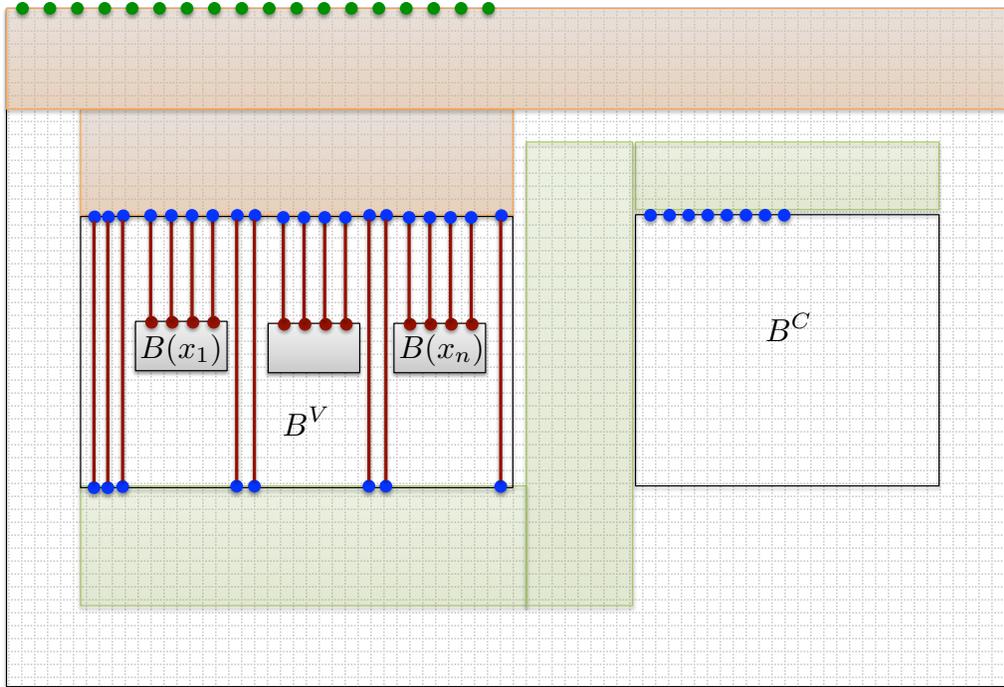


Figure 13: Routing the sets $\mathcal{P}_1, \mathcal{P}_2$ and \mathcal{P}_3 of paths inside $B(\mathcal{I})$. The paths in \mathcal{P}_2 are shown in red; the paths of \mathcal{P}_1 are routed inside the orange snake, and the paths of \mathcal{P}_3 are routed inside the green snake.

between these two sets. Let R' be the row of G_{i+1} containing the bottommost row of B^V , and let \mathcal{R}^V be the set of $N_{i+1} - 2$ consecutive rows of G_{i+1} lying below R' , including R' . Let R'' be the row of G_{i+1} containing the topmost row of B^C , and let \mathcal{R}^C be the set of $N_{i+1} - 2$ consecutive rows of G_{i+1} lying above R'' , including R'' . Finally, let \mathcal{R}^M be the set of rows lying between the top row of \mathcal{R}^C and the bottom row of \mathcal{R}^V , including these two rows. The snake \mathcal{Y}^2 consists of three corridors (see Figure 13). The first corridor, Υ_1^2 , is spanned by the set \mathcal{R}^V of rows and the set \mathcal{W}^V of columns. The second corridor, Υ_2^2 , is spanned by the set \mathcal{R}^M of rows and the set \mathcal{W}^M of columns. The third and the final corridor, Υ_3^2 , is spanned by the set \mathcal{R}^C of rows and the set \mathcal{W}^C of columns. We extend corridor Υ_1^2 by one column to the right, so that it intersects the boundary of Υ_2^2 , and similarly we extend corridor Υ_3^2 by one column to the left. It is now easy to verify that we obtain a snake of width at least $N_{i+1} - 2$, and so from Claim 5.3, there is a set \mathcal{P}_3 of node-disjoint paths, connecting every vertex of Γ'' to a distinct vertex of Γ''' . It is immediate to verify that the paths in \mathcal{P}_3 are order-preserving.

By combining the paths of $\mathcal{P}_0, \dots, \mathcal{P}_3$, we obtain a set \mathcal{P}'' of node-disjoint paths, that are internally disjoint from B^C . For every variable x_j of φ , set \mathcal{P}'' contains a set $\mathcal{P}(x_j)$ of paths connecting every vertex of $S(\hat{\mathcal{M}}(x_j))$ to a vertex of $A(x_j)$, so that the paths in $\mathcal{P}(x_j)$ are order-preserving. For every clause $C_q \in \mathcal{C}$, set \mathcal{P}'' contains a set $\mathcal{P}'(C_q)$ of node-disjoint order-preserving paths, connecting every vertex of $S(\hat{\mathcal{M}}(C_q))$ to a distinct vertex of Γ''' . We denote by $\Gamma(C_q) \subseteq \Gamma'''$ the set of vertices that serve as endpoints of the paths in $\mathcal{P}'(C_q)$. Note that for each clause $C_q \in \mathcal{C}$, the vertices of $\Gamma(C_q)$ appear consecutively in Γ''' . We define an ordering \mathcal{O}' of the clauses in \mathcal{C} as follows: clause C_q appears before clause $C_{q'}$ in this ordering iff the vertices of $\Gamma(C_q)$ appear to the left of the vertices of $\Gamma(C_{q'})$ on the top row of B^C .

Set \mathcal{P}_4 : In our final step, we construct, for each clause $C_q \in \mathcal{C}$, a set $\mathcal{P}''(C_q)$ of paths, connecting the vertices of $\Gamma(C_q)$ to the vertices of $A(C_q)$, so that the paths in $\mathcal{P}''(C_q)$ are order-preserving and the paths in $\bigcup_{C_q \in \mathcal{C}} \mathcal{P}''(C_q)$ are node-disjoint. We then let \mathcal{P}_4 be the union of these sets of paths.

Claim A.1 *For each clause $C_q \in \mathcal{C}$, there is a set $\mathcal{P}''(C_q) \subseteq B^C$ of paths, connecting every vertex of $\Gamma(C_q)$ to a distinct vertex of $A(C_q)$, such that the paths in set $\mathcal{P}_4 = \bigcup_{C_q \in \mathcal{C}} \mathcal{P}''(C_q)$ are mutually node-disjoint and internally disjoint from the boxes $\{B(C_1), \dots, B(C_m)\}$. Moreover, the paths in \mathcal{P}_4 are internally disjoint from the top boundary of B^C , and for each $1 \leq q \leq m$, the paths in $\mathcal{P}(C_q)$ are order-preserving.*

We obtain the final set \mathcal{P}' of paths by combining the paths of $\mathcal{P}_0, \dots, \mathcal{P}_4$. In order to complete the proof of Claim 5.4, it now remains to prove Claim A.1.

Proof: Consider some clause $C_q \in \mathcal{C}$, and let $\mathcal{W}(C_q)$ be the set of columns of G_{i+1} that intersect $B(C_q)$. We let $\mathcal{W}^L(C_q)$ be the set of N_{i+1} columns lying immediately to the left of $\mathcal{W}(C_q)$, and similarly we let $\mathcal{W}^R(C_q)$ be the set of N_{i+1} columns lying immediately to the right of C_q . Let \mathcal{R} be the set of all rows of G_{i+1} that intersect the box $B(C_q)$ (this set is the same for all clauses C_q). We let $B'(C_q)$ be the sub-grid of G_{i+1} spanned by the set $\mathcal{W}^L(C_q) \cup \mathcal{W}(C_q) \cup \mathcal{W}^R(C_q)$ of columns and the set \mathcal{R} of rows.

The idea is to define a snake-like routing, where the paths visit the boxes $B'(C_1), \dots, B'(C_m)$ in turn. For each clause C_q , only the paths corresponding to the clauses C_q, C_{q+1}, \dots, C_m will visit the box $B'(C_q)$. The paths corresponding to clause C_q will terminate on the top boundary of $B(C_q)$. For each $q+1 \leq q' \leq m$, if $C_{q'}$ appears before C_q in the ordering \mathcal{O}' , then the paths corresponding to $C_{q'}$ will traverse the box $B'(C_q)$ to the left of $B(C_q)$; otherwise, they will traverse the box $B'(C_q)$ to the right of $B(C_q)$. In order to implement this, we select a subset of vertices on the top and the bottom boundaries of $B'(C_q)$, through which the paths will enter and leave the box (see Figure 14).

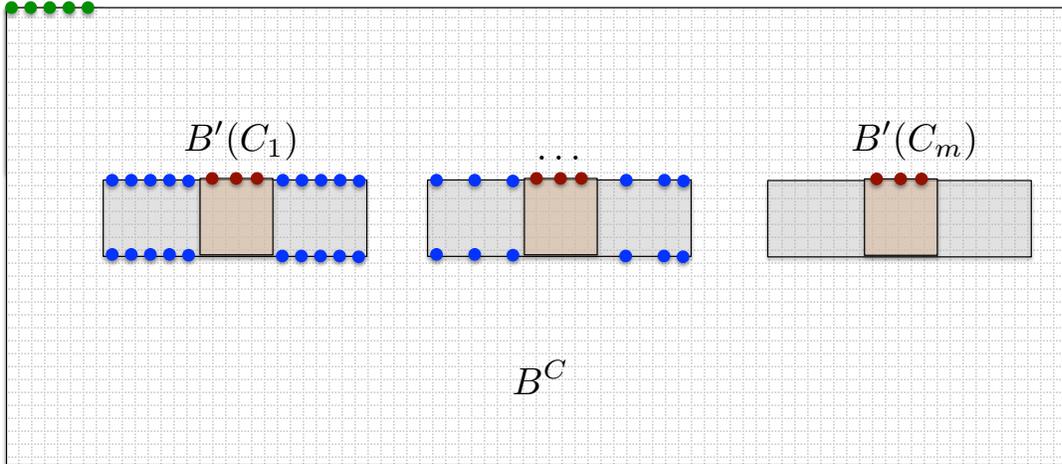


Figure 14: Selecting the sets $\Gamma(C_q)$ and $\Gamma'(C_q)$ of vertices for the clauses C_q . The vertices of Γ''' are shown in green.

Fix some clause C_q , for $1 \leq q \leq m$. Let n_q^L be the number of clauses $C_{q'}$, with $q < q' \leq m$, such that $C_{q'}$ appears before C_q in the ordering \mathcal{O}' , and let n_q^R be the number of such clauses with $C_{q'}$ appearing after C_q in \mathcal{O}' . Recall that for each clause $C_{q'}$, $|\hat{\mathcal{M}}(C_{q'})| = hc_{i+1}N_i$, and

$mhc_{i+1}N_i < N_{i+1} - 2$. We select an arbitrary set $\Gamma^L(C_q)$ of $n_q^L \cdot hc_{i+1}N_i$ vertices on the top boundary of $B'(C_q)$, strictly to the left of $B(C_q)$, and similarly, we select an arbitrary set $\Gamma^R(C_q)$ of $n_q^R \cdot hc_{i+1}N_i$ vertices on the top boundary of $B'(C_q)$, strictly to the right of $B(C_q)$. For each vertex $v \in \Gamma^L(C_q) \cup \Gamma^R(C_q)$, we let P_v be the column of $B'(C_q)$ that contains v , and we let v' be its other endpoint. For each vertex $v \in A(C_q)$, we let P_v be a path consisting of a single vertex v . We define $\Gamma(C_q) = \Gamma^L(C_q) \cup A(C_q) \cup \Gamma^R(C_q)$, a set of vertices on the top boundary of $B'(C_q)$. We define $\Gamma'(C_q) = \{v' \mid v \in \Gamma^L(C_q) \cup \Gamma^R(C_q)\}$, a set of vertices on the bottom boundary of $B'(C_q)$. Finally, we let $\tilde{\mathcal{P}}'_q = \{P_v \mid v \in \Gamma(C_q)\}$. It is easy to verify that for all $1 \leq q < m$, $|\Gamma'(C_q)| = |\Gamma(C_{q+1})|$.

In order to compute the routing, we use the following simple observation.

Observation A.2 *For all $1 \leq q < m$, there is a set $\tilde{\mathcal{P}}_q \subseteq B^C$ of paths, such that:*

- *Paths in $\tilde{\mathcal{P}}_1$ connect every vertex in Γ''' to a distinct vertex of $\Gamma(C_1)$ and are order-preserving;*
- *For $1 < q < m$, paths in $\tilde{\mathcal{P}}_q$ connect every vertex in $\Gamma'(C_q)$ to a distinct vertex of $\Gamma(C_{q+1})$ and are order-preserving;*
- *The paths in $\bigcup_q \tilde{\mathcal{P}}_q$ are mutually node-disjoint, and they are internally disjoint from all boxes in $\{B'(C_1), \dots, B'(C_m)\}$.*

Notice that combining the paths in sets $\bigcup_{q=1}^{m-1} \tilde{\mathcal{P}}_q$ and $\bigcup_{q=1}^m \tilde{\mathcal{P}}'_q$ finishes the proof of Claim A.1 and Claim 5.4. We now prove the observation.

Proof: We construct, for each $1 \leq q \leq m$, a snake $\mathcal{Y}(C_q)$, and route the set $\tilde{\mathcal{P}}_q$ of paths inside it (see Figure 15 for an illustration). Since $|\hat{\mathcal{M}}^C| < N_{i+1} - 2$, it is enough to ensure that each snake has width at least N_{i+1} .

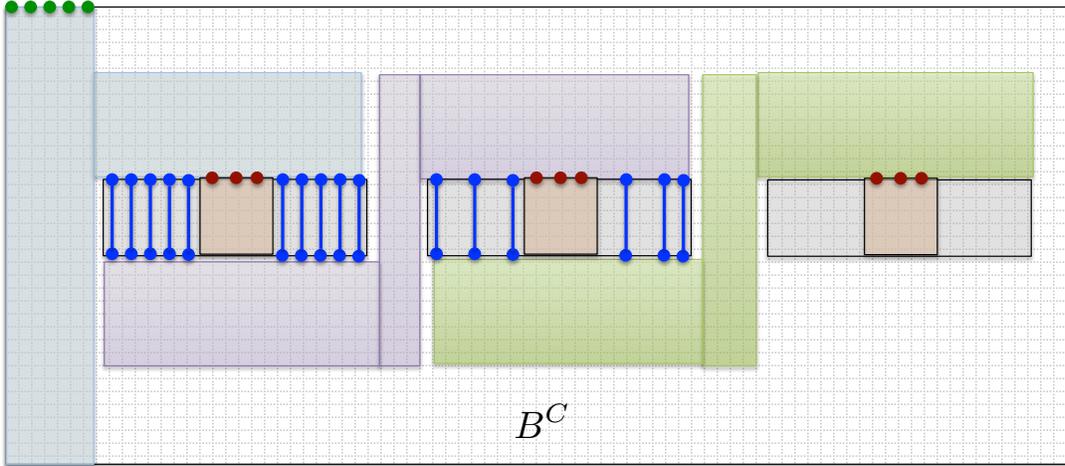


Figure 15: The snakes used to route the sets $\tilde{\mathcal{P}}_q$ of paths.

Recall that \mathcal{R} is the set of all rows of G_{i+1} that intersect the boxes $B(C_q)$. Let R' and R'' be the topmost and the bottommost rows of \mathcal{R} , respectively. Let \mathcal{R}' be the set of N_{i+1} consecutive rows lying above R' including R' , and let \mathcal{R}'' be the set of N_{i+1} consecutive rows lying below R'' , including R'' .

The first snake, $\mathcal{Y}(C_1)$ consists of two corridors. The first corridor, $\Upsilon_1(C_1)$, is spanned by the first N_{i+1} columns of B^C . For the second corridor, $\Upsilon_2(C_1)$, we let W be the last column of $\Upsilon_1(C_1)$, W' the last column of $\mathcal{W}^R(C_1)$, and \mathcal{W}' the set of all columns lying between W and W' , including these two columns. We then let $\Upsilon_2(C_1)$ be the corridor spanned by the columns of \mathcal{W}' and the rows of \mathcal{R}' .

For each $1 < q \leq m$, we construct a snake $\mathcal{Y}(C_q)$, that consists of three corridors. The first corridor, $\Upsilon_1(C_q)$ is spanned by the rows of \mathcal{R}'' and the columns of $\mathcal{W}^L(C_{q-1}) \cup \mathcal{W}(C_{q-1}) \cup \mathcal{W}^R(C_{q-1})$. The third corridor, $\Upsilon_3(C_q)$ is spanned by the rows of \mathcal{R}' and the columns of $\mathcal{W}^L(C_q) \cup \mathcal{W}(C_q) \cup \mathcal{W}^R(C_q)$. For the second corridor, $\Upsilon_2(C_q)$, we let \mathcal{R}^* be the set of consecutive rows starting from the top row of \mathcal{R}' and terminating at the bottom row of \mathcal{R}'' , including these two rows, and we let $\mathcal{W}'(C_q)$ be the set of all columns lying between $B'(C_{q-1})$ and $B'(C_q)$. We then let $\Upsilon_2(C_q)$ be the corridor spanned by the rows in \mathcal{R}^* and the columns in $\mathcal{W}'(C_q)$. In order to obtain a valid snake, we extend $\Upsilon_1(C_q)$ by one column to the right and $\Upsilon_3(C_q)$ by one column to the left (see Figure 15).

It is immediate to verify that for each $1 \leq q \leq m$, we obtain a valid snake $\mathcal{Y}(C_q)$ of width at least N_{i+1} ; the resulting snakes are mutually disjoint from each other, and each snake $\mathcal{Y}(C_q)$ intersects the boxes in $\{B(C_1), \dots, B(C_m)\}$ only on their boundaries. For each $1 \leq q < m$, we can now find the desired set $\mathcal{P}(C_q)$ of paths inside the snake $\mathcal{Y}(C_q)$, using Claim 5.3. It is immediate to verify that the paths are order-preserving. □ □

B Reducing the Maximum Vertex Degree and Hardness of EDP in Sub-Cubic Planar Graphs

In this section we prove Theorem 1.2, and show that Theorem 1.1 holds for sub-cubic planar graphs. We start with proving Theorem 1.2. Let $G = G^{\ell, h}$ be a grid of length ℓ and height h , where $\ell > 0$ is an even integer, and $h > 0$. For every column W_j of the grid, let e_1^j, \dots, e_{h-1}^j be the edges of W_j indexed in their top-to-bottom order. Let $E^*(G) \subseteq E(G)$ contain all edges e_z^j , where $z \neq j \pmod 2$, and let \hat{G} be the graph obtained from $G \setminus E^*(G)$, by deleting all degree-1 vertices from it. The resulting graph is called a *wall of length $\ell/2$ and height h* (see Figure 16). Consider the subgraph of \hat{G} induced by all horizontal edges of the grid G that belong to \hat{G} . This graph is a collection of h node-disjoint paths, that we refer to as the *rows* of \hat{G} , and denote them by R_1, \dots, R_h in this top-to-bottom order; notice that R_j is a sub-path of the j th row of G for all j . Graph \hat{G} contains a unique collection \mathcal{W} of $\ell/2$ node-disjoint paths that connect vertices of R_1 to vertices of R_h and are internally disjoint from R_1 and R_h . We refer to the paths in \mathcal{W} as the *columns* of \hat{G} , and denote them by $W_1, \dots, W_{\ell/2}$ in this left-to-right order. Paths $W_1, W_{\ell/2}, R_1$ and R_h are called the left, right, top and bottom boundary edges of \hat{G} , respectively, and the union of these paths is the boundary of \hat{G} .

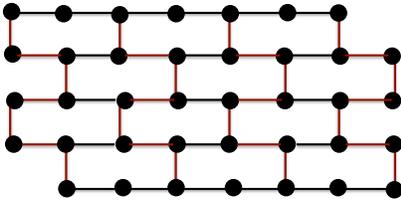


Figure 16: A wall of height 5 and length 4; the columns of the wall are shown in red.

Given a wall \hat{G} , a consecutive subset \mathcal{R}' of its rows, and a consecutive subset \mathcal{W}' of its columns, the sub-wall of \hat{G} spanned by the rows of \mathcal{R}' and the columns of \mathcal{W}' is the subgraph of \hat{G} induced by the set $\{v \mid \exists R \in \mathcal{R}', W \in \mathcal{W}' : v \in R \cap W\}$ of vertices. The first and the last columns of \mathcal{W}' serve as the left and the right boundary edges of the sub-wall, and the top and the bottom rows of \mathcal{R}' serve as its top and bottom boundary edges.

We perform a reduction from the 3SAT(5) problem. Assume that we are given an instance φ of 3SAT(5) on n variables and $m = 5n/3$ clauses. As before, our construction has $\Theta(\log n)$ levels. For every level $i \geq 0$, we define a family of instances of EDP. In order to construct a level- i instance \mathcal{I} , we define the parameters H_i, L_i and L'_i exactly as before, a path $Z(\mathcal{I})$ and a box $B(\mathcal{I})$ (which is a subgraph of the grid G_i of length L'_i and height H_i), a collection \mathcal{M} of demand pairs, and the mappings of the vertices of $S(\mathcal{M})$ to the vertices of $Z(\mathcal{I})$, and of the vertices of $T(\mathcal{M})$ to distinct vertices of the middle row of $B(\mathcal{I})$ exactly as before. In order to instantiate this instance, we select an arbitrary grid G_i of length $\ell \geq 2L_i + 2L'_i + 4H_i$, where ℓ is an even integer, and height $h \geq 3H_i$, map the vertices of $Z(\mathcal{I})$ to the vertices of the first row of G_i , and map the vertices of $B(\mathcal{I})$ to the vertices of a sub-grid G''_i of G_i exactly as before, obtaining an instantiation (G, \mathcal{M}) of the level- i instance \mathcal{I} of NDP. Our final step is to delete from G every edge of $E^*(G_i) \cap E(G)$, and then to delete all vertices that have degree 1 in the resulting graph. We also delete every other edge on the top row of G_i , and all horizontal edges that are incident to the vertices of $T(\mathcal{M})$, to ensure that the degree of every terminal is at most 2. The final graph, denoted by \hat{G} , is a subgraph of a wall of length $\ell/2$ and height h . We denote by $\hat{B}(\mathcal{I})$ the intersection of the image of $B(\mathcal{I})$ in G_i and the graph \hat{G} . This concludes the definition of the reduction. Since the resulting graph \hat{G} is a subgraph of G , the following observation is immediate.

Observation B.1 *If φ is a NO-INSTANCE, then for every level i , for every instantiation (G, \mathcal{M}) of the level- i instance \mathcal{I} of NDP, the corresponding instance (\hat{G}, \mathcal{M}) of EDP has optimal solution value at most N'_i .*

Proof: Assume otherwise. Consider an instantiation (G, \mathcal{M}) of a level- i instance \mathcal{I} of NDP, for some $i \geq 0$, and let (\hat{G}, \mathcal{M}) be the corresponding instance of EDP. Let \mathcal{P} be a set of edge-disjoint paths in \hat{G} , routing more than N'_i demand pairs. Since graph \hat{G} is sub-cubic, and the degree of every terminal in \hat{G} is at most 2, it is immediate to verify that the paths in \mathcal{P} are also vertex-disjoint. Since $\hat{G} \subseteq G$, there is a set \mathcal{P} of node-disjoint paths routing more than N'_i demand pairs in the instance (G, \mathcal{M}) of NDP, contradicting Theorem 6.1. \square

It is now enough to show that, if φ is a YES-INSTANCE, then for every level i , for every instantiation (G, \mathcal{M}) of the level- i instance \mathcal{I} of NDP, the corresponding instance (\hat{G}, \mathcal{M}) of EDP has a solution of value at least $N_i/2$. Before we do so, we need several definitions.

Suppose we are given some set \mathcal{P} of node-disjoint paths in some wall \hat{H} , and assume that every path in \mathcal{P} connects some vertex on a row R of \hat{H} to a vertex on a row R' of \hat{H} , where $R \neq R'$. As before, we say that the paths in \mathcal{P} are *order-preserving* iff the left-to-right ordering of their endpoints on R is the same as the left-to-right ordering of their endpoints on R' .

A subset U of the vertices lying on a row R of a wall \hat{H} is called *well-spread* iff U does not contain a pair of vertices connected by an edge in \hat{H} . Notice that if U is well-spread, then no two vertices of U may lie on the same column of \hat{H} .

We now define an analogue of box-respecting paths. Consider some level- i instance \mathcal{I} of NDP, for $i \geq 0$, and an instantiation (G, \mathcal{M}) of this instance. Let (\hat{G}, \mathcal{M}) be the corresponding instance of EDP, and let \mathcal{P} be a set of node-disjoint paths routing some subset of the demand pairs in \hat{G} . Let A

be the top boundary of $\hat{B}(\mathcal{I})$ (that is, A is the subgraph of G_i'' , that contains all the edges and the vertices of its top row that belong to $\hat{B}(\mathcal{I})$). We say that set \mathcal{P} of paths is *canonical with respect to the box $\hat{B}(\mathcal{I})$* iff for every path $P \in \mathcal{P}$, $P \cap A$ is a single edge, and the following holds. Denote $\mathcal{P} = (P_1, \dots, P_r)$, and denote, for every path P_i , its source vertex by s_i , and the unique edge of $P_i \cap A$ by e_i , such that s_1, \dots, s_r appear on the top row of \hat{G} in this left-to-right order. Then the edges of e_1, \dots, e_r must appear in this left-to-right order on A , and for each $1 \leq j \leq r$, e_j is the $(2j)$ th edge of A from the left.

Assume that φ is a YES-INSTANCE. Recall that for each $i \geq 0$, for every level- i instance \mathcal{I} of NDP, we have defined a collection $\mathcal{M}^*(\mathcal{I})$ of demand pairs, such that for every instantiation of \mathcal{I} , there is a set \mathcal{P} of node-disjoint paths, that respect the box $B(\mathcal{I})$, and route the set $\mathcal{M}^*(\mathcal{I})$ of demand pairs. Denote the pairs in $\mathcal{M}^*(\mathcal{I})$ by $(s_1, t_1), \dots, (s_p, t_p)$, and assume that the vertices s_1, \dots, s_p appear in this left-to-right order on $Z(\mathcal{I})$. We partition $\mathcal{M}^*(\mathcal{I})$ into two subsets: set $\mathcal{M}_1^*(\mathcal{I})$ contains all demand pairs (s_j, t_j) where j is odd, and set $\mathcal{M}_2^*(\mathcal{I})$ contains all remaining demand pairs. Notice that each of the sets $S(\mathcal{M}_1^*(\mathcal{I})), S(\mathcal{M}_2^*(\mathcal{I}))$ is well-spread, for any instantiation of \mathcal{I} and its corresponding instance of EDP. It is now enough to prove the following lemma.

Lemma B.2 *For all $i \geq 0$, for every level- i instance \mathcal{I} of NDP, for every instance (\hat{G}, \mathcal{M}) of EDP corresponding to an instantiation (G, \mathcal{M}) of \mathcal{I} , there is a set of node-disjoint paths in \hat{G} that route all demand pairs in $\mathcal{M}_1^*(\mathcal{I})$, such that the paths are canonical with respect to $\hat{B}(\mathcal{I})$, and the same holds for $\mathcal{M}_2^*(\mathcal{I})$.*

Proof: The proof is by induction on i . The lemma is clearly true for $i = 0$. We now assume that it holds for some $i \geq 0$ and prove it for a level- $(i + 1)$ instance \mathcal{I} of NDP. Let (G, \mathcal{M}) be any instantiation of \mathcal{I} , and let (\hat{G}, \mathcal{M}) be the corresponding instance of EDP. We show that the set $\mathcal{M}_1^*(\mathcal{I})$ of demand pairs can be routed in \hat{G} via node-disjoint paths, that are canonical with respect to $\hat{B}(\mathcal{I})$; the proof for set $\mathcal{M}_2^*(\mathcal{I})$ is identical.

Consider some level- i instance \mathcal{I}' of NDP that was used in the construction of the level- $(i + 1)$ instance \mathcal{I} . The current instantiation of \mathcal{I} also defines an instantiation of \mathcal{I}' . From the induction hypothesis, for all $z \in \{1, 2\}$, there is a set $\mathcal{P}_z(\mathcal{I}')$ of node-disjoint paths in \hat{G} , routing the demand pairs in $\mathcal{M}_z^*(\mathcal{I}')$, that respect the box $\hat{B}(\mathcal{I}')$. Moreover, the subset of all demand pairs in $\mathcal{M}_1^*(\mathcal{I})$ whose destinations lie in $\hat{B}(\mathcal{I}')$ is either equal to $\mathcal{M}_1^*(\mathcal{I}')$, or to $\mathcal{M}_2^*(\mathcal{I}')$.

The routing that we define for the set $\mathcal{M}_1^*(\mathcal{I})$ of demand pairs in graph \hat{G} is very similar to the one used in the proof of Lemma 5.2. However, we need to define corridors and snakes slightly differently.

Recall that the graph \hat{G} is a subgraph of a wall \hat{G}_i , that is obtained from G_i , by deleting the set $E^*(G_i)$ of its edges, the vertices of $U(B(\mathcal{I}))$, and all vertices whose degree in the resulting graph becomes 1. Given a sub-wall Υ of \hat{G}_i , spanned by a subset \mathcal{R} of its rows and a subset \mathcal{W} of its columns, we say that Υ is a *corridor* iff $\Upsilon \subseteq \hat{G}$. We say that two corridors Υ, Υ' are *internally disjoint* iff every vertex in $\Upsilon \cap \Upsilon'$ belongs to a single boundary edge of each corridor, and this boundary edge must be either the top or the bottom edge. We say that Υ and Υ' are *neighbors* iff they are internally disjoint and $\Upsilon \cap \Upsilon' \neq \emptyset$. As before, a snake of length ℓ is a sequence $\Upsilon_1, \dots, \Upsilon_\ell$ of corridors that are internally disjoint, such that for all $1 \leq j, j' \leq \ell$, corridors Υ_j and $\Upsilon_{j'}$ are neighbors iff $|j - j'| = 1$. We say that the width of the snake is w iff each of its corridors is spanned by at least w rows and at least w columns of \hat{G}_i , and for all $1 \leq j < \ell$, $\Upsilon_j \cap \Upsilon_{j+1}$ contains at least $2w$ vertices. Following is an analogue of Claim 5.3 for wall graphs; its proof is almost identical and is omitted here.

Claim B.3 *Let $\mathcal{Y} = (\Upsilon_1, \dots, \Upsilon_\ell)$ be a snake of width w , and let A, A' be two sets of vertices with $|A| = |A'| \leq w - 2$, such that the vertices of A either all lie on the top boundary edge of Υ_1 , or they all lie on the bottom boundary edge of Υ_1 , and the vertices of A' either all lie on the top boundary edge of Υ_ℓ , or they all lie on the bottom boundary edge of Υ_ℓ . Assume further that both A and A' are well-spread. Then there is a set \mathcal{Q} of node-disjoint paths contained in $\bigcup_{\ell'=1}^{\ell} \Upsilon_{\ell'}$, that connect every vertex of A to a distinct vertex of A' .*

We fix some assignment \mathcal{A} to the variables of φ that satisfies all clauses. For every clause $C_q \in \mathcal{C}$, we select a set $A(C_q)$ of vertices on the top boundary of $B(C_q)$ exactly as before, except that now we discard every other vertex from this set, in order to guarantee that these vertices are well-spread. Similarly, for every variable x of φ , we select the set $A(x)$ of vertices on the top boundary of $B(x)$ exactly as before, and then discard every other vertex from this set.

We select the sets Γ and Γ'' of vertices on the top and the bottom boundaries of B^V exactly as before, but now we are guaranteed that the vertices in each set are well-spread. The vertices of Z' and Γ''' are selected like before, except that now discard every other vertex in each set to ensure that the resulting sets are well-spread.

The remainder of the routing in Lemma 5.2 relies on the constructions of various snakes. For every snake $\mathcal{Y} = (\Upsilon_1, \dots, \Upsilon_r)$ that the construction uses, it is easy to modify the corresponding corridors, such that for all $1 \leq r' < r$, $\Upsilon_{r'} \cap \Upsilon_{r'+1}$ is contained in either the top or the bottom boundary edge of each corridor. This allows us to use Claim B.3 instead of the original Claim 5.3 in order to route inside the snakes. Even though each resulting snake in the wall graph will only be able to route half the number of paths it routed before, the amount of horizontal space that we left between the various boxes in the construction of the level- $(i+1)$ instance is sufficient to ensure that all demand pairs are routed. As observed above, for each level- i instance \mathcal{I}' used in the construction of \mathcal{I} , the subset of all demand pairs in $\mathcal{M}_1^*(\mathcal{I}')$, whose destinations lie in $\hat{B}(\mathcal{I}')$, is either equal to $\mathcal{M}_1^*(\mathcal{I}')$ or to $\mathcal{M}_2^*(\mathcal{I}')$, and from the induction hypothesis, each one of these sets of demand pairs has a set of node-disjoint paths routing it in \hat{G} , that is canonical with respect to $\hat{B}(\mathcal{I}')$. We use this fact to complete the routing inside the boxes $\hat{B}(\mathcal{I}')$ of the corresponding level- i instances \mathcal{I}' . \square

Hardness of NDP on Sub-Cubic Planar Graphs Consider the instances of EDP constructed above. Each such instance is defined on a sub-cubic planar graph, where the degree of every terminal is at most 2. It is easy to see that, if we are given a graph G with the above properties, and any set \mathcal{P} of paths whose endpoints are distinct terminals, the paths in \mathcal{P} are mutually edge-disjoint iff they are mutually node-disjoint. Therefore, the number of the demand pairs that can be routed in the YES-INSTANCE and the NO-INSTANCE via node-disjoint paths remains the same as for edge-disjoint paths. This completes the proof of Theorem 1.1.

C Maximum Vertex Degree in NDP Instances

Given an instance (G, \mathcal{M}) of the NDP problem, we denote by $\mathcal{T}(\mathcal{M})$ the set of all vertices that participate in the demand pairs in \mathcal{M} , and we refer to them as *terminals*. We start by defining equivalence between NDP instances.

Definition C.1 *We say that two instances $(G, \mathcal{M}), (G', \mathcal{M}')$ of the NDP problem are equivalent iff there is a bijection $f : \mathcal{T}(\mathcal{M}) \rightarrow \mathcal{T}(\mathcal{M}')$, such that, for every subset $\tilde{\mathcal{M}} \subseteq \mathcal{M}$ of demand pairs,*

there is a set \mathcal{P} of node-disjoint paths in graph G routing all pairs in $\tilde{\mathcal{M}}$ iff there is a set \mathcal{P}' of node-disjoint paths in graph G' routing all demand pairs in set $\tilde{\mathcal{M}}' = \{(f(s), f(t)) \mid (s, t) \in \tilde{\mathcal{M}}\}$.

The goal of this section is to prove the following theorem.

Theorem C.1 *For any integer $d > 3$, there is an instance (G, \mathcal{M}) of the NDP problem, where G is a planar graph, such that for every instance (G', \mathcal{M}') of NDP that is equivalent to (G, \mathcal{M}) , there is a vertex in G' of degree at least d .*

In a sense, the above theorem shows that the class of all planar graphs is strictly more general than the class of all planar graphs with maximum vertex degree at most 3 in the context of the NDP problem, and so our hardness result in Theorem 1.1 holds even for a restricted family of planar graphs.

Proof: Instance (G, \mathcal{M}) of NDP is defined as follows. Graph G is simply a star graph with $d + 2$ vertices (see Figure 17(a)), so $V(G) = \{v_1, \dots, v_{d+2}\}$, and $E(G) = \{e_1, \dots, e_{d+1}\}$, where for $1 \leq i \leq d + 1$, $e_i = (v_i, v_{d+2})$. We let $\mathcal{M} = \{(v_i, v_j) \mid 1 \leq i < j \leq d + 2\}$. It is easy to verify that a subset $\mathcal{M}' \subseteq \mathcal{M}$ of demand pairs can be routed via node-disjoint paths in G if and only if $|\mathcal{M}'| = 1$. Note that G is planar.

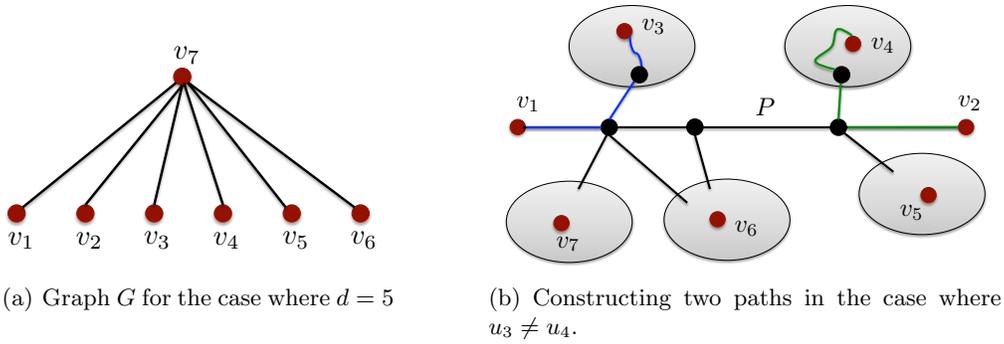


Figure 17: Illustration to the proof of Theorem C.1

Assume for contradiction that there is an instance (G', \mathcal{M}') of NDP that is equivalent to (G, \mathcal{M}) , such that the maximum vertex degree in G' is less than d . For convenience, for $1 \leq i \leq d + 2$, we denote the vertex $f(v_i)$ of G' by v_i . Let P be the shortest path in G' connecting a pair of terminals. We assume without loss of generality that P connects v_1 to v_2 ; so P does not contain any terminals as inner vertices. Consider the graph $H = G' \setminus P$. We claim that every vertex of $\{v_3, \dots, v_{d+2}\}$ must belong to a distinct connected component of H . Indeed, if, for example, v_3 and v_4 belong to the same connected component, then we can simultaneously route (v_3, v_4) and (v_1, v_2) in G' . Since we cannot do the same in G , the two instances are not equivalent. For $3 \leq i \leq d + 2$, let C_i be the connected component of H containing v_i , and let e_i be any edge connecting a vertex of C_i to a vertex of P (such an edge must exist as there must be a path connecting v_i to v_1 in G). Let u_i denote the endpoint of e_i that lies on P .

Assume first that for some $3 \leq i \neq j \leq d + 2$, $u_i \neq u_j$. Assume without loss of generality that u_i lies closer to v_1 on P than u_j . Then we can simultaneously route two pairs (v_i, v_1) and (v_j, v_2) via node-disjoint paths, as follows (see Figure 17(b)). The first path uses the segment of P from v_1 to u_i , the edge e_i , and some path connecting an endpoint of e_i to v_i inside C_i . The second path

similarly uses the segment of P from v_2 to u_j , the edge e_j , and some path connecting an endpoint of e_j to v_j inside C_j . Clearly, these two paths are disjoint. But we cannot route the pairs (v_i, v_1) and (v_j, v_2) simultaneously in G , and so the two instances are not equivalent.

We conclude that $u_3 = u_4 = \dots = u_{d+2}$ must hold. But then G' must contain a vertex of degree at least d , a contradiction. \square

D Maximum Vertex Degree in Planar EDP Instances

In this section we show that there is an instance (G, \mathcal{M}) of the EDP problem, where G is a planar graph, such that for every instance (G', \mathcal{M}') of EDP that is equivalent to (G, \mathcal{M}) , where G' is a planar graph, at least one vertex of G' must have degree at least 4. As before, given an instance (G, \mathcal{M}) of the EDP problem, the set of all vertices participating in the demand pairs is denoted by $\mathcal{T}(\mathcal{M})$, and the vertices in this set are called terminals. We say that a subset $\tilde{\mathcal{M}} \subseteq \mathcal{M}$ of demand pairs is *routable* in G iff there is a set \mathcal{P} of edge-disjoint paths in G , routing every demand pair in $\tilde{\mathcal{M}}$.

The equivalence between the EDP instances is defined similarly to the equivalence between the NDP instances, except that now we only consider restricted subsets of demand pairs: a subset $\tilde{\mathcal{M}} \subseteq \mathcal{M}$ of demand is called *restricted* iff every terminal participates in at most one demand pair of $\tilde{\mathcal{M}}$. We note that in the NDP problem, if a set $\tilde{\mathcal{M}}$ of demand pairs is routable via node-disjoint paths in a graph G , then $\tilde{\mathcal{M}}$ must be a restricted set; this is not necessarily true for EDP instances. We are now ready to define the equivalence between the instances of EDP.

Definition D.1 *Two instances (G, \mathcal{M}) and (G', \mathcal{M}') of the EDP problem are equivalent iff there is a bijection $f : \mathcal{T}(\mathcal{M}) \rightarrow \mathcal{T}(\mathcal{M}')$, such that for every restricted subset $\tilde{\mathcal{M}} \subseteq \mathcal{M}$ of demand pairs, $\tilde{\mathcal{M}}$ is routable in G iff the set $\tilde{\mathcal{M}}' = \{(f(s), f(t)) \mid (s, t) \in \tilde{\mathcal{M}}\}$ of demand pairs is routable in G' .*

The goal of this section is to prove the following theorem.

Theorem D.1 *There is an instance (G, \mathcal{M}) of the EDP problem, where G is a planar graph, such that for every instance (G', \mathcal{M}') of EDP that is equivalent to (G, \mathcal{M}) , where G' is planar, some vertex of G' has degree at least 4.*

We note that if we did not require that the sets $\tilde{\mathcal{M}}$ are restricted in the definition of the equivalence between EDP instances, then the theorem would be much easier to prove — in fact one could prove that for every d , there is an instance (G, \mathcal{M}) of EDP where G is planar, such that for every instance (G', \mathcal{M}') of EDP equivalent to (G, \mathcal{M}) , the maximum vertex degree in G' must be at least d . This can be done by letting the demand pairs in \mathcal{M} correspond to the edges of the star graph with d leaves.

Proof: We construct an instance (G, \mathcal{M}) of EDP with a set \mathcal{T} of terminals that consists of 25 vertices, partitioned into 5 subsets $\mathcal{T}_1, \dots, \mathcal{T}_5$ of 5 vertices each. For each $1 \leq i \leq 5$, we let $\mathcal{T}_i = \{t_j^i \mid 1 \leq j \leq 5\}$, and we let $\mathcal{T} = \bigcup_{i=1}^5 \mathcal{T}_i$. In order to construct the graph G , we start with the set \mathcal{T} of terminals. For each subset \mathcal{T}_i , for $1 \leq i \leq 5$, we add a vertex a_i , that connects to every terminal in \mathcal{T}_i with an edge. For all $1 \leq i < i' \leq 5$, we also add the edge $(a_i, a_{i'})$ to the graph. To summarize, so far, $V(G) = \mathcal{T} \cup \{a_1, \dots, a_5\}$, and $E(G) = \{(t, a_i) \mid 1 \leq i \leq 5, t \in \mathcal{T}_i\} \cup \{(a_i, a_j) \mid 1 \leq i < j \leq 5\}$. We fix a drawing of G in the plane, such that there is exactly one crossing

of the edges in $\{(a_i, a_j) \mid 1 \leq i < j \leq 5\}$ and no other crossings. We then replace this crossing with a vertex b , so that G becomes a planar graph (see Figure 18). Let \mathcal{M} contain all pairs (t, t') of terminals, where $t \neq t'$. This finishes the definition of the instance (G, \mathcal{M}) .

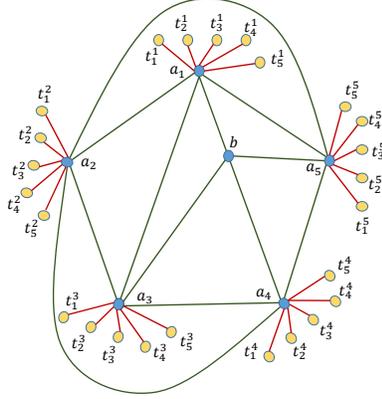


Figure 18: Graph G .

The following two observations are immediate.

Observation D.2 *Let $\hat{\mathcal{M}} \subseteq \mathcal{M}$ be a restricted set of demand pairs, such that for some $1 \leq i \leq 5$, each demand pair in $\hat{\mathcal{M}}$ contains a single distinct terminal of \mathcal{T}_i (and another terminal in $\mathcal{T} \setminus \mathcal{T}_i$). Then $\hat{\mathcal{M}}$ is routable in G iff $|\hat{\mathcal{M}}| \leq 4$.*

Observation D.3 *Let $\hat{\mathcal{M}} \subseteq \mathcal{M}$ be a set of demand pairs, defined as follows. Consider the graph $H = K_5$ (a clique on 5 vertices), and denote its vertices by v_1, \dots, v_5 . Set $\hat{\mathcal{M}}$ of demand pairs contains, for every edge $e = (v_i, v_j) \in E(H)$, a pair (s_e, t_e) , such that $s_e \in \mathcal{T}_i$ and $t_e \in \mathcal{T}_j$. The vertices s_e, t_e for all $e \in E(H)$ are chosen in way that ensures that $\hat{\mathcal{M}}$ is a restricted set of demand pairs. Then $\hat{\mathcal{M}}$ is routable in G .*

Assume now for contradiction that there is an instance (G', \mathcal{M}') of EDP that is equivalent to (G, \mathcal{M}) , such that G' is planar, and the maximum vertex degree in G' is at most 3.

For convenience, for every terminal $t \in \mathcal{T}$, we denote the corresponding terminal $f(t)$ of G' by t . We also define the partition $\{\mathcal{T}_i\}_{i=1}^5$ of the new terminals as before. Fix some index $1 \leq i \leq 5$, and let (A_i, B_i) be a partition of $V(G')$ with $\mathcal{T}_i \subseteq A_i$ and $\mathcal{T} \setminus \mathcal{T}_i \subseteq B_i$ that minimizes $|E(A_i, B_i)|$. Among all such partitions, choose the one where $|A_i|$ is the smallest. Since instances (G, \mathcal{M}) , (G', \mathcal{M}') are equivalent, from Observation D.2 and the max-flow min-cut theorem, $|E(A_i, B_i)| = 4$. We now claim that all sets A_i are mutually disjoint.

Claim D.4 *For all $1 \leq i < j \leq 5$, $A_i \cap A_j = \emptyset$.*

Proof: For every set $U \subseteq V(G')$ of vertices, let $\delta(U) = |E(U, V(G') \setminus U)|$. From the sub-modularity of cuts:

$$\delta(A_i) + \delta(A_j) \geq \delta(A_i \setminus A_j) + \delta(A_j \setminus A_i).$$

Recall that $\delta(A_i) = \delta(A_j) = 4$. From our definition of the sets $\{A_k\}_{k=1}^5$, $\mathcal{T}_i \subseteq A_i \setminus A_j$ and $\mathcal{T}_j \subseteq A_j \setminus A_i$, so $\delta(A_i \setminus A_j), \delta(A_j \setminus A_i) \geq 4$. This can only happen if both values are equal to 4, contradicting our choice of A_i and A_j . \square

We will show that G' contains the graph $H = K_5$ as a minor, which will lead to a contradiction, since G' is a planar graph. We will embed every vertex $v_i \in V(H)$ into the subgraph $G'[A_i]$ of G' . In order to be able to do so, we need to show each such subgraph is connected.

Observation D.5 *For all $1 \leq i \leq 5$, $G'[A_i]$ is connected.*

Proof: Assume otherwise, and let \mathcal{C} be the set of all connected components of $G'[A_i]$. By the definition of A_i , each such connected component must contain at least one terminal of \mathcal{T}_i . Assume for contradiction that $|\mathcal{C}| > 1$. Then there are two terminals of \mathcal{T}_i that lie in different connected components of \mathcal{T}_i . Assume w.l.o.g. that these terminals are t_1^i and t_2^i . Consider the following set \mathcal{M}'' of demand pairs: \mathcal{M}'' contains the pair (t_1^i, t_2^i) , and for all $3 \leq j \leq 5$, it contains a pair $(t_j^i, t_j^{i'})$ for some index $i' \neq i$. It is easy to verify that the set \mathcal{M}'' of demand pairs is restricted and it is routable in G , and so it must be routable in G' . Each of the pairs $\{(t_j^i, t_j^{i'})\}_{j=3}^5$ is routed on a path that must use at least one edge of $E(A_i, B_i)$. But the path routing the pair (t_1^i, t_2^i) must use two edges of $E(A_i, B_i)$, which is impossible since $|E(A_i, B_i)| = 4$. \square

Let $H = K_5$, and denote $V(H) = \{v_1, \dots, v_5\}$. We define a set $\hat{\mathcal{M}} \subseteq \mathcal{M}'$ of demand pairs exactly as in Observation D.3. Since the set $\hat{\mathcal{M}}$ of demand pairs is routable in G , it must also be routable in G' . Let \mathcal{P} be the set of edge-disjoint paths routing the set $\hat{\mathcal{M}}$ of demand pairs in G' . For each edge $e = (v_i, v_{i'}) \in E(H)$, we let $P_{i,i'}$ be the path in \mathcal{P} , routing the pair (s_e, t_e) . Assume that $i < i'$, and think of the path $P_{i,i'}$ as directed from a terminal in \mathcal{T}_i to a terminal in $\mathcal{T}_{i'}$. Let u be the last vertex of $P_{i,i'}$ that belongs to A_i , and let u' be the first vertex of $P_{i,i'}$ that appears after u on the path and belongs to $A_{i'}$. We then let $P'_{i,i'}$ be the sub-path of $P_{i,i'}$ from u to u' . Note that $P'_{i,i'}$ does not contain any vertex of $A_i \cup A_{i'}$ as an inner vertex. We need the following stronger claim.

Observation D.6 *Path $P_{i,i'}$ does not contain any vertex of $\bigcup_{k=1}^5 A_k$ as an inner vertex.*

Proof: Assume otherwise, and let $w \in \bigcup_{k=1}^5 A_k$ be some vertex that belongs to $P_{i,i'}$ as an inner vertex. Assume that $w \in A_k$. From the above discussion, $k \notin \{i, i'\}$. But there are four paths in \mathcal{P} that terminate at the vertices of \mathcal{T}_k : the paths routing the demand pairs corresponding to the edges of H that are incident to v_k . Each such path contains an edge of $E(A_k, B_k)$, and $|E(A_k, B_k)| = 4$. Therefore, path $P_{i,i'}$ cannot contain an edge of $E(A_k, B_k)$ and thus it cannot contain a vertex of A_k . \square

Let $\mathcal{P}' = \{P'_{i,i'} \mid P_{i,i'} \in \mathcal{P}\}$. Since the maximum vertex degree in G' is 3 and the paths in \mathcal{P} are edge-disjoint, they are almost vertex-disjoint: the only way for two paths in \mathcal{P} to share a vertex is when that vertex is an endpoint of one of these paths. It is then easy to see that whenever two paths in \mathcal{P}' share a vertex, that vertex must be an endpoint of each of these paths.

It is now immediate to show that H is a minor of G' : we map every vertex v_i of H to the connected subgraph $G'[A_i]$ of G' , and every edge e of H to the path P'_e of \mathcal{P}' . But G' is a planar graph and cannot contain a K_5 -minor, a contradiction. \square

References

- [ACG⁺10] Matthew Andrews, Julia Chuzhoy, Venkatesan Guruswami, Sanjeev Khanna, Kunal Talwar, and Lisa Zhang. Inapproximability of edge-disjoint paths and low congestion routing on undirected graphs. *Combinatorica*, 30(5):485–520, 2010.
- [ALM⁺98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM (JACM)*, 45(3):501–555, 1998.
- [And10] Matthew Andrews. Approximation algorithms for the edge-disjoint paths problem via Raecke decompositions. In *Proceedings of IEEE FOCS*, pages 277–286, 2010.
- [AR95] Yonatan Aumann and Yuval Rabani. Improved bounds for all optical routing. In *Proceedings of the sixth annual ACM-SIAM symposium on Discrete algorithms, SODA '95*, pages 567–576, Philadelphia, PA, USA, 1995. Society for Industrial and Applied Mathematics.
- [AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of np. *Journal of the ACM (JACM)*, 45(1):70–122, 1998.
- [AZ05] Matthew Andrews and Lisa Zhang. Hardness of the undirected edge-disjoint paths problem. In *STOC*, pages 276–283. ACM, 2005.
- [CC] Chandra Chekuri and Julia Chuzhoy. Half-integral all-or-nothing flow. Unpublished Manuscript.
- [CE13] Chandra Chekuri and Alina Ene. Poly-logarithmic approximation for maximum node disjoint paths with constant congestion. In *Proc. of ACM-SIAM SODA*, 2013.
- [Chu12] Julia Chuzhoy. Routing in undirected graphs with constant congestion. In *Proc. of ACM STOC*, pages 855–874, 2012.
- [CK15] Julia Chuzhoy and David H. K. Kim. On approximating node-disjoint paths in grids. In Naveen Garg, Klaus Jansen, Anup Rao, and José D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2015, August 24-26, 2015, Princeton, NJ, USA*, volume 40 of *LIPICs*, pages 187–211. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.
- [CKL16] Julia Chuzhoy, David H. K. Kim, and Shi Li. Improved approximation for node-disjoint paths in planar graphs. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016*, pages 556–569, New York, NY, USA, 2016. ACM.
- [CKS04] Chandra Chekuri, Sanjeev Khanna, and F Bruce Shepherd. Edge-disjoint paths in planar graphs. In *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on*, pages 71–80. IEEE, 2004.
- [CKS05] Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd. Multicommodity flow, well-linked terminals, and routing problems. In *Proc. of ACM STOC*, pages 183–192, 2005.

- [CKS06] Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd. An $O(\sqrt{n})$ approximation and integrality gap for disjoint paths and unsplittable flow. *Theory of Computing*, 2(1):137–146, 2006.
- [CL12] Julia Chuzhoy and Shi Li. A polylogarithmic approximation algorithm for edge-disjoint paths with congestion 2. In *Proc. of IEEE FOCS*, 2012.
- [EIS76] Shimon Even, Alon Itai, and Adi Shamir. On the complexity of timetable and multi-commodity flow problems. *SIAM J. Comput.*, 5(4):691–703, 1976.
- [Kar72] R. Karp. Reducibility among combinatorial problems. In R. Miller and J. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.
- [KK13] Ken-Ichi Kawarabayashi and Yusuke Kobayashi. An $O(\log n)$ -approximation algorithm for the edge-disjoint paths problem in Eulerian planar graphs. *ACM Trans. Algorithms*, 9(2):16:1–16:13, March 2013.
- [Kle05] Jon Kleinberg. An approximation algorithm for the disjoint paths problem in even-degree planar graphs. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '05, pages 627–636, Washington, DC, USA, 2005. IEEE Computer Society.
- [KS04] Stavros G. Kolliopoulos and Clifford Stein. Approximating disjoint-path problems using packing integer programs. *Mathematical Programming*, 99:63–87, 2004.
- [KT95] Jon M. Kleinberg and Éva Tardos. Disjoint paths in densely embedded graphs. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, pages 52–61, 1995.
- [KT98] Jon M. Kleinberg and Éva Tardos. Approximations for the disjoint paths problem in high-diameter planar networks. *J. Comput. Syst. Sci.*, 57(1):61–73, 1998.
- [KvL84] MR Kramer and Jan van Leeuwen. The complexity of wire-routing and finding minimum area layouts for arbitrary vlsi circuits. *Advances in computing research*, 2:129–146, 1984.
- [Lyn75] James F. Lynch. The equivalence of theorem proving and the interconnection problem. *SIGDA Newsl.*, 5(3):31–36, September 1975.
- [Räc02] Harald Räcke. Minimizing congestion in general networks. In *Proc. of IEEE FOCS*, pages 43–52, 2002.
- [RS90] N. Robertson and P. D. Seymour. Outline of a disjoint paths algorithm. In *Paths, Flows and VLSI-Layout*. Springer-Verlag, 1990.
- [RS95] Neil Robertson and Paul D Seymour. Graph minors. XIII. the disjoint paths problem. *Journal of Combinatorial Theory, Series B*, 63(1):65–110, 1995.
- [RT87] Prabhakar Raghavan and Clark D. Tompson. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7:365–374, December 1987.
- [RZ10] Satish Rao and Shuheng Zhou. Edge disjoint paths in moderately connected graphs. *SIAM J. Comput.*, 39(5):1856–1887, 2010.

- [SCS11] Loïc Seguin-Charbonneau and F. Bruce Shepherd. Maximum edge-disjoint paths in planar graphs with congestion 2. In *Proceedings of the 2011 IEEE 52Nd Annual Symposium on Foundations of Computer Science*, FOCS '11, pages 200–209, Washington, DC, USA, 2011. IEEE Computer Society.