# Almost Polynomial Hardness of Node-Disjoint Paths in Grids 

Julia Chuzhoy*<br>David H.K. Kim ${ }^{\dagger}$<br>Rachit Nimavat ${ }^{\ddagger}$

November 22, 2020


#### Abstract

In the classical Node-Disjoint Paths (NDP) problem, we are given an $n$-vertex graph $G=(V, E)$, and a collection $\mathcal{M}=\left\{\left(s_{1}, t_{1}\right), \ldots,\left(s_{k}, t_{k}\right)\right\}$ of pairs of its vertices, called source-destination, or demand pairs. The goal is to route as many of the demand pairs as possible, where to route a pair we need to select a path connecting it, so that all selected paths are disjoint in their vertices. The best current algorithm for NDP achieves an $O(\sqrt{n})$ approximation [Kolliopoulos, Stein, Mathematical Programming '04], while the best current negative result is a factor $2^{\Omega(\sqrt{\log n)}}$-hardness of approximation unless NP $\subseteq \operatorname{DTIME}\left(n^{O(\log n)}\right)$ [Chuzhoy, Kim, Nimavat, STOC '17], even if the underlying graph is a subgraph of a grid graph; unfortunately, this result does not extend to grid graphs.

The approximability of the NDP problem on grid graphs has remained a tantalizing open question, with the best current upper bound of $\tilde{O}\left(n^{1 / 4}\right)$, and the best current lower bound of APX-hardness [Chuzhoy, Kim, APPROX '15] only ruling out a ( $1+\boldsymbol{\delta}$ )-approximation algorithm for some fixed $\delta>0$, assuming that $\mathrm{P} \neq \mathrm{NP}$. In this paper we come close to


[^0]ACM Classification: F.2.2,
AMS Classification: 68Q17,68Q25,68W25
Key words and phrases: disjoint paths, graph routing, hardness of approximation

## Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat

resolving the approximability of NDP in general, and of NDP in grids in particular. Our main result is that NDP is $2^{\Omega\left(\log ^{1-\varepsilon} n\right)}$-hard to approximate for any constant $\varepsilon$, assuming that $\mathrm{NP} \nsubseteq \mathrm{DTIME}\left(n^{\text {poly } \log n}\right)$, and that it is $n^{\Omega\left(1 /(\log \log n)^{2}\right)}$-hard to approximate, assuming that for some constant $\delta>0, \operatorname{NP} \nsubseteq \operatorname{DTIME}\left(2^{n^{\delta}}\right)$. These results hold even for grid graphs and wall graphs, and extend to the closely related Edge-Disjoint Paths problem, even in wall graphs.

Our hardness proof performs a reduction from the $3 \mathrm{COL}(5)$ problem to NDP, using a new graph partitioning problem as a proxy. Unlike the more standard approach of employing Karp reductions to prove hardness of approximation, our proof is a Cook-type reduction, where, given an input instance of $3 \mathrm{COL}(5)$, we produce a large number of instances of NDP, and apply an approximation algorithm for NDP to each of them.

## 1 Introduction

We study the Node-Disjoint Paths (NDP) problem: given an undirected $n$-vertex graph $G$ and a collection $\mathcal{M}=\left\{\left(s_{1}, t_{1}\right), \ldots,\left(s_{k}, t_{k}\right)\right\}$ of pairs of its vertices, called source-destination, or demand pairs, we are interested in routing the demand pairs: in order to route a pair $\left(s_{i}, t_{i}\right)$, we need to select a path connecting $s_{i}$ to $t_{i}$. The goal is to route as many of the pairs as possible, subject to the constraint that the selected routing paths are mutually disjoint in their vertices and their edges. We let $S=\left\{s_{1}, \ldots, s_{k}\right\}$ be the set of the source vertices, $T=\left\{t_{1}, \ldots, t_{k}\right\}$ the set of the destination vertices, and we refer to the vertices of $S \cup T$ as terminals. We denote by NDP-Planar the special case of the problem where the graph $G$ is planar; by NDP-Grid the special case where $G$ is a square grid; and by NDP-Wall the special case where $G$ is a wall (see Figure 1 for an illustration of a wall and Section 2 for its formal definition).
NDP is a fundamental problem in the area of graph routing, that has been studied extensively. Robertson and Seymour [48, 49] showed, as part of their famous Graph Minors Series, an efficient algorithm for solving the problem if the number $k$ of the demand pairs is bounded by a constant. However, when $k$ is a part of input, the problem becomes NP-hard [29, 22], and it remains NP-hard even for planar graphs [40], and for grid graphs [37]. The best current upper bound on the approximability of NDP is $O(\sqrt{n})$, obtained by a simple greedy algorithm [36]. Until recently, the best known lower bound was an $\Omega\left(\log ^{1 / 2-\varepsilon} n\right)$-hardness of approximation for any constant $\varepsilon$, unless NP $\subseteq$ ZPTIME $\left(n^{\text {poly } \log n}\right)$ [5, 4]. For the special cases of NDP-Planar and NDP-Grid [17], the best known lower bound was APX-hardness, so that they do not have a $(1+\delta)$-approximation algorithm, for some fixed $\delta>0$, unless $\mathrm{P}=\mathrm{NP}$. In a recent paper [19], the authors have shown an improved $2^{\Omega(\sqrt{\log n})}$-hardness of approximation for NDP, assuming that $\mathrm{NP} \nsubseteq \mathrm{DTIME}\left(n^{O(\log n)}\right)$. This result holds even for planar graphs with maximum vertex degree 3 , where all source vertices lie on the boundary of a single face, and for subgraphs of grid graphs, with all source vertices lying on the boundary of the grid. We note that for general planar graphs, the $O(\sqrt{n})$ approximation algorithm of [36] was recently slightly improved to an $\tilde{O}\left(n^{9 / 19}\right)$-approximation [18].
The approximability status of NDP-Grid- the special case of NDP where the underlying graph is a square grid - remained a tantalizing open question. The study of this problem dates back to the 70's, and was initially motivated by applications to VLSI design. As grid graphs are extremely well-structured,

## Almost Polynomial Hardness of Node-Disjoint Paths in Grids

one would expect that good approximation algorithms can be designed for them, or that, at the very least, they should be easy to understand. However, establishing the approximability of NDP-Grid has been elusive so far. The simple greedy $O(\sqrt{n})$-approximation algorithm of [36] was only recently improved to a $\tilde{O}\left(n^{1 / 4}\right)$-approximation for NDP-Grid [17], while on the negative side only APX-hardness is known. In a very recent paper [20], the authors designed a $\left.2^{O(\sqrt{\log n} \cdot \log \log n}\right)$-approximation algorithm for a special case of NDP-Grid, where the source vertices appear on the grid boundary. This result can be seen as complementing the $2^{\Omega(\sqrt{\log n})}$-hardness of approximation of NDP on subgraphs of grids with all sources lying on the grid boundary [19] ${ }^{1}$. Furthermore, this result can be seen as suggesting that subpolynomial approximation algorithms may be achievable for NDP-Grid.

In this paper we show that this is unlikely to be the case, and come close to resolving the approximability status of NDP-Grid, and of NDP in general, by showing that NDP-Grid is $2^{\Omega\left(\log ^{1-\varepsilon} n\right)}$-hard to approximate for any constant $\varepsilon$, unless ${ }^{2}$ NP $\subseteq \operatorname{DTIME}\left(n^{\text {poly } \log n}\right)$. We further show that it is $n^{\Omega\left(1 /(\log \log n)^{2}\right)}$-hard to approximate, assuming that for some constant $\delta>0, \mathrm{NP} \nsubseteq \mathrm{DTIME}\left(2^{n^{\delta}}\right)$. The same hardness results also extend to NDP-Wall. These hardness results are stronger than the best currently known hardness for the general NDP problem, and should be contrasted with the $\left.2^{O(\sqrt{\log n} \cdot \log \log n}\right)$-approximation algorithm for NDP-Grid with all sources lying on the grid boundary [20].

Another basic routing problem that is closely related to NDP is Edge-Disjoint Paths (EDP). The input to this problem is the same as before: an undirected graph $G=(V, E)$ and a set $\mathcal{M}=\left\{\left(s_{1}, t_{1}\right), \ldots,\left(s_{k}, t_{k}\right)\right\}$ of demand pairs. The goal, as before, is to route the largest number of the demand pairs via paths. However, we now allow the paths to share vertices, and only require that they are mutually edge-disjoint. In general, it is easy to see that EDP is a special case of NDP. Indeed, given an EDP instance $(G, \mathcal{M})$, computing the line graph of the input graph $G$ transforms it into an equivalent instance of NDP. However, this transformation may inflate the number of the graph vertices, and so approximation factors that depend on $|V(G)|$ may no longer be preserved. Moreover, this transformation does not preserve planarity, and no such relationship is known between NDP and EDP in planar graphs. The approximability status of EDP is very similar to that of NDP: the best current approximation algorithm achieves an $O(\sqrt{n})$-approximation factor [13], and the recent $2^{\Omega(\sqrt{\log n})}$-hardness of approximation of [19], under the assumption that $\mathrm{NP} \nsubseteq \mathrm{DTIME}\left(n^{O(\log n)}\right)$, extends to EDP. Interestingly, EDP appears to be relatively easy on grid graphs, and has a constant-factor approximation for this special case [6, 35, 34]. The analogue of the grid graph in the setting of EDP seems to be the wall graph (see Figure 1): the approximability status of EDP on wall graphs is similar to that of NDP on grid graphs, with the best current upper bound of $\tilde{O}\left(n^{1 / 4}\right)$, and the best lower bound of APX-hardness [17]. The results of [19] extend to a $2^{\Omega(\sqrt{\log n})}$-hardness of approximation for EDP on subgraphs of wall graphs, with all source vertices lying on the wall boundary, under the same complexity assumption. We denote by EDP-Wall the special case of the EDP problem where the underlying graph is a wall. We show that our new almost polynomial hardness of approximation results also hold for EDP-Wall and for NDP-Wall.

[^1]
## Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat



Figure 1: A wall graph.

Other related work. Several other special cases of EDP are known to have reasonably good approximation algorithms. For example, for the special case of Eulerian planar graphs, Kleinberg [32] showed an $O\left(\log ^{2} n\right)$-approximation algorithm, while Kawarabayashi and Kobayashi [30] provide an improved $O(\log n)$-approximation for both Eulerian and 4-connected planar graphs. Polylogarithmic approximation algorithms are also known for bounded-degree expander graphs [38, 10, 9, 33, 26], and constant-factor approximation algorithms are known for trees [27, 15], and grids and grid-like graphs [6, 7, 35, 34]. Rao and Zhou [46] showed an efficient randomized $O$ (poly $\log n)$-approximation algorithm for the special case of EDP where the value of the global minimum cut in the input graph is $\Omega\left(\log ^{5} n\right)$. Recently, Fleszar et al. [25] designed an $O(\sqrt{r} \cdot \log (k r))$-approximation algorithm for EDP, where $r$ is the feedback vertex set number of the input graph $G=(V, E)$ - the smallest number of vertices that need to be deleted from $G$ in order to turn it into a forest.

A natural variation of NDP and EDP that relaxes the disjointness constraint by allowing a small vertexor edge-congestion has been a subject of extensive study. In the NDP with Congestion (NDPwC) problem, the input consists of an undirected graph and a set of demand pairs as before, and additionally a non-negative integer $c$. The goal is to route a maximum number of the demand pairs with congestion $c$, that is, each vertex may participate in at most $c$ paths in the solution. The EDP with Congestion problem (EDPwC) is defined similarly, except that now the congestion is measured on the graph edges and not vertices. The famous result of Raghavan and Thompson [43], that introduced the randomized LP-rounding technique, obtained a constant-factor approximation for NDPwC and EDPwC, for a congestion value $c=\Theta(\log n / \log \log n)$. A long sequence of work $[12,42,3,46,16,21,11]$ has led to an $O$ (poly $\log k$ )-approximation for EDPwC with congestion bound $c=2$, and for NDPwC with constant congestion. These results are essentially optimal, since it is known that for every constant $\varepsilon$, and for every congestion value $c=o(\log \log n / \log \log \log n)$, both problems are hard to approximate to within a factor $\Omega\left((\log n)^{\frac{1-\varepsilon}{c+1}}\right)$, unless NP $\subseteq \operatorname{ZPTIME}\left(n^{\text {poly } \log n}\right)$ [4]. When the input graph is planar, SeguinCharbonneau and Shepherd [50], improving on the result of Chekuri, Khanna and Shepherd [14], have shown a constant-factor approximation for EDPwC with congestion 2.

Our results and techniques. Our main result is the proof of the following two theorems.
Theorem 1.1. For every constant $\varepsilon>0$, there is no $2^{O\left(\log ^{1-\varepsilon} n\right)}$-approximation polynomial-time algorithm for NDP, assuming that NP $\nsubseteq \mathrm{DTIME}\left(n^{\mathrm{poly} \log n}\right)$. Moreover, there is no $n^{O\left(1 /(\log \log n)^{2}\right)}$-approximation polynomial-time algorithm for NDP, assuming that for some constant $\delta>0, \mathrm{NP} \nsubseteq \operatorname{DTIME}\left(2^{n^{\delta}}\right)$. These

## Almost Polynomial Hardness of Node-Disjoint Paths in Grids

results hold even when the input graph is a grid graph or a wall graph.
Theorem 1.2. For every constant $\varepsilon>0$, there is no $2^{O\left(\log ^{1-\varepsilon} n\right)}$-approximation polynomial-time algorithm for EDP , assuming that $\mathrm{NP} \nsubseteq \mathrm{DTIME}\left(n^{\text {poly } \log n}\right)$. Moreover, there is no $n^{O\left(1 /(\log \log n)^{2}\right)}$-approximation polynomial-time algorithm for EDP, assuming that for some constant $\delta>0, \operatorname{NP} \nsubseteq \operatorname{DTIME}\left(2^{n^{\delta}}\right)$. These results hold even when the input graph is a wall graph.

We now provide a high-level overview of our techniques. The starting point of our hardness of approximation proof is $3 \mathrm{COL}(5)$ - a special case of the 3 -coloring problem, where the underlying graph is 5 -regular. We define a new graph partitioning problem, that we refer to as $(r, h)$-Graph Partitioning, and denote by $(r, h)$-GP. In this problem, we are given a bipartite graph $\tilde{G}=\left(V_{1}, V_{2}, E\right)$ and two integral parameters $r, h>0$. A solution to the problem is a partition $\left(W_{1}, W_{2}, \ldots, W_{r}\right)$ of $V_{1} \cup V_{2}$ into $r$ subsets, and for each $1 \leq i \leq r$, a subset $E_{i} \subseteq E\left(W_{i}\right)$ of edges, so that $\left|E_{i}\right| \leq h$ holds, and the goal is to maximize $\sum_{i=1}^{r}\left|E_{i}\right|$. A convenient intuitive way to think about this problem is that we would like to partition $\tilde{G}$ into a large number of subgraphs, in a roughly balanced way (with respect to the number of edges), so as to preserve as many of the edges as possible. We show that NDP-Grid is at least as hard to approximate as the ( $\mathrm{r}, \mathrm{h}$ )-GP problem (to within polylogarithmic factors). Our reduction exploits the fact that routing in grids is closely related to graph drawing, and that graphs with small crossing number have small balanced separators. The ( $r, h$ )-GP problem itself appears similar in flavor to the Densest $k$-Subgraph problem (DkS). In the DkS problem, we are given a graph $G=(V, E)$ and a parameter $k$, and the goal is to find a subset $U \subseteq V$ of $k$ vertices, that maximizes the number of edges in the induced graph $G[U]$. Intuitively, in the (r,h)-GP problem, the goal is to partition the graph into many dense subgraphs, and so in order to prove that ( $\mathrm{r}, \mathrm{h}$ )-GP is hard to approximate, it is natural to employ techniques used in hardness of approximation proofs for DkS. The best current approximation algorithm for DkS achieves a $n^{1 / 4+\varepsilon}$-approximation for any constant $\varepsilon$ [8]. Even though the problem appears to be very hard to approximate, its hardness of approximation proof has been elusive until recently: only constant-factor hardness results were known for DkS under various worst-case complexity assumptions, and $2^{\Omega\left(\log ^{2 / 3} n\right)}$-hardness under average-case assumptions [23, 1, 31, 44]. In a recent breakthrough, Manurangsi [41] has shown that for some constant $c$, DkS is hard to approximate to within a factor $n^{1 /(\log \log n)^{c}}$, under the Exponential Time Hypothesis. Despite our feeling that $(r, h)-G P$ is somewhat similar to DkS , we were unable to extend the techniques of [41] to this problem, or to prove its hardness of approximation via other techniques.
We overcome this difficulty as follows. First, we define a graph partitioning problem that is slightly more general than ( $\mathrm{r}, \mathrm{h}$ )-GP. The definition of this problem is somewhat technical and is deferred to Section 3. This problem is specifically designed so that the reduction to NDP-Grid still goes through, but it is somewhat easier to control its solutions and prove hardness of approximation for it. Furthermore, instead of employing a standard Karp-type reduction (where an instance of 3COL(5) is reduced to a single instance of NDP-Grid, while using the graph partitioning problem as a proxy), we employ a sequential Cook-type reduction. We assume for contradiction that an $\alpha$-approximation algorithm $\mathcal{A}$ for NDP-Grid exists, where $\alpha$ is the hardness of approximation factor we are trying to prove. Our reduction is iterative. In every iteration $j$, we reduce the $3 \mathrm{COL}(5)$ instance to a collection $\mathcal{J}_{j}$ of instances of NDP-Grid, and apply the algorithm $\mathcal{A}$ to each of them. If the 3COL(5) instance is a Yes-Instance, then we are guaranteed that each resulting instance of NDP-Grid has a large solution, and so all solutions returned by
$\mathcal{A}$ are large. If the $3 \operatorname{COL}(5)$ instance is a No-Instance, then unfortunately it is still possible that the resulting instances of NDP-Grid will have large solutions. However, we can use these resulting solutions in order to further refine our reduction, and construct a new collection $\mathcal{J}_{j+1}$ of instances of NDP-Grid. While in the Yes-Instance case we will continue to obtain large solutions to all NDP-Grid instances that we construct, we can show that in the No-Instance case, in some iteration of the algorithm, we will fail to find such a large solution. Our reduction is crucially sequential, and we exploit the solutions returned by algorithm $\mathcal{A}$ in previous iterations in order to construct new instances of NDP-Grid for subsequent iterations. It is interesting whether these techniques may be helpful in obtaining new hardness of approximation results for DkS.

We note that our approach is completely different from the previous hardness of approximation proof of [19]. The proof in [19] proceeded by performing a reduction from 3SAT(5). Initially, a simple reduction from 3SAT(5) to the NDP problem on a subgraph of the grid graph is used in order to produce a constant hardness gap. The resulting instance of NDP is called a level-1 instance. The reduction then employs a boosting technique, that, in order to obtain a level- $i$ instance, combines a number of level- $(i-1)$ instances with a single level- 1 instance. The hardness gap grows by a constant factor from iteration to iteration, until the desired hardness of approximation bound is achieved. All source vertices in the constructed instances appear on the grid boundary, and a large number of vertices are carefully removed from the grid in order to create obstructions to routing, and to force the routing paths to behave in a prescribed way. The reduction itself is a Karp-type reduction, and eventually produces a single instance of NDP with a large gap between the Yes-Instance and No-Instance solutions.

Organization. We start with Preliminaries in Section 2, and introduce the new graph partitioning problems in Section 3. We introduce two main tools that we use in the hardness of approximation proof in Section 4. The tools are a reduction from 3COL to the graph partitioning problem, and a reduction from the graph partitioning problem to NDP-Grid. The hardness of approximation proof for NDP-Grid appears in Section 5, with the details of the reduction from the graph partitioning problem to NDP-Grid deferred to Section 6. Finally, we extend our hardness results to NDP-Wall and EDP-Wall in Section 7.

## 2 Preliminaries

All graphs in this paper are simple. We use standard graph-theoretic notation. Given a graph $G$ and a subset $W \subseteq V(G)$ of its vertices, $E(W)$ denotes the set of all edges of $G$ that have both their endpoints in $W$. Given a path $P$ and a subset $U$ of vertices of $G$, we say that $P$ is internally disjoint from $U$ iff every vertex in $P \cap U$ is an endpoint of $P$. Similarly, $P$ is internally disjoint from a subgraph $G^{\prime}$ of $G$ iff $P$ is internally disjoint from $V\left(G^{\prime}\right)$. Given a subset $\mathcal{N}^{\prime} \subseteq \mathcal{M}$ of the demand pairs in $G$, we denote by $S\left(\mathcal{N}^{\prime}\right)$ and $T\left(\mathcal{M}^{\prime}\right)$ the sets of the source and the destination vertices of the demand pairs in $\mathcal{M}^{\prime}$, respectively. We let $\mathcal{T}\left(\mathcal{M}^{\prime}\right)=S\left(\mathcal{M}^{\prime}\right) \cup T\left(\mathcal{M}^{\prime}\right)$ denote the set of all terminals participating as a source or a destination in $\mathcal{M}^{\prime}$. Given a graph $G$ and a vertex $v \in V(G)$, we denote by $d_{G}(v)$ the degree of $v$ in $G$. All logarithms in this paper are to the base of 2 .

## Almost Polynomial Hardness of Node-Disjoint Paths in Grids

Grid graphs. For a pair $h, \ell>0$ of integers, we let $G^{h, \ell}$ denote the grid of height $h$ and length $\ell$. The set of its vertices is $V\left(G^{h, \ell}\right)=\{v(i, j) \mid 1 \leq i \leq h, 1 \leq j \leq \ell\}$, and the set of its edges is the union of two subsets: the set $E^{H}=\{(v(i, j), v(i, j+1)) \mid 1 \leq i \leq h, 1 \leq j<\ell\}$ of horizontal edges and the set $E^{V}=\{(\nu(i, j), v(i+1, j)) \mid 1 \leq i<h, 1 \leq j \leq \ell\}$ of vertical edges. The subgraph of $G^{h, \ell}$ 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 $\ell$ 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. Given a vertex $v=v(i, j)$ of the grid, we denote by $\operatorname{row}(v)$ and $\operatorname{col}(v)$ the row and the column of the grid, respectively, that contain $v$. We say that $G^{h, \ell}$ is a square grid iff $h=\ell$. The boundary of the grid is $R_{1} \cup R_{h} \cup W_{1} \cup W_{\ell}$. We sometimes refer to $R_{1}$ and $R_{h}$ as the top and the bottom boundary edges of the grid respectively, and to $W_{1}$ and $W_{\ell}$ as the left and the right boundary edges of the grid.
Given a subset $\mathcal{R}^{\prime}$ of consecutive rows of $G$ and a subset $\mathcal{W}^{\prime}$ of consecutive columns of $G$, the subgrid of $G$ spanned by the rows in $\mathcal{R}^{\prime}$ and the columns in $\mathcal{W}^{\prime}$ is the subgraph of $G$ induced by the set $\left\{v \mid \operatorname{row}(v) \in \mathcal{R}^{\prime}, \operatorname{col}(v) \in \mathcal{W}^{\prime}\right\}$ of its vertices.

Given two vertices $u=v(i, j)$ and $u^{\prime}=v\left(i^{\prime}, j^{\prime}\right)$ of a grid $G$, the shortest-path distance between them is denoted by $d\left(u, u^{\prime}\right)$. Given two vertex subsets $X, Y \subseteq V(G)$, the distance between them is $d(X, Y)=$ $\min _{u \in X, u^{\prime} \in Y}\left\{d\left(u, u^{\prime}\right)\right\}$. When $H, H^{\prime}$ are subgraphs of $G$, we use $d\left(H, H^{\prime}\right)$ to denote $d\left(V(H), V\left(H^{\prime}\right)\right)$.

Wall graphs. Let $G=G^{\ell, h}$ be a grid of length $\ell$ and height $h$. Assume that $\ell>0$ is an even integer, and that $h>0$. For every column $W_{j}$ of the grid, let $e_{1}^{j}, \ldots, e_{h-1}^{j}$ 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 $z \neq j \bmod 2$, and let $\hat{G}$ be the graph obtained from $G \backslash E^{*}(G)$, by deleting all degree-1 vertices from it. Graph $\hat{G}$ is called a wall of length $\ell / 2$ and height $h$ (see Figure 1). 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 $R_{1}, \ldots, R_{h}$ in this top-to-bottom order; notice that $R_{j}$ is also the $j$ th row of the grid $G$ 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 $W_{1}, \ldots, W_{\ell / 2}$ 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 their union is the boundary of $\hat{G}$.

Semiregular bipartite graphs. We will use the following notion of semiregular bipartite graphs.
Definition 2.1. Let $G=(A, B, E)$ be a bipartite graph. We say that $G$ is $\left(d, d^{\prime}\right)$-semiregular, for integers $d, d^{\prime}>0$ if the following hold:

- For every vertex $v \in A, d \leq d_{G}(v)<2 d$;
- For every vertex $u \in B, d_{G}(u)<2 d^{\prime}$; and
- $|E| \geq d^{\prime}|B| /(4 \log |B|)$ if $|B|>1$, and $|E|=d^{\prime}$ otherwise.


## Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat

We say that $G$ is semiregular if it is $\left(d, d^{\prime}\right)$-semiregular for any integers $d, d^{\prime}>0$.
Claim 2.2. There is a polynomial-time algorithm, that, given any bipartite graph $G=(A, B, E)$ with $|A|,|B|>1$, computes a semiregular subgraph $G^{\prime}=\left(A^{\prime}, B^{\prime}, E^{\prime}\right)$ of $G$, with $A^{\prime} \subseteq A, B^{\prime} \subseteq B$, and $\left|E^{\prime}\right| \geq$ $\Omega(|E| /(\log |A| \log |B|))$.
(we note that, if $|A|=1$ or $|B|=1$, then the graph is semiregular by definition).
Proof. We transform the graph $G$ into a semiregular graph $G^{\prime}$ in two steps. In the first step, we regularize the degrees of the vertices of $B$. We partition the vertices in $B$ into $r=\lceil\log |A|\rceil+1$ groups $Y_{1}, \ldots, Y_{r}$, as follows: group $Y_{i}$ contains all vertices $u \in B$ with $2^{i} \leq d_{G}(u)<2^{i+1}$. If a vertex $u \in B$ belongs to group $Y_{i}$, then we say that all edges that are incident to $u$ also belong to group $Y_{i}$. Clearly, there is some index $1 \leq i \leq r$, such that at least $|E| / r=\Omega(|E| / \log |A|)$ edges of $E$ belong to group $Y_{i}$. We let $d^{\prime}=2^{i}$, and we define a new bipartite graph $\hat{G}$. The set of vertices consists of the set $A$, and of a set $B^{\prime} \subseteq B$, containing all vertices that belong to group $Y_{i}$. The set of edges, that we denote by $\hat{E}$, contains all edges of $E$ that are incident to the vertices of $B^{\prime}$. Notice that $|\hat{E}| \geq d^{\prime}\left|B^{\prime}\right|$.

In the second step, we similarly regularize the degrees of the vertices in $A$. We partition the vertices of $A$ into $r^{\prime}=\left\lceil\log \left|B^{\prime}\right|\right\rceil+1$ groups $X_{1}, \ldots, X_{r^{\prime}}$, as follows: group $X_{i}$ contains all vertices $v \in A$ with $2^{j} \leq d_{\hat{G}}(v)<2^{j+1}$. If a vertex $v \in A$ belongs to group $X_{j}$, then we say that all edges that are incident to $v$ in $\hat{G}$ also belong to group $X_{j}$. Clearly, there is some index $1 \leq j \leq r^{\prime}$, such that at least $|\hat{E}| / r^{\prime} \geq$ $d^{\prime}\left|B^{\prime}\right| /\left(\left\lceil\log \left|B^{\prime}\right|\right\rceil+1\right) \geq d^{\prime} B^{\prime} /\left(4 \log \left|B^{\prime}\right|\right) \geq \Omega(|E| /(\log |A| \log |B|))$ edges of $\hat{E}$ belong to group $X_{j}$. We set $d=2^{j}$. We now define the final bipartite graph $G^{\prime}=\left(A^{\prime}, B^{\prime}, E^{\prime}\right)$. The set $B^{\prime}$ of vertices remains the same as before. The set $A^{\prime}$ of vertices contains every vertex of $A$ that belongs to class $X_{j}$. The set $E^{\prime}$ of edges contains every edge of $\hat{E}$ that is incident to a vertex in $A^{\prime}$. It is now immediate to verify that for every vertex $v \in A^{\prime}, d \leq d_{G^{\prime}}(v)<2 d$, and for every vertex $u \in B^{\prime}, d_{G^{\prime}}(u)<2 d^{\prime}$. From the above discussion, $\left|E^{\prime}\right| \geq \Omega(|E| / \log |A| \log |B|)$. Moreover, $\left|E^{\prime}\right| \geq d^{\prime}\left|B^{\prime}\right| /\left(4 \log \left|B^{\prime}\right|\right)$ as required.

The 3COL(5) problem. The starting point of our reduction is the 3COL(5) problem. In this problem, we are given a 5-regular graph $G=(V, E)$. Note that, if $n=|V|$ and $m=|E|$, then $m=5 n / 2$. We are also given a set $\mathcal{C}=\{r, b, g\}$ of 3 colors. A coloring $\chi: V \rightarrow \mathcal{C}$ is an assignment of a color in $\mathcal{C}$ to every vertex in $V$. We say that an edge $e=(u, v)$ is satisfied by the coloring $\chi$ iff $\chi(u) \neq \chi(v)$. The coloring $\chi$ is valid iff it satisfies every edge. We say that $G$ is a Yes-Instance iff there is a valid coloring $\chi: V \rightarrow \mathcal{C}$. We say that it is a No-Instance with respect to some given parameter $\varepsilon$, iff for every coloring $\chi: V \rightarrow \mathcal{C}$, at most a $(1-\varepsilon)$-fraction of the edges are satisfied by $\chi$. We use the following theorem of Feige et al. [24]:
Theorem 2.3. [Proposition 15 in [24]] There is some constant $\varepsilon$, such that distinguishing between the Yes-Instances and the No-Instances (with respect to $\varepsilon$ ) of 3 COL(5) is NP-hard.

A two-prover protocol. We use the following two-prover protocol. The two provers are called an edge-prover and a vertex-prover. Given a 5 -regular graph $G$, the verifier selects an edge $e=(u, v)$ of $G$ uniformly at random, and then selects a random endpoint (say $v$ ) of this edge. It then sends edge $e$ to the edge-prover and vertex $v$ to the vertex-prover. The edge-prover must return an assignment of colors from
$\mathcal{C}$ to $u$ and $v$, such that the two colors are distinct; the vertex-prover must return an assignment of a color from $\mathcal{C}$ to $v$. The verifier accepts iff both provers assign the same color to $v$. Given a 2-prover game $\mathcal{G}$, its value is the maximum acceptance probability of the verifier over all possible strategies of the provers.

Notice that, if $G$ is a Yes-Instance, then there is a strategy for both provers that guarantees acceptance with probability 1 : the provers fix a valid coloring $\chi: V \rightarrow \mathcal{C}$ of $G$ and respond to the queries according to this coloring.
We claim that if $G$ is a No-Instance, then for any strategy of the two provers, the verifier accepts with probability at most ( $1-\varepsilon / 2$ ). Note first that we can assume without loss of generality that the strategies of the provers are deterministic. Indeed, if the provers have a probability distribution over the answers to each query $Q$, then the edge-prover, given a query $Q^{\prime}$, can return an answer that maximizes the acceptance probability of the verifier under the random strategy of the vertex-prover. This defines a deterministic strategy for the edge-prover that does not decrease the acceptance probability of the verifier. The vertexprover in turn, given any query $Q$, can return an answer that maximizes the acceptance probability of the verifier, under the new deterministic strategy of the edge-prover. The acceptance probability of this final deterministic strategy of the two provers is at least as high as that of the original randomized strategy. The deterministic strategy of the vertex-prover defines a coloring of the vertices of $G$. This coloring must dissatisfy at least $\varepsilon m$ edges. The probability that the verifier chooses one of these edge is at least $\varepsilon$. The response of the edge-prover on such an edge must differ from the response of the vertex-prover on at least one endpoint of the edge. The verifier chooses this endpoint with probability at least $\frac{1}{2}$, and so overall the verifier rejects with probability at least $\varepsilon / 2$. Therefore, if $G$ is a Yes-Instance, then the value of the corresponding game is 1 , and if it is a No-InSTANCE, then the value of the game is at most $(1-\varepsilon / 2)$.

Parallel repetition. We perform $\ell$ rounds of parallel repetition of the above protocol, for some integer $\ell>0$, that may depend on $n=|V(G)|$. Specifically, the verifier chooses a sequence ( $e_{1}, \ldots, e_{\ell}$ ) of $\ell$ edges, where each edge $e_{i}$ is selected independently uniformly at random from $E(G)$. For each chosen edge $e_{i}$, one of its endpoints $v_{i}$ is then chosen independently at random. The verifier sends ( $e_{1}, \ldots, e_{\ell}$ ) to the edge-prover, and $\left(v_{1}, \ldots, v_{\ell}\right)$ to the vertex-prover. The edge-prover returns a coloring of both endpoints of each edge $e_{i}$. This coloring must satisfy the edge (so the two endpoints must be assigned different colors), but it need not be consistent across different edges. In other words, if two edges $e_{i}$ and $e_{j}$ share the same endpoint $v$, the edge-prover may assign different colors to each occurrence of $v$. The vertex-prover returns a coloring of the vertices in $\left(v_{1}, \ldots, v_{\ell}\right)$. Again, if some vertex $v_{i}$ repeats twice, the coloring of the two occurrences need not be consistent. The verifier accepts iff for each $1 \leq i \leq \ell$, the coloring of the vertex $v_{i}$ returned by both provers is consistent. (No verification of consistency is performed across different $i$ 's. So, for example, if $v_{i}$ is an endpoint of $e_{j}$ for $i \neq j$, then it is possible that the two colorings do not agree and the verifier still accepts). We say that a pair $\left(A, A^{\prime}\right)$ of answers to the two queries $\left(e_{1}, \ldots, e_{\ell}\right)$ and $\left(v_{1}, \ldots, v_{\ell}\right)$ is matching, or consistent, iff it causes the verifier to accept. We let $\mathcal{G}^{\ell}$ denote this 2 -prover protocol with $\ell$ repetitions.

Theorem 2.4 (Parallel Repetition). [47, 28, 45] There is some constant $0<\gamma^{\prime}<1$, such that for each 2 -prover game $\tilde{\mathcal{G}}$, if the value of $\tilde{\mathcal{G}}$ is $x$, then the value of the game $\tilde{\mathcal{G}}^{\ell}$, obtained from $\ell>0$ parallel repetitions of $\tilde{\mathcal{G}}$, is at most $\gamma^{\gamma}$.

Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat

Corollary 2.5. There is some constant $0<\gamma<1$, such that, if $G$ is $a$ Yes-Instance, then $\mathcal{G}^{\ell}$ has value 1 , and if $G$ is $a$ No-Instance, then $\mathcal{G}^{\ell}$ has value at most $2^{-\gamma \ell}$.

We now summarize the parameters and introduce some basic notation:

- Let $Q^{E}$ denote the set of all possible queries to the edge-prover, so each query is an $\ell$-tuple of edges. Then $\left|Q^{E}\right|=m^{\ell}=(5 n / 2)^{\ell}$. Each query has $6^{\ell}$ possible answers -6 colorings per edge. The set of feasible answers is the same for each edge-query, and we denote it by $\mathcal{A}^{E}$.
- Let $Q^{V}$ denote the set of all possible queries to the vertex-prover, so each query is an $\ell$-tuple of vertices, and $\left|Q^{V}\right|=n^{\ell}$. Each query has $3^{\ell}$ feasible answers -3 colorings per vertex. The set of feasible answers is the same for each vertex-query, and we denote it by $\mathcal{A}^{V}$.
- We think about the verifier as choosing a number of random bits, that determine the choices of the queries $Q \in Q^{E}$ and $Q^{\prime} \in Q^{V}$ that it sends to the provers. We sometimes call each such random choice a "random string". The set of all such choices is denoted by $\mathcal{R}$, where for each $R \in \mathcal{R}$, we denote $R=\left(Q, Q^{\prime}\right)$, with $Q \in Q^{E}, Q^{\prime} \in Q^{V}$ - the two queries sent to the two provers when the verifier chooses $R$. Then $|\mathcal{R}|=(2 m)^{\ell}=(5 n)^{\ell}$, and each random string $R \in \mathcal{R}$ is chosen with the same probability.
- It is important to note that each query $Q=\left(e_{1}, \ldots, e_{\ell}\right) \in Q^{E}$ of the edge-prover participates in exactly $2^{\ell}$ random strings (one random string for each choice of one endpoint per edge of $\left\{e_{1}, \ldots, e_{\ell}\right\}$ ), while each query $Q^{\prime}=\left(v_{1}, \ldots, v_{\ell}\right)$ of the vertex-prover participates in exactly $5^{\ell}$ random strings (one random string for each choice of an edge incident to each of $v_{1}, \ldots, v_{\ell}$ ).

A function $f: Q^{E} \cup Q^{V} \rightarrow \mathcal{A}^{E} \cup \mathcal{A}^{V}$ is called a global assignment of answers to queries iff for every query $Q \in Q^{E}$ to the edge-prover, $f(Q) \in \mathcal{A}^{E}$, and for every query $Q^{\prime} \in Q^{V}$ to the vertex-prover, $f\left(Q^{\prime}\right) \in \mathcal{A}^{V}$. We say that $f$ is a perfect global assignment iff for every random string $R=\left(Q^{E}, Q^{V}\right),\left(f\left(Q^{E}\right), f\left(Q^{V}\right)\right)$ is a matching pair of answers. The following simple theorem, whose proof appears in Section A of the Appendix, shows that in the Yes-Instance case, there are many perfect global assignments, that neatly partition all answers to the edge-queries.

Theorem 2.6. Assume that $G$ is a Yes-Instance. Then there are $6^{\ell}$ perfect global assignments $f_{1}, \ldots, f_{6^{\ell}}$ of answers to queries, such that:

- for each query $Q \in Q^{E}$ to the edge-prover, for each possible answer $A \in \mathcal{A}^{E}$, there is exactly one index $1 \leq i \leq 6^{\ell}$ with $f_{i}(Q)=A$; and
- for each query $Q^{\prime} \in Q^{V}$ to the vertex-prover, for each possible answer $A^{\prime} \in \mathcal{A}^{V}$ to $Q^{\prime}$, there are exactly $2^{\ell}$ indices $1 \leq i \leq 6^{\ell}$, for which $f_{i}\left(Q^{\prime}\right)=A^{\prime}$.


## Almost Polynomial Hardness of Node-Disjoint Paths in Grids

Constraint graph $H$. Given a 3COL(5) instance $G$ with $|V(G)|=n$, and an integer $\ell>0$, we associate a graph $H$, that we call the constraint graph, with it. For every query $Q \in Q^{E} \cup Q^{V}$, there is a vertex $v(Q)$ in $H$, while for each random string $R=\left(Q, Q^{\prime}\right)$, there is an edge $e(R)=\left(v(Q), v\left(Q^{\prime}\right)\right)$. Notice that $H$ is a bipartite graph. We denote by $U^{E}$ the set of its vertices corresponding to the edge-queries, and by $U^{V}$ the set of its vertices corresponding to the vertex-queries. Recall that $\left|U^{E}\right|=(5 n / 2)^{\ell},\left|U^{V}\right|=n^{\ell}$; the degree of every vertex in $U^{E}$ is $2^{\ell}$; the degree of every vertex in $U^{V}$ is $5^{\ell}$, and $|E(H)|=|\mathcal{R}|=(5 n)^{\ell}$.

Assignment graphs $L\left(H^{\prime}\right)$. Assume now that we are given some subgraph $H^{\prime} \subseteq H$ of the constraint graph. We build a bipartite graph $L\left(H^{\prime}\right)$ associated with it (this graph takes into account the answers to the queries; it may be convenient for now to think that $H^{\prime}=H$, but later we will use smaller subgraphs of $H)$. The vertices of $L\left(H^{\prime}\right)$ are partitioned into two subsets:

- For each edge-query $Q \in Q^{E}$ with $v(Q) \in H^{\prime}$, for each possible answer $A \in \mathcal{A}^{E}$ to $Q$, we introduce a vertex $v(Q, A)$. We denote by $S(Q)$ the set of these $6^{\ell}$ vertices corresponding to $Q$, and we call them a group representing $Q$. We denote by $\hat{U}^{E}$ the resulting set of vertices:

$$
\hat{U}^{E}=\left\{v(Q, A) \mid\left(Q \in Q^{E} \text { and } v(Q) \in H^{\prime}\right), A \in \mathcal{A}^{E}\right\} .
$$

- For each vertex-query $Q^{\prime} \in Q^{V}$ with $v\left(Q^{\prime}\right) \in H^{\prime}$, for each possible answer $A^{\prime} \in \mathcal{A}^{V}$ to $Q^{\prime}$, we introduce $2^{\ell}$ vertices $v_{1}\left(Q^{\prime}, A^{\prime}\right), \ldots, v_{2^{\ell}}\left(Q^{\prime}, A^{\prime}\right)$. We call all these vertices the copies of answer $A^{\prime}$ to query $Q^{\prime}$. We denote by $S\left(Q^{\prime}\right)$ the set of all vertices corresponding to $Q^{\prime}$ :

$$
S\left(Q^{\prime}\right)=\left\{v_{i}\left(Q^{\prime}, A^{\prime}\right) \mid A^{\prime} \in \mathcal{A}^{V}, 1 \leq i \leq 2^{\ell}\right\},
$$

so $\left|S\left(Q^{\prime}\right)\right|=6^{\ell}$. We call $S\left(Q^{\prime}\right)$ the group representing $Q^{\prime}$. We denote by $\hat{U}^{V}$ the resulting set of vertices:

$$
\hat{U}^{V}=\left\{v_{i}\left(Q^{\prime}, A^{\prime}\right) \mid\left(Q^{\prime} \in Q^{V} \text { and } v\left(Q^{\prime}\right) \in H^{\prime}\right), A^{\prime} \in \mathcal{A}^{V}, 1 \leq i \leq 2^{\ell}\right\} .
$$

The final set of vertices of $L\left(H^{\prime}\right)$ is $\hat{U}^{E} \cup \hat{U}^{V}$. We define the set of edges of $L\left(H^{\prime}\right)$ as follows. For each random string $R=\left(Q^{E}, Q^{V}\right)$ whose corresponding edge $e(R)$ belongs to $H^{\prime}$, for every answer $A \in \mathcal{A}^{E}$ to $Q^{E}$, let $A^{\prime} \in \mathcal{A}^{V}$ be the unique answer to $Q^{V}$ consistent with $A$. For each copy $v_{i}\left(Q^{V}, A^{\prime}\right)$ of answer $A^{\prime}$ to query $Q^{V}$, we add an edge $\left(v\left(Q^{E}, A\right), v_{i}\left(Q^{V}, A^{\prime}\right)\right)$. Let

$$
E(R)=\left\{\left(v\left(Q^{E}, A\right), v_{i}\left(Q^{V}, A^{\prime}\right)\right) \mid A \in \mathcal{A}^{E}, A^{\prime} \in \mathcal{A}^{V}, A \text { and } A^{\prime} \text { are consistent answers to } R, 1 \leq i \leq 2^{\ell}\right\}
$$

be the set of the resulting edges, so $|E(R)|=6^{\ell} \cdot 2^{\ell}=12^{\ell}$. We denote by $\hat{E}$ the set of all edges of $L\left(H^{\prime}\right)$ the union of the sets $E(R)$ for all random strings $R$ with $e(R) \in H^{\prime}$.

Recall that we have defined a partition of the set $\hat{U}^{E}$ of vertices into groups $S(Q)$ - one group for each query $Q \in Q^{E}$ with $v(Q) \in H^{\prime}$. We denote this partition by $U_{1}$. Similarly, we have defined a partition of $\hat{U}^{V}$ into groups, that we denote by $\mathcal{U}_{2}=\left\{S\left(Q^{\prime}\right) \mid Q^{\prime} \in Q^{V}\right.$ and $\left.v\left(Q^{\prime}\right) \in H^{\prime}\right\}$. Recall that for each group $U \in \mathcal{U}_{1} \cup \mathcal{U}_{2},|U|=6^{\ell}$.

Finally, we need to define bundles of edges in graph $L\left(H^{\prime}\right)$. For every vertex $v \in \hat{U}^{E} \cup \hat{U}^{V}$, we define a partition $\mathcal{B}(v)$ of the set of all edges incident to $v$ in $L\left(H^{\prime}\right)$ into bundles, as follows. Fix some group $U \in \mathcal{U}_{1} \cup \mathcal{U}_{2}$ that we have defined. If there is at least one edge of $L\left(H^{\prime}\right)$ connecting $v$ to the vertices of $U$, then we define a bundle containing all edges connecting $v$ to the vertices of $U$, and add this bundle to $\mathcal{B}(v)$. Therefore, if $v \in S(Q)$, then for each random string $R$ in which $Q$ participates, with $e(R) \in H^{\prime}$, we have defined one bundle of edges in $\mathcal{B}(v)$. For each vertex $v \in \hat{U}^{E} \cup \hat{U}^{V}$, the set of all edges incident to $v$ is thus partitioned into a collection of bundles, that we denote by $\mathcal{B}(v)$, and we denote $\beta(v)=|\mathcal{B}(v)|$. Note that, if $v \in S(Q)$ for some query $Q \in Q^{E} \cup Q^{V}$, then $\beta(v)$ is exactly the degree of the vertex $v(Q)$ in graph $H^{\prime}$. Note also that $\bigcup_{v \in V\left(H^{\prime}\right)} \mathcal{B}(v)$ does not define a partition of the edges of $\hat{E}$, as each such edge belongs to two bundles. However, each of $\bigcup_{v \in \hat{U}^{E}} \mathcal{B}(v)$ and $\bigcup_{v \in \hat{U}^{V}} \mathcal{B}(v)$ does define a partition of $\hat{E}$. It is easy to verify that every bundle that we have defined contains exactly $2^{\ell}$ edges. We use this fact later.

## 3 The ( $r, h$ )-Graph Partitioning problem

We will use a graph partitioning problem as a proxy in order to reduce the $3 \mathrm{COL}(5)$ problem to NDP-Grid. The specific graph partitioning problem is somewhat complex. We first define a simpler variant of this problem, and then provide the intuition and the motivation for the more complex variant that we eventually use.
In the basic $(r, h)$-Graph Partitioning problem, that we denote by $(r, h)$-GP, we are given a bipartite graph $\tilde{G}=\left(V_{1}, V_{2}, E\right)$, and two integral parameters $h, r>0$. A solution consists of a partition $\left(W_{1}, \ldots, W_{r}\right)$ of $V_{1} \cup V_{2}$ into $r$ subsets, and for each $1 \leq i \leq r$, a subset $E_{i} \subseteq E\left(W_{i}\right)$ of edges, such that $\left|E_{i}\right| \leq h$. The goal is to maximize $\sum_{i}\left|E_{i}\right|$.
One intuitive way to think about the ( $\mathrm{r}, \mathrm{h}$ )-GP problem is that we would like to partition the vertices of $\tilde{G}$ into $r$ clusters, that are roughly balanced (in terms of the number of edges in each cluster). However, unlike the standard balanced partitioning problems, that attempt to minimize the number of edges connecting the different clusters, our goal is to maximize the total number of edges that remain in the clusters. We suspect that the $(r, h)$-GP problem is very hard to approximate; in particular it appears to be somewhat similar to the Densest $k$-Subgraph problem (DkS). Like in the DkS problem, we are looking for dense subgraphs of $\tilde{G}$ (the subgraphs $\tilde{G}\left[W_{i}\right]$ ), but unlike DkS, where we only need to find one such dense subgraph, we would like to partition all vertices of $\tilde{G}$ into a prescribed number of dense subgraphs. We can prove that NDP-Grid is at least as hard as ( $r, h$ )-GP (to within polylogarithmic factors; see below), but unfortunately we could not prove strong hardness of approximation results for ( $r, h$ )-GP. In particular, known hardness proofs for DkS do not seem to generalize to this problem. To overcome this difficulty, we define a slightly more general problem, and then use it as a proxy in our reduction. Before defining the more general problem, we start with intuition.

## Almost Polynomial Hardness of Node-Disjoint Paths in Grids

Intuition: Given a 3COL(5) instance $G$, we can construct the constraint graph $H$, and the assignment graph $L(H)$, as described in Section 2. We can then view $L(H)$ as an instance of $(r, h)-G P$, with $r=6^{\ell}$ and $h=|\mathcal{R}|$. Assume that $G$ is a Yes-Instance. Then we can use the perfect global assignments $f_{1}, \ldots, f_{r}$ of answers to the queries, given by Theorem 2.6, in order to partition the vertices of $L(H)$ into $r=6^{\ell}$ clusters $W_{1}, \ldots, W_{r}$, as follows. Fix some $1 \leq i \leq r$. For each query $Q \in Q^{E}$ to the edge-prover, set $W_{i}$ contains a single vertex $v(Q, A) \in S(Q)$, where $A=f_{i}(Q)$. For each query $Q^{\prime} \in Q^{V}$ to the vertex-prover, set $W_{i}$ contains a single vertex $v_{j}\left(Q^{\prime}, A^{\prime}\right)$, where $A^{\prime}=f_{i}\left(Q^{\prime}\right)$, and the indices $j$ are chosen so that every vertex $v_{j}\left(Q^{\prime}, A^{\prime}\right)$ participates in exactly one cluster $W_{i}$. From the construction of the assignment graph $L(H)$ and the properties of the assignments $f_{i}$ guaranteed by Theorem 2.6 , we indeed obtain a partition $W_{1}, \ldots, W_{r}$ of the vertices of $L(H)$. For each $1 \leq i \leq r$, we then set $E_{i}=E\left(W_{i}\right)$. Notice that for every query $Q \in Q^{E} \cup Q^{V}$, exactly one vertex of $S(Q)$ participates in each cluster $W_{i}$. Therefore, for each group $U \in \mathcal{U}_{1} \cup \mathcal{U}_{2}$, each cluster $W_{i}$ contains exactly one vertex from this group. It is easy to verify that for each $1 \leq i \leq r$, for each random string $R \in \mathcal{R}$, set $E_{i}$ contains exactly one edge of $E(R)$, and so $\left|E_{i}\right|=|\mathcal{R}|=h$, and the solution value is $h \cdot r$. Unfortunately, in the No-InStance, we may still obtain a solution of a high value, as follows: instead of distributing, for each query $Q \in Q^{E} \cup Q^{V}$, the vertices of $S(Q)$ to different clusters $W_{i}$, we may put all vertices of $S(Q)$ into a single cluster. While in our intended solution to the ( $\mathrm{r}, \mathrm{h}$ )-GP problem instance each cluster can be interpreted as an assignment of answers to the queries, and the number of edges in each cluster is bounded by the number of random strings satisfied by this assignment, we may no longer use this interpretation with this new type of solutions ${ }^{3}$. Moreover, unlike in the Yes-Instance solutions, if we now consider some cluster $W_{i}$, and some random string $R \in \mathcal{R}$, we may add several edges of $E(R)$ to $E_{i}$, which will further allow us to accumulate a high solution value. One way to get around this problem is to impose additional restrictions on the feasible solutions to the ( $\mathrm{r}, \mathrm{h}$ )-GP problem, which are consistent with our Yes-Instance solution, and thereby obtain a more general (and hopefully more difficult) problem. But while doing so we still need to ensure that we can prove that NDP-Grid remains at least as hard as the newly defined problem. Recall the definition of bundles in graph $L(H)$. It is easy to verify that in our intended solution to the Yes-Instance, every bundle contributes at most one edge to the solution. This motivates our definition of a slight generalization of the $(r, h)$-GP problem, that we call $(r, h)$-Graph Partitioning with Bundles, or $(r, h)-G P w B$.
The input to (r,h)-GPwB problem is almost the same as before: we are given a bipartite graph $\tilde{G}=$ $\left(V_{1}, V_{2}, E\right)$, and two integral parameters $h, r>0$. Additionally, we are given a partition $\mathcal{U}_{1}$ of $V_{1}$ into groups, and a partition $\mathcal{U}_{2}$ of $V_{2}$ into groups, so that for each group $U \in \mathcal{U}_{1} \cup \mathcal{U}_{2},|U|=r$. Using these groups, we define bundles of edges as follows: for every vertex $v \in V_{1}$, for each group $U \in \mathcal{U}_{2}$, such that some edge of $E$ connects $v$ to a vertex of $U$, the set of all edges that connect $v$ to the vertices of $U$ defines a single bundle. Similarly, for every vertex $v \in V_{2}$, for each group $U \in \mathcal{U}_{1}$, such that some edge of $E$ connects $v$ to a vertex of $U$, all edges that connect $v$ to the vertices of $U$ define a bundle. We denote, for each vertex $v \in V_{1} \cup V_{2}$, by $\mathcal{B}(v)$ the set of all bundles into which the edges incident to $v$ are partitioned, and we denote by $\beta(v)=|\mathcal{B}(v)|$ the number of such bundles. We also denote by $\mathcal{B}=\bigcup_{v \in V_{1} \cup V_{2}} \mathcal{B}(v)$ - the set of all bundles. Note that as before, $\mathcal{B}$ is not a partition of $E$, but every edge of $E$ belongs to exactly two bundles: one bundle in $\bigcup_{v \in V_{1}} \mathcal{B}(v)$, and one bundle in $\bigcup_{v \in V_{2}} \mathcal{B}(v)$. As before, we need to compute a partition $\left(W_{1}, \ldots, W_{r}\right)$ of $V_{1} \cup V_{2}$ into $r$ subsets, and for each $1 \leq i \leq r$, select a subset $E_{i} \subseteq E\left(W_{i}\right)$ of

[^2]Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat

edges, such that $\left|E_{i}\right| \leq h$. But now there is an additional restriction: we require that for each $1 \leq i \leq r$, for every bundle $B \in \mathcal{B}, E_{i}$ contains at most one edge $e \in B$. As before, the goal is to maximize $\sum_{i}\left|E_{i}\right|$.

Valid instances, regular instances, and perfect solutions. Given an instance $\mathcal{J}=\left(\tilde{G}=\left(V_{1}, V_{2}, E\right), \mathcal{U}_{1}, \mathcal{U}_{2}, h, r\right)$ of $(r, h)$-GPwB, let $\beta^{*}(\mathcal{J})=\sum_{v \in V_{1}} \beta(v)$. Note that for any solution to $\mathcal{J}$, the solution value must be bounded by $\beta^{*}(\mathcal{J})$, since for every vertex $v \in V_{1}$, for every bundle $B \in \mathcal{B}(v)$, at most one edge from the bundle may contribute to the solution value. In all instances of $(r, h)$-GPwB that we consider, we always set $h=\beta^{*}(\mathcal{J}) / r$. Next, we define valid instances; they are defined so that the instances that we obtain when reducing from $3 \mathrm{COL}(5)$ are always valid, as we show later.

Definition 3.1. We say that instance $\mathcal{J}$ of $(r, h)-G P w B$ is valid iff $h=\beta^{*}(\mathcal{J}) / r$ and $h \geq \max _{v \in V_{1} \cup V_{2}}\{\beta(v)\}$.

Recall that for every group $U \in \mathcal{U}_{1} \cup \mathcal{U}_{2},|U|=r$. We now define perfect solutions to the ( $r, h$ )-GPwB problem. We will ensure that our intended solutions in the Yes-Instance are always perfect, as we show later.

Definition 3.2. We say that a solution $\left(\left(W_{1}, \ldots, W_{r}\right),\left(E_{1}, \ldots, E_{r}\right)\right)$ to a valid $(\mathrm{r}, \mathrm{h})$-GPwB instance $\mathcal{J}$ is perfect iff:

- For each group $U \in \mathcal{U}_{1} \cup \mathcal{U}_{2}$, exactly one vertex of $U$ belongs to each cluster $W_{i}$; and
- For each $1 \leq i \leq r,\left|E_{i}\right|=h$.

Note that the value of a perfect solution to a valid instance $\mathcal{J}$ is $h \cdot r=\beta^{*}(\mathcal{J})$, and this is the largest value that any solution can achieve.
Lastly, we say that an instance $\mathcal{J}=\left(\tilde{G}=\left(V_{1}, V_{2}, E\right), \mathcal{U}_{1}, \mathcal{U}_{2}, h, r\right)$ of the $(r, h)$-GPwB problem is $\left(d_{1}, d_{2}, b\right)$ regular iff the graph $\tilde{G}$ is $\left(d_{1}, d_{2}\right)$-semiregular, and for every vertex $v \in V_{1} \cup V_{2}$, every bundle $B \in \mathcal{B}(v)$ has cardinality exactly $b$.

## 4 Tools for the hardness proof

In this section we provide two basic tools that will be used in the hardness of approximation proof: a reduction from the $3 \mathrm{COL}(5)$ problem to the $(r, h)$-GPwB problem, and a reduction from the $(r, h)$-GPwB problem to the NDP-Grid problem. We also establish some properties of these reductions that will be used later.

### 4.1 From 3COL(5) to ( $\mathrm{r}, \mathrm{h}$ )-GPwB

Suppose we are given an instance $G$ of the 3COL(5) problem, and an integral parameter $\ell>0$ (the number of repetitions). Consider the corresponding constraint graph $H$, and suppose we are given some
connected subgraph $H^{\prime} \subseteq H$. We define an instance $\mathcal{J}\left(H^{\prime}\right)$ of $(r, h)$-GPwB, as follows. First, we use Claim 2.2 in order to compute a semiregular subgraph $H^{\prime \prime} \subseteq H^{\prime}$, with $\left|E\left(H^{\prime \prime}\right)\right| \geq \Omega\left(\left|E\left(H^{\prime}\right)\right| / \log ^{2} \mid E\left(H^{\prime} \mid\right)\right.$. In our reduction, we will work with the graph $H^{\prime \prime}$ instead of the graph $H^{\prime}$ from now on. We now define the instance $\mathcal{J}\left(H^{\prime}\right)$ of ( $\mathrm{r}, \mathrm{h}$ )-GPwB.

- The underlying graph is $L\left(H^{\prime \prime}\right)=\left(\hat{U}^{E}, \hat{U}^{V}, \hat{E}\right)$;
- The parameters are $r=6^{\ell}$ and $h=\left|E\left(H^{\prime \prime}\right)\right|$;
- The partition $\mathcal{U}_{1}$ of $\hat{U}^{E}$ is the same as before: the vertices of $\hat{U}^{E}$ are partitioned into groups $S(Q)$ - one group for each query $Q \in Q^{E}$ with $v(Q) \in V\left(H^{\prime \prime}\right)$. Similarly, the partition $\mathcal{U}_{2}$ of $\hat{U}^{V}$ into groups is also defined exactly as before, and contains, for each query $Q^{\prime} \in Q^{V}$ with $v\left(Q^{\prime}\right) \in V\left(H^{\prime \prime}\right)$, a group $S\left(Q^{\prime}\right)$. (Recall that for all $Q \in Q^{E} \cup Q^{V}$ with $v(Q) \in H^{\prime},|S(Q)|=6^{\ell}$ ).

Claim 4.1. Let $G$ be an instance of the $3 C O L(5)$ problem, $\ell>0$ an integral parallel repetition parameter, and $H^{\prime} \subseteq H$ a connected subgraph of the corresponding constraint graph. Consider the corresponding instance $\mathcal{J}\left(H^{\prime}\right)$ of $(r, \mathrm{~h})-\mathrm{GPwB}$. Then $\mathcal{J}\left(H^{\prime}\right)$ is a valid instance, it is a $\left(d^{\prime}, d^{\prime \prime}, b\right)$-regular instance for some $d^{\prime}, d^{\prime \prime}>0$ and $b=2^{\ell}$, and moreover, if $G$ is a Yes-Instance, then there is a perfect solution to $\mathcal{J}\left(H^{\prime}\right)$.

Proof. We first verify that $\mathcal{J}\left(H^{\prime}\right)$ is a valid instance of $(r, h)$-GPwB. Recall that for a query $Q \in Q^{E}$ to the edge-prover and an answer $A \in \mathcal{A}^{E}$, the number of bundles incident to vertex $v(Q, A)$ in $L\left(H^{\prime}\right)$ is exactly the degree of the vertex $v(Q)$ in graph $H^{\prime \prime}$. The total number of bundles incident to the vertices of $S(Q)$ is then the degree of $v(Q)$ in $H^{\prime \prime}$ times $\left|\mathcal{A}^{E}\right|$. Therefore, $\beta^{*}(\mathcal{J})=\sum_{v(Q, A) \in \hat{U}^{E}}|\beta(v)|=\left|E\left(H^{\prime \prime}\right)\right| \cdot\left|\mathcal{A}^{E}\right|=$ $h \cdot r$. It is now immediate to verify that $h=\beta^{*}(\mathcal{J}) / r$. Similarly, for a vertex $v=v_{j}\left(Q^{\prime}, A^{\prime}\right) \in U^{V}$, the number of bundles incident to $v$ is exactly the degree of $v\left(Q^{\prime}\right)$ in $H^{\prime \prime}$. Since $h=\left|E\left(H^{\prime \prime}\right)\right|$, we get that $h \geq \max _{v \in V_{1} \cup V_{2}}\{\beta(v)\}$, and so $\mathcal{J}\left(H^{\prime}\right)$ is a valid instance.
Next, we show that $\mathcal{J}\left(H^{\prime}\right)$ is a $\left(d^{\prime}, d^{\prime \prime}, b\right)$-regular instance, for some $d^{\prime}, d^{\prime \prime}>0$ and $b=2^{\ell}$. Denote $\mathcal{J}\left(H^{\prime}\right)=\left(\tilde{G}=\left(V_{1}, V_{2}, E\right), \mathcal{U}_{1}, \mathcal{U}_{2}, h, r\right)$. Recall that we have defined a graph $H^{\prime \prime} \subseteq H^{\prime}$, that is $\left(d_{1}, d_{2}\right)$ semiregular, for some $d_{1}, d_{2}>0$. Recall also that, from the discussion at the end of Section 2, for every vertex $v \in V_{1} \cup V_{2}$, every bundle $B \in \mathcal{B}(v)$ contains exactly $2^{\ell}$ edges. Moreover, if $v \in S(Q)$ for any query $Q$, then the number of its bundles $\beta(v)$ is exactly the degree of the vertex $v(Q)$ in graph $H^{\prime \prime}$. Therefore, $d_{\tilde{G}}(v)=2^{\ell} \cdot d_{H^{\prime \prime}}(v(Q))$. It follows that for every vertex $v \in V_{1}, d_{1} \cdot 2^{\ell} \leq d_{\tilde{G}}(v)<2 d_{1} \cdot 2^{\ell}$, and for every vertex $v^{\prime} \in V_{2}, d_{\tilde{G}}\left(v^{\prime}\right)<2 d_{2} \cdot 2^{\ell}$. Moreover, $|E(\tilde{G})|=12^{\ell} \cdot\left|E\left(H^{\prime \prime}\right)\right|$.
Recall that $V\left(H^{\prime}\right)=U^{E} \cup U^{V}$, where the vertices of $U^{E}$ represent the edge-queries, and the vertices of $U^{V}$ represent the vertex-queries. Since, from the definition of semiregular graphs, we are guaranteed that $\left|E\left(H^{\prime \prime}\right)\right| \geq d_{2}\left|U^{V} \cap V\left(H^{\prime \prime}\right)\right| /\left(4 \log \left|U^{V} \cap V\left(H^{\prime \prime}\right)\right|\right)$, and $\left|V_{2}\right|=6^{\ell} \cdot\left|U^{V} \cap V\left(H^{\prime \prime}\right)\right|$, we get that:

$$
\begin{aligned}
|E(\tilde{G})| & =12^{\ell} \cdot\left|E\left(H^{\prime \prime}\right)\right| \\
& \geq \frac{12^{\ell} \cdot d_{2}\left|U^{V} \cap V\left(H^{\prime \prime}\right)\right|}{4 \log \left|U^{V} \cap V\left(H^{\prime \prime}\right)\right|} \\
& \geq \frac{\left(d_{2} \cdot 2^{\ell}\right) \cdot\left|V_{2}\right|}{4 \log \left|U^{V} \cap V\left(H^{\prime \prime}\right)\right|} \\
& \geq \frac{\left(d_{2} \cdot 2^{\ell}\right) \cdot\left|V_{2}\right|}{4 \log \left|V_{2}\right|}
\end{aligned}
$$

We conclude that instance $\mathcal{J}\left(H^{\prime}\right)$ is $\left(d^{\prime}, d^{\prime \prime}, b\right)$-regular, for $d^{\prime}=2^{\ell} \cdot d_{1}, d^{\prime \prime}=2^{\ell} \cdot d_{2}$, and $b=2^{\ell}$.
Assume now that $G$ is a Yes-Instance. We define a perfect solution $\left(\left(W_{1}, \ldots, W_{r}\right),\left(E_{1}, \ldots, E_{r}\right)\right)$ to instance $\mathcal{J}\left(H^{\prime}\right)$ of the $(r, h)$-GPwB problem. Let $\left\{f_{1}, f_{2}, \ldots, f_{6^{\ell}}\right\}$ be the collection of perfect global assignments of answers to the queries, given by Theorem 2.6. Recall that $\mathcal{U}_{1}=\left\{S(Q) \mid Q \in Q^{E}, v(Q) \in H^{\prime \prime}\right\}$ and $\mathcal{U}_{2}=\left\{S\left(Q^{\prime}\right) \mid Q^{\prime} \in Q^{V}, v\left(Q^{\prime}\right) \in H^{\prime \prime}\right\}$, where each group in $\mathcal{U}_{1} \cup \mathcal{U}_{2}$ has cardinality $r=6^{\ell}$. We now fix some $1 \leq i \leq r$, and define the set $W_{i}$ of vertices. For each query $Q \in Q^{E}$ to the edge-prover with $v(Q) \in V\left(H^{\prime \prime}\right)$, if $A=f_{i}(Q)$, then we add the vertex $v(Q, A)$ to $W_{i}$. For each query $Q^{\prime} \in Q^{V}$ to the vertex-prover with $v\left(Q^{\prime}\right) \in V\left(H^{\prime \prime}\right)$, if $A^{\prime}=f_{i}\left(Q^{\prime}\right)$, then we select some index $1 \leq j \leq 2^{\ell}$, and add the vertex $v_{j}\left(Q^{\prime}, A^{\prime}\right)$ to $W_{i}$. The indices $j$ are chosen so that every vertex $v_{j}\left(Q^{\prime}, A^{\prime}\right)$ participates in at most one cluster $W_{i}$. From the construction of the graph $L\left(H^{\prime \prime}\right)$ and the properties of the assignments $f_{i}$ guaranteed by Theorem 2.6 , it is easy to verify that $W_{1}, \ldots, W_{r}$ partition the vertices of $L\left(H^{\prime \prime}\right)$, and moreover, for each group $S(Q) \in \mathcal{U}_{1} \cup \mathcal{U}_{2}$, each set $W_{i}$ contains exactly one vertex of $S(Q)$.
Finally, for each $1 \leq i \leq r$, we set $E_{i}=E\left(W_{i}\right)$. We claim that for each bundle $B \in \mathcal{B}$, set $E_{i}$ may contain at most one edge of $B$. Indeed, let $v \in W_{i}$ be some vertex, let $U \in \mathcal{U}_{1} \cup \mathcal{U}_{2}$ be some group, and let $B$ be the bundle containing all edges that connect $v$ to the vertices of $U$. Since $W_{i}$ contains exactly one vertex of $U$, at most one edge of $B$ may belong to $E_{i}$.

It now remains to show that $\left|E_{i}\right|=h$ for all $i$. Fix some $1 \leq i \leq r$. It is easy to verify that for each random string $R=\left(Q, Q^{\prime}\right)$ with $e(R) \in H^{\prime \prime}$, set $W_{i}$ contains a pair of vertices $v(Q, A), v_{j}\left(Q^{\prime}, A^{\prime}\right)$, where $A$ and $A^{\prime}$ are matching answers to $Q$ and $Q^{\prime}$ respectively, and so the corresponding edge connecting this pair of vertices in $L\left(H^{\prime \prime}\right)$ belongs to $E_{i}$. Therefore, $\left|E_{i}\right|=\left|E\left(H^{\prime \prime}\right)\right|=h$.

Claim 4.1 shows that, if an instance $G$ of the 3COL(5) problem is a Yes-Instance, then the corresponding instance $\mathcal{J}(H)$ of ( $r, h$ )-GPwB has a solution of a high value (a perfect solution). Ideally, we would like to show that, if $G$ is a No-Instance, then any solution to the corresponding $(r, h)$-GPwB instance has a low value. Unfortunately, we cannot do this directly, and this is the reason that our hardness of approximation proof, provided in Section 5, employs a Cook reduction and not a Karp reduction. As we will show in Section 5, if instance $G$ of 3COL(5) is a No-InSTANCE, but the corresponding instance of $(r, h)$-GPwB has a high solution value, then we can exploit the solution to the $(r, h)$-GPwB problem instance in order to compute a partition of the constraint graph $H$ into smaller subgraphs, and then employ the same reduction on each one of the resulting subgraphs; see Section 5 for more details.

### 4.2 From (r,h)-GPwB to NDP

The following definition will be useful for us later, when we extend our results to NDP and EDP on wall graphs.

Definition 4.2. Let $\mathcal{P}$ be a set of paths in a grid $\hat{G}$. We say that $\mathcal{P}$ is a spaced-out set iff for each pair $P, P^{\prime} \in \mathcal{P}$ of paths, $d\left(V(P), V\left(P^{\prime}\right)\right) \geq 2$, and all paths in $\mathcal{P}$ are internally disjoint from the boundaries of the grid $\hat{G}$.

Note that if $\mathcal{P}$ is a set of paths that is spaced-out, then all paths in $\mathcal{P}$ are mutually node-disjoint. The following theorem is central to our hardness of approximation proof.

Theorem 4.3. There is a constant $c^{*}>0$, and there is an algorithm, that, given a valid and a $\left(d, d^{\prime}, b\right)$ regular instance $\mathcal{J}=\left(\tilde{G}, \mathcal{U}_{1}, \mathcal{U}_{2}, h, r\right)$ of $(\mathrm{r}, \mathrm{h})-\mathrm{GPwB}$, for some $d, d^{\prime}, b>0$, with $|V(\tilde{G})|=N$ and $|E(\tilde{G})|=$ $M$, in time poly $(N M)$ constructs an instance $\hat{\mathcal{J}}=(\hat{G}, \mathcal{M})$ of NDP-Grid with $|V(\hat{G})|=O\left(M^{4} \log ^{2} M\right)$, such that the following hold:

- If $\mathcal{J}$ has a perfect solution (of value $\beta^{*}=\beta^{*}(\mathcal{J})$ ), then instance $\hat{\mathcal{J}}$ has a solution $\mathcal{P}$ that routes at least $\frac{\beta^{*}}{c^{*} \log ^{2} M}$ demand pairs, such that the paths in $\mathcal{P}$ are spaced-out; and
- There is an algorithm with running time poly $(N M)$, that, given a solution $\mathcal{P}^{*}$ to the NDP-Grid problem instance $\hat{\mathcal{J}}$, constructs a solution to the $(\mathrm{r}, \mathrm{h})-\mathrm{GPwB}$ instance $\mathcal{J}$, of value at least $\frac{\left|\mathcal{P}^{*}\right|}{\mathcal{C}^{*} \log ^{3} M}$.

We note that the theorem is slightly stronger than what is needed in order to prove hardness of approximation of NDP-Grid: if $\mathcal{J}$ has a perfect solution, then it is sufficient to ensure that the set $\mathcal{P}$ of paths in the corresponding NDP-Grid instance $\hat{\mathcal{J}}$ is node-disjoint. But we will use the stronger guarantee that it is spaced-out when extending our results to NDP and EDP in wall graphs. Note that in the second assertion we are only guaranteed that the paths in $\mathcal{P}^{*}$ are node-disjoint. The proof of the theorem is somewhat technical and is deferred to Section 6.

### 4.3 Combining the tools for the Yes-Instance

Assume now that we are given an instance $G$ of $3 \mathrm{COL}(5)$, an integral parallel repetitions parameter $\ell>0$, and a connected subgraph $H^{\prime} \subseteq H$ of the corresponding constraint graph. Recall that we have constructed a corresponding instance $\mathcal{J}\left(H^{\prime}\right)$ of $(r, h)$-GPwB that is valid and $\left(d, d^{\prime}, b\right)$-regular for some $d, d^{\prime}, b>0$. We can then use Theorem 4.3 to construct an instance of NDP-Grid, that we denote by $\hat{\mathcal{J}}\left(H^{\prime}\right)$. Note that $\left|E\left(L\left(H^{\prime \prime}\right)\right)\right| \leq 2^{O(\ell)} \cdot|E(H)| \leq n^{O(\ell)}$. Let $\hat{c}$ be a constant, such that $\left|E\left(L\left(H^{\prime \prime}\right)\right)\right| \leq n^{\hat{c} \ell}$ for any connected subgraph $H^{\prime} \subseteq H$; we can assume w.l.o.g. that $\hat{c}>1$. We can also assume w.l.o.g. that $c^{*} \geq 1$, where $c^{*}$ is the constant from Theorem 4.3, and we assume that the algorithm from Claim 2.2 returns a subgraph $G^{\prime}=\left(A^{\prime}, B^{\prime}, E^{\prime}\right)$ of $G$ with $\left|E^{\prime}\right| \geq|E(G)| /\left(c^{*} \log ^{2}|E(G)|\right)$, if $G$ is a connected graph. Lastly, we denote $c_{\mathrm{YI}}=\left(\hat{c} \cdot c^{*}\right)^{4}$. We obtain the following immediate corollary of Theorem 4.3:

Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat

Corollary 4.4. Suppose we are given 3COL(5) instance $G$ that is a Yes-Instance, an integral parallel repetition parameter $\ell>0$, and a connected subgraph $H^{\prime} \subseteq H$ of the corresponding constraint graph. Then instance $\hat{\mathcal{J}}\left(H^{\prime}\right)$ of NDP-Grid has a solution of value at least $\frac{\mid E\left(H^{\prime}\right) \cdot \cdot^{\ell}}{c_{1} \ell^{4} \log ^{4} n}$, where $n=|V(G)|$.

Proof. From Claim 4.1, instance $\mathcal{J}\left(H^{\prime}\right)=\left(L\left(H^{\prime \prime}\right), \mathcal{U}_{1}, \mathcal{U}_{2}, r, h\right)$ of $(r, h)$-GPwB is a valid instance, it is a $\left(d^{\prime}, d^{\prime \prime}, b\right)$-regular instance for some $d^{\prime}, d^{\prime \prime}, b>0$, and it has a perfect solution, whose value is $\beta^{*}=\beta^{*}\left(\mathcal{J}\left(H^{\prime \prime}\right)\right)=h \cdot r=\left|E\left(H^{\prime \prime}\right)\right| \cdot 6^{\ell}$. Recall that $H^{\prime \prime}$ is a semiregularized graph, that was obtained from $H^{\prime}$ by applying Claim 2.2 to it, so $\left|E\left(H^{\prime \prime}\right)\right| \geq\left|E\left(H^{\prime}\right)\right| /\left(c^{*} \log ^{2}\left|E\left(H^{\prime}\right)\right|\right)$. From Theorem 4.3, instance $\hat{\mathcal{J}}\left(H^{\prime}\right)$ of NDP-Grid has a solution of value at least $\frac{\left|E\left(H^{\prime \prime}\right)\right| \cdot 6^{\ell}}{c^{*} \log ^{2}\left|E\left(H^{\prime \prime}\right)\right|} \geq \frac{\left|E\left(H^{\prime}\right)\right| \cdot 6^{\ell}}{\left(c^{*}\right)^{2} \log ^{4} M}$, where $M=\left|E\left(L\left(H^{\prime \prime}\right)\right)\right|$. Since $\log M \leq \hat{c} \ell \log n$, the corollary follows.

As already mentioned before, if $G$ is a No-InSTANCE, it is still possible that the corresponding instance of ( $\mathrm{r}, \mathrm{h}$ )-GPwB, and hence the resulting instance of NDP-Grid, will have a high solution value. We get around this problem by employing a Cook reduction, as shown in Section 5.

## 5 The hardness proof

Let $G$ be an input instance of 3COL(5). Recall that $\gamma$ is the absolute constant from the Parallel Repetition Theorem (Corollary 2.5). We will set the value of the parallel repetition parameter $\ell$ later, ensuring that $\ell>\log ^{2} n$, where $n=|V(G)|$. Let $\alpha^{*}=2^{\Theta(\ell / \log n)}$ be the hardness of approximation factor that we are trying to achieve.
Given the tools developed in the previous sections, a standard way to prove hardness of NDP-Grid would work as follows. Given an instance $G$ of $3 \mathrm{COL}(5)$ and the chosen parameter $\ell$, construct the corresponding graph $H$ (the constraint graph), together with the assignment graph $L(H)$. We then construct an instance $\mathcal{J}(H)$ of $(r, h)$-GPwB as described in the previous section, and convert it into an instance $\hat{\mathcal{J}}(H)$ of NDP-Grid.
Note that, if $G$ is a Yes-Instance, then from Corollary 4.4, there is a solution to $\hat{\mathcal{J}}(H)$ of value at least $\frac{|\mathcal{R}| \cdot 6^{\ell}}{c_{\mathrm{y}} \ell^{4} \log ^{4} n}$. Assume now that $G$ is a No-Instance. If we could show that any solution to the corresponding $(r, h)$-GPwB instance $\mathcal{J}(H)$ has value less than $\frac{|\mathcal{R}| \cdot 6^{\ell}}{c_{\mathrm{YI}}^{2} \cdot \alpha^{*} \ell^{7} \log ^{7} n}$, we would be done. Indeed, in such a case, from Theorem 4.3, every solution to the NDP-Grid instance $\hat{\mathcal{J}}(H)$ routes fewer than $\frac{|\mathcal{R}| \cdot 6^{\ell}}{c_{\mathrm{YI}} \alpha^{*} \ell^{4} \log ^{4} n}$ demand pairs. If we assume for contradiction that an $\alpha^{*}$-approximation algorithm exists for NDP-Grid, then, if $G$ is a YES-Instance, the algorithm would have to return a solution to $\hat{\mathcal{J}}(H)$ routing at least $\frac{|\mathcal{R}| \cdot 6^{6}}{c_{\mathrm{YI}} \alpha^{*} \ell^{4} \log ^{4} n}$ demand pairs, while, if $G$ is a No-Instance, no such solution would exist. Therefore, we could use the $\alpha^{*}$-approximation algorithm for NDP-Grid to distinguish between the Yes-Instances and the No-Instances of 3COL(5).
Unfortunately, we are unable to prove this directly. For simplicity, assume that $H$ is a semiregular graph. Our intended solution to the $(r, h)$-GPwB instance $\mathcal{J}(H)$, defined over the graph $L(H)$, for each query
$Q \in Q^{E} \cup Q^{V}$, places every vertex of $S(Q)$ into a distinct cluster. Any such solution will indeed have a low value in the No-Instance. But a cheating solution may place many vertices from the same set $S(Q)$ into some cluster $W_{j}$. Such a solution may end up having a high value, but it may not translate into a good strategy for the two provers, that satisfies a large fraction of the random strings. In an extreme case, for each query $Q$, we may place all vertices of $S(Q)$ into a single cluster $W_{j}$. The main idea in our reduction is to overcome this difficulty by noting that such a cheating solution can be used to compute a partition of the constraint graph $H$. The graph is partitioned into as many as $6^{\ell}$ pieces, each of which is significantly smaller than the original graph $H$. At the same time, a large fraction of the edges of $H$ will survive the partitioning procedure. Intuitively, if we now restrict ourselves to only those random strings $R \in \mathcal{R}$, whose corresponding edges have survived the partitioning procedure, then the problem does not become significantly easier, and we can recursively apply the same argument to the resulting subgraphs of $H$. We make a significant progress in each such iteration, since the sizes of the resulting subgraphs of $H$ decrease very fast. The main tool that allows us to execute this plan is the following theorem.

Theorem 5.1. Suppose we are given an instance $G$ of the 3COL(5) problem with $|V(G)|=n$, and an integral parallel repetition parameter $\ell>\log ^{2} n$, together with some connected subgraph $H^{\prime} \subseteq H$ of the corresponding constraint graph $H$. Consider the corresponding instance $\mathcal{J}\left(H^{\prime}\right)$ of $(\mathrm{r}, \mathrm{h})$-GPwB, and assume that we are given a solution to this instance of value at least $\left|E\left(H^{\prime}\right)\right| \cdot 6^{\ell} / \alpha$, where $\alpha=$ $c_{\mathrm{YI}}^{2} \cdot \alpha^{*} \cdot \ell^{7} \log ^{7} n$. Then there is an algorithm, whose running time is $O\left(n^{O(\ell)}\right)$, that returns one of the following:

- Either a strategy for the two provers that satisfies more than a $2^{-\gamma \ell / 2}$-fraction of the constraints $R \in \mathcal{R}$ with $e(R) \in E\left(H^{\prime}\right)$; or
- A collection $\mathcal{H}$ of disjoint connected subgraphs of $H^{\prime}$, such that for each $\tilde{H} \in \mathcal{H},|E(\tilde{H})| \leq$ $\left|E\left(H^{\prime}\right)\right| / 2^{\gamma \ell / 16}$, and $\sum_{\tilde{H} \in \mathcal{H}}|E(\tilde{H})| \geq \frac{c^{\prime}\left|E\left(H^{\prime}\right)\right|}{\ell^{2} \alpha^{2}}$, for some universal constant $c^{\prime}$.

We postpone the proof of the theorem to the following subsection, after we complete the hardness proof for NDP-Grid. We assume for contradiction that we are given a factor- $\alpha^{*}$ approximation algorithm $\mathcal{A}$ for NDP-Grid (recall that $\alpha^{*}=2^{\Theta(\ell / \log n)}$ ). We will use this algorithm to distinguish between the Yes-Instances and the No-Instances of $3 \operatorname{COL}(5)$ with $|V(G)|=n$. Suppose we are given an instance $G$ of 3COL(5).
For an integral parallel repetition parameter $\ell>\log ^{2} n$, let $H$ be the constraint graph corresponding to $G$ and $\ell$. We next show an algorithm, that uses $\mathcal{A}$ as a subroutine, in order to determine whether $G$ is a Yes-Instance or a No-Instance. The running time of the algorithm is $n^{O(\ell)}$.
Throughout the algorithm, we maintain a collection $\mathcal{H}$ of disjoint connected subgraphs of $H$, that we sometimes call clusters. Set $\mathcal{H}$ is in turn partitioned into two subsets: set $\mathcal{H}_{1}$ of active clusters, and set $\mathcal{H}_{2}$ of inactive clusters. Consider now some inactive cluster $H^{\prime} \in \mathcal{H}_{2}$. This cluster defines a 2-prover game $\mathcal{G}\left(H^{\prime}\right)$, where the queries to the two provers are $\left\{Q^{E} \in Q^{E} \mid v\left(Q^{E}\right) \in V\left(H^{\prime}\right)\right\}$, and $\left\{Q^{V} \in Q^{V} \mid v\left(Q^{V}\right) \in V\left(H^{\prime}\right)\right\}$ respectively, and the constraints of the verifier are:

$$
\mathcal{R}\left(H^{\prime}\right)=\left\{R \in \mathcal{R} \mid e(R) \in E\left(H^{\prime}\right)\right\} .
$$

For each inactive cluster $H^{\prime} \in \mathcal{H}_{2}$, we will store a strategy of the two provers for game $\mathcal{G}\left(H^{\prime}\right)$, that satisfies at least a $2^{-\gamma \ell / 2}$-fraction of the constraints in $\mathcal{R}\left(H^{\prime}\right)$.
At the beginning, the clusters in $\mathcal{H}$ are the connected components of $H$, and every cluster is active. The algorithm is executed while $\mathcal{H}_{1} \neq \emptyset$, and its execution is partitioned into phases. In every phase, we process each of the clusters that belongs to $\mathcal{H}_{1}$ at the beginning of the phase. Each phase is then in turn is partitioned into iterations, where in every iteration we process a distinct active cluster $H^{\prime} \in \mathcal{H}_{1}$. We describe an iteration when an active cluster $H^{\prime} \in \mathcal{H}_{1}$ is processed in Figure 2 (see also the flowchart in Figure 3).

## Iteration for Processing a Cluster $H^{\prime} \in \mathcal{H}_{1}$

1. Construct an instance $\mathcal{J}\left(H^{\prime}\right)$ of $(r, h)-G P w B$.
2. Use Theorem 4.3 to construct an instance $\hat{\mathcal{J}}\left(H^{\prime}\right)$ of the NDP-Grid problem.
3. Run the $\alpha^{*}$-approximation algorithm $\mathcal{A}$ on this instance. If the resulting solution routes fewer than $\frac{\left|E\left(H^{\prime}\right)\right| 6^{\ell}}{c_{\mathrm{YI}} \alpha^{*} \ell^{4} \log ^{4} n}$ demand pairs, halt and return " $G$ is a No-InSTANCE".
4. Otherwise, consider the returned solution routing at least $\frac{\left|E\left(H^{\prime}\right)\right| 6^{\ell}}{c_{11} \alpha^{*} \ell^{4} \log ^{4} n}$ demand pairs. Denote $\left|E\left(L\left(H^{\prime \prime}\right)\right)\right|=M$, and recall that $M \leq n^{\hat{c} \ell}$. Use Theorem 4.3 to compute a solution $\left(\left(W_{1}, \ldots, W_{r}\right),\left(E_{1}, \ldots, E_{r}\right)\right)$ to the instance $\mathcal{J}\left(H^{\prime}\right)$ of the $(\mathrm{r}, \mathrm{h})$-GPwB problem, of value at least:

$$
\frac{\left|E\left(H^{\prime}\right)\right| \cdot 6^{\ell}}{\left(c_{\mathrm{Y} 1} \alpha^{*} \ell^{4} \log ^{4} n\right)\left(c^{*} \log ^{3} M\right)} \geq \frac{\left|E\left(H^{\prime}\right)\right| \cdot 6^{\ell}}{c_{\mathrm{Y} 1} *^{*} \hat{c}^{3} \alpha^{*} \ell^{7} \log ^{7} n} \geq \frac{\left|E\left(H^{\prime}\right)\right| \cdot 6^{\ell}}{c_{\mathrm{Y} 1}^{2} \alpha^{*} \ell^{7} \log ^{7} n}=\frac{\left|E\left(H^{\prime}\right)\right| \cdot 6^{\ell}}{\alpha} .
$$

5. Apply the algorithm from Theorem 5.1 to this solution.
(a) If the outcome is a strategy for the provers satisfying more than a $2^{-\gamma / / 2}$-fraction of constraints $R \in \mathcal{R}$ with $e(R) \in E\left(H^{\prime}\right)$, then declare cluster $H^{\prime}$ inactive and move it from $\mathcal{H}_{1}$ to $\mathscr{H}_{2}$. Store the resulting strategy of the provers.
(b) Otherwise, let $\tilde{\mathcal{H}}$ be the collection of subgraphs of $H^{\prime}$ returned by the algorithm. Remove $H^{\prime}$ from $\mathcal{H}_{1}$ and add all graphs of $\tilde{\mathcal{H}}$ to $\mathcal{H}_{1}$.

Figure 2: An iteration description

If the algorithm terminates with $\mathcal{H}$ containing only inactive clusters, then we return " $G$ is a YesInstance".


Figure 3: Flowchart for an iteration execution

Correctness. We establish the correctness of the algorithm in the following two lemmas.
Lemma 5.2. If $G$ is $a$ Yes-Instance, then the algorithm always returns " $G$ is a Yes-Instance".
Proof. Consider an iteration of the algorithm when an active cluster $H^{\prime}$ is processed. Notice that the algorithm may only determine that $G$ is a No-Instance in Step (3).
From Corollary 4.4, an instance $\hat{\mathcal{J}}\left(H^{\prime}\right)$ of NDP-Grid is guaranteed to have a solution of value at least $\frac{\left|E\left(H^{\prime}\right)\right| \cdot 6^{4}}{c_{\mathrm{Y} l^{4}} \log ^{4} n}$, and our $\alpha^{*}$-approximation algorithm to NDP-Grid must then return a solution of value at least $\frac{\left|E\left(H^{\prime}\right)\right| \cdot 6^{\ell}}{c_{\mathrm{Y} 1} \alpha^{*} \ell^{4} \log ^{4} n}$. Therefore, our algorithm will never return " $G$ is a No-InSTANCE".

Lemma 5.3. If $G$ is a NO-InStance, then the algorithm always returns " $G$ is a NO-InSTANCE".
Proof. From Corollary 2.5, it is enough to show that, whenever the algorithm classifies $G$ as a YesInSTANCE, there is a strategy for the two provers, that satisfies more than a $2^{-\gamma \ell}$-fraction of the constraints in $\mathcal{R}$.
Note that the original graph $H$ has at most $n^{\hat{\imath}}$ edges. In every phase, the number of edges in each active graph decreases by a factor of at least $2^{\gamma \ell / 16}$. Therefore, the number of phases is bounded by $O(\log n)$. If the algorithm classifies $G$ as a Yes-Instance, then it must terminate when no active clusters remain. In
every phase, the number of edges in $\bigcup_{H^{\prime} \in \mathcal{H}} E\left(H^{\prime}\right)$ goes down by at most a factor $\ell^{2} \alpha^{2} / c^{\prime}$. Therefore, at the end of the algorithm:

$$
\sum_{H^{\prime} \in \mathcal{H}_{2}}\left|E\left(H^{\prime}\right)\right| \geq \frac{|\mathcal{R}|}{\left(\ell^{2} \alpha^{2} / c^{\prime}\right)^{O(\log n)}}=\frac{|\mathcal{R}|}{\left(\ell^{16} \cdot\left(\alpha^{*}\right)^{2} \cdot \log ^{14} n\right)^{O(\log n)}}=\frac{|\mathcal{R}|}{\left(\alpha^{*}\right)^{O(\log n)}} .
$$

By appropriately setting $\alpha^{*}=2^{\Theta(\ell / \log n)}$, we will ensure that the number of edges remaining in the inactive clusters $H^{\prime} \in \mathcal{H}_{2}$ is at least $|\mathcal{R}| / 2^{\gamma \ell / 4}$. Each such edge corresponds to a distinct random string $R \in \mathcal{R}$. Recall that for each inactive cluster $H^{\prime}$, there is a strategy for the provers in the corresponding game $\mathcal{G}\left(H^{\prime}\right)$ that satisfies at least $\left|E\left(H^{\prime}\right)\right| / 2^{\gamma / / 2}$ of its constraints. Taking the union of all these strategies, we can satisfy more than $|\mathcal{R}| / 2^{\gamma \ell}$ constraints of $\mathcal{R}$, contradicting the fact that $G$ is a No-Instance.

Running time and the hardness factor. It is easy to see that the algorithm has at most $n^{O(\ell)}$ iterations, where in every iteration it processes a distinct active cluster $H^{\prime} \subseteq H$. The corresponding graph $L\left(H^{\prime \prime}\right)$ has at most $n^{O(\ell)}$ edges, and so the resulting instance of NDP-Grid contains at most $n^{O(\ell)}$ vertices. Therefore, the overall running time of the algorithm is $n^{O(\ell)}$. From the above analysis, the algorithm always correctly classifies $G$ as a Yes-Instance or as a No-Instance. The hardness factor that we obtain is $\alpha^{*}=2^{\Theta(\ell / \log n)}$, while we only apply our approximation algorithm to instances of NDP-Grid containing at most $N=n^{O(\ell)}$ vertices.
Setting $\ell=\log ^{p} n$ for a large enough integer $p$, we obtain $\alpha^{*}=2^{\Theta\left((\log N)^{1-2 /(p+1)}\right)}$, giving us a $2^{(\log n)^{(1-\varepsilon)}}-$ hardness of approximation for NDP-Grid for any constant $\varepsilon$, assuming NP $\nsubseteq$ DTIME ( $\left.n^{\text {polylog } n}\right)$.
Setting $\ell=n^{\delta}$ for some constant $\delta$, we get that $N=2^{O\left(n^{\delta} \log n\right)}$ and $\alpha^{*}=2^{\Theta\left(n^{\delta} / \log n\right)}$, giving us a $n^{\Omega\left(1 /(\log \log n)^{2}\right)}$-hardness of approximation for NDP-Grid, assuming that $\operatorname{NP} \nsubseteq \operatorname{DTIME}\left(2^{n^{\delta}}\right)$ for some constant $\delta>0$.

Note that the approximation factor $\alpha^{*}$ is in fact eventually expressed as a function of the size of the NDPGrid instance, and that we apply the approximation algorithm to instances of NDP-Grid of various sizes. However, the size of each resulting instance of NDP-Grid is bounded by $N$, and so the approximation factor achieved for each such instance should be at most $\alpha^{*}$ (as defined with respect to $N$ ).

### 5.1 Proof of Theorem 5.1

Recall that each edge of graph $H^{\prime}$ corresponds to some constraint $R \in \mathcal{R}$. Let $\mathcal{R}^{\prime} \subseteq \mathcal{R}$ be the set of all constraints $R$ with $e(R) \in E\left(H^{\prime}\right)$. Let $\mathcal{R}^{\prime \prime} \subseteq \mathcal{R}^{\prime}$ be the set of all constraints $R$ with $e(R) \in E\left(H^{\prime \prime}\right)$. Denote the solution to the $(r, \mathrm{~h})-\mathrm{GPwB}$ instance $\mathcal{J}\left(H^{\prime}\right)$ by $\left(\left(W_{1}, \ldots, W_{r}\right),\left(E_{1}, \ldots, E_{r}\right)\right)$, and let $E^{\prime}=\bigcup_{i=1}^{r} E_{i}$. Recall that for each random string $R \in \mathcal{R}^{\prime \prime}$, there is a set $E(R)$ of $12^{\ell}$ edges in graph $L\left(H^{\prime \prime}\right)$ representing $R$. Due to the way these edges are partitioned into bundles, at most $6^{\ell}$ edges of $E(R)$ may belong to $E^{\prime}$. We say that a random string $R \in \mathcal{R}^{\prime \prime}$ is good iff $E^{\prime}$ contains at least $6^{\ell} /(2 \alpha)$ edges of $E(R)$, and we say that it is bad otherwise.

Observation 5.4. At least $\left|\mathcal{R}^{\prime}\right| /(2 \alpha)$ random strings of $\mathcal{R}^{\prime \prime}$ are good.

## Almost Polynomial Hardness of Node-Disjoint Paths in Grids

Proof. A good random string contributes at most $6^{\ell}$ edges to $E^{\prime}$, while a bad random string contributes at most $6^{\ell} /(2 \alpha)$ edges. If the number of good random strings is less than $\left|\mathcal{R}^{\prime}\right| /(2 \alpha)$, then a simple accounting shows that $\left|E^{\prime}\right|<\left|\mathcal{R}^{\prime}\right| \cdot \sigma^{\ell} / \alpha=\left|E\left(H^{\prime}\right)\right| \cdot 6^{\ell} / \alpha$, a contradiction.

Consider some random string $R \in \mathcal{R}^{\prime \prime}$, and assume that $R=\left(Q^{E}, Q^{V}\right)$. We denote by $E^{\prime}(R)=E(R) \cap E^{\prime}$. Intuitively, say that a cluster $W_{i}$ is a terrible cluster for $R$ if the number of edges of $E(R)$ that lie in $E_{i}$ is much smaller than $\left|Q^{E} \cap W_{i}\right|$ or $\left|Q^{V} \cap W_{i}\right|$. We now give a formal definition of a terrible cluster.

Definition 5.5. Given a random string $R \in \mathcal{R}^{\prime \prime}$ and an index $1 \leq i \leq 6^{\ell}$, we say that a cluster $W_{i}$ is a terrible cluster for $R$, if:

- either $\left|E^{\prime}(R) \cap E_{i}\right|<\left|W_{i} \cap S\left(Q^{V}\right)\right| /(8 \alpha)$; or
- $\left|E^{\prime}(R) \cap E_{i}\right|<\left|W_{i} \cap S\left(Q^{E}\right)\right| /(8 \alpha)$.

We say that an edge $e \in E^{\prime}(R)$ is a terrible edge if it belongs to the set $E_{i}$, where $W_{i}$ is a terrible cluster for $R$.

Observation 5.6. For each good random string $R \in \mathcal{R}^{\prime \prime}$, at most $6^{\ell} /(4 \alpha)$ edges of $E^{\prime}(R)$ are terrible.
Proof. Assume for contradiction that more than $6^{\ell} /(4 \alpha)$ edges of $E^{\prime}(R)$ are terrible. Denote $R=$ $\left(Q^{E}, Q^{V}\right)$. Consider some such terrible edge $e \in E^{\prime}(R)$, and assume that $e \in E_{i}$ for some cluster $W_{i}$, that is terrible for $R$. We say that $e$ is a type-1 terrible edge iff $\left|E^{\prime}(R) \cap E_{i}\right|<\left|W_{i} \cap S\left(Q^{V}\right)\right| /(8 \alpha)$, and it is a type-2 terrible edge otherwise, in which case $\left|E^{\prime}(R) \cap E_{i}\right|<\left|W_{i} \cap S\left(Q^{E}\right)\right| /(8 \alpha)$ must hold. Let $E^{1}(R)$ and $E^{2}(R)$ be the sets of all terrible edges of $E^{\prime}(R)$ of types 1 and 2 , respectively. Then either $\left|E^{1}(R)\right|>6^{\ell} /(8 \alpha)$, or $\left|E^{2}(R)\right|>6^{\ell} /(8 \alpha)$ must hold.
Assume first that $\left|E^{1}(R)\right|>6^{\ell} /(8 \alpha)$. Fix some index $1 \leq i \leq 6^{\ell}$, such that $W_{i}$ is a cluster that is terrible for $R$, and $\left|E(R) \cap E_{i}\right|<\left|W_{i} \cap S\left(Q^{V}\right)\right| /(8 \alpha)$. We assign, to each edge $e \in E_{i} \cap E^{1}(R)$, a set of $8 \alpha$ vertices of $W_{i} \cap S\left(Q^{V}\right)$ arbitrarily, so that every vertex is assigned to at most one edge; we say that the corresponding edge is responsible for the vertex. Every edge of $E_{i} \cap E^{1}(R)$ is now responsible for $8 \alpha$ distinct vertices of $W_{i} \cap S\left(Q^{V}\right)$. Once we finish processing all such clusters $W_{i}$, we will have assigned, to each edge of $E^{1}(R)$, a set of $8 \alpha$ distinct vertices of $S\left(Q^{V}\right)$. We conclude that $\left|S\left(Q^{V}\right)\right| \geq 8 \alpha\left|E^{1}(R)\right|>6^{\ell}$. But $\left|S\left(Q^{V}\right)\right|=6^{\ell}$, a contradiction.
The proof for the second case, where $\left|E^{2}(R)\right|>6^{\ell} /(8 \alpha)$ is identical, and relies on the fact that $\left|S\left(Q^{E}\right)\right|=$ 6 .

We will use the following simple observation.
Observation 5.7. Let $R \in \mathcal{R}^{\prime \prime}$ be a good random string, with $R=\left(Q^{E}, Q^{V}\right)$, and let $1 \leq i \leq 6^{\ell}$ be an index, such that $W_{i}$ is not terrible for $R$. Then $\left|W_{i} \cap S\left(Q^{V}\right)\right| \geq\left|W_{i} \cap S\left(Q^{E}\right)\right| /(8 \alpha)$ and $\left|W_{i} \cap S\left(Q^{E}\right)\right| \geq$ $\left|W_{i} \cap S\left(Q^{V}\right)\right| /(8 \alpha)$.

Proof. Assume first for contradiction that $\left|W_{i} \cap S\left(Q^{V}\right)\right|<\left|W_{i} \cap S\left(Q^{E}\right)\right| /(8 \alpha)$. Consider the edges of $E(R) \cap E_{i}$. Each such edge must be incident to a distinct vertex of $S\left(Q^{V}\right)$. Indeed, if two edges $e, e^{\prime} \in E(R) \cap E_{i}$ are incident to the same vertex $v_{j}\left(Q^{V}, A\right) \in S\left(Q^{V}\right)$, then, since the other endpoint of each such edge lies in $S\left(Q^{E}\right)$, the two edges belong to the same bundle, a contradiction. Therefore, $\left|E_{i} \cap E(R)\right| \leq\left|W_{i} \cap S\left(Q^{V}\right)\right|<\left|W_{i} \cap S\left(Q^{E}\right)\right| /(8 \alpha)$, contradicting the fact that $W_{i}$ is not a terrible cluster for $R$.

The proof for the second case, where $\left|W_{i} \cap S\left(Q^{E}\right)\right|<\left|W_{i} \cap S\left(Q^{V}\right)\right| /(8 \alpha)$ is identical. As before, each edge of $E(R) \cap E_{i}$ must be incident to a distinct vertex of $S\left(Q^{E}\right)$, as otherwise, a pair $e, e^{\prime} \in E(R)$ of edges that are incident on the same vertex $v\left(Q^{E}, A\right) \in S\left(Q^{E}\right)$ belong the same bundle. Therefore, $\left|E_{i} \cap E(R)\right| \leq\left|W_{i} \cap S\left(Q^{E}\right)\right|<\left|W_{i} \cap S\left(Q^{V}\right)\right| /(8 \alpha)$, contradicting the fact that $W_{i}$ is not a terrible cluster for $R$.

For each good random string $R \in \mathcal{R}^{\prime \prime}$, we discard the terrible edges from set $E^{\prime}(R)$, so $\left|E^{\prime}(R)\right| \geq 6^{\ell} /(4 \alpha)$ still holds.

Let $z=2^{\gamma \ell / 8}$. We say that cluster $W_{i}$ is heavy for a random string $R=\left(Q^{E}, Q^{V}\right) \in \mathcal{R}^{\prime \prime}$ iff $\left|W_{i} \cap S\left(Q^{E}\right)\right|, \mid W_{i} \cap$ $S\left(Q^{V}\right) \mid>z$. We say that an edge $e \in E^{\prime}(R)$ is heavy iff it belongs to set $E_{i}$, where $W_{i}$ is a heavy cluster for $R$. Finally, we say that a random string $R \in \mathcal{R}^{\prime \prime}$ is heavy iff at least half of the edges in $E^{\prime}(R)$ are heavy. Random strings and edges that are not heavy are called light. We now consider two cases. The first case happens if at least half of the good random strings are light. In this case, we compute a strategy for the provers to choose assignments to the queries, so that at least a $2^{-\gamma \ell / 2}$-fraction of the constraints in $\mathcal{R}^{\prime}$ are satisfied. In the second case, at least half of the good random strings are heavy. We then compute a partition $\mathcal{H}$ of $H^{\prime}$ as desired. We now analyze the two cases. Note that if $\left|E\left(H^{\prime}\right)\right|<z /(8 \alpha)$, then Case 2 cannot happen. This is since $h=\left|E\left(H^{\prime}\right)\right|<z /(8 \alpha)$ in this case. Therefore, for all $1 \leq i \leq r$, $\left|E_{i}^{\prime}\right| \leq h<z /(8 \alpha)$ must hold, and so no random strings may be heavy. Therefore, if $H^{\prime}$ is small enough, we will return a strategy of the provers that satisfies a large fraction of the constraints in $\mathcal{R}^{\prime}$.

Case 1. This case happens if at least half of the good random strings are light. Let $\mathcal{L} \subseteq \mathcal{R}^{\prime \prime}$ be the set of the good light random strings, so $|\mathcal{L}| \geq\left|\mathcal{R}^{\prime}\right| /(4 \alpha)$. For each such random string $R \in \mathcal{L}$, we let $E^{L}(R) \subseteq E^{\prime}(R)$ be the set of all light edges corresponding to $R$, so $\left|E^{L}(R)\right| \geq 6^{\ell} /(8 \alpha)$. We now describe a randomized algorithm to choose an answer to every query $Q \in Q^{E} \cup Q^{V}$ with $v(Q) \in H^{\prime}$. Our algorithm chooses a random index $1 \leq i \leq r$. For every query $Q \in Q^{E} \cup Q^{V}$ with $v(Q) \in H^{\prime}$, we consider the set $\mathcal{A}(Q)$ of all answers $A$, such that some vertex $v(Q, A)$ belongs to $W_{i}$ (for the case where $Q \in Q^{V}$, the vertex is of the form $v_{j}(Q, A)$ ). We then choose one of the answers from $\mathcal{A}(Q)$ uniformly at random, and assign it to $Q$. If $\mathcal{A}(Q)=\emptyset$, then we choose an arbitrary answer to $Q$.

We claim that the expected number of satisfied constraints of $\mathcal{R}^{\prime}$ is at least $\left|\mathcal{R}^{\prime}\right| / 2^{\gamma \ell / 2}$. Since $\mathcal{L} \geq\left|\mathcal{R}^{\prime}\right| /(4 \alpha)$, it is enough to show that the expected fraction the good light constraints that are satisfied is at least $4 \alpha|\mathcal{L}| / 2^{\gamma \ell / 2}$, and for that it is sufficient to show that each light constraint $R \in \mathcal{L}$ is satisfied with probability at least $4 \alpha / 2^{\gamma \ell / 2}$.

Fix one such constraint $R=\left(Q^{E}, Q^{V}\right) \in \mathcal{L}$, and consider an edge $e \in E^{L}(R)$. Assume that $e$ connects a vertex $v\left(Q^{E}, A\right)$ to a vertex $v_{j}\left(Q^{V}, A^{\prime}\right)$, and that $e \in E_{i}$. We say that edge $e$ is happy iff our algorithm chose
the index $i$, the answer $A$ to query $Q^{E}$, and the answer $A^{\prime}$ to query $Q^{V}$. Notice that due to our construction of bundles, at most one edge $e \in E^{L}(R)$ may be happy with any choice of the algorithm; moreover, if any edge $e \in E^{L}(R)$ is happy, then the constraint $R$ is satisfied. The probability that a fixed edge $e$ is happy is at least $1 /\left(8 \cdot 6^{\ell} z^{2} \alpha\right)$. Indeed, we choose the correct index $i$ with probability $1 / 6^{\ell}$. Since $e$ belongs to $E_{i}, W_{i}$ is a light cluster for $R$, and so either $\left|S\left(Q^{E}\right) \cap W_{i}\right| \leq z$, or $\left|S\left(Q^{V}\right) \cap W_{i}\right| \leq z$. Assume without loss of generality that it is the former; the other case is symmetric. Then, since $e$ is not terrible, from Observation 5.7, $\left|S\left(Q^{V}\right) \cap W_{i}\right| \leq 8 \alpha z$, and so $\left|\mathcal{A}\left(Q^{V}\right)\right| \leq 8 \alpha z$, while $\left|\mathcal{A}\left(Q^{E}\right)\right| \leq z$. Therefore, the probability that we choose answer $A$ to $Q^{E}$ and answer $A^{\prime}$ to $A^{V}$ is at least $1 /\left(8 \alpha z^{2}\right)$, and overall, the probability that a fixed constraint $R \in \mathcal{L}$ is satisfied is at least $\left|E^{L}(R)\right| /\left(8 \cdot 6^{\ell} z^{2} \alpha\right) \geq 1 /\left(64 z^{2} \alpha^{2}\right) \geq 4 \alpha / 2^{\gamma \ell / 2}$, since $z=2^{\gamma \ell / 8}$, and $\alpha<2^{\gamma / / 32}$.

So far, we have defined a randomized strategy for the provers that satisfies at least a $2^{-\gamma \ell / 2}$-fraction of the constraints in $\mathcal{R}^{\prime}$ in expectation. We can turn this strategy into a deterministic one, that satisfies at least a $2^{-\gamma / / 2}$-fraction of the constraints in $\mathcal{R}^{\prime}$, using standard techniques: first, for every cluster $W_{i}$, we can compute the expected number of satisfied constraints if the index $i$ is chosen, and then fix the cluster $W_{i}$ that maximizes this expectation. The randomized strategy of the provers outlined above can then be turned into a deterministic one using the method of conditional expectation, as sketched in Section 2.

Case 2. This case happens if at least half of the good random strings are heavy. Let $\mathcal{R}^{H} \subseteq \mathcal{R}^{\prime \prime}$ be the set of all random strings that are good and heavy, so $\left|\mathcal{R}^{H}\right| \geq\left|\mathcal{R}^{\prime}\right| /(4 \alpha)$. For each such random string $R \in \mathcal{R}^{H}$, we let $E^{H}(R) \subseteq E^{\prime}(R)$ be the set of all heavy edges corresponding to $R$. Recall that $\left|E^{H}(R)\right| \geq 6^{\ell} /(8 \alpha)$.
Fix some heavy random string $R \in \mathcal{R}^{H}$ and assume that $R=\left(Q^{E}, Q^{V}\right)$. For each $1 \leq i \leq r$, let $E_{i}(R)=$ $E^{H}(R) \cap E_{i}$. Recall that, if $E_{i}(R) \neq \emptyset$, then $\left|W_{i} \cap S\left(Q^{E}\right)\right|,\left|W_{i} \cap S\left(Q^{V}\right)\right| \geq z$ must hold, and, from the definition of terrible clusters, $\left|E_{i}(R)\right| \geq z /(8 \alpha)$. It is also immediate that $\left|E_{i}(R)\right| \leq\left|E^{\prime}(R)\right| \leq 6^{\ell}$.

In the remainder of the algorithm and its analysis, it would be convenient for us if the values $\left|E_{i}(R)\right|$ were roughly the same for all indices $1 \leq i \leq r$ and random strings $R \in \mathcal{R}^{H}$ with $E_{i}(R) \neq \emptyset$. This can be achieved by discarding all but a polylogarithmic fraction of edges in $\bigcup_{R \in \mathcal{R}^{H}} E^{H}(R)$, using standard geometric grouping and averaging techniques, which we do next.
Consider some heavy random string $R \in \mathcal{R}^{H}$. We partition the set $\left\{1, \ldots, 6^{\ell}\right\}$ of indices into at most $\log \left(\left|E^{H}(R)\right|\right) \leq \log \left(6^{\ell}\right)$ classes, where index $1 \leq y \leq 6^{\ell}$ belongs to class $\mathcal{C}_{j}(R)$ iff $2^{j-1}<\left|E^{H}(R) \cap E_{y}\right| \leq$ $2^{j}$. Then there is some index $j_{R}$, so that $\sum_{y \in \mathcal{C}_{j_{R}}(R)}\left|E^{H}(R) \cap E_{y}\right| \geq\left|E^{H}(R)\right| / \log \left(6^{\ell}\right)$. We say that $R$ chooses the index $j_{R}$. Notice that:

$$
\sum_{y \in \mathcal{C}_{j_{R}(R)}}\left|E^{H}(R) \cap E_{y}\right| \geq \frac{\left|E^{H}(R)\right|}{\log \left(6^{\ell}\right)} \geq \frac{6^{\ell}}{8 \ell \alpha \log 6} .
$$

Moreover,

$$
\begin{equation*}
\left|\mathcal{C}_{j_{R}}(R)\right| \geq \frac{\left|E^{H}(R)\right|}{\log \left(6^{\ell}\right) \cdot 2^{j_{R}}} \geq \frac{6^{\ell}}{8 \cdot 2^{j_{R}} \cdot \ell \alpha \log 6} \tag{5.1}
\end{equation*}
$$

## Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat

Let $j^{*}$ be the index that was chosen by at least $\left|\mathcal{R}^{H}\right| / \log \left(6^{\ell}\right)$ random strings, and let $\mathcal{R}^{*} \subseteq \mathcal{R}^{H}$ be the set of all random strings that chose $j^{*}$.
We now proceed in two steps. In the first step, we show a randomized algorithm for computing a collection $\mathcal{H}$ of subgraphs of $H^{\prime}$. In the second step, we derandomize this algorithm using the method of conditional expectation.

Randomized algorithm. We now show a randomized algorithm to construct a collection $\mathcal{H}=\left\{H_{1}, \ldots, H_{6}\right.$ e $\}$ of subgraphs of $H^{\prime}$. We first define the sets of vertices in these subgraphs, and then the sets of edges. Choose a random ordering of the clusters $W_{1}, \ldots, W_{6}$; re-index the clusters according to this ordering. For each query $Q \in Q^{E} \cup Q^{V}$ with $v(Q) \in H^{\prime \prime}$, add the vertex $v(Q)$ to set $V\left(H_{i}\right)$, where $i$ is the smallest index for which $W_{i}$ contains at least $2^{j^{*}-1}$ vertices of $S(Q)$; if no such index $i$ exists, then we do not add $v(Q)$ to any set.
In order to define the edges of each graph $H_{i}$, for every random string $R=\left(Q^{E}, Q^{V}\right) \in \mathcal{R}^{*}$, if $i \in \mathcal{C}_{j^{*}}(R)$, and both $v\left(Q^{E}\right)$ and $v\left(Q^{V}\right)$ belong to $V\left(H_{i}\right)$, then we add the corresponding edge $e(R)$ to $E\left(H_{i}\right)$. This completes the definition of the family $\mathcal{H}=\left\{H_{1}, \ldots, H_{6^{6}}\right\}$ of subgraphs of $H^{\prime}$. We now analyze their properties. It is immediate to verify that the graphs in $\mathcal{H}$ are disjoint.

Claim 5.8. For each $1 \leq i \leq 6^{\ell},\left|E\left(H_{i}\right)\right| \leq\left|E\left(H^{\prime}\right)\right| / 2^{\gamma \ell / 16}$.

Proof. Fix some index $1 \leq i \leq 6^{\ell}$. An edge $e(R)$ may belong to $H_{i}$ only if $R \in \mathcal{R}^{*}$, and $i \in \mathcal{C}_{j^{*}}(R)$. In that case, $E_{i}$ contained at least $z /(8 \alpha)$ edges of $E(R)$ (since $W_{i}$ must be heavy for $R$ and it is not terrible for $R$ ). Therefore, the number of edges in $H_{i}$ is bounded by $\left|E_{i}\right| \cdot 8 \alpha / z \leq 8 \alpha h / z \leq 8 \alpha\left|E\left(H^{\prime \prime}\right)\right| / 2^{\gamma \ell / 8} \leq$ $8 \alpha\left|E\left(H^{\prime}\right)\right| / 2^{\gamma \ell / 8} \leq\left|E\left(H^{\prime}\right)\right| / 2^{\gamma \ell / 16}$, since $\alpha \leq 2^{\gamma / / 32}$.

Claim 5.9. E $\left[\sum_{i=1}^{r}\left|E\left(H_{i}\right)\right|\right] \geq \frac{\left|E\left(H^{\prime}\right)\right|}{128 \ell^{2} \alpha^{2} \log ^{2} 6}$.

Proof. Recall that $\left|\mathcal{R}^{*}\right| \geq\left|\mathcal{R}^{H}\right| / \log \left(6^{\ell}\right) \geq\left|\mathcal{R}^{\prime}\right| /\left(4 \alpha \log \left(6^{\ell}\right)\right)$. We now fix $R \in \mathcal{R}^{*}$ and analyze the probability that $e(R) \in \bigcup_{i=1}^{r} E\left(H_{i}\right)$. Assume that $R=\left(Q^{E}, Q^{V}\right)$. Let $J$ be the set of indices $1 \leq y \leq 6^{\ell}$, such that $\left|W_{y} \cap S\left(Q^{V}\right)\right| \geq 2^{j^{*}-1}$. Clearly, $|J| \leq 6^{\ell} / 2^{j^{*}-1}$, and $v\left(Q^{V}\right)$ may only belong to graph $H_{i}$ if $i \in J$. Similarly, let $J^{\prime}$ be the set of indices $1 \leq y \leq 6^{\ell}$, such that $\left|W_{y} \cap S\left(Q^{E}\right)\right| \geq 2^{j^{*}-1}$. As before, $\left|J^{\prime}\right| \leq 6^{\ell} / 2^{j^{*}-1}$, and $v\left(Q^{E}\right)$ may only belong to graph $H_{i}$ if $i \in J$. Observe that every index $y \in \mathcal{C}_{j^{*}}(R)$ must belong to $J \cap J^{\prime}$, and, since $j^{*}=j_{R}$, from Equation (5.1), $\left|\mathcal{C}_{j^{*}}(R)\right| \geq \frac{6^{\ell}}{8 \cdot 2 j^{*} \cdot \ell \alpha \log 6}$.
Let $y \in J \cup J^{\prime}$ be the first index that occurs in our random ordering. If $y \in \mathcal{C}_{j^{*}}(R)$, then edge $e(R)$ is added to $H_{y}$. The probability of this happening is:

$$
\frac{\left|\bigodot_{j^{*}}(R)\right|}{\left|J \cup J^{\prime}\right|} \geq \frac{6^{\ell} /\left(8 \cdot 2^{j^{*}} \cdot \ell \alpha \log 6\right)}{2 \cdot 6^{\ell} / 2^{j^{*}-1}}=\frac{1}{32 \ell \alpha \log 6} .
$$

Overall, the expectation of $\sum_{i=1}^{r}\left|E\left(H_{i}\right)\right|$ is at least:

$$
\frac{\left|\mathcal{R}^{*}\right|}{32 \ell \alpha \log 6} \geq \frac{\left|\mathcal{R}^{\prime}\right|}{128 \ell^{2} \alpha^{2} \log ^{2} 6}=\frac{\left|E\left(H^{\prime}\right)\right|}{128 \ell^{2} \alpha^{2} \log ^{2} 6} .
$$

Deterministic algorithm. We obtain a deterministic algorithm for computing a collection $\mathcal{H}$ of subgraphs of $H^{\prime}$ with the desired properties, by applying the method of conditional expectation. Specifically, we perform $r$ iterations, where in the $i$ th iteration we choose one of the clusters $W_{j}$, for $1 \leq j \leq r$, to be the $i$ th cluster in the ordering. Notice that Claim 5.8 holds regardless of the ordering of the clusters that we choose. Therefore, it is sufficient to choose the ordering in such a way that, for the resulting subgraphs $H_{1}, \ldots, H_{r}, \sum_{i=1}^{r}\left|E\left(H_{i}\right)\right| \geq \frac{\left|E\left(H^{\prime}\right)\right|}{128 \ell^{2} \alpha^{2} \log ^{2} 6}$.
Consider the first iteration. We process every cluster $W_{j}$ in turn, and compute the expected value of $\sum_{i=1}^{r}\left|E\left(H_{i}\right)\right|$, if $W_{j}$ is chosen to be the first cluster in the ordering, and the order of the remaining clusters is random. Once $W_{j}$ is chosen to be the first cluster, the graph $H_{1}$ if fixed, so we can compute $\left|E\left(H_{1}\right)\right|$ directly. For every random string $R \in \mathcal{R}^{*}$ with $e(R) \notin E\left(H_{1}\right)$, we can calculate the probability that $e(R)$ is included in $\bigcup_{i>1} E\left(H_{i}\right)$ using the same analysis as in Claim 5.9. We then choose a cluster $W_{j}$ that gives the largest expectation of $\sum_{i=1}^{r}\left|E\left(H_{i}\right)\right|$, make $W_{j}$ the first cluster in the ordering, and continue to the next iteration. It is easy to verify that, if $\left\{H_{1}, \ldots, H_{r}\right\}$ is the final collection of subgraphs of $H^{\prime}$, then $\sum_{i=1}^{r}\left|E\left(H_{i}\right)\right| \geq \frac{\left|E\left(H^{\prime}\right)\right|}{128 \ell^{2} \alpha^{2} \log ^{2} 6}$.
Lastly, the final set $\mathcal{H}$ of clusters contains all connected components of every graph in $\left\{H_{1}, \ldots, H_{r}\right\}$. It is now easy to verify that $\mathcal{H}$ has all required properties.

## 6 Reduction from (r,h)-GPwB to NDP-Grid

In this section we prove Theorem 4.3, by providing a reduction from ( $r, h$ )-GPwB to NDP-Grid. We assume that we are given an instance $\mathcal{J}=\left(\tilde{G}=\left(V_{1} \cup V_{2}, E\right), \mathcal{U}_{1}, \mathcal{U}_{2}, h, r\right)$ of $(\mathrm{r}, \mathrm{h})$-GPwB. Let $\left|V_{1}\right|=$ $N_{1},\left|V_{2}\right|=N_{2},|E|=M$, and $N=N_{1}+N_{2}$. We assume that $\mathcal{J}$ is a valid instance, so, if we denote by $\beta^{*}=\beta^{*}(\mathcal{J})=\sum_{v \in V_{1}} \beta(v)$, then $h=\beta^{*} / r$, and $h \geq \max _{v \in V_{1} \cup V_{2}}\{\beta(v)\}$. We also assume that $\mathcal{J}$ is a $\left(d, d^{\prime}, b\right)$-regular instance, for some $d, d^{\prime}, b>0$, so for every vertex $v \in V_{1} \cup V_{2}$, every bundle $B \in \mathcal{B}(v)$ has cardinality $b$, and graph $\tilde{G}$ is $\left(d, d^{\prime}\right)$-semiregular. We start by describing the construction of the NDP-Grid instance from the instance $\mathcal{J}$ of $(r, h)-G P w B$.

### 6.1 Proof intuition

In this section, we assume that we are given an instance $\mathcal{J}$ of $(r, h)$-GPwB as described above, and we show how to construct a corresponding instance of the NDP-Grid problem. Before we provide a formal description of the construction, we provide intuition for it. In our construction, the graph $\hat{G}$ associated with the NDP-Grid instance that we are constructing is simply a large enough square grid. We select two rows of this grid, denoted by $R^{\prime}$ and $R^{\prime \prime}$, that are far enough from each other and from the top and bottom

## Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat

grid boundaries. Consider now the graph $\tilde{G}=\left(V_{1} \cup V_{2}, E\right)$ associated with the input instance $\mathcal{J}$ of the $(r, h)-G P w B$ problem. We will associate, with every vertex $v_{i} \in V_{1}$, a contiguous subpath $K_{i}$ of the row $R^{\prime}$, that we call a block representing $v_{i}$, such that all resulting subpaths $K_{i}$ are far from each other, and far from the grid boundaries. Similarly, we associate, with every vertex $v_{j}^{\prime} \in V_{2}$, a subpath $K_{j}^{\prime}$ of row $R^{\prime \prime}$, such that all resulting subpaths $K_{j}^{\prime}$ are far from each other and far from the grid boundaries. Intuitively, for every edge $e=\left(v_{i}, v_{j}^{\prime}\right) \in E$ of graph $\tilde{G}$, with $v_{i} \in V_{1}$ and $v_{j}^{\prime} \in V_{2}$, we would like to create a demand pair $s(e), t(e)$, where $s(e)$ is a vertex lying on $K_{i}$, and $t(e)$ is a vertex lying on $K_{j}^{\prime}$. Let us assume for now that we create these demand pairs in such a way that they do not share source or destination vertices. If we now consider some block $K_{i}$ corresponding to some vertex $v_{i} \in V_{1}$, then we need to designate some of its vertices as the source vertices $s(e)$ for all edges $e \in E(\tilde{G})$ that are incident to $v_{i}$. For convenience, denote the set of all edges incident to $v_{i}$ in $\tilde{G}$ by $\left\{e_{1}, \ldots, e_{z}\right\}$. Notice that for each such edge $e_{z^{\prime}}$, its destination vertex $t\left(e_{z^{\prime}}\right)$ lies in a distinct block $K_{j_{z}}^{\prime}$. Therefore, the ordering of the destination vertices is fixed, once we fix the ordering of the blocks $\left\{K_{j}^{\prime}\right\}_{j=1}^{N_{2}}$ along the row $R^{\prime \prime}$. In order to make the routing of these demand pairs easier, it is convenient to place the source vertices $s\left(e_{1}\right), \ldots, s\left(e_{z}\right)$ on the path $K_{i}$ in exactly the same order as the order of their destination vertices on the row $R^{\prime \prime}$. Also, we need to select the spacing between these source vertices carefully: if the sources are too far apart, then it is easy to show that we can always route all demand pairs, regardless of whether the corresponding $(r, h)$-GPwB instance has a high-value solution or not. On the other hand, if the source vertices are too close to each other, then routing the demand pairs, in the case where ( $\mathrm{r}, \mathrm{h}$ )-GPwB has a perfect solution can be challenging. The placement of the destination vertices in blocks $K_{j}^{\prime}$ is done similarly.
At a high level, if we assume that the resulting NDP-Grid instance has a large solution value, then we can define a large subgraph $G^{\prime} \subseteq \tilde{G}$ that can be drawn in the plane with relatively few crossings. Graph $G^{\prime}$ will contain all edges $e \in E(\tilde{G})$, whose corresponding demand pair is routed by the solution to the NDP-Grid problem instance. It is well known that graphs that can be drawn in the plane with few crossings have balanced cuts of a small value (that is, a partition of the vertices of the graph into two subsets of similar size with only a small number of vertices in each subset that have a neighbor in the other subset). We exploit this in order to iteratively compute balanced cuts in $G^{\prime}$, until we obtain a partition of $G^{\prime}$ into $r$ clusters, thus obtaining a solution to the underlying ( $\mathrm{r}, \mathrm{h}$ )-GPwB problem. There are two subtleties with this part of the proof. First, we can only ensure that a drawing of $G^{\prime}$ with few crossings exists if the distances between the source vertices within each block $K_{i}$ are small enough, and the same holds for the destination vertices. Second, we are not guaranteed that the resulting solution to the ( $\mathrm{r}, \mathrm{h}$ )-GPwB problem obeys the bundle condition: that is, it is possible that several edges that belong to the same bundle are added to the solution. In order to overcome the latter problem, we slightly modify the construction of the NDP-Grid instance, by unifying some source vertices and some destination vertices, so that, whenever two edges lie in the same bundle, their corresponding demand pairs share their source or destination vertex. This ensures that no two edges of the graph $G^{\prime}$ induced by the solution to the NDP-Grid problem instance lie in the same bundle, and that the final solution to the ( $\mathrm{r}, \mathrm{h}$ )-GPwB problem instance is indeed a feasible solution.

Lastly, we show that, if the instance $\mathcal{J}$ of $(r, h)$-GPwB has a perfect solution, then there is a solution to the NDP-Grid problem that routes a large enough number of demand pairs. Intuitively, given a perfect
solution $\left(\left(W_{1}, \ldots, W_{r}\right),\left(E_{1}, \ldots, E_{r}\right)\right)$ to the $(r, h)$-GPwB problem instance, we will route a large subset of the demand pairs corresponding to the edges in the sets $E_{1}, \ldots, E_{r}$. Due to the large distances between the blocks $\left\{K_{i}\right\}_{v_{i} \in V_{1}}$ and between the blocks $\left\{K_{j}^{\prime}\right\}_{v_{j}^{\prime} \in V_{2}}$, in a sense, we can think about routing $r$ different sets of demand pairs (where each set corresponds to the edges in a distinct set $E_{i}$ ) independently, and exploit the large distances between the blocks in order to combine these routings together. There are two crucial properties that the routing uses. First, we use the fact that within each block $K_{i}$, the source vertices are ordered in the same order as their corresponding destination vertices on row $R^{\prime \prime}$, and the same holds for each block $K_{j}^{\prime} \subseteq R^{\prime \prime}$. Second, we exploit a "distance property", that, intuitively (though somewhat imprecisely) says that, if $\mathcal{N}^{\prime}$ is the set of the demand pairs that we attempt to route, then for any pair $(s(e), t(e)),\left(s\left(e^{\prime}\right), t\left(e^{\prime}\right)\right) \in \mathcal{M}^{\prime}$ of demand pairs, the total number of destination vertices of the pairs in $\mathcal{M}^{\prime}$ lying between $t(e)$ and $t\left(e^{\prime}\right)$ is at most the distance between $s(e)$ and $s\left(e^{\prime}\right)$ on row $R^{\prime}$. In order to achieve this distance property, we will carefully select an ordering of the blocks $K_{1}^{\prime}, \ldots, K_{N_{2}}^{\prime}$ that correspond to the vertices of $V_{2}$ along the row $R^{\prime \prime}$. More precisely, we will fix an arbitrary ordering $\rho$ of the groups in $\mathcal{U}_{1}$, and we will carefully select an ordering $\rho^{\prime}$ of the groups in $\mathcal{U}_{2}$. The former will be used in order to select an ordering of the blocks $K_{i}$ along the row $R^{\prime}$, while the latter will be used to determine an ordering of the blocks $K_{j}^{\prime}$ along the row $R^{\prime \prime}$.

### 6.2 The construction

We now proceed to describe the reduction formally. We assume for now that we are given some ordering $\rho$ of the groups in $\mathcal{U}_{1}$, and some ordering $\rho^{\prime}$ of the groups in $\mathcal{U}_{2}$; we show later how to select these orderings. Using the ordering $\rho$, we define an ordering $\sigma$ of the vertices of $V_{1}$, as follows. The vertices that belong to the same group $U \in \mathcal{U}_{1}$ are placed consecutively in the ordering $\sigma$, in an arbitrary order. The ordering between the groups in $\mathcal{U}_{1}$ is the same as their ordering in $\rho$. We assume that $V_{1}=\left\{v_{1}, v_{2}, \ldots, v_{N_{1}}\right\}$, where the vertices are indexed according to their order in $\sigma$. We then define an ordering $\sigma^{\prime}$ of the vertices of $V_{2}$ in exactly the same manner, using the ordering $\rho^{\prime}$ of $\mathcal{U}_{2}$. We assume that $V_{2}=\left\{v_{1}^{\prime}, v_{2}^{\prime}, \ldots, v_{N_{2}}^{\prime}\right\}$, where the vertices are indexed according to their ordering in $\sigma^{\prime}$.
Consider some vertex $v \in V_{1}$. Recall that $\mathcal{B}(v)$ denotes the partition of the edges incident to $v$ in $\tilde{G}$ into bundles, where every bundle is a non-empty subsets of edges, and that $\beta(v)=|\mathcal{B}(v)|$. Recall also that, since graph $\tilde{G}$ is $\left(d, d^{\prime}\right)$-semiregular, $d \leq d_{\tilde{G}}(v)<2 d$. Moreover, since every bundle has cardinality $b$, we get that for every vertex $v \in V_{1}$ :

$$
\begin{equation*}
d / b \leq \beta(v)<2 d / b . \tag{6.1}
\end{equation*}
$$

Recall also that, since instance $\mathcal{J}$ is valid, $\sum_{v \in V_{1}} \beta(v)=\beta^{*}(\mathcal{J})=h r$. As every bundle contains exactly $b$ edges, we get that:

$$
\begin{equation*}
|E(\tilde{G})|=\sum_{v \in V_{1}} d(v)=\sum_{v \in V_{1}} \beta(v) \cdot b=h r b . \tag{6.2}
\end{equation*}
$$

Consider again some vertex $v \in V_{1}$. Recall that each bundle $B \in \mathcal{B}(v)$ corresponds to a single group

## Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat

$U(B) \in \mathcal{U}_{2}$, and contains all edges that connect $v$ to the vertices of $U(B)$. For every bundle $B \in \mathcal{B}(v)$, we say that the corresponding group $U(B) \in \mathcal{U}_{2}$ covers $v$.
The ordering $\rho^{\prime}$ of the groups in $\mathcal{U}_{2}$ naturally induces an ordering of the bundles in $\mathcal{B}(v)$, where $B$ appears before $B^{\prime}$ in the ordering iff $U(B)$ appears before $U\left(B^{\prime}\right)$ in $\rho^{\prime}$. We denote $\mathcal{B}(v)=\left\{B_{1}(v), B_{2}(v), \ldots, B_{\beta(v)}(v)\right\}$, where the bundles are indexed according to this ordering.

Similarly, for a vertex $v^{\prime} \in V_{2}$, every bundle $B \in \mathcal{B}\left(v^{\prime}\right)$ corresponds to a group $U(B) \in \mathcal{U}_{1}$, and contains all edges that connect $v^{\prime}$ to the vertices of $U(B)$. As before, the ordering $\rho$ of the groups in $\mathcal{U}_{1}$ naturally defines an ordering of the bundles in $\mathcal{B}\left(v^{\prime}\right)$. We denote $\mathcal{B}\left(v^{\prime}\right)=\left\{B_{1}\left(v^{\prime}\right), B_{2}\left(v^{\prime}\right), \ldots, B_{\beta\left(v^{\prime}\right)}\left(v^{\prime}\right)\right\}$, and we assume that the bundles are indexed according to this ordering.

Lastly, since graph $\tilde{G}$ is $\left(d, d^{\prime}\right)$-semiregular, we get that $|E(\tilde{G})| \geq d^{\prime} N_{2} /\left(4 \log N_{2}\right) \geq d^{\prime} N_{2} /(4 \log M)$. Since every group in $\mathcal{U}_{2}$ contains exactly $r$ vertices of $V_{2}$, we get that:

$$
\begin{equation*}
\left|\mathcal{U}_{2}\right|=\frac{N_{2}}{r} \leq \frac{4|E(\tilde{G})| \log M}{d^{\prime} r} . \tag{6.3}
\end{equation*}
$$

We are now ready to define the instance $\hat{\mathcal{J}}=(\hat{G}, \mathcal{M})$ of NDP-Grid, from the input instance $(\tilde{G}=$ $\left.\left(V_{1}, V_{2}, E\right), \mathcal{U}_{1}, \mathcal{U}_{2}, h, r\right)$ of $(r, h)-G P w B$. Let $\lambda=2048 \cdot\left\lceil M^{2} \cdot \log M\right\rceil$. The graph $\hat{G}$ is simply the $(\lambda \times \lambda)$ grid, so $V(\hat{G})=O\left(M^{4} \log ^{2} M\right)$ as required. We now turn to define the set $\mathcal{M}$ of the demand pairs. We first define the set $\mathcal{M}$ itself, without specifying the locations of the corresponding vertices in $\hat{G}$, and later specify a mapping of all vertices participating in the demand pairs to $V(\hat{G})$.
Consider the underlying graph $\tilde{G}=\left(V_{1}, V_{2}, E\right)$ of the ( $r, h$ )-GPwB problem instance. Initially, for every edge $e=(u, v) \in E$, with $u \in V_{1}, v \in V_{2}$, we define a demand pair $(s(e), t(e))$ representing $e$, and add it to $\mathcal{M}$, so that the vertices participating in the demand pairs are all distinct. Next, we process the vertices $v \in V_{1} \cup V_{2}$ one-by-one. Consider first some vertex $v \in V_{1}$, and some bundle $B \in \mathcal{B}(v)$. Assume that $B=\left\{e_{1}, \ldots, e_{z}\right\}$. Recall that for each $1 \leq i \leq z$, set $\mathcal{M}$ currently contains a demand pair $\left(s\left(e_{i}\right), t\left(e_{i}\right)\right)$ representing $e_{i}$. We unify all vertices $s\left(e_{1}\right), \ldots, s\left(e_{z}\right)$ into a single vertex $s_{B}$. We then replace the demand pairs $\left(s\left(e_{1}\right), t\left(e_{1}\right)\right), \ldots,\left(s\left(e_{z}\right), t\left(e_{z}\right)\right)$ with the demand pairs $\left(s_{B}, t\left(e_{1}\right)\right), \ldots,\left(s_{B}, t\left(e_{z}\right)\right)$. Once we finish processing all vertices in $V_{1}$, we perform the same procedure for every vertex of $V_{2}$ : given a vertex $v^{\prime} \in V_{2}$, for every bundle $B^{\prime} \in \mathcal{B}\left(v^{\prime}\right)$, we unify all destination vertices $t(e)$ with $e \in B^{\prime}$ into a single destination vertex, that we denote by $t_{B^{\prime}}$, and we update $\mathcal{M}$ accordingly. This completes the definition of the set $\mathcal{M}$ of the demand pairs.

Observe that each edge of $e \in E$ still corresponds to a unique demand pair in $\mathcal{M}$, that we will denote by $\left(s_{B(e)}, t_{B^{\prime}(e)}\right)$, where $B(e)$ and $B^{\prime}(e)$ are the two corresponding bundles containing $e$. Given a subset $E^{\prime} \subseteq E$ of edges of $\tilde{G}$, we denote by $\mathcal{M}\left(E^{\prime}\right)=\left\{\left(s_{B(e)}, t_{B^{\prime}(e)}\right) \mid e \in E^{\prime}\right\}$ the set of all demand pairs corresponding to the edges of $E^{\prime}$.
In order to complete the reduction, we need to show a mapping of all source and all destination vertices of $\mathcal{M}$ to the vertices of $\hat{G}$. Let $R^{\prime}$ and $R^{\prime \prime}$ be two rows of the grid $\hat{G}$, lying at a distance at least $\lambda / 4$ from each other and from the top and the bottom boundaries of the grid, with $R^{\prime}$ lying above $R^{\prime \prime}$. We will map all vertices of $S(\mathcal{M})$ to $R^{\prime}$, and all vertices of $T(\mathcal{M})$ to $R^{\prime \prime}$, as follows.

Locations of the sources. Let $K_{1}, K_{2}, \ldots, K_{N_{1}}$ be a collection of $N_{1}$ disjoint subpaths of $R^{\prime}$, where each subpath contains $1024 \cdot\lceil h \cdot \log M\rceil$ vertices; the subpaths are indexed according to their left-to-right ordering on $R^{\prime}$, and every consecutive pair of the paths is separated by at least $10 M$ vertices from each other and from the left and the right boundaries of $\hat{G}$. Observe that the width $\lambda$ of the grid is large enough to allow this, as $h \leq M$ must hold. For all $1 \leq i \leq N_{1}$, we call $K_{i}$ the block representing the vertex $v_{i} \in V_{1}$. We now fix some $1 \leq i \leq N_{1}$ and consider the block $K_{i}$ representing the vertex $v_{i}$. We map the source vertices $s_{B_{1}\left(v_{i}\right)}, s_{B_{2}\left(v_{i}\right)}, \ldots, s_{B_{\beta\left(v_{i}\right)}\left(v_{i}\right)}$ to vertices of $K_{i}$, so that they appear on $K_{i}$ in this order, and every consecutive pair of sources is separated by exactly $512 \cdot\left\lceil h \cdot \log M / \beta\left(v_{i}\right)\right\rceil$ vertices.

Locations of the destinations. Similarly, we let $K_{1}^{\prime}, K_{2}^{\prime}, \ldots, K_{N_{2}}^{\prime}$ be a collection of $N_{2}$ disjoint subpaths of $R^{\prime \prime}$, each of which contains $1024 \cdot\lceil h \cdot \log M\rceil$ vertices, so that the subpaths are indexed according to their left-to-right ordering on $R^{\prime \prime}$, and every consecutive pair of the paths is separated by at least 10 M vertices from each other and from the left and the right boundaries of $\hat{G}$. We call $K_{i}^{\prime}$ the block representing the vertex $v_{i}^{\prime} \in V_{2}$. We now fix some $1 \leq i \leq N_{2}$ and consider the block $K_{i}^{\prime}$ representing the vertex $v_{i}^{\prime}$. We map the destination vertices $t_{B_{1}\left(v_{i}^{\prime}\right)}, t_{B_{2}\left(v_{i}^{\prime}\right)}, \ldots, t_{B_{\beta\left(v_{i}^{\prime}\right)}\left(v_{i}^{\prime}\right)}$ to vertices of $K_{i}^{\prime}$, so that they appear on $K_{i}^{\prime}$ in this order, and every consecutive pair of destinations is separated by exactly $512 \cdot\left\lceil h \cdot \log M / \beta\left(v_{i}^{\prime}\right)\right\rceil$ vertices. This concludes the definition of the instance $\hat{\mathcal{J}}=(\hat{G}, \mathcal{M})$ of NDP-Grid, except for the procedure for selecting the orderings $\rho$ and $\rho^{\prime}$ of the groups in $\mathcal{U}_{1}$ and $\mathcal{U}_{2}$, respectively. The following immediate observation will be useful to us.

Observation 6.1. Consider a vertex $v_{j} \in V_{1}$, and let $\mathcal{N} \subseteq \mathcal{N}$ be any subset of demand pairs, whose sources are all distinct and lie on $K_{j}$. Assume that $\mathcal{N}_{j}=\left\{\left(s_{1}, t_{1}\right), \ldots,\left(s_{y}, t_{y}\right)\right\}$, where the demand pairs are indexed according to the left-to-right ordering of their source vertices on $K_{j}$. Then $t_{1}, \ldots, t_{y}$ appear in this left-to-right order on $R^{\prime \prime}$.

Choosing the orderings of the groups. We now show an algorithm for selecting the ordering $\rho$ of $\mathcal{U}_{1}$ and $\rho^{\prime}$ of $\mathcal{U}_{2}$ that are used in the reduction. We let $\rho$ be an arbitrary ordering of the groups in $\mathcal{U}_{1}$. Before we describe an algorithm for computing the ordering $\rho^{\prime}$ of $\mathcal{U}_{2}$, we first define the properties that we need from it.

Assume that we are given some ordering $\rho^{\prime}$ over the groups in $\mathcal{U}_{2}$. Consider first some vertex $v_{i} \in V_{1}$, and recall that we have denoted its bundles by $B_{1}\left(v_{i}\right), \ldots, B_{\beta\left(v_{i}\right)}$. Recall that every bundle $B_{j}\left(v_{i}\right)$ contains all edges connecting $v_{i}$ to a single group of $\mathcal{U}_{2}$, that we denote, for convenience, by $U_{j}^{i}$. We assume that the bundles are indexed so that $U_{1}^{i}, \ldots, U_{\beta_{i}\left(v_{i}\right)}^{i}$ appear in $\rho^{\prime}$ in this order. Recall that, from Equation 6.1, $d / b \leq \beta\left(v_{i}\right)<2 d / b$. Intuitively, we will only focus on the first $\lceil d /(2 b)\rceil$ bundles of $v_{i}$, and we will never attempt to route demand pairs that correspond to the remaining bundles. We say that vertex $v_{i}$ is happy with the ordering $\rho^{\prime}$ of $\mathcal{U}_{2}$, if, for all $1 \leq j<\lceil d /(2 b)\rceil$, the total number of groups that appear strictly between $U_{j}^{i}$ and $U_{j}^{i+1}$ in $\rho^{\prime}$ is at most $O\left(\frac{\left|\mathcal{U}_{2}\right| b \log ^{2} M}{d}\right)$ (notice that the indexing of the bundles $B_{j}\left(v_{i}\right)$ and of the corresponding groups $U_{j}^{i}$ depends on $\rho^{\prime}$ ). We say that ordering $\rho^{\prime}$ is good iff every vertex $v_{i} \in V_{1}$ is happy with this ordering. Intuitively, choosing an ordering $\rho^{\prime}$ that is good will ensure that the distance property mentioned above holds, which will allow us to route a large set of the demand pairs in the case

## Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat

where there is a perfect solution to the $(r, h)-G P w B$ instance. We now show that we can compute such an ordering $\rho^{\prime}$ efficiently.

Lemma 6.2. There is an algorithm whose running time is polynomial in $|V(\tilde{G})|$, that computes a good ordering $\rho^{\prime}$ of $\mathcal{U}_{2}$.

Proof. The algorithm will compute a partition $\mathcal{U}^{1}, \mathcal{U}^{2}, \ldots, \mathcal{U}^{q}$ of $\mathcal{U}_{2}$ into subsets, each of which has cardinality $O\left(\frac{\| \mathcal{U}_{2} \mid b \log ^{2} M}{d}\right)$. In the final ordering $\rho^{\prime}$, groups of $\mathcal{U}_{2}$ that lie within a single set $\mathcal{U}^{i}$ are placed consecutively, in an arbitrary order, while groups lying in different sets are ordered according to the ordering $U^{1}, \mathfrak{U}^{2}, \ldots, \mathfrak{U}^{q}$ of these sets.
The algorithm consists of a number of iterations, where in the $i$ th iteration we compute the subset $\mathcal{U}^{i} \subseteq \mathcal{U}_{2}$ of groups. We now describe a single iteration. Recall that for every vertex $v \in V_{1}$, we have defined a collection of groups that cover $v$ - these are all groups $U \in \mathcal{U}_{2}$ such that there is a bundle $B \in \mathcal{B}(v)$, connecting $v$ to vertices of $U$. The total number of groups that cover $v$ is $\beta(v) \geq d / b$ from Equation 6.1. We say that vertex $v$ is active in iteration $i$ iff the number of groups that cover $u$ and belong to $\mathcal{U}^{1} \cup \cdots \cup \bigcup^{i-1}$ is less than $\lceil d /(2 b)\rceil$. Let $\mathcal{U}^{\prime}=\mathcal{U}_{2} \backslash\left(\mathcal{U}^{1} \cup \cdots \cup \chi^{i-1}\right)$, and let $V^{\prime} \subseteq V_{1}$ be the set of all vertices that are active in iteration $i$. Recall that by the definition of active vertices, for each $v \in V^{\prime}$, there are at least $d /(2 b)$ groups in $U^{\prime}$ that cover $v$. We next show that we can efficiently compute a subset $\mathcal{U}^{i} \subseteq \mathcal{U}_{2}$ of $O\left(\frac{\left|U_{2}\right| b \log ^{2} M}{d}\right)$ groups, such that for every active vertex $v \in \mathcal{U}^{\prime}$, at least one group in $\mathcal{U}^{i}$ covers $v$. Note that, if we manage to select the groups $\mathcal{U}^{1}, \mathcal{U}^{2}, \ldots, \mathcal{U}^{q}$ with this property, such that $\bigcup_{i} \mathcal{U}^{i}=\mathcal{U}_{2}$, then it is immediate to verify that the resulting ordering $\rho^{\prime}$ of $\mathcal{U}_{2}$ is a good ordering. This is since, for every vertex $v \in V_{1}$, if $v$ was active for $i$ iterations, then $U^{1} \cup \cdots \cup U^{i}$ contain at least $\lceil d /(2 b)\rceil$ groups that cover $v$, and every set $u^{i^{\prime}}$ for $1 \leq i^{\prime} \leq i$ contains at least one such group. The following claim will then complete the proof of the lemma.

Claim 6.3. There is an efficient algorithm, that, given an integer $i>0$, collections $\chi^{1}, \ldots, \chi^{i-1}$ of subsets of $\mathcal{U}_{2}$, and the set $V^{\prime} \subseteq V_{1}$ of vertices that are active in iteration $i$, such that $V^{\prime} \neq \emptyset$, computes a collection $\mathcal{U}^{i} \subseteq \mathcal{U}_{2} \backslash\left(\mathcal{U}^{1} \cup \cdots \cup U^{i-1}\right)$ of $O\left(\frac{\left|\mathcal{U}_{2}\right| b \log M}{d}\right)$ groups, such that for every active vertex $v \in \mathcal{U}^{\prime}$, at least one group in $U^{i}$ covers $v$.

Proof. For convenience, we denote $\mathcal{U}^{\prime}=\mathcal{U}_{2} \backslash\left(\mathcal{U}^{1} \cup \cdots \cup \mathcal{U}^{i-1}\right)$. The algorithm exploits the Set Cover problem. In this problem, we are given as input a collection $\{1, \ldots, n\}$ of elements, and a collection $\mathcal{S}=\left\{S_{1}, \ldots, S_{m}\right\}$ of subsets of elements, so for all $1 \leq j \leq m, S_{j} \subseteq\{1, \ldots, m\}$. The goal is to compute a minimum cardinality collection $\mathcal{S}^{\prime} \subseteq \mathcal{S}$ of subsets, such that every element in $\{1, \ldots, n\}$ belongs to at least one subset in $\mathcal{S}^{\prime}$. We use the following observation, whose proof appears in Appendix ${ }^{4}$.

Observation 6.4. There is an efficient algorithm, that, given an instance of the Set Cover problem with $n$ elements and $m$ sets, such that every element belongs to at least $m / t$ sets, computes a solution containing at most $O(t \log n)$ sets.

[^3]In order to complete the proof of Claim 6.3, we set up an instance of the Set Cover problem, where every group in $U^{\prime}$ represents a set, every vertex in $V^{\prime}$ is an element, and element $v \in V^{\prime}$ belongs to a set $S_{U}$ representing a group $U \in U^{\prime}$ if and only if $U$ covers $v$. Recall that we are guaranteed that every element belongs to at least $d /(2 b)$ sets. Setting $t=\left|U^{\prime}\right| \cdot(2 b) / d$, and using the algorithm from Observation 6.4, we obtain a solution to the resulting instance of the Set Cover problem of value $O\left(\frac{\left|\mathcal{U}_{2}\right| b \log N_{1}}{d}\right) \leq O\left(\frac{\left|\mathcal{U}_{2}\right| b \log M}{d}\right)$, which immediately defines the desired set $U^{i} \subseteq \mathcal{U}^{\prime}$ of groups.

Our algorithm employs Claim 6.3 to construct the sets $U^{1}, U^{2}, \ldots$ of groups, until the set $V^{\prime}$ of vertices that are active becomes empty. Each subsequent set $\mathcal{U}^{i}$ then contains an arbitrary subset of $O\left(\frac{\left|\mathcal{U}_{2}\right| b \log M}{d}\right)$ groups of $\mathcal{U}_{2} \backslash\left(\mathcal{U}^{1} \cup \cdots \cup \mathfrak{U}^{i-1}\right)$.

This concludes the description of the reduction from (r,h)-GPwB to NDP-Grid. In order to complete the proof of Theorem 4.3, it is now enough to prove the following two theorems, which is done in the next two subsections.

Theorem 6.5. Suppose we are given a valid and $\left(d, d^{\prime}, b\right)$-regular instance $\mathcal{J}=\left(\tilde{G}=\left(V_{1}, V_{2}, E\right), \mathcal{U}_{1}, \mathcal{U}_{2}, h, r\right)$ of $(\mathrm{r}, \mathrm{h})-\mathrm{GPwB}$, such that $\mathcal{J}$ has a perfect solution. Then there is a solution to instance $\hat{\mathcal{J}}$ of NDP-Grid, routing $\Omega\left(\beta^{*}(\mathcal{J}) / \log ^{2} M\right)$ demand pairs via a set of spaced-out paths.

Theorem 6.6. There is an efficient algorithm, that, given a valid and ( $d, d^{\prime}, b$ )-regular instance $\mathcal{J}=$ $\left(\tilde{G}=\left(V_{1}, V_{2}, E\right), \mathcal{U}_{1}, \mathcal{U}_{2}, h, r\right)$ of the $(\mathrm{r}, \mathrm{h})$-GPwB problem with $|E|=M$, the corresponding instance $\hat{J}$ of NDP-Grid, and a solution $\mathcal{P}^{*}$ to $\hat{\mathcal{J}}$, computes a solution to the $(\mathrm{r}, \mathrm{h})-\mathrm{GPwB}$ instance $\mathcal{J}$ of value at least $\Omega\left(\left|\mathcal{P}^{*}\right| / \log ^{3} M\right)$.

### 6.3 From partitioning to routing: proof of Theorem 6.5

The goal of this subsection is to prove Theorem 6.5.
Let $\left(\left(W_{1}, \ldots, W_{r}\right),\left(E_{1}, \ldots, E_{r}\right)\right)$ be a perfect solution to instance $\mathcal{J}$ of $(r, h)-G P w B$. Recall that by the definition of a perfect solution, for each group $U \in \mathcal{U}_{1} \cup \mathcal{U}_{2}$, every set $W_{i}$ contains exactly one vertex of $U$, and moreover, for each $1 \leq i \leq r,\left|E_{i}\right|=h=\beta^{*}(\mathcal{J}) / r$.
We let $E^{0}=\bigcup_{i=1}^{r} E_{i}$, so $\left|E^{0}\right|=\beta^{*}(\mathcal{J})=h r$. Let $\mathcal{N}^{0} \subseteq \mathcal{M}$ be the set of all demand pairs corresponding to the edges of $E^{0}$. Note that we are guaranteed that no two demand pairs in $\mathcal{M}^{0}$ share a source or a destination, since no two edges of $E^{0}$ belong to the same bundle. We now define a slightly smaller subset $E^{1} \subseteq E^{0}$ of edges, and a corresponding set $\mathcal{M}^{1} \subseteq \mathcal{M}^{0}$ of demand pairs. Since the value of the solution $\left(\left(W_{1}, \ldots, W_{r}\right),\left(E_{1}, \ldots, E_{r}\right)\right)$ of the $(r, \mathrm{~h})$-GPwB problem instance is $h r=\beta^{*}(\mathcal{J})=\sum_{v \in V_{1}} \beta(v)$, for every vertex $v \in V_{1}$, there are exactly $\beta(v)$ edges in $E^{0}$ - one edge from each bundle $B \in \mathcal{B}(v)$. Recall that the total number of such bundles is $\beta(v)<2 d / b$ from Equation 6.1. Recall that every bundle $B \in \mathcal{B}(v)$ corresponds to a distinct group $U(B) \in \mathcal{U}_{2}$, that covers $v$. Moreover, we have used the ordering $\rho^{\prime}$ of the groups in $U_{2}$ in order to define an ordering of the bundles in $\mathcal{B}(v)$. Let $\left\{e_{1}^{i}, \ldots, e_{\beta\left(v_{i}\right)}^{i}\right\}$ be the set of all edges of $E^{0}$ that are incident to vertex $v_{i}$. We assume that these edges are indexed according to the
ordering of their corresponding bundles. We discard all but the first $\lceil d /(2 b)\rceil$ edges from this set. We then denote by $E^{1} \subseteq E^{0}$ the resulting set of edges, obtained after processing every vertex in $V_{1}$. It is easy to verify that $\left|E^{1}\right|=\Omega\left(\left|E^{0}\right|\right)$. We let $\mathcal{M}^{1} \subseteq \mathcal{M}^{0}$ be the set of all demand pairs corresponding to the edges of $E^{1}$.

Assume now that we are given any subset $\mathcal{N}^{\prime} \subseteq \mathcal{N}^{1}$ of demand pairs. We define an ordering $\eta\left(\mathcal{N}^{\prime}\right)$ of the demand pairs in $\mathcal{M}^{\prime}$, by exploiting the perfect solution $\left(\left(W_{1}, \ldots, W_{r}\right),\left(E_{1}, \ldots, E_{r}\right)\right)$ to instance $\mathcal{J}$ of $(r, h)-G P w B$, as follows. Recall that every demand pair $(s, t) \in \mathcal{M}^{1}$ corresponds to some edge $e \in E_{1} \cup \cdots \cup E_{r}$, that we denote by $e(s, t)$. First, we partition the demand pairs in $\mathcal{M}^{\prime}$ into subsets $\mathcal{M}_{1}^{\prime}, \ldots, \mathcal{M}_{r}^{\prime}$, as follows: demand pair $(s, t) \in \mathcal{M}^{\prime}$ is added to subset $\mathcal{M}_{i}^{\prime}$ iff the corresponding edge $e(s, t)$ lies in $E_{i}$. The demand pairs that lie in different sets $\mathcal{M}_{1}^{\prime}, \ldots, \mathcal{M}_{r}^{\prime}$ are then ordered according to the natural order of these sets. Consider now some set $\mathcal{M}_{i}^{\prime}$ of demand pairs. We order the demand pairs in $\mathcal{M}_{i}^{\prime}$ according to the left-to-right order of their destination vertices on row $R^{\prime \prime}$. Let $\eta\left(\mathcal{M}^{\prime}\right)$ be the resulting ordering of the demand pairs in $\mathcal{N}^{\prime}$. Next, we define a property of a subset of the demand pairs, called a distance property. We later show that, if we are given a subset $\mathcal{M}^{\prime} \subseteq \mathcal{M}^{1}$ of demand pairs that has the distance property, then all demand pairs in $\mathcal{M}^{\prime}$ can be routed via spaced-out paths. We also show that there is a large subset $\mathcal{M}^{\prime} \subseteq \mathcal{M}^{1}$ of the demand pairs that has the the distance property.

We now proceed to define a distance property. We assume that we are given a set $\mathcal{M}^{\prime} \subseteq \mathcal{H}^{1}$ of demand pairs, and consider the corresponding ordering $\eta=\eta\left(\mathcal{M}^{\prime}\right)$ of these demand pairs. We denote $\mathcal{M}^{\prime}=$ $\left\{\left(s_{1}, t_{1}\right), \ldots,\left(s_{\left|\mathcal{M}^{\prime}\right|}, t_{\left|\mathcal{M}^{\prime}\right|}\right)\right\}$, where the demand pairs are indexed according to their ordering in $\eta$. Notice that this ordering is not necessarily the same as the ordering of the set $S\left(\mathcal{M}^{\prime}\right)$ of source vertices on row $R^{\prime}$, or the ordering of the set $T\left(\mathcal{M}^{\prime}\right)$ of the destination vertices on row $R^{\prime \prime}$.

Consider some pair $s, s^{\prime} \in S\left(\mathcal{M}^{\prime}\right)$ of source vertices, and assume that $s=s_{i}$ and $s^{\prime}=s_{j}$. We denote $N_{\mathcal{M}^{\prime}}\left(s, s^{\prime}\right)=|i-j|-1$ - the number of the demand pairs in the ordering $\eta$, that lie between $\left(s_{i}, t_{i}\right)$ and $\left(s_{j}, t_{j}\right)$.
Definition 6.7. Given a set $\mathcal{M}^{\prime} \subseteq \mathcal{M}^{1}$ of demand pairs, we say that two source vertices $s, s^{\prime} \in S\left(\mathcal{M}^{\prime}\right)$ are consecutive with respect to $\mathcal{M}^{\prime}$, iff $s \neq s^{\prime}$, and no vertex of $S\left(\mathcal{M}^{\prime}\right)$ lies strictly between $s$ and $s^{\prime}$ on row $R^{\prime}$. We say that $\mathcal{M}^{\prime}$ has the distance property iff for every pair $s, s^{\prime} \in S\left(\mathcal{M}^{\prime}\right)$ of source vertices that are consecutive with respect to $\mathcal{M}^{\prime}, N_{\mathcal{M}^{\prime}}\left(s, s^{\prime}\right)<d\left(s, s^{\prime}\right) / 4$ holds, where $d\left(s, s^{\prime}\right)$ is the distance between $s$ and $s^{\prime}$ in graph $\hat{G}$.

In order to complete the proof of Theorem 6.5, we proceed in two steps. First, we show that there is a large enough subset $\mathcal{N}^{\prime} \subseteq \mathcal{M}^{1}$ of demand pairs that has the distance property. Next, we show that any such set $\mathcal{N}^{\prime}$ of demand pairs can be routed via spaced-out paths.

Lemma 6.8. There is a subset $\mathcal{N}^{\prime} \subseteq \mathcal{M}^{1}$ of demand pairs that has the distance property, and $\left|\mathcal{M}^{\prime}\right|=$ $\Omega\left(\left|\mathcal{M}^{1}\right| / \log ^{2} M\right)=\Omega\left(\left|\mathcal{M}^{0}\right| / \log ^{2} M\right)$.
Lemma 6.9. Assume that $\mathcal{N}^{\prime} \subseteq \mathcal{M}^{1}$ is a subset of demand pairs that has the distance property. Then there is a spaced-out set $\mathcal{P}$ of paths routing all pairs of $\mathcal{M}^{\prime}$ in graph $\hat{G}$.

The above two lemmas finish the proof of Theorem 6.5, since $\left|\mathcal{M}^{0}\right|=\beta^{*}(\mathcal{J})$. We prove these lemmas in the following two subsections.

### 6.3.1 Proof of Lemma 6.8

Let $\eta^{1}=\eta\left(\mathcal{M}^{1}\right)$ be the ordering of the demand pairs in $\mathcal{M}^{1}$. We show that this ordering almost has the distance property: if $s, s^{\prime}$ are two consecutive sources with respect to $\mathcal{M}^{1}$, then $N_{\mathcal{M}^{1}} \leq O\left(d\left(s, s^{\prime}\right) \log ^{2} M\right)$ holds. We then use a simple procedure in order to select a subset $\mathcal{N}^{\prime} \subseteq \mathcal{M}^{1}$ of $\Omega\left(\left|\mathcal{M}^{1}\right| / \log ^{2} M\right)$ demand pairs that has the distance property.

Analyzing ordering $\eta^{1}$. Recall that ordering $\eta^{1}$ of the demand pairs in $\mathcal{N}^{1}$ was defined using the perfect solution $\left(\left(W_{1}, \ldots, W_{r}\right),\left(E_{1}, \ldots, E_{r}\right)\right)$ to instance $\mathcal{J}$ of $(r, h)-G P w B$. We have partitioned the demand pairs in $\mathcal{M}^{1}$ into subsets $\mathcal{M}_{1}, \ldots, \mathcal{M}_{r}$, by using the partition $\left\{E_{1}, \ldots, E_{r}\right\}$ of their corresponding edges. The demand pairs that lie in different sets $\mathcal{M}_{1}, \ldots, \mathcal{M}_{r}$ are then ordered according to the natural order of these sets, while the demand pairs in each set $\mathcal{M}_{i}$ are ordered according to the left-to-right order of their destination vertices on row $R^{\prime \prime}$. We need the following claim.

Claim 6.10. Let $s, s^{\prime} \in S\left(\mathcal{N}^{1}\right)$ be any pair of source vertices that are consecutive with respect to $\mathcal{N}^{1}$. Then $N_{\mathcal{M}^{1}}\left(s, s^{\prime}\right) \leq O\left(d\left(s, s^{\prime}\right) \log ^{2} M\right)$.

Proof. First, if $s$ and $s^{\prime}$ belong to distinct blocks $K_{i}$ on row $R^{\prime}$, then, since the distance between every pair of blocks is at least $M$, while $\left|\mathcal{M}^{1}\right| \leq M$, the assertion clearly holds. Therefore, we can assume that there is some block $K_{i}$, with $s, s^{\prime} \in K_{i}$. Recall that block $K_{i}$ represents vertex $v_{i} \in V_{1}$. Assume w.l.o.g. that $v_{i} \in W_{j}$, in the solution to the ( $\mathrm{r}, \mathrm{h}$ )-GPwB problem.
Let $(s, t),\left(s^{\prime}, t^{\prime}\right) \in \mathcal{M}^{1}$ be the demand pairs corresponding to these source vertices; the demand pairs are uniquely defined, since the pairs in $\mathcal{N}^{1}$ do not share their source or their destination vertices. Let $e$ be the edge of $\tilde{G}$ corresponding to the demand pair $(s, t)$, and let $e^{\prime}$ be the edge of $\tilde{G}$ corresponding to the demand pair $\left(s^{\prime}, t^{\prime}\right)$. Then $e$ and $e^{\prime}$ are both incident to the vertex $v_{i}$, and they lie in distinct bundles of $\mathcal{B}\left(v_{i}\right)$. Recall that we have used the ordering $\rho^{\prime}$ of the groups in $\mathcal{U}_{2}$ in order to define an ordering of the bundles in $\mathcal{B}\left(v_{i}\right)$, which we used in turn in order to define an ordering $\left\{e_{i}^{1}, \ldots, e_{i}^{\beta\left(v_{i}\right)}\right\}$ of the edges incident to $v_{i}$ that belong to $E^{0}$. Only the first $\lceil d /(2 b)\rceil$ of these edges were added to $E^{1}$. Moreover, since the solution to instance $\mathcal{J}$ of $(r, h)-G P w B$ that we have considered is perfect, every bundle incident to $v_{i}$ contributed exactly one edge to $E^{0}$. Therefore, if we denote $e=e_{i}^{z}$, then $e^{\prime}=e_{i}^{z+1}$ (or the other way around; we assume it is the former), and $z+1 \leq\lceil d /(2 b)\rceil$. Let $U \in \mathcal{U}_{2}$ be the group corresponding to the bundle of $e$; that is, $e$ connects $v_{i}$ to a vertex of $U$. Let $U^{\prime}$ be defined similarly to edge $e^{\prime}$. From our definition of a good ordering, there are at most $O\left(\frac{\left|U_{2}\right| b \log ^{2} M}{d}\right)$ groups that lie between $U$ and $U^{\prime}$ in the ordering $\rho^{\prime}$. We will use this fact in order to bound $N_{\mathcal{M}^{1}}\left(s, s^{\prime}\right)$, that we denote, for convenience, by $N\left(s, s^{\prime}\right)$.
Let $\left(s^{\prime \prime}, t^{\prime \prime}\right) \in \mathcal{M}^{1}$ be any demand pair that lies between $(s, t)$ and $\left(s^{\prime}, t^{\prime}\right)$ in the ordering $\eta^{1}$. Then the edge $e^{\prime \prime} \in E^{1}$ corresponding to this demand pair must also lie in set $E_{j}$ that corresponds to the cluster $W_{j}$ in the solution to instance $\mathcal{J}$ of $(r, h)-G P w B$, and moreover, $t^{\prime \prime}$ must lie between $t$ and $t^{\prime}$ on row $R^{\prime \prime}$. If $t^{\prime \prime}$ lies in some block $K_{x}^{\prime}$, that represents a vertex $v_{x}^{\prime} \in V_{2}$, then the group $U^{\prime \prime}$ to which $v_{x}^{\prime}$ belongs must lie between $U$ and $U^{\prime}$ in the ordering $\rho^{\prime}$. Recall that, since $\rho^{\prime}$ is a good ordering of the groups in $\mathcal{U}_{2}$, there are at most

Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat

$O\left(\frac{\left|u_{2}\right| b \log ^{2} M}{d}\right)$ groups that lie between $U$ and $U^{\prime}$ in the ordering $\rho^{\prime}$. Let us now consider any such group $U^{\prime \prime}$, that consists of $r$ vertices. Exactly one vertex of $U^{\prime \prime}$ lies in $W_{j}$; we denote it by $v^{\prime}\left(U^{\prime \prime}\right)$. This vertex has degree at most $d^{\prime}$ in $\tilde{G}$, and so $\mathcal{B}\left(v^{\prime}\left(U^{\prime \prime}\right)\right)$ contains at most $d^{\prime} / b$ bundles. Therefore, at most $d^{\prime} / b$ edges that are incident to $v^{\prime}\left(U^{\prime \prime}\right)$ belong to $E^{1}$. Altogether, we conclude that:

$$
\begin{aligned}
N\left(s, s^{\prime}\right) & \leq O\left(\frac{\left|\mathcal{U}_{2}\right| b \log ^{2} M}{d}\right) \cdot \frac{d^{\prime}}{b} \\
& \leq O\left(\frac{\left|\mathcal{U}_{2}\right| d^{\prime} \log ^{2} M}{d}\right) \\
& \leq O\left(\frac{|E(\tilde{G})| \log ^{3} M}{d r}\right) \\
& \leq O\left(\frac{h b \log ^{3} M}{d}\right)
\end{aligned}
$$

(we have used Equation 6.3 for the penultimate inequality and Equation 6.2 for the last inequality.)
Recall that, from our construction, $d\left(s, s^{\prime}\right)=512 \cdot\left\lceil h \cdot \log M / \beta\left(v_{i}\right)\right\rceil \geq 256 h b \log M / d$, since, from Equation 6.1, $\beta(v)<2 d / b$. We conclude that $N_{\mathcal{M}^{1}}\left(s, s^{\prime}\right) \leq O\left(d\left(s, s^{\prime}\right) \log ^{2} M\right)$

We obtain the following immediate corollary of Claim 6.10.
Corollary 6.11. There is a constant $c$, such that, for any pair $s, s^{\prime} \in S\left(\mathcal{M}^{1}\right)$ of source vertices (that are not necessarily consecutive with respect to $\left.\mathcal{M}^{1}\right), N_{\mathcal{M}^{1}}\left(s, s^{\prime}\right) \leq c d\left(s, s^{\prime}\right) \log ^{2} M$ holds.

Defining $\mathcal{M}^{\prime}$. We denote $\mathcal{M}^{1}=\left\{\left(s_{1}, t_{1}\right), \ldots,\left(s_{\left|\mathcal{M}^{1}\right|}, t_{\left|\mathcal{M}^{1}\right|}\right)\right\}$, where the demand pairs are indexed according to the ordering $\eta^{1}$. We now let:

$$
\mathcal{M}^{\prime}=\left\{\left(s_{i}, t_{i}\right) \mid i=1 \quad \bmod \left\lceil 8 c \log ^{2} M\right\rceil\right\} .
$$

We let $\eta$ be the ordering of the demand pairs in $\mathcal{M}^{\prime}$, induced by $\eta^{1}$. It is immediate to verify that $\left|\mathcal{M}^{\prime}\right|=\Omega\left(\left|\mathcal{M}^{1}\right| / \log ^{2} M\right)=\Omega\left(\left|\mathcal{M}^{0}\right| / \log ^{2} M\right)$, and that $\eta=\eta\left(\mathcal{M}^{\prime}\right)$. Moreover, if $s, s^{\prime} \in S\left(\mathcal{M}^{\prime}\right)$ is any pair of source vertices, then $N_{\mathcal{M}^{\prime}}\left(s, s^{\prime}\right) \leq N_{\mathcal{M}^{1}} /\left\lceil 8 c \log ^{2} M\right\rceil \leq d\left(s, s^{\prime}\right) / 4$ from Corollary 6.11. Therefore, $\mathcal{M}^{\prime}$ has the distance property.

### 6.3.2 Proof of Lemma 6.9

Recall that $R^{\prime}$ and $R^{\prime \prime}$ are the rows of $\hat{G}$ containing the vertices of $S(\mathcal{M})$ and $T(\mathcal{M})$ respectively, and that $R^{\prime}$ and $R^{\prime \prime}$ lie at distance at least $\lambda / 4$ form each other and from the top and the bottom boundaries of $\hat{G}$,
where $\lambda>M^{2}$ is the dimension of the grid. Let $R$ be any row lying between $R^{\prime}$ and $R^{\prime \prime}$, within distance at least $\lambda / 16$ from each of them.

We denote $M^{\prime}=\left|\mathcal{M}^{\prime}\right|$, and $\mathcal{M}^{\prime}=\left\{\left(s_{1}, t_{1}\right), \ldots,\left(s_{M^{\prime}}, t_{M^{\prime}}\right)\right\}$, where the pairs are indexed according to their ordering in $\eta=\eta\left(\mathcal{M}^{\prime}\right)$. To recap, the ordering $\eta$ of $\mathcal{M}^{\prime}$ was defined as follows. We have defined a partition $\mathcal{M}_{1}, \ldots, \mathcal{M}_{r}$ of the demand pairs in $\mathcal{M}^{\prime}$, according to the partition $E_{1}, \ldots, E_{r}$ of their corresponding edges. For each $1 \leq i \leq r$, the demand pairs in $\mathcal{M}_{i}$ appear consecutively in $\eta$, ordered according to the left-to-right ordering of their destination vertices on row $R^{\prime \prime}$, while the order of the demand pairs lying in different sets $\mathcal{M}_{i}$ follows the indexing of these subsets.
Let $X=\left\{x_{i} \mid 1 \leq i \leq M^{\prime}\right\}$ be a set of vertices of $R$, where $x_{i}$ is the (2i)th vertex of $R$ from the left. We will construct a set $\mathcal{P}^{*}$ of spaced-out paths routing all demand pairs in $\mathcal{M}^{\prime}$, so that the path $P_{i}$ routing $\left(s_{i}, t_{i}\right)$ intersects the row $R$ at exactly one vertex - the vertex $x_{i}$. Notice that row $R$ partitions the grid $\hat{G}$ into two subgrids: a top subgrid $G^{t}$ spanned by all rows that appear above $R$ (including $R$ ), and a bottom subgrid $G^{b}$ spanned by all rows that appear below $R$ (including $R$ ).
It is now enough to show that there are two sets of spaced-out paths: set $\mathcal{P}^{1}$ routing all pairs in $\left\{\left(s_{i}, x_{i}\right) \mid 1 \leq i \leq M^{\prime}\right\}$ in the top grid $G^{t}$, and set $\mathcal{P}^{2}$ routing all pairs in $\left\{\left(t_{i}, x_{i}\right) \mid 1 \leq i \leq M^{\prime}\right\}$ in the bottom grid $G^{b}$, so that both sets of paths are internally disjoint from $R$.

Routing in the top grid. Consider some vertex $v_{j} \in V_{1}$, and the corresponding block $K_{j}$. We construct a subgrid $\hat{K}_{j}$ of $\hat{G}$, containing $K_{j}$, that we call a box, as follows. Let $\mathcal{C}_{j}$ be the set of all columns of $\hat{G}$ intersecting $K_{j}$. We augment $\mathcal{C}_{j}$ by adding $2 M$ columns lying immediately to the left and $2 M$ columns lying immediately to the right of $\mathcal{C}_{j}$, obtaining a set $\hat{\mathcal{C}}_{j}$ of columns. Let $\hat{\mathcal{R}}$ contain three rows: row $R^{\prime}$; the row lying immediately above $R^{\prime}$; and the row lying immediately below $R^{\prime}$. Then box $\hat{K}_{j}$ is the subgrid of $\hat{G}$ spanned by the rows in $\hat{\mathcal{R}}$ and the columns in $\hat{\mathcal{C}}_{j}$. Since every block is separated by at least $10 M$ columns from every other block, as well as from the left and the right boundaries of $\hat{G}$, the resulting boxes are all disjoint, and every box is separated by at least $2 M$ columns of $\hat{G}$ from every other box, and from the left and the right boundaries of $\hat{G}$.
We will initially construct a set $\mathcal{P}^{1}$ of spaced-out paths in $G^{t}$, such that each path $P_{i} \in \mathcal{P}^{1}$, for $1 \leq i \leq M^{\prime}$, originates from the vertex $x_{i}$, and visits the boxes $\hat{K}_{1}, \hat{K}_{2}, \ldots, \hat{K}_{N_{1}}$ in this order. We will ensure that each such path $P_{i}$ contains the corresponding source vertex $s_{i}$. Eventually, by suitably truncating each such path $P_{i}$, we will ensure that it connects $x_{i}$ to $s_{i}$.

Claim 6.12. Consider some vertex $v_{j} \in V_{1}$, and the corresponding box $\hat{K}_{j}$. Denote $Y_{j}=S\left(\mathcal{M}^{\prime}\right) \cap V\left(K_{j}\right)$, $M_{j}=\left|Y_{j}\right|$, and assume that $Y_{j}=\left\{s_{i_{1}}, s_{i_{2}}, \ldots, s_{i_{M_{j}}}\right\}$, where the indexing of the vertices of $Y_{j}$ is consistent with the indexing of the vertices of $S\left(\mathcal{M}^{\prime}\right)$ induced by the ordering $\eta$ of the demand pairs in $\mathcal{N}^{\prime}$, and $i_{1}<i_{2}<\ldots<i_{M_{j}}$. Then there is a set $\mathcal{W}_{j}$ of $M^{\prime}$ columns of the box $\hat{K}_{j}$, such that:

- set $\mathcal{W}_{j}$ does not contain a pair of consecutive columns; and
- for each $1 \leq z \leq M_{j}$, the $i_{z}$ th column of $\mathcal{W}_{j}$ from the left contains the source vertex $s_{i_{z}}$.


## Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat

Proof. Observe that from Observation 6.1, the vertices $s_{i_{1}}, s_{i_{2}}, \ldots, s_{i_{j}}$ must appear in this left-to-right order on $R^{\prime}$, while the vertices $x_{i_{1}}, x_{i_{2}}, \ldots, x_{i_{M_{j}}}$ appear in this left-to-right order on $R$. Moreover, for all $1 \leq z<M_{j}, i_{z+1}-i_{z}-1=N_{\mathcal{M}^{\prime}}\left(s_{i_{z}}, s_{i_{z+1}}\right) \leq d\left(s_{i_{z}}, s_{i_{z+1}}\right) / 4$, from the distance property of $\mathcal{M}_{j}$. We add to $\mathcal{W}_{j}$ all columns of $\hat{K}_{j}$ where the vertices of $Y_{j}$ lie. For each $1 \leq z<M_{j}$, we also add to $\mathcal{W}_{j}$ an arbitrary set of $\left(i_{z+1}-i_{z}-1\right)$ columns lying between the column of $s_{i_{z}}$ and the column of $s_{i_{z+1}}$, so that no pair of columns in $\mathcal{W}_{j}$ is consecutive. Finally, we add to $\mathcal{W}_{j}\left(i_{1}-1\right)$ columns that lie to the left of the column of $s_{i_{1}}$, and $\left(M^{\prime}-i_{M_{j}}\right)$ columns that lie to the right of the column of $s_{i_{j}}$. We make sure that no pair of columns in $\mathcal{W}_{j}$ is consecutive - it is easy to see that there are enough columns to ensure that.

Claim 6.13. There is a set $\mathcal{P}^{1}=\left\{P_{1}, \ldots, P_{M^{\prime}}\right\}$ of spaced-out paths in $G^{t}$, that are internally disjoint from $R$, such that for each $1 \leq i \leq M^{\prime}$, path $P_{i}$ originates from vertex $x_{i}$, and for all $1 \leq j \leq N_{1}$, it contains the ith column of $\mathcal{W}_{j}$; in particular, it contains vertex $s_{i}$.

We defer the proof of the claim to Section C of the appendix; see Figure 4 for an illustration of the routing.


Figure 4: Routing in the top grid.

Routing in the bottom grid. Consider some vertex $v_{j}^{\prime} \in V_{2}$, and the corresponding block $K_{j}^{\prime}$. We construct a box $\hat{K}_{j}^{\prime}$ containing $K_{j}^{\prime}$ exactly as before. As before, the resulting boxes are all disjoint, and every box is separated by at least $2 M$ columns of $\hat{G}$ from every other box, and from the left and the right boundaries of $\hat{G}$.
As before, we will initially construct a set $\mathcal{P}^{2}$ of spaced-out paths in $G^{b}$, such that each path $P_{i}^{\prime} \in \mathcal{P}^{2}$, for $1 \leq i \leq M^{\prime}$, originates from the vertex $x_{i}$, and visits the boxes $\hat{K}_{1}^{\prime}, \hat{K}_{2}^{\prime}, \ldots, \hat{K}_{N_{2}}^{\prime}$ in turn. We will ensure that each such path $P_{i}^{\prime}$ contains the corresponding destination vertex $t_{i}$. Eventually, by suitably truncating each such path $P_{i}^{\prime}$, we will ensure that it connects $x_{i}$ to $t_{i}$.

Claim 6.14. Consider some vertex $v_{j}^{\prime} \in V_{2}$, and the corresponding box $\hat{K}_{j}^{\prime}$. Denote $Y_{j}^{\prime}=T\left(\mathcal{M}^{\prime}\right) \cap V\left(K_{j}^{\prime}\right)$, $M_{j}^{\prime}=\left|Y_{j}^{\prime}\right|$, and assume that $Y_{j}^{\prime}=\left\{t_{i_{1}}, t_{i_{2}}, \ldots, t_{M_{M_{j}^{\prime}}}\right\}$, where the indexing of the vertices of $Y_{j}^{\prime}$ is consistent
with the ordering $\eta=\eta\left(\mathcal{K}^{\prime}\right)$ of the demand pairs in $\mathcal{M}^{\prime}$, and $i_{1}<i_{2}<\ldots<i_{M_{j}^{\prime}}$. Then there is a set $\mathcal{W}_{j}^{\prime}$ of $M^{\prime}$ columns of the box $\hat{K}_{j}^{\prime}$, such that:

- set $\mathcal{W}_{j}^{\prime}$ does not contain a pair of consecutive columns; and
- for each $1 \leq z \leq M_{j}^{\prime}$, the $i_{z}$ th column of $\mathcal{W}_{j}^{\prime}$ from the left contains the destination vertex $t_{i_{z}}$.

Proof. From the definition of the ordering $\eta$ of $\mathcal{M}^{\prime}$, the demand pairs $\left(s_{i_{1}}, t_{i_{1}}\right),\left(s_{i_{2}}, t_{i_{2}}\right), \ldots,\left(s_{i_{M_{j}}}, t_{i_{M_{j}}}\right)$ appear consecutively in this order in $\eta$. Moreover, from the construction of the NDP-Grid instance $(\hat{G}, \mathcal{M})$, every pair of destination vertices in $T\left(\mathcal{M}^{\prime}\right)$ is separated by at least one vertex. We add to $\mathcal{W}_{j}^{\prime}$ all columns in which the vertices $t_{i_{1}}, t_{i_{2}}, \ldots, t_{i_{M_{j}}}$ lie. We also add to $\mathcal{W}_{j}^{\prime} i_{1}-1$ columns that lie to the left of the column of $t_{i_{1}}$, and $M^{\prime}-i_{M_{j}^{\prime}}$ columns that lie to the right of the column of $t_{i_{M_{j}}}$ in $\hat{K}_{j}^{\prime}$. We make sure that no pair of columns in $\mathcal{W}_{j}^{\prime}$ is consecutive - it is easy to see that there are enough columns to ensure that.

The proof of the following claim is identical to the proof of Claim 6.13 and is omitted here.
Claim 6.15. There is a set $\mathcal{P}^{2}=\left\{P_{1}^{\prime}, \ldots, P_{M^{\prime}}^{\prime}\right\}$ of spaced-out paths in $G^{b}$, that are internally disjoint from $R$, such that for each $1 \leq i \leq M^{\prime}$, path $P_{i}^{\prime}$ originates from vertex $x_{i}$, and for all $1 \leq j \leq N_{2}$, it contains the ith column of $\mathcal{W}_{j}^{\prime}$; in particular it contains $t_{i}$.

By combining the paths in sets $\mathcal{P}^{1}$ and $\mathcal{P}^{2}$, we obtain a new set $\mathcal{P}^{*}=\left\{P_{1}^{*}, \ldots, P_{M^{\prime}}^{*}\right\}$ of spaced-out paths, such that for all $1 \leq i \leq M^{\prime}$, path $P_{i}^{*}$ contains $s_{i}, x_{i}$ and $t_{i}$. By suitably truncating each such path, we obtain a collection of spaced-out paths routing all demand pairs in $\mathcal{N}^{\prime}$.

### 6.4 From routing to partitioning: proof of Theorem 6.6

In this subsection we conclude the proof of Theorem 4.3, by proving Theorem 6.6. Let $\mathcal{N}^{*} \subseteq \mathcal{M}$ be the set of the demand pairs routed by the solution $\mathcal{P}^{*}$, and let $E^{*} \subseteq E$ be the set of all edges $e$, whose corresponding demand pair belongs to $\mathcal{M}^{*}$. Let $\tilde{G}^{\prime} \subseteq \tilde{G}$ be the subgraph of $\tilde{G}$ induced by the edges in $E^{*}$. Notice that whenever two edges of $\tilde{G}$ belong to the same bundle, their corresponding demand pairs share a source or a destination. Since all paths in $\mathcal{P}^{*}$ are node-disjoint, all demand pairs in $\mathcal{M}^{*}$ have distinct sources and destinations, and so no two edges in $E^{*}$ belong to the same bundle.
Note that, if $\left|\mathcal{P}^{*}\right| \leq 2^{64} h \log ^{3} M$, then we can return the solution $\left(\left(W_{1}, \ldots, W_{r}\right),\left(E_{1}, \ldots, E_{r}\right)\right)$, where $W_{1}=V(\tilde{G})$ and $W_{2}=W_{3}=\cdots=W_{r}=\emptyset$; set $E_{1}$ contains an arbitrary subset of $\left\lceil\frac{\left|\mathcal{P}^{*}\right|}{2^{64} \log ^{3} M}\right\rceil \leq h$ edges of $E^{*}$, and all other sets $E_{i}$ are empty. Since no two edges of $E^{*}$ belong to the same bundle, we obtain a feasible solution to the $(\mathrm{r}, \mathrm{h})$-GPwB problem instance of value $\Omega\left(\left|\mathcal{P}^{*}\right| / \log ^{3} M\right)$. Therefore, from now on, we assume that $\left|\mathcal{P}^{*}\right|>2^{64} h \log ^{3} M$.
Our algorithm computes a solution to the $(r, h)-G P w B$ instance $\mathcal{J}$ by repeatedly partitioning $\tilde{G}^{\prime}$ into smaller and smaller subgraphs, by employing suitably defined balanced cuts.

Recall that, given a graph $\mathbf{H}$, a cut in $\mathbf{H}$ is a bi-partition $(A, B)$ of its vertices. We denote by $E_{\mathbf{H}}(A, B)$ the set of all edges with one endpoint in $A$ and another in $B$, and by $E_{\mathbf{H}}(A)$ and $E_{\mathbf{H}}(B)$ the sets of all edges with both endpoints in $A$ and in $B$, respectively. Given a cut $(A, B)$ of $\mathbf{H}$, the value of the cut is $\left|E_{\mathbf{H}}(A, B)\right|$. We will omit the subscript $\mathbf{H}$ when clear from context.

Definition 6.16. Given a graph $\mathbf{H}$ and a parameter $0<\rho<1$, a cut $(A, B)$ of $\mathbf{H}$ is called a $\rho$-edge-balanced cut iff $|E(A)|,|E(B)| \geq \rho \cdot|E(\mathbf{H})|$.

The following theorem is central to the proof of Theorem 6.6.
Theorem 6.17. There is an efficient algorithm, that, given a vertex-induced subgraph $\mathbf{H}$ of $\tilde{G}^{\prime}$ with $|E(\mathbf{H})|>22^{64} h \log ^{3} M$, computes a $1 / 32$-edge-balanced cut of $\mathbf{H}$, of value at most $\frac{|E(\mathbf{H})|}{64 \log M}$.

We prove Theorem 6.17 below, after we complete the proof of Theorem 6.6 using it. Our algorithm maintains a collection $\mathcal{G}$ of disjoint vertex-induced subgraphs of $\tilde{G}^{\prime}$, and consists of a number of phases. The input to the first phase is the collection $\mathcal{G}$ containing a single graph - the graph $\tilde{G}^{\prime}$. The algorithm continues as long as $\mathcal{G}$ contains a graph $\mathbf{H} \in \mathcal{G}$ with $|E(\mathbf{H})|>2^{64} \cdot h \log ^{3} M$; if no such graph $\mathbf{H}$ exists, then the algorithm terminates. Each phase is executed as follows. We process every graph $\mathbf{H} \in \mathcal{G}$ with $|E(\mathbf{H})|>2^{64} \cdot h \log ^{3} M$ one-by-one. When graph $\mathbf{H}$ is processed, we apply Theorem 6.17 to it, obtaining a 1/32-edge-balanced cut $(A, B)$ of $\mathbf{H}$, of value at most $\frac{|E(\mathbf{H})|}{64 \log M}$. We then remove $\mathbf{H}$ from $\mathcal{G}$, and add $\mathbf{H}[A]$ and $\mathbf{H}[B]$ to $\mathcal{G}$ instead. This completes the description of the algorithm. We use the following claim to analyze it.
Claim 6.18. Let $\mathcal{G}^{\prime}$ be the final set of disjoint subgraphs of $\tilde{G}^{\prime}$ obtained at the end of the algorithm. Then $\sum_{\mathbf{H} \in \mathcal{G}^{\prime}}|E(\mathbf{H})| \geq \Omega\left(\left|E\left(\tilde{G}^{\prime}\right)\right|\right)$, and $\left|\mathcal{G}^{\prime}\right| \leq r$.

Proof. We construct a binary partitioning tree $\tau$ of graph $\tilde{G}^{\prime}$, that simulates the graph partitions computed by the algorithm. For every subgraph $\mathbf{H} \subseteq \tilde{G}^{\prime}$ that belonged to $\mathcal{G}$ over the course of the algorithm, tree $\tau$ contains a node $v(\mathbf{H})$. The root of the tree is the node $v\left(\tilde{G}^{\prime}\right)$. If, over the course of our algorithm, we have partitioned the graph $\mathbf{H}$ into two disjoint vertex-induced subgraphs $\mathbf{H}^{\prime}$ and $\mathbf{H}^{\prime \prime}$, then we add edges $\left(v(\mathbf{H}), v\left(\mathbf{H}^{\prime}\right)\right)$ and $\left(v(\mathbf{H}), v\left(\mathbf{H}^{\prime \prime}\right)\right)$ to the tree, and $v\left(\mathbf{H}^{\prime}\right), v\left(\mathbf{H}^{\prime \prime}\right)$ become the children of $v(\mathbf{H})$ in $\tau$.
The level of a node $v(\mathbf{H})$ in the tree is the length of the path connecting $v(\mathbf{H})$ to the root of the tree; so the root of the tree is at level 0 . The depth of the tree, that we denote by $\Delta$, is the length of the longest leaf-to-root path in the tree. Since the cuts computed over the course of the algorithm are $1 / 32$-edge-balanced, $\Delta \leq \frac{\log M}{\log (32 / 31)}$. Consider now some level $0 \leq i \leq \Delta$, and let $V_{i}$ be the set of all nodes of the tree $\tau$ lying at level $i$. Let $V_{i}^{\prime}$ be a set of nodes consisting of all nodes of $V_{i}$, and all nodes that lie at levels $0, \ldots,(i-1)$ that are leaves of the tree. Let $\hat{E}_{i}=\bigcup_{v(\mathbf{H}) \in V_{i}^{\prime}} E(\mathbf{H})$ be the set of all edges contained in all subgraphs $\mathbf{H}$ of $\tilde{G}^{\prime}$, whose corresponding node $v(\mathbf{H})$ lies in $V_{i}^{\prime}$. Finally, let $m_{i}=\left|\hat{E}_{i}\right|$. Then $m_{0}=\left|E\left(\tilde{G}^{\prime}\right)\right|$, and for all $1 \leq i \leq \Delta$, the number of edges discarded over the course of phase $i$ is $m_{i-1}-m_{i} \leq m_{i-1} /(64 \log M) \leq m_{0} /(64 \log M)$ - this is since, whenever we partition a graph $\mathbf{H}$ into two subgraphs, we lose at most $|E(\mathbf{H})| /(64 \log M)$ of its edges. Overall, we get that:

$$
m_{0}-m_{\Delta} \leq \frac{\Delta m_{0}}{64 \log M} \leq \frac{\log M}{\log 32 / 31} \cdot \frac{m_{0}}{64 \log M} \leq \frac{m_{0}}{2}
$$

and so $\sum_{\mathbf{H} \in \mathcal{G}^{\prime}}|E(\mathbf{H})|=m_{\Delta} \geq m_{0} / 2$. This finishes the proof of the first assertion. We now turn to prove the second assertion. Recall that no two edges of $E^{*}$ may belong to the same bundle. Since $h=\beta^{*}(\mathcal{J}) / r=\left(\sum_{v \in V_{1}} \beta(v)\right) / r$, for any subset $E^{\prime} \subseteq E^{*}$ of edges, $\left|E^{\prime}\right| \leq \sum_{v \in V_{1}} \beta(v) \leq h r$ must hold, and in particular, $\left|E^{*}\right| \leq h r$. It is now enough to prove that for every leaf vertex $v(\mathbf{H})$ of $\tau,|E(\mathbf{H})| \geq h$ holds - since all graphs in $\mathcal{G}^{\prime}$ are mutually disjoint, and each such graph corresponds to a distinct leaf of $\tau$, this would imply that $\left|\mathcal{G}^{\prime}\right| \leq r$.
Consider now some leaf vertex $v(\mathbf{H})$ of $\tau$, and let $v\left(\mathbf{H}^{\prime}\right)$ be its parent. Then $\left|E\left(\mathbf{H}^{\prime}\right)\right| \geq 2^{64} h \log ^{3} M$, and, since the partition of $\mathbf{H}^{\prime}$ that we have computed was 1/32-balanced, $|E(\mathbf{H})| \geq\left|E\left(\mathbf{H}^{\prime}\right)\right| / 32 \geq h$. We conclude that $\left|\mathcal{G}^{\prime}\right| \leq r$.

We are now ready to define the solution $\left(\left(W_{1}, \ldots, W_{r}\right),\left(E_{1}, \ldots, E_{r}\right)\right)$ to the $(r, \mathrm{~h})$-GPwB problem instance $\mathcal{J}$. Let $\mathcal{G}^{\prime}$ be the set of the subgraphs of $\tilde{G}^{\prime}$ obtained at the end of our algorithm, and denote $\mathcal{G}^{\prime}=$ $\left\{\mathbf{H}_{1}, \mathbf{H}_{2}, \ldots, \mathbf{H}_{z}\right\}$. Recall that from Claim $6.18, z \leq r$. For $1 \leq i \leq z$, we let $W_{i}=V\left(\mathbf{H}_{i}\right)$. If $\left|E\left(\mathbf{H}_{i}\right)\right| \leq h$, then we let $E_{i}=E\left(\mathbf{H}_{i}\right)$; otherwise, we let $E_{i}$ contain any subset of $h$ edges of $E\left(\mathbf{H}_{i}\right)$. Since $\left|E\left(\mathbf{H}_{i}\right)\right| \leq$ $2^{64} h \log ^{3} M$, in either case, $\left|E_{i}\right| \geq \Omega\left(\left|E\left(\mathbf{H}_{i}\right)\right| / \log ^{3} M\right)$. For $i>z$, we set $W_{i}=\emptyset$ and $E_{i}=\emptyset$. Since, as observed before, no pair of edges of $E^{*}$ belongs to the same bundle, it is immediate to verify that we obtain a feasible solution to the $(r, h)$-GPwB problem instance. The value of the solution is:

$$
\sum_{i=1}^{r}\left|E_{i}\right| \geq \sum_{i=1}^{r} \Omega\left(\left|E\left(\mathbf{H}_{i}\right)\right| / \log ^{3} M\right)=\Omega\left(\left|E\left(\tilde{G}^{\prime}\right)\right| / \log ^{3} M\right)=\Omega\left(\left|\mathcal{P}^{*}\right| / \log ^{3} M\right)
$$

from Claim 6.18. In order to complete the proof of Theorem 6.6, it now remains to prove Theorem 6.17.
Proof of Theorem 6.17. Let $\mathbf{H}$ be a vertex-induced subgraph of $\tilde{G}^{\prime}$ with $|E(\mathbf{H})|>2^{64} h \log ^{3} M$. Our proof consists of two parts. First, we show an efficient algorithm to compute a drawing of $\mathbf{H}$ with relatively few crossings. Next, we show how to exploit this drawing in order to compute a $1 / 32$-edge-balanced cut of $\mathbf{H}$ of small value. We start by defining a drawing of a given graph in the plane and the crossings in this drawing.

Definition 6.19. A drawing of a given graph $\mathbf{H}^{\prime}$ in the plane is a mapping, in which every vertex of $\mathbf{H}$ is mapped to a point in the plane, and every edge to a continuous curve connecting the images of its endpoints, such that whenever a point $x$ belongs to three such curves, $x$ is an endpoint of all three curves; no curve intersects itself; and no curve contains an image of any vertex other than its endpoints. A crossing in such a drawing is a point contained in the images of two edges.

Lemma 6.20. There is an efficient algorithm that, given a solution $\mathcal{P}^{*}$ to the NDP-Grid problem instance $\hat{\jmath}$, and a vertex-induced subgraph $\mathbf{H} \subseteq \tilde{G}^{\prime}$, computes a drawing of $\mathbf{H}$ with at most $2048 \cdot|E(\mathbf{H})|\lceil h \log M\rceil$ crossings.

Proof. We consider the natural embedding of the grid $\hat{G}$ into the plane. We will use this drawing of $\hat{G}$ in order to guide our drawing of the graph $\mathbf{H}$. Notice that this embedding of $\hat{G}$ naturally defines a drawing of every path $P \in \mathcal{P}^{*}$, obtained by taking the union of the images of the edges of $P$. We will map vertices of $\mathbf{H}$ to points that lie inside some faces defined by the cells of the grid. Consider a vertex $v_{i} \in V(\mathbf{H}) \cap V_{1}$,

## Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat

and let $K_{i}$ be the block representing this vertex. Let $\kappa_{i}$ be any cell of the grid $\hat{G}$ that has a vertex of $K_{i}$ on its boundary. We map the vertex $v_{i}$ to a point $p_{i}$ lying in the middle of the cell $\kappa_{i}$ (it is sufficient that $p_{i}$ is far enough from the boundaries of the cell). For every vertex $v_{j}^{\prime} \in V(\mathbf{H}) \cap V_{2}$, we select a cell $\kappa_{j}^{\prime}$ whose boundary contains a vertex of the corresponding block $K_{j}^{\prime}$, and map $v_{j}^{\prime}$ to a point $p_{j}^{\prime}$ lying in the middle of $\kappa_{j}^{\prime}$ similarly.
Next, we define the drawings of the edges of $E(\mathbf{H})$. Consider any such edge $e=\left(v_{i}, v_{j}^{\prime}\right) \in E(\mathbf{H})$, with $v_{i} \in V_{1}$ and $v_{j}^{\prime} \in V_{2}$, and let $(s, t) \in \mathcal{M}^{*}$ be its corresponding demand pair. Let $K_{i}$ and $K_{j}^{\prime}$ be the blocks containing $s$ and $t$ respectively, and let $P \in \mathcal{P}^{*}$ be the path routing the demand pair $(s, t)$ in our solution to the NDP-Grid problem instance $\hat{J}$. The drawing of the edge $e$ is a concatenation of the following three segments: (i) the image of the path $P$, that we refer to as a type-1 segment; (ii) a straight line connecting $s$ to the image of $v_{i}$, that we refer to as a type-2 segment; and (iii) a straight line connecting $t$ to the image of $v_{j}^{\prime}$, that we refer to as a type- 3 segment. If the resulting curve has any self-loops, then we delete them.
We now bound the number of crossings in the resulting drawing of $\mathbf{H}$. Since the paths in $\mathcal{P}^{*}$ are nodedisjoint, whenever the images of two edges $e$ and $e^{\prime}$ cross, the crossing must be between the type-1 segment of $e$, and either the type- 2 or the type- 3 segment of $e^{\prime}$ (or the other way around).
Consider now some edge $e=\left(v_{i}, v_{j}^{\prime}\right) \in E(\mathbf{H})$ with $v_{i} \in V_{1}$ and $v_{j}^{\prime} \in V_{2}$, and let $K_{i}$ and $K_{j}^{\prime}$ be the blocks representing $v_{i}$ and $v_{j}^{\prime}$ respectively. Let $(s, t) \in \mathcal{M}^{*}$ be the demand pair corresponding to $e$. Assume that a type- 1 segment of some edge $e^{\prime}$ crosses a type- 2 segment of $e$. This can only happen if the path $P^{\prime} \in \mathcal{P}^{*}$ routing the demand pair $\left(s^{\prime}, t^{\prime}\right)$ corresponding to the edge $e^{\prime}$ contains a vertex of $K_{i}$. Since $\left|V\left(K_{i}\right)\right| \leq 1024 \cdot\lceil h \log M\rceil$, at most $1024 \cdot\lceil h \log M\rceil$ type-1 segments of other edges may cross the type-2 segment of $e$. The same accounting applies to the type- 3 segment of $e$. Overall, the number of crossings in the above drawing is bounded by:

$$
\sum_{e \in E(\mathbf{H})} 2 \cdot 1024 \cdot\lceil h \log M\rceil=2048 \cdot|E(\mathbf{H})|\lceil h \log M\rceil .
$$

Next, we show an efficient algorithm that computes a small balanced partition of any given graph $\mathbf{H}^{\prime}$, as long as its maximum vertex degree is suitably bounded, and we are given a drawing of $\mathbf{H}^{\prime}$ with a small number of crossings.

Lemma 6.21. There is an efficient algorithm that, given any graph $\mathbf{H}$ with $|E(\mathbf{H})|=m$ and maximum vertex degree at most $d$, and a drawing $\varphi$ of $\mathbf{H}$ with at most $\mathrm{cr} \leq m d \alpha$ crossings for some $\alpha>1$, such that $m>2^{20} d \alpha$, computes a $1 / 32$-edge-balanced cut $(A, B)$ of $\mathbf{H}$ of value $|E(A, B)| \leq 64 \sqrt{8 m d \alpha}$.

Before we complete the proof of the lemma, we show that the proof of Theorem 6.17 follows from it. Let $m=|E(\mathbf{H})|$. Note that the maximum vertex degree in graph $\mathbf{H}$ is bounded by $\max _{v \in V_{1} \cup V_{2}}\{\beta(v)\}$, as $E(\mathbf{H})$ cannot contain two edges that belong to the same bundle. From the definition of valid instances, $h \geq \beta(v)$ for all $v \in V_{1} \cup V_{2}$, and so the maximum vertex degree in $\mathbf{H}$ is bounded by $d=h$. From Lemma 6.20, the drawing $\varphi$ of $\mathbf{H}$ has at most $2048|E(\mathbf{H})|\lceil h \log M\rceil \leq 2^{12} m d \log M$ crossings. Setting $\alpha=2^{12} \log M$, the number of crossings cr in $\varphi$ is bounded by $m d \alpha$. Moreover, since $m>2^{64} h \log ^{3} M$, we get that

## Almost Polynomial Hardness of Node-Disjoint Paths in Grids



Figure 5: Grid $Q_{v}$ obtained from $v$
$m>2^{20} d \alpha$. We can now apply Lemma 6.21 to graph $\mathbf{H}$ to obtain a $1 / 32$-edge-balanced cut $(A, B)$ with $|E(A, B)| \leq 64 \sqrt{8 m d \alpha} \leq 64 \sqrt{2^{15} m h \log M} \leq 64 \sqrt{m^{2} /\left(2^{49} \log ^{2} M\right)} \leq \frac{|E(\mathbf{H})|}{64 \log M}$, since we have assumed that $|E(\mathbf{H})|=m>2^{64} h \log ^{3} M$. It now remains to prove Lemma 6.21.

Proof of Lemma 6.21. We start by providing some intuition for the proof. In general, it is well known that a graph $G$ that can be drawn in the plane with at most cr crossings has a balanced separator (a vertex-cut), containing $O(\sqrt{|V(G)|+\mathrm{cr}})$ vertices. However, in our case, we need to compute an edge-cut, and since some vertex degrees may be quite large, using this bound and then deleting all edges incident to the vertices in the separator is not sufficient to us. In order to overcome this difficulty, we slightly transform our graph into a constant-degree graph, and then work with the transformed graph directly.
For each vertex $v \in V(\mathbf{H})$, we denote the degree of $v$ in $\mathbf{H}$ by $d_{v}$. We assume without loss of generality that for all $v \in V(\mathbf{H}), d_{v} \geq 1$ : otherwise, we can remove all isolated vertices from $\mathbf{H}$ and then apply our algorithm to compute a $1 / 32$-edge-balanced cut $(A, B)$ in the remaining graph. At the end we can add the isolated vertices to $A$ or $B$, while ensuring that the cut remains 1/32-balanced, and without increasing its value.

We construct a new graph $\hat{\mathbf{H}}$ from graph $\mathbf{H}$ as follows. For every vertex $v \in V(\mathbf{H})$, we add a $\left(d_{v} \times d_{v}\right)$-grid $Q_{v}$ to $\hat{\mathbf{H}}$, so that the resulting grids are mutually disjoint. We call the edges of the resulting grids regular edges. Let $e_{1}(v), \ldots, e_{d_{v}}(v)$ be the edges of $\mathbf{H}$ incident to $v$, indexed in the clockwise order of their entering the vertex $v$ in the drawing $\varphi$ of $\mathbf{H}$. We denote by $\Pi(v)=\left\{p_{1}(v), \ldots, p_{d_{v}}(v)\right\}$ the set of vertices on the top boundary of $Q_{v}$, where the vertices are indexed in the clock-wise order of their appearance on the boundary of $Q_{v}$. We refer to the vertices of $\Pi(v)$ as the portals of $Q_{v}$ (see Figure 5). Let $\Pi=\bigcup_{v \in V(\mathbf{H})} \Pi(v)$ be the set of all portals. For every edge $e=(u, v) \in E(\mathbf{H})$, we add a new special edge to graph $\hat{\mathbf{H}}$, as follows. Assume that $e=e_{i}(v)=e_{j}(u)$. Then we add an edge $\left(p_{i}(v), p_{j}(u)\right)$ to $\hat{\mathbf{H}}$. We think of this edge as the special edge representing $e$. This finishes the definition of the graph $\hat{\mathbf{H}}$. It is immediate to see that the drawing $\varphi$ of $\mathbf{H}$ can be extended to a drawing $\varphi^{\prime}$ of $\hat{\mathbf{H}}$ without introducing any new crossings, that is, the number of crossings in $\varphi^{\prime}$ remains at most cr. Note that every portal vertex is incident to exactly one special edge, and the maximum vertex degree in $\hat{\mathbf{H}}$ is 4 . We will use the following bound on $|V(\hat{\mathbf{H}})|$ :

Observation 6.22. $|V(\hat{\mathbf{H}})| \leq 2 m d$.

Proof. Clearly, $|V(\hat{\mathbf{H}})|=\sum_{v \in V(\mathbf{H})} d_{v}^{2} \leq \sum_{v \in V(\mathbf{H})} d \cdot d_{v} \leq 2 m d$.

Let $\hat{\mathbf{H}}^{\prime}$ be the graph obtained from $\hat{\mathbf{H}}$ by replacing every intersection point in the drawing $\varphi^{\prime}$ of $\hat{\mathbf{H}}$ with a vertex. Then $\hat{\mathbf{H}}^{\prime}$ is a planar graph with $\left|V\left(\hat{\mathbf{H}}^{\prime}\right)\right| \leq|V(\hat{\mathbf{H}})|+\mathrm{cr} \leq 2 m d+m d \alpha \leq 4 m d \alpha$, as $\alpha \geq 1$. We assign weights to the vertices of $\hat{\boldsymbol{H}}^{\prime}$ as follows: every vertex of $\Pi$ is assigned the weight 1 , and every other vertex is assigned the weight 0 . Note that the weight of a vertex is exactly the number of special edges incident to it, and the total weight of all vertices is $W=|\Pi|=2 \mathrm{~m}$. We will use the following version of the planar separator theorem [51, 39, 2].
Theorem 6.23 ([39]). There is an efficient algorithm, that, given a planar graph $G=(V, E)$ with $n$ vertices, and an assignment $w: V \rightarrow R^{+}$of non-negative weights to the vertices of $G$, with $\sum_{v \in V} w(v)=W$, computes a partition $(A, X, B)$ of $V(G)$, such that:

- no edge connecting a vertex of $A$ to a vertex of $B$ exists in $G$;
- $\sum_{v \in A} w(v), \sum_{v \in B} w(v) \leq 2 W / 3$; and
- $|X| \leq 2 \sqrt{2 n}$.

We apply Theorem 6.23 to graph $\hat{\mathbf{H}}^{\prime}$, to obtain a partition $(A, X, B)$ of $V\left(\hat{\mathbf{H}}^{\prime}\right)$, with $|X| \leq 2 \sqrt{2\left|V\left(\hat{\mathbf{H}}^{\prime}\right)\right|} \leq$ $2 \sqrt{8 m d \alpha}$. Since $W=\sum_{v \in V\left(\hat{\mathbf{H}}^{\prime}\right)} w(v)=2 m$, we get that $|A \cap \Pi|=\sum_{v \in A} w(v) \leq 2 W / 3 \leq 4 m / 3$, and similarly $|B \cap \Pi| \leq 4 m / 3$. Assume without loss of generality that $|A \cap \Pi| \leq|B \cap \Pi|$. We obtain a bipartition $\left(A^{\prime}, B^{\prime}\right)$ of $V\left(\hat{\mathbf{H}}^{\prime}\right)$ by setting $A^{\prime}=A \cup X$ and $B^{\prime}=B$. Since $|X| \leq 2 \sqrt{8 m d \alpha} \leq m / 3\left(\right.$ as $\left.m>22^{20} d \alpha\right)$, we are guaranteed that $\left|A^{\prime} \cap \Pi\right|,\left|B^{\prime} \cap \Pi\right| \leq 4 m / 3$ holds. Moreover, as all vertex degrees in $\hat{\mathbf{H}}^{\prime}$ are at most $4,\left|E_{\hat{\mathbf{H}}^{\prime}}\left(A^{\prime}, B^{\prime}\right)\right| \leq 4|X| \leq 8 \sqrt{8 m d \alpha}$.
Note that cut $\left(A^{\prime}, B^{\prime}\right)$ in graph $\hat{\mathbf{H}}^{\prime}$ naturally defines a cut $\left(A^{\prime \prime}, B^{\prime \prime}\right)$ in graph $\hat{\mathbf{H}}$, where a vertex $v \in$ $V(\hat{\mathbf{H}})$ is added to $A^{\prime \prime}$ if it lies in the set $A^{\prime}$, and it is added to $B^{\prime \prime}$ otherwise. It is easy to verify that $\left|A^{\prime \prime} \cap \Pi\right|,\left|B^{\prime \prime} \cap \Pi\right| \leq 4 m / 3$ still holds. Moreover, the number of edges in the cut may only decrease, that is, $\left|E_{\hat{\mathbf{H}}}\left(A^{\prime \prime}, B^{\prime \prime}\right)\right| \leq\left|E_{\hat{\mathbf{H}}^{\prime}}\left(A^{\prime}, B^{\prime}\right)\right| \leq 8 \sqrt{8 m d \alpha}$. For convenience, we will denote the cut $\left(A^{\prime \prime}, B^{\prime \prime}\right)$ by $\left(A^{\prime}, B^{\prime}\right)$ from now on, and we will only continue working with this cut.
Unfortunately, the cut $\left(A^{\prime}, B^{\prime}\right)$ of $\hat{H}$ does not directly translate into a balanced cut in $\mathbf{H}$, since for some vertices $v \in V(\mathbf{H})$, the corresponding grid $Q_{v}$ may be split between $A^{\prime}$ and $B^{\prime}$. We now show how to overcome this difficulty, by moving each such grid entirely to one of the two sides. Before we proceed, we state a simple fact about grid graphs.
Observation 6.24. Let $z>1$ be an integer, and let $Q$ be the $(z \times z)$-grid. Let $U$ be the set of vertices lying on the top row of $Q$, and let $(X, Y)$ be a bi-partition of $V(Q)$. Then $\left|E_{Q}(X, Y)\right| \geq \min \{|U \cap X|,|U \cap Y|\}$.

Proof. It is easy to verify that for any bi-partition $\left(X^{\prime}, Y^{\prime}\right)$ of $U$ into two disjoint subsets, there is a set $\mathcal{P}$ of $\min \left\{\left|X^{\prime}\right|,\left|Y^{\prime}\right|\right\}$ node-disjoint paths in $Q$ connecting vertices of $X^{\prime}$ to vertices of $Y^{\prime}$. The observation follows from the maximum flow - minimum cut theorem.

We say that a vertex $v \in V(\mathbf{H})$ is split by the cut $\left(A^{\prime}, B^{\prime}\right)$ iff $V\left(Q_{v}\right) \cap A^{\prime}$ and $V\left(Q_{v}\right) \cap B^{\prime} \neq \emptyset$. We say that it is split evenly iff $\left|\Pi(v) \cap A^{\prime}\right|,\left|\Pi(v) \cap B^{\prime}\right| \geq d_{v} / 8$; otherwise we say that it is split unevenly. We modify the cut $\left(A^{\prime}, B^{\prime}\right)$ in the following two steps, to ensure that no vertex of $V(\mathbf{H})$ remains split.

Step 1 [Unevenly split vertices]. We process each vertex $v \in V(\mathbf{H})$ that is unevenly split one-by-one. Consider any such vertex $v$. If $\left|\Pi(v) \cap A^{\prime}\right|>\left|\Pi(v) \cap B^{\prime}\right|$, then we move all vertices of $Q_{v}$ to $A^{\prime}$; otherwise we move all vertices of $Q_{v}$ to $B^{\prime}$. Assume without loss of generality that the former happens. Notice that from Observation 6.24, $\left|E\left(A^{\prime}, B^{\prime}\right)\right|$ does not increase, since $E\left(Q_{v}\right)$ contributed at least $\left|\Pi(v) \cap B^{\prime}\right|$ regular edges to the cut before the current iteration. Moreover, $\left|\Pi(v) \cap A^{\prime}\right|$ increases by the factor of at most 8/7. Therefore, at the end of this procedure, once all unevenly split vertices of $\mathbf{H}$ are processed, $\left|A^{\prime} \cap \Pi\right|,\left|B^{\prime} \cap \Pi\right| \leq \frac{8}{7} \cdot \frac{4}{3} m=\frac{32}{21} m$ and $\left|E\left(A^{\prime}, B^{\prime}\right)\right| \leq 8 \sqrt{8 m d \alpha}$.

Step 2 [Evenly split vertices]. In this step, we process each vertex $v \in V(\mathbf{H})$ that is evenly split one-byone. Consider an iteration where some such vertex $v \in V(\mathbf{H})$ is processed. If $\left|A^{\prime} \cap \Pi\right| \leq\left|B^{\prime} \cap \Pi\right|$, then we move all vertices of $Q_{v}$ to $A^{\prime}$; otherwise we move them to $B^{\prime}$. Assume without loss of generality that the former happened. Then before the current iteration $\left|A^{\prime} \cap \Pi\right| \leq|\Pi| / 2 \leq m$, and, since $|\Pi(v)| \leq d<m / 21$, $\left|A^{\prime} \cap \Pi\right| \leq \frac{32}{21} m$, while $\left|B^{\prime} \cap \Pi\right| \leq \frac{32}{21} m$ as before. Moreover, from Observation 6.24, before the current iteration, the regular edges of $Q_{v}$ contributed at least $d(v) / 8$ edges to $E\left(A^{\prime}, B^{\prime}\right)$, and after the current iteration, no regular edges of $Q_{v}$ contribute to the cut, but we may have added up to $d(v)$ new special edges to it. Therefore, after all vertices of $\mathbf{H}$ that are evenly split are processed, $\left|E\left(A^{\prime}, B^{\prime}\right)\right|$ grows by the factor of at most 8 , and remains at most $64 \sqrt{8 m d \alpha}$.
We are now ready to define the final cut $\left(A^{*}, B^{*}\right)$ in graph $\mathbf{H}$. We let $A^{*}$ contain all vertices $v \in$ $V(\mathbf{H})$ with $V\left(Q_{v}\right) \subseteq A^{\prime}$, and we let $B^{*}$ contain all remaining vertices of $V(\mathbf{H})$. Clearly, $\left|E_{\mathbf{H}}\left(A^{*}, B^{*}\right)\right| \leq$ $\left|E_{\hat{\mathbf{H}}^{\prime}}\left(A^{\prime}, B^{\prime}\right)\right| \leq 64 \sqrt{8 m d \alpha}$. It remains to show that $\left|E_{\mathbf{H}}\left(A^{*}\right)\right|,\left|E_{\mathbf{H}}\left(B^{*}\right)\right| \geq|E(\mathbf{H})| / 32$. We show that $\left|E_{\mathbf{H}}\left(A^{*}\right)\right| \geq|E(\mathbf{H})| / 32$; the proof that $\left|E_{\mathbf{H}}\left(B^{*}\right)\right| \geq|E(\mathbf{H})| / 32$ is symmetric. Observe that $\sum_{v \in B^{*}} d_{v}=$ $\left|B^{\prime} \cap \Pi\right| \leq \frac{32 m}{21}$, while $\sum_{v \in V(\mathbf{H})} d_{v}=2 m$. Therefore, $\sum_{v \in A^{*}} d_{v} \geq 2 m-\frac{32 m}{21}=\frac{10 m}{21}$. But $\left|E_{\mathbf{H}}\left(A^{*}, B^{*}\right)\right| \leq$ $64 \sqrt{8 m d \alpha} \leq 64 \sqrt{m^{2} / 2^{17}}<m / 4$ (since $m>2^{20} d \alpha$ ). Therefore,

$$
\left|E_{\mathbf{H}}\left(A^{*}\right)\right|=\frac{\sum_{v \in A^{*}} d_{v}-\left|E_{\mathbf{H}}\left(A^{*}, B^{*}\right)\right|}{2} \geq \frac{5 m}{21}-\frac{m}{8} \geq \frac{m}{32} .
$$

## 7 Hardness of NDP and EDP on wall graphs

In this section we extend our results to NDP and EDP on wall graphs, completing the proofs of Theorem 1.1 and Theorem 1.2. We first prove hardness of NDP-Wall, and show later how to extend it to EDP-Wall. Let $\hat{G}=G^{\ell, h}$ be a grid of length $\ell$ and height $h$, where $\ell>0$ is an even integer, and $h>0$. We denote by $\hat{G}^{\prime}$ the wall corresponding to $\hat{G}$, as defined in Section 2. We prove the following analogue of Theorem 4.3.

## Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat

Theorem 7.1. There is a constant $c^{*}>0$, and there is an algorithm, that, given a valid and a $\left(d, d^{\prime}, b\right)$ regular instance $\mathcal{J}=\left(\tilde{G}, \mathcal{U}_{1}, \mathcal{U}_{2}, h, r\right)$ of $(\mathrm{r}, \mathrm{h})-\mathrm{GPwB}$, for some $d, d^{\prime}, b>0$, with $|V(\tilde{G})|=N$ and $|E(\tilde{G})|=$ $M$, in time poly $(N M)$ constructs an instance $\hat{\jmath}^{\prime}=\left(\hat{G}^{\prime}, \mathcal{M}\right)$ of NDP-Wall with $\left|V\left(\hat{G}^{\prime}\right)\right|=O\left(M^{4} \log ^{2} M\right)$, such that the following hold:

- If $\mathcal{J}$ has a perfect solution (of value $\beta^{*}=\beta^{*}(\mathcal{J})$ ), then instance $\hat{\jmath}^{\prime}$ has a solution $\mathcal{P}^{\prime}$ that routes at least $\frac{\beta^{*}}{c^{*} \log ^{2} M}$ demand pairs via node-disjoint paths; and
- There is an algorithm with running time poly $(N M)$, that, given a solution $\mathcal{P}^{*}$ to the NDP-Wall problem instance $\hat{\mathcal{J}}^{\prime}$, constructs a solution to the $(\mathrm{r}, \mathrm{h})-\mathrm{GPwB}$ instance $\mathcal{J}$, of value at least $\frac{\left|\mathcal{P}^{*}\right|}{c^{*} \cdot \log ^{3} M}$.

Notice that plugging Theorem 7.1 into the hardness of approximation proof instead of Theorem 4.3, we extend the hardness result to the NDP problem on wall graphs and complete the proof of Theorem 1.1.
Proof of Theorem 7.1. Let $\hat{J}=(\hat{G}, \mathcal{M})$ be the instance of NDP-Grid constructed in Theorem 4.3. In order to obtain an instance $\hat{\mathcal{J}}^{\prime}$ of NDP-Wall, we replace the grid $\hat{G}$ with the corresponding wall $\hat{G}^{\prime}$ as described above; the set of the demand pairs remains unchanged. We now prove the two assertions about the resulting instance $\hat{\mathcal{J}}^{\prime}$, starting from the second one.
Suppose we are given a solution $\mathcal{P}^{*}$ to the NDP-Wall problem instance $\hat{J}^{\prime}$. Since $\hat{G}^{\prime} \subseteq \hat{G}$, and the set of demand pairs in instances $\hat{\mathcal{J}}$ and $\hat{\mathcal{J}}^{\prime}$ is the same, $\mathcal{P}^{*}$ is also a feasible solution to the NDP-Grid problem instance $\hat{\mathcal{J}}$, and so we can use the efficient algorithm from Theorem 4.3 to construct a solution to the ( $r$ r,h)-GPwB instance $\mathcal{J}$, of value at least $\frac{\left|\mathcal{P}^{*}\right|}{c^{*} \log ^{3} M}$.
It now remains to prove the first assertion. Assume that $\mathcal{J}$ has a perfect solution. From Theorem 4.3, instance $\hat{\mathcal{J}}$ of NDP-Grid has a solution $\mathcal{P}$ that routes at least $\frac{\beta^{*}}{c^{*} \log ^{2} M}$ demand pairs via paths that are spaced-out. It is now enough to show that there is a solution of value $\frac{\beta^{*}}{c^{*} \log ^{2} M}$ to the corresponding instance $\hat{\jmath}^{\prime}$ of NDP-Wall.
Consider the spaced-out set $\mathcal{P}$ of paths in $\hat{G}$. Recall that for every pair $P, P^{\prime}$ of paths, $d\left(V(P), V\left(P^{\prime}\right)\right) \geq 2$, and all paths in $\mathcal{P}$ are internally disjoint from the boundaries of the grid $\hat{G}$. For each path $P \in \mathcal{P}$, we will slightly modify $P$ to obtain a new path $P^{\prime}$ contained in the wall $\hat{G}^{\prime}$, so that the resulting set $\mathcal{P}^{\prime}=\left\{P^{\prime} \mid P \in \mathcal{P}\right\}$ of paths is node-disjoint.
For all $1 \leq i<\ell, 1 \leq j \leq \ell$, let $e_{i}^{j}$ denote the $i$ th edge from the top lying in column $W_{j}$ of the grid $\hat{G}$, so that $e_{j}^{i}=(v(i, j), v(i+1, j))$. Let $E^{*}=E(\hat{G}) \backslash E\left(\hat{G}^{\prime}\right)$ be the set of edges that were deleted from the grid $\hat{G}$ when constructing the wall $\hat{G}^{\prime}$. We call the edges of $E^{*}$ bad edges. Notice that only vertical edges may be bad, and, if $e_{i}^{j} \in E\left(W_{j}\right)$ is a bad edge, for $1<j<\ell$, then $e_{i}^{j+1}$ is a good edge. Consider some bad edge $e_{i}^{j}=(v(i, j), v(i+1, j))$, such that $1<j<\ell$, so $e_{i}^{j}$ does not lie on the boundary of $\hat{G}$. Let $Q_{i}^{j}$ be the path $(v(i, j), v(i, j+1), v(i+1, j+1), v(i+1, j))$. Clearly, path $Q_{i}^{j}$ is contained in the wall $\hat{G}^{\prime}$. For every path $P \in \mathcal{P}$, we obtain the new path $P^{\prime}$ by replacing every bad edge $e_{i}^{j} \in P$ with the corresponding path $Q_{i}^{j}$. It is easy to verify that $P^{\prime}$ is a path with the same endpoints as $P$, and that it is contained in the wall $\hat{G}^{\prime}$. Moreover, since the paths in $\mathcal{P}$ are spaced-out, the paths in the resulting set $\mathcal{P}^{\prime}=\left\{P^{\prime} \mid P \in \mathcal{P}\right\}$ are node-disjoint.

This completes the proof of Theorem 1.1. In order to prove Theorem 1.2, we show an approximationpreserving reduction from NDP-Wall to EDP-Wall.

Claim 7.2. Let $\mathcal{J}=(G, \mathcal{M})$ be an instance of NDP-Wall, and let $\mathcal{J}^{\prime}$ be the instance of EDP-Wall consisting of the same graph $G$ and the same set $\mathcal{M}$ of demand pairs. Let OPT and $\mathrm{OPT}^{\prime}$ be the optimal solution values for $\mathcal{J}$ and $\mathcal{J}^{\prime}$, respectively. Then $\mathrm{OPT}^{\prime} \geq \mathrm{OPT}$, and there is an efficient algorithm, that, given any solution $\mathcal{P}^{\prime}$ to instance $\mathcal{J}^{\prime}$ of EDP-Wall, computes a solution $\mathcal{P}$ to instance $\mathcal{J}$ of $N D P-W a l l$ of value $\Omega\left(\left|\mathcal{P}^{\prime}\right|\right)$.

Notice that, from Claim 7.2, if there is an $\alpha$-approximation algorithm for EDP-Wall with running time $f(n)$, for $\alpha>1$ that may be a function of the graph size $n$, then there is an $O(\alpha)$-approximation algorithm for NDP-Wall with running time $f(n)+\operatorname{poly}(n)$. Therefore, the proof of Claim 7.2 will complete the proof of Theorem 1.2

Proof. The assertion that $\mathrm{OPT}^{\prime} \geq$ OPT is immediate, as any set $\mathcal{P}$ of node-disjoint paths in the wall $G$ is also a set of edge-disjoint paths.

Assume now that we are given a set $\mathcal{P}^{\prime}$ of edge-disjoint paths in $G$. We show an efficient algorithm to compute a subset $\mathcal{P} \subseteq \mathcal{P}^{\prime}$ of $\Omega\left(\left|\mathcal{P}^{\prime}\right|\right)$ paths that are node-disjoint. Since the maximum vertex degree in $G$ is 3 , the only way for two paths $P, P^{\prime} \in \mathcal{P}^{\prime}$ to share a vertex $x$ is when $x$ is an endpoint of at least one of these two paths. If $x$ is an endpoint of $P$, and $x \in V\left(P^{\prime}\right)$, then we say that $P$ has a conflict with $P^{\prime}$.

We construct a directed graph $H$, whose vertex set is $\left\{v_{P} \mid P \in \mathcal{P}^{\prime}\right\}$, and there is an edge $\left(v_{P}, v_{P^{\prime}}\right)$ iff $P$ has a conflict with $P^{\prime}$. It is immediate to verify that the maximum out-degree of any vertex in $H$ is at most 4 , as each of the two endpoints of a path $P$ may be shared by at most two additional paths. Therefore, every subgraph $H^{\prime} \subseteq H$ of $H$ contains a vertex of total degree at most 8 . We construct a set $U$ of vertices of $H$, such that no two vertices of $U$ are connected by an edge, using a standard greedy algorithm: while $H \neq \emptyset$, select a vertex $v \in V(H)$ with total degree at most 8 and add it to $U$; remove $v$ and all its neighbors from $H$. It is easy to verify that at the end of the algorithm, $|U|=\Omega(|V(H)|)=\Omega\left(\left|\mathcal{P}^{\prime}\right|\right)$, and no pair of vertices in $U$ is connected by an edge. Let $\mathcal{P}=\left\{P \mid v_{P} \in U\right\}$. Then the paths in $\mathcal{P}$ are node-disjoint, and $|\mathcal{P}|=\Omega\left(\left|\mathcal{P}^{\prime}\right|\right)$.

## 8 Acknowledgment

We thank the anonymous reviewers for carefully reading this manuscript and for their many helpful suggestions on simplifying and improving the presentation, especially for the suggestion on how to derandomize the original randomized reduction.

## Appendix

## A Proof of Theorem 2.6

Suppose $G$ is a Yes-Instance, and let $\chi$ be a valid coloring of $V(G)$. Let $\pi_{1}, \ldots, \pi_{6}$ be 6 different permutations of $\{r, g, b\}$. For each $1 \leq i \leq 6$, permutation $\pi_{i}$ defines a valid coloring $\chi_{i}$ of $G$ : for every vertex $v \in V(G)$, if $v$ is assigned a color $c \in\{\mathcal{C}\}$ by $\chi$, then $\chi_{i}$ assigns the color $\pi_{i}(c)$ to $v$. Notice that for each vertex $v$ and for each color $c \in \mathcal{C}$, there are exactly two indices $i \in\{1, \ldots, 6\}$, such that $\chi_{i}$ assigns the color $c$ to $v$. Notice also that for each edge $(u, v)$, if $c, c^{\prime} \in \mathcal{C}$ is any pair of distinct colors, then there is exactly one index $i \in\{1, \ldots, 6\}$, such that $u$ is assigned the color $c$ and $v$ is assigned the color $c^{\prime}$ by $\chi_{i}$.
Let $B$ be the set of all vectors of length $\ell$, whose entries belong to $\{1, \ldots, 6\}$, so that $|B|=6^{\ell}$. For each such vector $b \in B$, we define a perfect global assignment $f_{b}$ of answers to the queries, as follows. Let $Q \in Q^{E}$ be a query to the edge-player, and assume that $Q=\left(e_{1}, \ldots, e_{\ell}\right)$. Fix some index $1 \leq j \leq \ell$, and assume that $e_{j}=\left(v_{j}, u_{j}\right)$. Assume that $b_{j}=z$, for some $1 \leq z \leq 6$. We assign to $v_{j}$ the color $\chi_{z}\left(v_{j}\right)$, and we assign to $u_{j}$ the color $\chi_{z}\left(u_{j}\right)$. Since $\chi_{z}$ is a valid coloring of $V(G)$, the two colors are distinct. This defines an answer $A \in \mathcal{A}^{E}$ to the query $Q$, that determines $f_{b}(Q)$.
Consider now some query $Q^{\prime} \in Q^{V}$ to the vertex-player, and assume that $Q^{\prime}=\left(v_{1}, \ldots, v_{\ell}\right)$. Fix some index $1 \leq j \leq \ell$, and assume that $b_{j}=z$, for some $1 \leq z \leq 6$. We assign to $v_{j}$ the color $\chi_{z}\left(v_{j}\right)$. This defines an answer $A^{\prime} \in \mathcal{A}^{V}$ to the query $Q^{\prime}$, that determines $f_{b}\left(Q^{\prime}\right)$. Notice that for each $1 \leq j \leq \ell$, the answers that we choose for the $j$ th coordinate of each query are consistent with the valid coloring $\chi_{b_{j}}$ of $G$. Therefore, it is immediate to verify that for each $b \in B, f_{b}$ is a perfect global assignment.

We now fix some query $Q \in Q^{E}$ of the edge-prover, and some answer $A \in \mathcal{A}^{E}$ to it. Assume that $Q=\left(e_{1}, \ldots, e_{\ell}\right)$, where for $1 \leq j \leq \ell, e_{j}=\left(v_{j}, u_{j}\right)$. Let $c_{j}, c_{j}^{\prime}$ are the assignments to $v_{j}$ and $u_{j}$ given by the $j$ th coordinate of $A$, so that $c_{j} \neq c_{j}^{\prime}$. Recall that there is exactly one index $z_{j} \in\{1, \ldots, 6\}$, such that $\chi_{z_{j}}$ assigns the color $c_{j}$ to $v_{j}$ and the color $c_{j}^{\prime}$ to $u_{j}$. Let $b^{*} \in B$ be the vector, where for $1 \leq j \leq \ell, b_{j}^{*}=z_{j}$. Then $f_{b^{*}}(Q)=A$, and for all $b \neq b^{*}, f_{b}(Q) \neq A$.
Finally, fix some query $Q^{\prime} \in \mathcal{Q}^{V}$ of the vertex-prover, and some answer $A^{\prime} \in \mathcal{A}^{V}$ to it. Let $Q^{\prime}=\left(v_{1}, \ldots, v_{\ell}\right)$. Assume that for each $1 \leq j \leq \ell$, the $j$ th coordinate of $A^{\prime}$ contains the color $c_{j}$. Recall that there are exactly two indices $z \in\{1, \ldots, 6\}$, such that $\chi_{z}$ assigns the color $c_{j}$ to $v_{j}$. Denote this set of two indices by $Z_{j} \subseteq\{1, \ldots, 6\}$. Consider now some vector $b \in B$. If, for all $1 \leq j \leq \ell, b_{j} \in Z_{j}$, then $f_{b}\left(Q^{\prime}\right)=A^{\prime}$; otherwise, $f_{b}\left(Q^{\prime}\right) \neq A^{\prime}$. Therefore, the total number of vectors $b \in B$, for which $f_{b}\left(Q^{\prime}\right)=A^{\prime}$ is exactly $2^{\ell}$.

## B Proof of Observation 6.4

We use the standard greedy algorithm. Start with $\mathcal{S}^{\prime}=\emptyset$, and then iterate. In each iteration $i$, we consider the current partial solution $\mathcal{S}^{\prime}$, and let $J_{i}$ be the set of all elements $j$, such that $j \notin \bigcup_{S \in \mathcal{S}^{\prime}} S$. We let $S_{i} \in \mathcal{S}$ be the set containing largest number of elements in $J_{i}$, add $S_{i}$ to $\mathcal{S}^{\prime}$, and continue to the next iteration. The algorithm terminates once $J_{i}=\emptyset$.

## Almost Polynomial Hardness of Node-Disjoint Paths in Grids

In order to analyze the algorithm, consider the $i$ th iteration. Notice that, since every element belongs to at least $m / t$ sets of $\mathcal{S}$, there must be at least one set $S \in \mathcal{S}$ that contains at least $\left|J_{i}\right| / t$ elements of $J_{i}$. Therefore, set $S_{i}$ that the algorithm chooses covers at least $\left|J_{i}\right| / t$ elements of $J_{i}$, and $\left|J_{i+1}\right| \leq(1-1 / t)\left|J_{i}\right|$. Overall, we get that for all $i,\left|J_{i}\right| \leq(1-1 / t)^{i-1} n$, and so both $\left|\mathcal{S}^{\prime}\right|$, and the number of iterations are bounded by $O(t \log n)$.

## C Proof of Claim 6.13

The goal of this section is to prove Claim 6.13. We show an efficient algorithm to construct a set $\mathcal{P}^{1}=\left\{P_{1}, \ldots, P_{M^{\prime}}\right\}$ of spaced-out paths, that originate at the vertices of $X$ on $R$, and traverse the boxes $\hat{K}_{j}$ in a snake-like fashion (see Figure 4). We will ensure that for each $1 \leq i \leq M^{\prime}$ and $1 \leq j \leq N_{1}$, the intersection of the path $P_{i}$ with the box $\hat{K}_{j}$ is the $i$ th column of $\mathcal{W}_{j}$, and that $P_{i}$ contains the vertex $x_{i}$.
Fix some index $1 \leq j \leq N_{1}$, and consider the box $\hat{K}_{j}$. Let $I_{j}^{\prime}$ and $I_{j}^{\prime \prime}$ denote the top and the bottom boundaries of $\hat{K}_{j}$, respectively. For each $1 \leq i \leq M^{\prime}$, let $W_{j}^{i}$ denote the $i$ th column of $\mathcal{W}_{j}$, and let $x^{\prime}(j, i)$ and $x^{\prime \prime}(j, i)$ denote the topmost and the bottommost vertices of $W_{j}^{i}$, respectively.
For convenience, for each $1 \leq i \leq M^{\prime}$ we denote $x^{\prime \prime}(0, i)=x_{i}$, and we let $I_{0}^{\prime \prime} \subseteq R$ be a subpath of $R$ containing exactly $2 M$ vertices, such that $x_{1}, \ldots, x_{M^{\prime}} \in I_{0}^{\prime \prime}$. The following claim is central to the proof.
Claim C.1. There is an efficient algorithm to construct, for each $1 \leq j \leq N_{1}$, a set $\mathcal{P}_{j}=\left\{P_{j}^{1}, \ldots, P_{j}^{M^{\prime}}\right\}$ of paths, such that:

- For each $1 \leq j \leq N_{1}$ and $1 \leq i \leq M^{\prime}$, path $P_{j}^{i}$ connects $x^{\prime \prime}(j-1, i)$ to $x^{\prime}(j, i)$;
- The set $\bigcup_{j=1}^{N_{1}} \mathcal{P}_{j}$ of paths is spaced-out; and
- All paths in $\bigcup_{j=1}^{N_{1}} \mathcal{P}_{j}$ are contained in $G^{t}$, and are internally disjoint from $R$.

Notice that by combining the paths in $\bigcup_{j=1}^{N_{1}} \mathcal{P}_{j}$ with the set $\bigcup_{j=1}^{N_{1}} \mathcal{W}_{j}$ of columns of the boxes, we obtain the desired set $\mathcal{P}^{1}$ of paths, completing the proof of Claim 6.13. In the remainder of this section we focus on proving Claim C.1.
Recall that each box $\hat{K}_{j}$ is separated by at least $2 M$ columns from every other such box, and from the left and right boundaries of $\hat{G}$. It is also separated by at least $4 M$ rows from the top boundary of $\hat{G}$ and from the row $R$. We exploit this spacing in order to construct the paths, by utilizing a special structure called a snake, that we define next.
Given a set $\mathcal{L}$ of consecutive rows of $\hat{G}$ and a set $\mathcal{W}$ of consecutive columns of $\hat{G}$, we denote by $\Upsilon(\mathcal{L}, \mathcal{W})$ the subgraph of $\hat{G}$ spanned by the rows in $\mathcal{L}$ and the columns in $\mathcal{W}$; we refer to such a graph as a corridor. Let $\Upsilon=\Upsilon(\mathcal{L}, \mathcal{W})$ be any such corridor. Let $L^{\prime}$ and $L^{\prime \prime}$ be the top and the bottom row of $\mathcal{L}$ respectively, and let $W^{\prime}$ and $W^{\prime \prime}$ be the first and the last column of $\mathcal{W}$ respectively. The four paths $\mathrm{r} \cap L^{\prime}, \mathrm{r} \cap L^{\prime \prime}, \mathrm{r} \cap W^{\prime}$ and $\Upsilon \cap W^{\prime \prime}$ are called the top, bottom, left and right boundaries of $\Upsilon$ respectively, and their union is called

## Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat

the boundary of $\Upsilon$. The width of the corridor $\Upsilon$ is $w(\Upsilon)=\min \{|\mathcal{L}|,|\mathcal{W}|\}$. We say that two corridors $\Upsilon, \Upsilon^{\prime}$ are internally disjoint, iff every vertex $v \in \Upsilon \cap \mathrm{\Upsilon}^{\prime}$ belongs to the boundaries of both corridors. We say that two internally disjoint corridors $\Upsilon, \Upsilon^{\prime}$ are neighbors iff $\Upsilon \cap \Upsilon^{\prime} \neq \emptyset$. We are now ready to define snakes.

A snake $y$ of length $z$ is a sequence $\left(\Upsilon_{1}, \Upsilon_{2}, \ldots, \Upsilon_{z}\right)$ of $z$ corridors that are pairwise internally disjoint, such that for all $1 \leq z^{\prime}, z^{\prime \prime} \leq z, \Upsilon_{z^{\prime}}$ is a neighbor of $\Upsilon_{z^{\prime \prime}}$ iff $\left|z^{\prime}-z^{\prime \prime}\right|=1$. The width of the snake is defined to be the minimum of two quantities: (i) $\min _{1 \leq z^{\prime}<z}\left\{\left|\Upsilon_{z^{\prime}} \cap \Upsilon_{z^{\prime}+1}\right|\right\}$; and (ii) $\min _{1 \leq z^{\prime} \leq z}\left\{w\left(\Upsilon_{z^{\prime}}\right)\right\}$. Notice that, given a snake $y$, there is a unique simple cycle $\sigma(y)$ contained in $\bigcup_{z^{\prime}=1}^{z} \Upsilon_{z^{\prime}}$, such that, if $D$ denotes the disc on the plane whose boundary is $\sigma(y)$, then every vertex of $\bigcup_{z^{\prime}=1}^{z} \Upsilon_{z^{\prime}}$ lies in $D$, while every other vertex of $\hat{G}$ lies outside $D$. We call $\sigma(y)$ the boundary of $y$. We say that a vertex $u$ belongs to a snake $y$, and denote $u \in \mathcal{y}$, iff $u \in \bigcup_{z^{\prime}=1}^{z} \Upsilon_{z^{\prime}}$. We use the following simple claim, whose proof can be found, e.g. in [19].
Claim C.2. Let $y=\left(\Upsilon_{1}, \ldots, \Upsilon_{z}\right)$ be a snake of width $w$, and let $A$ and $A^{\prime}$ be two sets of vertices with $|A|=\left|A^{\prime}\right| \leq w-2$, such that the vertices of A lie on a single boundary edge of $\Upsilon_{1}$, and the vertices of $A^{\prime}$ lie on a single boundary edge of $\Upsilon_{z}$. There is an efficient algorithm, that, given the snake $y$, and the sets $A$ and $A^{\prime}$ of vertices as above, computes a set $Q$ of node-disjoint paths contained in $y$, that connect every vertex of $A$ to a distinct vertex of $A^{\prime}$.

Let $L, L^{\prime}$ be any pair of rows of $G$. Let $Q$ be a set of node-disjoint paths connecting some set of vertices $B \subseteq L$ to $B^{\prime} \subseteq L^{\prime}$. We say that the paths in $Q$ are order-preserving iff the left-to-right ordering of their endpoints on $L$ is same as that of their endpoints on $L^{\prime}$.
Corollary C.3. Let $y=\left(\Upsilon_{1}, \ldots, \Upsilon_{z}\right)$ be a snake of width $w$, and let $B$ and $B^{\prime}$ be two sets of $r \leq\lfloor w / 2\rfloor-1$ vertices each, such that the vertices of $B$ lie on the bottom boundary edge of $\Upsilon_{1}$, the vertices of $B^{\prime}$ lie on the top boundary edge of $\Upsilon_{z}$ and for every pair $v, v^{\prime} \in B \cup B^{\prime}$ of vertices, $d_{\hat{G}}\left(v, v^{\prime}\right) \geq 2$. There is an efficient algorithm, that, given the snake $y$, and the sets $B$ and $B^{\prime}$ of vertices as above, computes a set $\hat{Q}$ of spaced-out order-preserving paths contained in $y$.

Proof. Let $B=\left\{b_{1}, b_{2}, \ldots, b_{r}\right\}$ and $B^{\prime}=\left\{b_{1}^{\prime}, b_{2}^{\prime}, \ldots, b_{r}^{\prime}\right\}$. Assume that the vertices in both sets are indexed according to their left-to-right ordering on their corresponding rows of the grid. Since set $B$ does not contain a pair of neighboring vertices, we can augment it to a larger set $A$, by adding a vertex between every consecutive pair of vertices of $B$. In other words, we obtain a vertex set $A=\left\{a_{1}, \ldots, a_{2 r-1}\right\}$, such that for all $1 \leq i \leq r, a_{2 i-1}=b_{i}$, and the vertices of $A$ are indexed according to their left-to-right ordering on the bottom boundary of $\Upsilon_{1}$. Similarly, we can augment the set $B^{\prime}$ to a set $A^{\prime}=\left\{a_{1}^{\prime}, \ldots, a_{2 r-1}^{\prime}\right\}$ of vertices, such that for all $1 \leq i \leq r, a_{2 i-1}^{\prime}=b_{i}^{\prime}$, and the vertices of $A^{\prime}$ are indexed according to their left-to-right ordering on the top boundary of $\Upsilon_{z}$.
We apply Claim C. 2 to the sets $A, A^{\prime}$ of vertices, obtaining a set $Q$ of node-disjoint paths, that are contained in $y$, and connect every vertex of $A$ to a distinct vertex of $A^{\prime}$. For all $1 \leq i \leq r$, let $Q_{i}$ be the path originating from $a_{i}$. We claim that $Q$ is an order-preserving set of paths. Indeed, assume for contradiction that some path $Q_{i} \in Q$ connects $a_{i}$ to $a_{i^{\prime}}^{\prime}$, for $i \neq i^{\prime}$. Notice that the path $Q_{i}$ partitions the snake $y$ into two subgraphs: one containing $(i-1)$ vertices of $A$ and $\left(i^{\prime}-1\right)$ vertices of $A^{\prime}$; and the other containing the remaining vertices of $A$ and $A^{\prime}$ (excluding the endpoints of $Q_{i}$ ). Since $i \neq i^{\prime}$, there must be a path $Q_{i^{\prime \prime}} \in Q$ intersecting the path $Q_{i}$, a contradiction to the fact that $Q$ is a set of node-disjoint paths.

## Almost Polynomial Hardness of Node-Disjoint Paths in Grids

Similarly, it is easy to see that for all $1 \leq i<r, d\left(Q_{2 i-1}, Q_{2 i+1}\right) \geq 2$. This is since the removal of the path $Q_{2 i}$ partitions the snake $y$ into two disjoint subgraphs, with path $Q_{2 i-1}$ contained in one and path $Q_{2 i+1}$ contained in the other.

Our final set of path is $\hat{Q}=\left\{Q_{2 i-1}: 1 \leq i \leq r\right\}$. From the above discussion, it is a spaced-out set of paths contained in $y$, and for each $1 \leq i \leq r$, path $Q_{2 i-1} \in \hat{Q}$ connects $b_{i}$ to $b_{i}^{\prime}$.

In order to complete the proof, we need the following easy observation.
Observation C.4. There is an efficient algorithm that constructs, for each $1 \leq j \leq N_{1}$, a snake $y_{j}$ of width at least $2 M$ in $G^{t}$, such that all resulting snakes are mutually disjoint, and for each $1 \leq j \leq N_{1}$ :

- the bottom boundary of the first corridor of $y_{j}$ contains $I_{j-1}^{\prime \prime}$;
- the top boundary of the last corridor of $y_{j}$ contains $I_{j}^{\prime}$; and
- all snakes are disjoint from $R$, except for $y_{1}$, that contains $I_{0}^{\prime \prime} \subseteq R$ as part of is boundary, and does not contain any other vertices of $R$.

The construction of the snakes is immediate and exploits the ample space between the boxes $\hat{K}_{j}$; (see Figure 4 for an illustration). From Corollary C.3, for each $1 \leq j \leq N_{1}$, we obtain a set $\mathcal{P}_{j}$ of spacedout paths contained in $y_{j}$, such that for each $1 \leq i \leq M^{\prime}$, there is a path $P_{j}^{i} \in \mathcal{P}_{j}$ connecting vertex $x^{\prime \prime}(j-1, i)$ to vertex $x^{\prime}(j, i)$. For each $1 \leq i \leq M^{\prime}$, let $P_{i}$ be the path obtained by concatenating the paths $\left\{P_{1}^{i}, \mathcal{W}_{1}^{i}, P_{2}^{i}, \ldots, P_{N_{1}}^{i}, \mathcal{W}_{N_{1}}^{i}\right\}$. The final set of paths is $\mathcal{P}^{1}=\left\{P_{1}, \ldots, P_{M^{1}}\right\}$.

## References

[1] Noga Alon, Sanjeev Arora, Rajsekar Manokaran, Dana Moshkovitz, and Omri WEINSTEIN: Inapproximabilty of densest k-subgraph from average case hardness, 2011. Manuscript. 5
[2] Noga Alon, Paul Seymour, and Robin Thomas: Planar separators. SIAM Journal on Discrete Mathematics, 7(2):184-193, 1994. 44
[3] Matthew Andrews: Approximation algorithms for the edge-disjoint paths problem via raecke decompositions. In 51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA, pp. 277-286, 2010. [doi:10.1109/FOCS.2010.33] 4
[4] 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. 2, 4
[5] Matthew Andrews and Lisa Zhang: Logarithmic hardness of the undirected edge-disjoint paths problem. J. ACM, 53(5):745-761, 2006. [doi:10.1145/1183907.1183910] 2

## Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat

[6] 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, pp. 567-576, Philadelphia, PA, USA, 1995. Society for Industrial and Applied Mathematics. 3, 4
[7] Baruch Awerbuch, Rainer Gawlick, Tom Leighton, and Yuval Rabani: On-line admission control and circuit routing for high performance computing and communication. In Proceedings 35th Annual Symposium on Foundations of Computer Science, pp. 412-423, Nov 1994. [doi:10.1109/SFCS.1994.365675] 4
[8] Aditya Bhaskara, Moses Charikar, Eden Chlamtac, Uriel Feige, and Aravindan ViJAYARAGHAVAN: Detecting high log-densities: an $O\left(n^{1 / 4}\right)$ approximation for densest $k$-subgraph. In Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010, pp. 201-210, 2010. [doi:10.1145/1806689.1806718] 5
[9] Andrei Z. Broder, Alan M. Frieze, Stephen Suen, and Eli Upfal: Optimal construction of edge-disjoint paths in random graphs. SIAM Journal on Computing, 28(2):541-573, 1998. [doi:10.1137/S0097539795290805, doi.org10.1137S0097539795290805] 4
[10] Andrei Z. Broder, Alan M. Frieze, and Eli Upfal: Existence and construction of edge disjoint paths on expander graphs. In Proceedings of the Twenty-fourth Annual ACM Symposium on Theory of Computing, STOC '92, pp. 140-149, New York, NY, USA, 1992. ACM. [doi:10.1145/129712.129727] 4
[11] Chandra Chekuri and Alina Ene: Poly-logarithmic approximation for maximum node disjoint paths with constant congestion. In Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013, pp. 326-341, 2013. [doi:10.1137/1.9781611973105.24] 4
[12] Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd: Multicommodity flow, well-linked terminals, and routing problems. In Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005, pp. 183-192, 2005. [doi:10.1145/1060590.1060618] 4
[13] 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. 3
[14] Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd: Edge-disjoint paths in planar graphs with constant congestion. SIAM Journal on Computing, 39(1):281-301, 2009. 4
[15] Chandra Chekuri, Marcelo Mydlarz, and F. Bruce Shepherd: Multicommodity demand flow in a tree and packing integer programs. ACM Trans. Algorithms, 3(3), August 2007. [doi:10.1145/1273340.1273343] 4
[16] Julia Chuzhoy: Routing in undirected graphs with constant congestion. SIAM J. Comput., 45(4):1490-1532, 2016. [doi:10.1137/130910464] 4

## Almost Polynomial Hardness of Node-Disjoint Paths in Grids

[17] Julia Chuzhoy and David H. K. Kim: On approximating node-disjoint paths in grids. In Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2015, August 24-26, 2015, Princeton, NJ, USA, pp. 187-211, 2015. [doi:10.4230/LIPIcs.APPROX-RANDOM.2015.187] 2, 3
[18] 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, pp. 556-569, New York, NY, USA, 2016. ACM. [doi:10.1145/2897518.2897538] 2
[19] Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat: New hardness results for routing on disjoint paths. In Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017, pp. 86-99, 2017. [doi:10.1145/3055399.3055411] 2, 3, 6, 50
[20] Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat: Improved Approximation for Node-Disjoint Paths in Grids with Sources on the Boundary. In 45th International Colloquium on Automata, Languages, and Programming (ICALP 2018), volume 107 of Leibniz International Proceedings in Informatics (LIPIcs), pp. 38:1-38:14, 2018. [doi:10.4230/LIPIcs.ICALP.2018.38] 3
[21] Julia Chuzhoy and Shi Li: A polylogarithmic approximation algorithm for edge-disjoint paths with congestion 2. J. ACM, 63(5):45:1-45:51, 2016. 4
[22] Shimon Even, Alon Itai, and Adi Shamir: On the complexity of timetable and multicommodity flow problems. SIAM J. Comput., 5(4):691-703, 1976. [doi:10.1137/0205048] 2
[23] URIEL FEIGE: Relations between average case complexity and approximation complexity. In Proceedings of the Thiry-fourth Annual ACM Symposium on Theory of Computing, STOC '02, pp. 534-543, New York, NY, USA, 2002. ACM. [doi:10.1145/509907.509985] 5
[24] Uriel Feige, Magnús M. Halldórsson, Guy Kortsarz, and Aravind SriniVASAN: Approximating the domatic number. SIAM J. Comput., 32(1):172-195, January 2003. [doi:10.1137/S0097539700380754] 8
[25] Krzysztof Fleszar, Matthias Mnich, and Joachim Spoerhase: New Algorithms for Maximum Disjoint Paths Based on Tree-Likeness. In 24th Annual European Symposium on Algorithms (ESA 2016), volume 57 of Leibniz International Proceedings in Informatics (LIPIcs), pp. 42:1-42:17, 2016. [doi:10.4230/LIPIcs.ESA.2016.42] 4
[26] Alan M. FriEze: Edge-disjoint paths in expander graphs. In Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '00, pp. 717-725, Philadelphia, PA, USA, 2000. Society for Industrial and Applied Mathematics. 4
[27] Naveen Garg, Vijay V. VaZirani, and Mihalis Yannakakis: Primal-dual approximation algorithms for integral flow and multicut in trees. Algorithmica, 18(1):3-20, 1997. [doi:10.1007/BF02523685] 4

## Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat

[28] Thomas Holenstein: Parallel repetition: Simplifications and the no-signaling case. In Proceedings of the Thirty-ninth Annual ACM Symposium on Theory of Computing, STOC '07, pp. 411-419, New York, NY, USA, 2007. ACM. [doi:10.1145/1250790.1250852] 9
[29] Richard Karp: On the complexity of combinatorial problems. Networks, 5:45-68, 1975. 2
[30] 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):16:116:13, March 2013. [doi:10.1145/2438645.2438648] 4
[31] Subhash Кнот: Ruling out PTAS for graph min-bisection, dense k-subgraph, and bipartite clique. SIAM J. Comput., 36(4):1025-1071, 2006. [doi:10.1137/S0097539705447037] 5
[32] 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, pp. 627-636, Washington, DC, USA, 2005. IEEE Computer Society. [doi:10.1109/SFCS.2005.18] 4
[33] Jon Kleinberg and Ronitt Rubinfeld: Short paths in expander graphs. In Proceedings of the 37th Annual Symposium on Foundations of Computer Science, FOCS '96, pp. 86-95, Washington, DC, USA, 1996. IEEE Computer Society. 4
[34] Jon M. Kleinberg and Éva Tardos: Disjoint paths in densely embedded graphs. In Proceedings of the 36th Annual Symposium on Foundations of Computer Science, pp. 52-61, 1995. 3, 4
[35] Jon M. Kleinberg and Éva Tardos: Approximations for the disjoint paths problem in highdiameter planar networks. J. Comput. Syst. Sci., 57(1):61-73, 1998. 3, 4
[36] Stavros G. Kolliopoulos and Clifford Stein: Approximating disjoint-path problems using packing integer programs. Mathematical Programming, 99:63-87, 2004. [doi:10.1007/s10107-002-0370-6] 2, 3
[37] Mark R. 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. 2
[38] Tom Leighton and Satish Rao: Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. J. ACM, 46(6):787-832, November 1999. [doi:10.1145/331524.331526] 4
[39] Richard J Lipton and Robert Endre Tarjan: A separator theorem for planar graphs. SIAM Journal on Applied Mathematics, 36(2):177-189, 1979. 44
[40] James F. Lynch: The equivalence of theorem proving and the interconnection problem. SIGDA Newsl., 5(3):31-36, September 1975. [doi:10.1145/1061425.1061430] 2
[41] Pasin Manurangsi: Almost-polynomial ratio eth-hardness of approximating densest k-subgraph. In Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017, pp. 954-961, 2017. [doi:10.1145/3055399.3055412] 5
[42] Harald Räcke: Minimizing congestion in general networks. In 43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings, pp. 43-52, 2002. [doi:10.1109/SFCS.2002.1181881] 4
[43] Prabhakar Raghavan and Clark D. Thompson: Randomized rounding: a technique for provably good algorithms and algorithmic proofs. Combinatorica, 7:365-374, December 1987. [doi:10.1007/BF02579324] 4
[44] Prasad Raghavendra and David Steurer: Graph expansion and the unique games conjecture. In Proceedings of the Forty-second ACM Symposium on Theory of Computing, STOC '10, pp. 755-764, New York, NY, USA, 2010. ACM. [doi:10.1145/1806689.1806792] 5
[45] Anup RaO: Parallel repetition in projection games and a concentration bound. In Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, STOC '08, pp. 1-10, New York, NY, USA, 2008. ACM. [doi:10.1145/1374376.1374378] 9
[46] Satish Rao and Shuheng Zhou: Edge disjoint paths in moderately connected graphs. SIAM J. Comput., 39(5):1856-1887, 2010. 4
[47] Ran Raz: A parallel repetition theorem. SIAM J. Comput., 27(3):763-803, June 1998. [doi:10.1137/S0097539795280895] 9
[48] Neil Robertson and Paul D. Seymour: Outline of a disjoint paths algorithm. In Paths, Flows and VLSI-Layout. Springer-Verlag, 1990. 2
[49] Neil Robertson and Paul D Seymour: Graph minors. XIII. The disjoint paths problem. Journal of Combinatorial Theory, Series B, 63(1):65-110, 1995. 2
[50] 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, pp. 200-209, Washington, DC, USA, 2011. IEEE Computer Society. [doi:10.1109/FOCS.2011.30] 4
[51] Peter Ungar: A theorem on planar graphs. Journal of the London Mathematical Society, 1(4):256-262, 1951. 44

## AUTHORS

Julia Chuzhoy<br>Professor<br>Toyota Technological Institute at Chicago<br>Chicago, IL, USA<br>cjulia@ttic.edu<br>https://ttic.uchicago.edu/~cjulia

David H. K. Kim
Software Engineer
Google
Sunnyvale, CA, USA
dave.hongk@gmail.com

## Rachit Nimavat

Student
Toyota Technological Institute at Chicago
Chicago, IL, USA
nimavat@ttic.edu

## ABOUT THE AUTHORS

Julia Chuzhoy is a Professor at the Toyota Technological Institute at Chicago. She completed her Ph.D. in Technion, Israel, and spend three years as a postdoctoral scholar at MIT, University of Pennsylvania and Institute for Advanced Studies in Princeton. She mainly works on algorithms for graph problems. In her spare time she likes to read books, learns to play piano and studies French.

David H. K. Kim received his Ph.D. in Computer Science in May 2018 from the University of Chicago, where he studied approximation algorithms for network routing problems under the supervision of Julia Chuzhoy. He also worked on scheduling problems and their applications on real systems with Henry Hoffmann. He now works as a software engineer at Google and enjoys spending his free time hiking and listening to music.


[^0]:    An extended abstract of this paper appeared in the Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing
    *Supported in part by NSF grants CCF-1318242 and CCF-1616584.
    ${ }^{\dagger}$ Supported in part by NSF grant CCF-1318242 and CCF-1616584.
    ${ }^{\ddagger}$ Supported in part by NSF grant CCF-1318242 and CCF-1616584.

[^1]:    ${ }^{1}$ Note that the results are not strictly complementary: the algorithm only works for grid graphs, while the hardness result is only valid for subgraphs of grids.
    ${ }^{2}$ The original version of this paper contained a randomized reduction, with a somewhat stronger complexity assumptions that NP $\nsubseteq \operatorname{RTIME}\left(n^{\text {poly } \log n}\right)$ and NP $\nsubseteq \operatorname{RTIME}\left(2^{n^{\delta}}\right)$, respectively. The derandomization strategy that we have used was proposed to us by a reviewer.

[^2]:    ${ }^{3}$ We note that a similar problem arises if one attempts to design naive hardness of approximation proofs for DkS.

[^3]:    ${ }^{4}$ This proof was suggested to us by one of the reviewers of the paper. A previous version of the paper contained a slightly weaker proof.

