

# Improved Approximation Algorithm for Node-Disjoint Paths in Grid Graphs with Sources on Grid Boundary

Julia Chuzhoy\*

David H. K. Kim<sup>†</sup>

Rachit Nimavat<sup>‡</sup>

May 24, 2018

## Abstract

We study the classical Node-Disjoint Paths (NDP) problem: given an undirected  $n$ -vertex graph  $G$ , together with a set  $\{(s_1, t_1), \dots, (s_k, t_k)\}$  of pairs of its vertices, called source-destination, or demand pairs, find a maximum-cardinality set  $\mathcal{P}$  of mutually node-disjoint paths that connect the demand pairs. The best current approximation for the problem is achieved by a simple greedy  $O(\sqrt{n})$ -approximation algorithm. Until recently, the best negative result was an  $\Omega(\log^{1/2-\epsilon} n)$ -hardness of approximation, for any fixed  $\epsilon$ , under standard complexity assumptions. A special case of the problem, where the underlying graph is a grid, has been studied extensively. The best current approximation algorithm for this special case achieves an  $\tilde{O}(n^{1/4})$ -approximation factor. On the negative side, a recent result by the authors shows that NDP is hard to approximate to within factor  $2^{\Omega(\sqrt{\log n})}$ , even if the underlying graph is a subgraph of a grid, and all source vertices lie on the grid boundary. In a very recent follow-up work, the authors further show that NDP in grid graphs is hard to approximate to within factor  $\Omega(2^{\log^{1-\epsilon} n})$  for any constant  $\epsilon$  under standard complexity assumptions, and to within factor  $n^{\Omega(1/(\log \log n)^2)}$  under randomized ETH.

In this paper we study the NDP problem in grid graphs, where all source vertices  $\{s_1, \dots, s_k\}$  appear on the grid boundary. Our main result is an efficient randomized  $2^{O(\sqrt{\log n} \cdot \log \log n)}$ -approximation algorithm for this problem. Our result in a sense complements the  $2^{\Omega(\sqrt{\log n})}$ -hardness of approximation for sub-graphs of grids with sources lying on the grid boundary, and should be contrasted with the above-mentioned almost polynomial hardness of approximation of NDP in grid graphs (where the sources and the destinations may lie anywhere in the grid). Much of the work on approximation algorithms for NDP relies on the multicommodity flow relaxation of the problem, which is known to have an  $\Omega(\sqrt{n})$  integrality gap, even in grid graphs, with all source and destination vertices lying on the grid boundary. Our work departs from this paradigm, and uses a (completely different) linear program only to select the pairs to be routed, while the routing itself is computed by other methods. We generalize this result to instances where the source vertices lie within a prescribed distance from the grid boundary.

---

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

<sup>†</sup>Computer Science Department, University of Chicago. Email: [hongk@cs.uchicago.edu](mailto:hongk@cs.uchicago.edu). Supported in part by NSF grants CCF-1318242 and CCF-1616584.

<sup>‡</sup>Toyota Technological Institute at Chicago. Email: [nimavat@ttic.edu](mailto:nimavat@ttic.edu). Supported in part by NSF grant CCF-1318242.

# 1 Introduction

We study the classical Node-Disjoint Paths (NDP) problem, where the input consists of an undirected  $n$ -vertex graph  $G$  and a collection  $\mathcal{M} = \{(s_1, t_1), \dots, (s_k, t_k)\}$  of pairs of its vertices, called *source-destination* or *demand* pairs. We say that a path  $P$  *routes* a demand pair  $(s_i, t_i)$  iff the endpoints of  $P$  are  $s_i$  and  $t_i$ . The goal is to compute a maximum-cardinality set  $\mathcal{P}$  of node-disjoint paths, where each path  $P \in \mathcal{P}$  routes a distinct demand pair in  $\mathcal{M}$ . We denote by NDP-Planar the special case of the problem when the underlying graph  $G$  is planar, and by NDP-Grid the special case where  $G$  is a square grid<sup>1</sup>. We refer to the vertices in set  $S = \{s_1, \dots, s_k\}$  as *source vertices*; to the vertices in set  $T = \{t_1, \dots, t_k\}$  as *destination vertices*, and to the vertices in set  $S \cup T$  as *terminals*.

NDP is a fundamental graph routing problem that has been studied extensively in both graph theory and theoretical computer science communities. Robertson and Seymour [RS90, RS95] explored the problem in their Graph Minor series, providing an efficient algorithm for NDP when the number  $k$  of the demand pairs is bounded by a constant. But when  $k$  is a part of input, the problem becomes NP-hard [Kar75, EIS76], even in planar graphs [Lyn75], and even in grid graphs [KvL84]. The best current approximation factor of  $O(\sqrt{n})$  for NDP is achieved by a simple greedy algorithm [KS04]. Until recently, this was also the best approximation algorithm for NDP-Planar and NDP-Grid. A natural way to design approximation algorithms for NDP is via the multicommodity flow relaxation: instead of connecting each routed demand pair with a path, send maximum possible amount of (possibly fractional) flow between them. The optimal solution to this relaxation can be computed via a standard linear program. The  $O(\sqrt{n})$ -approximation algorithm of [KS04] can be cast as an LP-rounding algorithm of this relaxation. Unfortunately, it is well-known that the integrality gap of this relaxation is  $\Omega(\sqrt{n})$ , even when the underlying graph is a grid, with all terminals lying on its boundary. In a recent work, Chuzhoy and Kim [CK15] designed an  $\tilde{O}(n^{1/4})$ -approximation for NDP-Grid, thus bypassing this integrality gap barrier. Their main observation is that, if all terminals lie close to the grid boundary (say within distance  $O(n^{1/4})$ ), then a simple dynamic programming-based algorithm yields an  $O(n^{1/4})$ -approximation. On the other hand, if, for every demand pair, either the source or the destination lies at a distance at least  $\Omega(n^{1/4})$  from the grid boundary, then the integrality gap of the multicommodity flow relaxation improves, and one can obtain an  $\tilde{O}(n^{1/4})$ -approximation via LP-rounding. A natural question is whether the integrality gap improves even further, if all terminals lie further away from the grid boundary. Unfortunately, the authors show in [CK15] that the integrality gap remains at least  $\Omega(n^{1/8})$ , even if all terminals lie within distance  $\Omega(\sqrt{n})$  from the grid boundary. The  $\tilde{O}(n^{1/4})$ -approximation algorithm for NDP-Grid was later extended and generalized to an  $\tilde{O}(n^{9/19})$ -approximation algorithm for NDP-Planar [CKL16].

On the negative side, until recently, only an  $\Omega(\log^{1/2-\epsilon} n)$ -hardness of approximation was known for the general version of NDP, for any constant  $\epsilon$ , unless  $\text{NP} \subseteq \text{ZPTIME}(n^{\text{poly} \log n})$  [AZ06, ACG<sup>+</sup>10], and only APX-hardness was known for NDP-Planar and NDP-Grid [CK15]. In a recent work [CKN17b], the authors have shown that NDP is hard to approximate to within a  $2^{\Omega(\sqrt{\log n})}$  factor unless  $\text{NP} \subseteq \text{DTIME}(n^{O(\log n)})$ , even if the underlying graph is a planar graph with maximum vertex degree at most 3, and all source vertices lie on the boundary of a single face. The result holds even when the input graph  $G$  is a vertex-induced subgraph of a grid, with all sources lying on the grid boundary. In a very recent work [CKN17a], the authors show that NDP-Grid is  $2^{\Omega(\log^{1-\epsilon} n)}$ -hard to approximate for any constant  $\epsilon$  assuming  $\text{NP} \not\subseteq \text{BPTIME}(n^{\text{poly} \log n})$ , and moreover, assuming randomized ETH, the hardness of approximation factor becomes  $n^{\Omega(1/(\log \log n)^2)}$ . We note that the instances constructed in these latter hardness proofs require all terminals to lie far from the grid boundary.

---

<sup>1</sup>We use the standard convention of denoting  $n = |V(G)|$ , and so the grid has dimensions  $(\sqrt{n} \times \sqrt{n})$ ; we assume that  $\sqrt{n}$  is an integer.

In this paper we explore NDP-Grid. This important special case of NDP was initially motivated by applications in VLSI design, and has received a lot of attention since the 1960's. We focus on a restricted version of NDP-Grid, that we call Restricted NDP-Grid: here, in addition to the graph  $G$  being a square grid, we also require that all source vertices  $\{s_1, \dots, s_k\}$  lie on the grid boundary. We do not make any assumptions about the locations of the destination vertices, that may appear anywhere in the grid. The best current approximation algorithm for Restricted NDP-Grid is the same as that for the general NDP-Grid, and achieves a  $\tilde{O}(n^{1/4})$ -approximation [CK15]. Our main result is summarized in the following theorem.

**Theorem 1.1** *There is an efficient randomized  $2^{O(\sqrt{\log n} \cdot \log \log n)}$ -approximation algorithm for Restricted NDP-Grid.*

This result in a sense complements the  $2^{\Omega(\sqrt{\log n})}$ -hardness of approximation of NDP on sub-graphs of grids with all sources lying on the grid boundary of [CKN17b]<sup>2</sup>, and should be contrasted with the recent almost polynomial hardness of approximation of [CKN17a] for NDP-Grid mentioned above. Our algorithm departs from previous work on NDP in that it does not use the multicommodity flow relaxation. Instead, we define sufficient conditions that allow us to route a subset  $\mathcal{M}'$  of demand pairs via disjoint paths, and show that there exists a subset of demand pairs satisfying these conditions, whose cardinality is at least  $\text{OPT}/2^{O(\sqrt{\log n} \cdot \log \log n)}$ , where  $\text{OPT}$  is the value of the optimal solution. It is then enough to compute a maximum-cardinality subset of the demand pairs satisfying these conditions. We write an LP-relaxation for this problem and design a  $2^{O(\sqrt{\log n} \cdot \log \log n)}$ -approximation LP-rounding algorithm for it. We emphasize that the linear program is only used to select the demand pairs to be routed, and not to compute the routing itself.

We then generalize the result to instances where the source vertices lie within a prescribed distance from the grid boundary.

**Theorem 1.2** *For every integer  $\delta \geq 1$ , there is an efficient randomized  $(\delta \cdot 2^{O(\sqrt{\log n} \cdot \log \log n)})$ -approximation algorithm for the special case of NDP-Grid where all source vertices lie within distance at most  $\delta$  from the grid boundary.*

We note that for instances of NDP-Grid where both the sources and the destinations are within distance at most  $\delta$  from the grid boundary, it is easy to obtain an efficient  $O(\delta)$ -approximation algorithm (see, e.g. [CK15]).

A problem closely related to NDP is the Edge-Disjoint Paths (EDP) problem. It is defined similarly, except that now the paths chosen to route the demand pairs may share vertices, and are only required to be edge-disjoint. The approximability status of EDP is very similar to that of NDP: there is an  $O(\sqrt{n})$ -approximation algorithm [CKS06], and an  $\Omega(\log^{1/2-\epsilon} n)$ -hardness of approximation for any constant  $\epsilon$ , unless  $\text{NP} \subseteq \text{ZPTIME}(n^{\text{poly} \log n})$  [AZ06, ACG<sup>+</sup>10]. As in the NDP problem, we can use the standard multicommodity flow LP-relaxation of the problem, in order to obtain the  $O(\sqrt{n})$ -approximation algorithm, and the integrality gap of the LP-relaxation is  $\Omega(\sqrt{n})$  even in planar graphs. Recently, Fleszar et al. [FMS16] designed an  $O(\sqrt{r} \cdot \log(kr))$ -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.

Several special cases of EDP have better approximation algorithms: an  $O(\log^2 n)$ -approximation is known for even-degree planar graphs [CKS05, CKS04, Kle05], and an  $O(\log n)$ -approximation is known for nearly-Eulerian uniformly high-diameter planar graphs, and nearly-Eulerian densely embedded

---

<sup>2</sup>Note that the two results are not strictly complementary: our algorithm only applies to grid graphs, while the hardness result is only valid for sub-graphs of grids.

graphs, including grid graphs [AR95, KT98, KT95]. Furthermore, an  $O(\log n)$ -approximation algorithm is known for EDP on 4-edge-connected planar, and Eulerian planar graphs [KK13]. It appears that the restriction of the graph  $G$  to be Eulerian, or near-Eulerian, makes the EDP problem on planar graphs significantly simpler, and in particular improves the integrality gap of the standard multicommodity flow LP-relaxation.

The analogue of the grid graph for the EDP problem is the wall graph (see Figure 1): the integrality gap of the multicommodity flow relaxation for EDP on wall graphs is  $\Omega(\sqrt{n})$ . The  $\tilde{O}(n^{1/4})$ -approximation algorithm of [CK15] for NDP-Grid extends to EDP on wall graphs, and the  $2^{\Omega(\sqrt{\log n})}$ -hardness of approximation of [CKN17b] for NDP-Planar also extends to EDP on sub-graphs of walls, with all sources lying on the top boundary of the wall. The recent hardness result of [CKN17a] for NDP-Grid also extends to an  $2^{\Omega(\log^{1-\epsilon} n)}$ -hardness of EDP on wall graphs, assuming  $\text{NP} \not\subseteq \text{BPTIME}(n^{\text{poly} \log n})$ , and to  $n^{\Omega(1/(\log \log n)^2)}$ -hardness assuming randomized ETH. We extend our results to EDP and NDP on wall graphs:

**Theorem 1.3** *There is an efficient randomized  $2^{O(\sqrt{\log n} \cdot \log \log n)}$ -approximation algorithm for EDP and for NDP on wall graphs, when all source vertices lie on the wall boundary.*

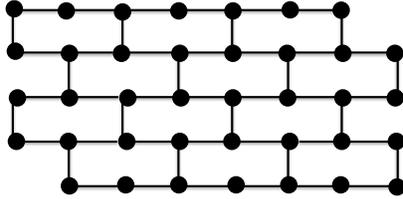


Figure 1: A wall graph.

**Other related work.** Cutler and Shiloach [CS78] studied an even more restricted version of NDP-Grid, where all source vertices lie on the top row  $R^*$  of the grid, and all destination vertices lie on a single row  $R'$  of the grid, far enough from its top and bottom boundaries. They considered three different settings of this special case. In the packed-packed setting, all sources appear consecutively on  $R^*$ , and all destinations appear consecutively on  $R'$  (but both sets may appear in an arbitrary order). They show a necessary and a sufficient condition for all demand pairs to be routable via node-disjoint paths in this setting. The second setting is the packed-spaced setting. Here, the sources again appear consecutively on  $R^*$ , but all destinations are at a distance at least  $d$  from each other. For this setting, the authors show that if  $d \geq k$ , then all demand pairs can be routed. We note that [CK15] extended their algorithm to a more general setting, where the destination vertices may appear anywhere in the grid, as long as the distance between any pair of the destination vertices, and any destination vertex and the boundary of the grid, is at least  $\Omega(k)$ . Robertson and Seymour [RS88] provided sufficient conditions for the existence of node-disjoint routing of a given set of demand pairs in the more general setting of graphs drawn on surfaces, and they designed an algorithm whose running time is  $\text{poly}(n) \cdot f(k)$  for finding the routing, where  $f(k)$  is at least exponential in  $k$ . Their result implies the existence of the routing in grids, when the destination vertices are sufficiently far from each other and from the grid boundaries, but it does not provide an efficient algorithm to compute such a routing. The third setting studied by Cutler and Shiloach is the spaced-spaced setting, where the distances between every pair of source vertices, and every pair of destination vertices are at least  $d$ . The authors note that they could not come up with a better algorithm for this setting, than the one provided for the packed-spaced case. Aggarwal, Kleinberg, and Williamson [AKW00] considered a special case of NDP-Grid, where the set of the demand pairs is a permutation: that is, every vertex

of the grid participates in exactly one demand pair. They show that  $\Omega(\sqrt{n}/\log n)$  demand pairs are routable in this case via node-disjoint paths. They further show that if all terminals are at a distance at least  $d$  from each other, then at least  $\Omega(\sqrt{nd}/\log n)$  pairs are routable.

A variation of the NDP and EDP problems, where small congestion is allowed, has been a subject of extensive study, starting with the classical paper of Raghavan and Thompson [RT87] that introduced the randomized rounding technique. We say that a set  $\mathcal{P}$  of paths causes congestion  $c$ , if at most  $c$  paths share the same vertex or the same edge, for the NDP and the EDP settings respectively. A recent line of work [CKS05, Rác02, And10, RZ10, Chu16, CL16, CE13, CC] has led to an  $O(\text{poly log } k)$ -approximation for both NDP and EDP problems with congestion 2. For planar graphs, a constant-factor approximation with congestion 2 is known [SCS11].

**Organization.** We start with a high-level intuitive overview of our algorithm in Section 2. We then provide Preliminaries in Section 3 and the algorithm for Restricted NDP-Grid in Section 4, with parts of the proof being deferred to Sections 5–9. We extend our algorithm to EDP and NDP on wall graphs in Section 10. We generalize our algorithm to the setting where the sources are within some prescribed distance from the grid boundary in Section 11.

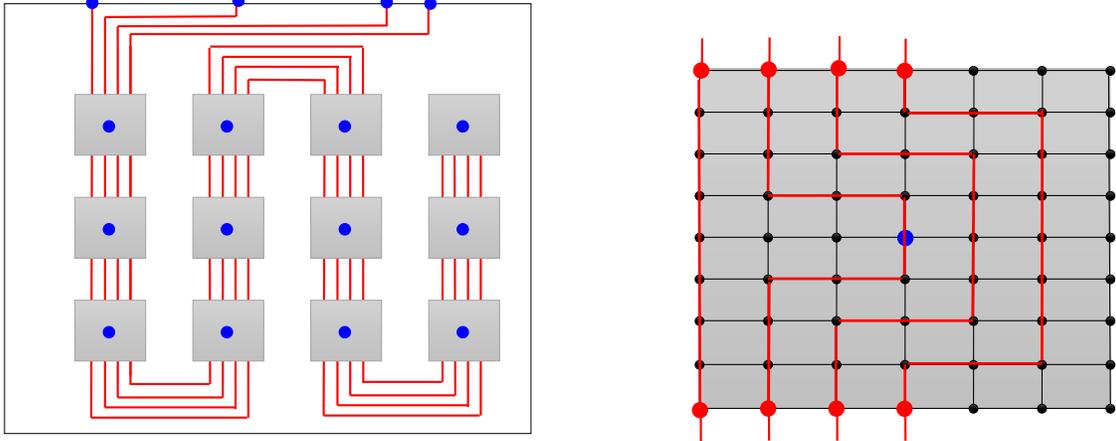
## 2 High-Level Overview of the Algorithm

The goal of this section is to provide an informal high-level overview of the main result of the paper – the proof of Theorem 1.1. With this goal in mind, the values of various parameters are given imprecisely in this section, in a way that best conveys the intuition. The following sections contain a formal description of the algorithm and the precise settings of all parameters.

We first consider an even more restricted special case of NDP-Grid, where all source vertices appear on the top boundary of the grid, and all destination vertices appear far enough from the grid boundary, and design an efficient randomized  $2^{O(\sqrt{\log n \cdot \log \log n})}$ -approximation algorithm  $\mathcal{A}$  for this problem. We later show how to reduce Restricted NDP-Grid to this special case of the problem; we focus on the description of the algorithm  $\mathcal{A}$  for now.

We assume that our input graph  $G$  is the  $(\ell \times \ell)$ -grid, and we denote by  $n = \ell^2$  the number of its vertices. We further assume that the set of the demand pairs is  $\mathcal{M} = \{(s_1, t_1), \dots, (s_k, t_k)\}$ , with the vertices in set  $S = \{s_1, \dots, s_k\}$  called source vertices; the vertices in set  $T = \{t_1, \dots, t_k\}$  called destination vertices; and the vertices in  $S \cup T$  called terminals. Let  $\text{OPT}$  denote the value of the optimal solution to the NDP instance  $(G, \mathcal{M})$ . We assume that the vertices of  $S$  lie on the top boundary of the grid, that we denote by  $R^*$ , and the vertices of  $T$  lie sufficiently far from the grid boundary – say, at a distance at least  $\text{OPT}$  from it. For a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of the demand pairs, we denote by  $S(\mathcal{M}')$  and  $T(\mathcal{M}')$  the sets of the source and the destination vertices of the demand pairs in  $\mathcal{M}'$ , respectively. As our starting point, we consider a simple observation of Chuzhoy and Kim [CK15], that generalizes the results of Cutler and Shiloach [CS78]. Suppose we are given an instance of NDP-Grid with  $k$  demand pairs, where the sources lie on the top boundary of the grid, and the destination vertices may appear anywhere in the grid, but the distance between every pair of the destination vertices, and every destination vertex and the boundary of the grid, is at least  $(8k + 8)$  – we call such instances *spaced-out instances*. In this case, all demand pairs in  $\mathcal{M}$  can be efficiently routed via node-disjoint paths, as follows. Consider, for every destination vertex  $t_i \in T$ , a square sub-grid  $B_i$  of  $G$ , of size  $(2k \times 2k)$ , such that  $t_i$  lies roughly at the center of  $B_i$ . We construct a set  $\mathcal{P}$  of  $k$  node-disjoint paths, that originate at the vertices of  $S$ , and traverse the sub-grids  $B_i$  one-by-one in a snake-like fashion (see a schematic view on Figure 2(a)). We call this part of the routing *global routing*. The *local routing* needs to specify how

the paths in  $\mathcal{P}$  traverse each box  $B_i$ . This is done in a straightforward manner, while ensuring that the unique path originating at vertex  $s_i$  visits the vertex  $t_i$  (see Figure 2(b)). By suitably truncating the final set  $\mathcal{P}$  of paths, we obtain a routing of all demand pairs in  $\mathcal{M}$  via node-disjoint paths.



(a) Global routing. In this figure, the sub-grids  $B_i$  are aligned vertically and horizontally. A similar (but somewhat more complicated) routing can be performed even if they are not aligned. For convenience we did not include all source vertices and all paths.

(b) Local routing inside  $B_i$

Figure 2: Schematic view of routing of spaced-out instances.

Unfortunately, in our input instance  $(G, \mathcal{M})$ , the destination vertices may not be located sufficiently far from each other. We can try to select a large subset  $\mathcal{M}' \subseteq \mathcal{M}$  of the demand pairs, so that every pair of destination vertices in  $T(\mathcal{M}')$  appear at a distance at least  $\Omega(|\mathcal{M}'|)$  from each other; but in some cases the largest such set  $\mathcal{M}'$  may only contain  $O(\text{OPT}/\sqrt{k})$  demand pairs (for example, suppose all destination vertices lie consecutively on a single row of the grid). One of our main ideas is to generalize this simple algorithm to a number of recursive levels.

For simplicity, let us first describe the algorithm with just two recursive levels. Suppose we partition the top row of the grid into  $z$  disjoint intervals,  $I_1, \dots, I_z$ . Let  $\mathcal{M}' \subseteq \mathcal{M}$  be a set of demand pairs that we would like to route. Denote  $|\mathcal{M}'| = k'$ , and assume that we are given a collection  $\mathcal{Q}$  of square sub-grids of  $G$ , of size  $(4k' \times 4k')$  each (that we call *squares*), such that every pair  $Q, Q' \in \mathcal{Q}$  of distinct squares is at a distance at least  $4k'$  from each other. Assume further that each such sub-grid  $Q \in \mathcal{Q}$  is assigned a color  $\chi(Q) \in \{c_1, \dots, c_z\}$ , such that, if  $Q$  is assigned the color  $c_j$ , then all demand pairs  $(s, t) \in \mathcal{M}'$  whose destination  $t$  lies in  $Q$  have their source  $s \in I_j$  (so intuitively, each color  $c_j$  represents an interval  $I_j$ ). Let  $\mathcal{M}'_j \subseteq \mathcal{M}'$  be the set of all demand pairs  $(s, t) \in \mathcal{M}'$  with  $s \in I_j$ . We would like to ensure that  $|\mathcal{M}'_j|$  is roughly  $k'/z$ , and that all destination vertices of  $T(\mathcal{M}'_j)$  are at a distance at least  $|\mathcal{M}'_j|$  from each other. We claim that if we could find the collection  $\{I_1, \dots, I_z\}$  of the intervals of the first row, a collection  $\mathcal{Q}$  of sub-grids of  $G$ , a coloring  $\chi : \mathcal{Q} \rightarrow \{c_1, \dots, c_z\}$ , and a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of the demand pairs with these properties, then we would be able to route all demand pairs in  $\mathcal{M}'$ .

In order to do so, for each square  $Q \in \mathcal{Q}$ , we construct an augmented square  $Q^+$ , by adding a margin of  $k'$  rows and columns around  $Q$ . Our goal is to construct a collection  $\mathcal{P}$  of node-disjoint paths routing the demand pairs in  $\mathcal{M}'$ . We start by constructing a global routing, where all paths in  $\mathcal{P}$  originate from the vertices of  $S(\mathcal{M}')$  and then visit the squares in  $\{Q^+ \mid Q \in \mathcal{Q}\}$  in a snake-like fashion, just like we did for the spaced-out instances described above (see Figure 2(a)). Consider now some square  $Q \in \mathcal{Q}$  and the corresponding augmented square  $Q^+$ . Assume that  $\chi(Q) = c_j$ , and let  $\mathcal{P}_j \subseteq \mathcal{P}$  be the set of paths originating at the source vertices that lie in  $I_j$ . While traversing the square  $Q^+$ , we

ensure that only the paths in  $\mathcal{P}_j$  enter the square  $Q$ ; the remaining paths use the margins on the left and on the right of  $Q$  in order to traverse  $Q^+$ . This can be done because the sources of the paths in  $\mathcal{P}_j$  appear consecutively on  $R^*$ , relatively to the sources of all paths in  $\mathcal{P}$ . In order to complete the local routing inside the square  $Q$ , observe that the destination vertices appear far enough from each other, and so we can employ the simple algorithm for spaced-out instances inside  $Q$ .

In order to optimize the approximation factor that we achieve, we extend this approach to  $\rho = O(\sqrt{\log n})$  recursive levels. Let  $\eta = 2^{\lceil \sqrt{\log n} \rceil}$ . We define auxiliary parameters  $d_1 > d_2 > \dots > d_\rho > d_{\rho+1}$ . Roughly speaking, we can think of  $d_{\rho+1}$  as being a constant (say 16), of  $d_1$  as being comparable to OPT, and for all  $1 \leq h \leq \rho$ ,  $d_{h+1} = d_h/\eta$ . The setup for the algorithm consists of three ingredients: (i) a hierarchical decomposition  $\tilde{\mathcal{H}}$  of the grid into square sub-grids (that we refer to as squares); (ii) a hierarchical partition  $\mathcal{I}$  of the first row  $R^*$  of the grid into intervals; and (iii) a hierarchical coloring  $f$  of the squares in  $\tilde{\mathcal{H}}$  with colors that correspond to the intervals of  $\mathcal{I}$ , together with a selection of a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of the demand pairs to route. We define sufficient conditions on the hierarchical system  $\tilde{\mathcal{H}}$  of squares, the hierarchical partition  $\mathcal{I}$  of  $R^*$  into intervals, the coloring  $f$  and the subset  $\mathcal{M}'$  of the demand pairs, under which a routing of all pairs in  $\mathcal{M}'$  exists and can be found efficiently. For a fixed hierarchical system  $\tilde{\mathcal{H}}$  of squares, a triple  $(\mathcal{I}, f, \mathcal{M}')$  satisfying these conditions is called a *good ensemble*. We show that a good ensemble with a large enough set  $\mathcal{M}'$  of demand pairs exists, and then design an approximation algorithm for computing a good ensemble maximizing  $|\mathcal{M}'|$ . We now describe each of these ingredients in turn.

## 2.1 A Hierarchical System of Squares

A hierarchical system  $\tilde{\mathcal{H}}$  of squares consists of a sequence  $\mathcal{Q}_1, \mathcal{Q}_2, \dots, \mathcal{Q}_\rho$  of sets of sub-grids of  $G$ . For each  $1 \leq h \leq \rho$ ,  $\mathcal{Q}_h$  is a collection of disjoint sub-grids of  $G$  (that we refer to as *level- $h$  squares*); every such square  $Q \in \mathcal{Q}_h$  has size  $(d_h \times d_h)$ , and every pair of distinct squares  $Q, Q' \in \mathcal{Q}_h$  are within distance at least  $d_h$  from each other (see Figure 3). We require that for each  $1 < h \leq \rho$ , for every square  $Q \in \mathcal{Q}_h$ , there is a unique square  $Q' \in \mathcal{Q}_{h-1}$  (called the *parent-square* of  $Q$ ) that contains  $Q$ . We say that a demand pair  $(s, t)$  *belongs* to the hierarchical system  $\tilde{\mathcal{H}} = (\mathcal{Q}_1, \mathcal{Q}_2, \dots, \mathcal{Q}_\rho)$  of squares iff  $t \in \bigcup_{Q \in \mathcal{Q}_\rho} Q$ . We show a simple efficient algorithm to construct  $2^{O(\sqrt{\log n})}$  such hierarchical systems of squares, so that every demand pair belongs to at least one of them. Each such system  $\tilde{\mathcal{H}}$  of squares induces an instance of NDP—the instance is defined over the same graph  $G$ , and the set  $\tilde{\mathcal{M}} \subseteq \mathcal{M}$  of demand pairs that belong to the system  $\tilde{\mathcal{H}}$ . It is then enough to obtain a factor  $2^{O(\sqrt{\log n} \cdot \log \log n)}$ -approximation algorithm for each resulting instance  $(G, \tilde{\mathcal{M}})$  separately. From now on we fix one such hierarchical system  $\tilde{\mathcal{H}} = (\mathcal{Q}_1, \mathcal{Q}_2, \dots, \mathcal{Q}_\rho)$  of squares, together with the set  $\tilde{\mathcal{M}} \subseteq \mathcal{M}$  of demand pairs, containing all pairs  $(s, t)$  that belong to  $\tilde{\mathcal{H}}$ , and focus on designing an  $2^{O(\sqrt{\log n} \cdot \log \log n)}$ -approximation algorithm for instance  $(G, \tilde{\mathcal{M}})$ .

## 2.2 A Hierarchical Partition of the Top Grid Boundary

Recall that  $R^*$  denotes the first row of the grid. A hierarchical partition  $\mathcal{I}$  of  $R^*$  is a sequence  $\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_\rho$  of sets of sub-paths of  $R^*$ , such that for each  $1 \leq h \leq \rho$ , the paths in  $\mathcal{I}_h$  (that we refer to as *level- $h$  intervals*) partition the vertices of  $R^*$ . We also require that for all  $1 < h \leq \rho$ , every level- $h$  interval  $I \in \mathcal{I}_h$  is contained in a unique level- $(h-1)$  interval  $I' \in \mathcal{I}_{h-1}$ , that we refer to as the *parent-interval* of  $I$ . For every level  $1 \leq h \leq \rho$ , we define a collection  $\chi_h$  of colors, containing one color  $c_h(I)$  for each level- $h$  interval  $I \in \mathcal{I}_h$ . If  $I' \in \mathcal{I}_h$  is a parent-interval of  $I \in \mathcal{I}_{h+1}$ , then we say that color  $c_h(I')$  is a *parent-color* of  $c_{h+1}(I)$ .

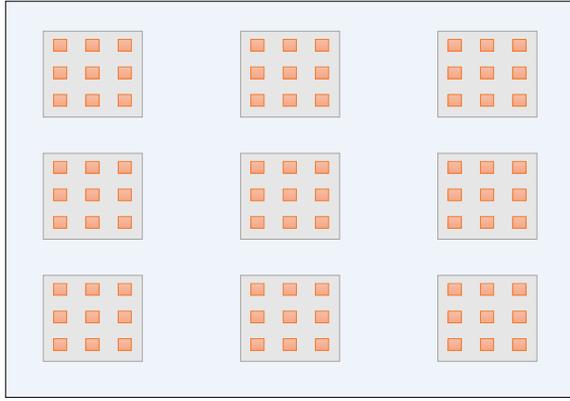


Figure 3: A schematic view of a hierarchical system of squares with 2 levels.

### 2.3 Coloring the Squares and Selecting Demand Pairs to Route

The third ingredient of our algorithm is an assignment  $f$  of colors to the squares, and a selection of a subset of the demand pairs to be routed. For every level  $1 \leq h \leq \rho$ , for every level- $h$  square  $Q \in \mathcal{Q}_h$ , we would like to assign a single level- $h$  color  $c_h(I) \in \chi_h$  to  $Q$ , denoting  $f(Q) = c_h(I)$ . Intuitively, if color  $c_h(I)$  is assigned to  $Q$ , then the only demand pairs  $(s, t)$  with  $t \in Q$  that we may route are those whose source vertex  $s$  lies on the level- $h$  interval  $I$ . We require that the coloring is consistent across levels: that is, for all  $1 < h \leq \rho$ , if a level- $h$  square is assigned a level- $h$  color  $c_h$ , and its parent-square is assigned a level- $(h-1)$  color  $c_{h-1}$ , then  $c_{h-1}$  must be a parent-color of  $c_h$ . We call such a coloring  $f$  a *valid coloring* of  $\tilde{\mathcal{H}}$  with respect to  $\mathcal{I}$ .

Finally, we would like to select a subset  $\mathcal{M}' \subseteq \tilde{\mathcal{M}}$  of the demand pairs to route. Consider some demand pair  $(s, t)$  and some level  $1 \leq h \leq \rho$ . Let  $I_h$  be the level- $h$  interval to which  $s$  belongs. Then we say that  $s$  has the level- $h$  color  $c_h(I_h)$ . Therefore, for each level  $1 \leq h \leq \rho$ , vertex  $s$  is assigned the unique level- $h$  color  $c_h(I_h)$ , and for  $1 \leq h < \rho$ ,  $c_h(I_h)$  is the parent-color of  $c_{h+1}(I_{h+1})$ . Let  $Q_\rho \in \mathcal{Q}_\rho$  be the level- $\rho$  square to which  $t$  belongs. We may only add  $(s, t)$  to  $\mathcal{M}'$  if the level- $\rho$  color of  $Q_\rho$  is  $c_\rho(I_\rho)$  (that is, it is the same as the level- $\rho$  color of  $s$ ). Notice that in particular, this means that for every level  $1 \leq h \leq \rho$ , if  $Q_h$  is the level- $h$  square containing  $t$ , and it is assigned the color  $c_h(I_h)$ , then  $s$  is assigned the same level- $h$  color, and so  $s \in I_h$ . Finally, we require that for all  $1 \leq h \leq \rho$ , for every level- $h$  color  $c_h$ , the total number of all demand pairs  $(s, t) \in \mathcal{M}'$ , such that the level- $h$  color of  $s$  is  $c_h$ , is no more than  $d_{h+1}/16$  (if  $h = \rho$ , then the number is no more than 1). If  $\mathcal{M}'$  has all these properties, then we say that it *respects the coloring*  $f$ . We say that  $(\mathcal{I}, f, \mathcal{M}')$  is a *good ensemble* iff  $\mathcal{I}$  is a hierarchical partition of  $R^*$  into intervals;  $f$  is a valid coloring of the squares in  $\tilde{\mathcal{H}}$  with respect to  $\mathcal{I}$ ; and  $\mathcal{M}' \subseteq \tilde{\mathcal{M}}$  is a subset of the demand pairs that respects the coloring  $f$ . The *size* of the ensemble is  $|\mathcal{M}'|$ .

### 2.4 The Routing

We show that, if we are given a good ensemble  $(\mathcal{I}, f, \mathcal{M}')$ , then we can route all demand pairs in  $\mathcal{M}'$ . The routing itself follows the high-level idea outlined above. We gradually construct a collection  $\mathcal{P}$  of node-disjoint paths routing the demand pairs in  $\mathcal{M}'$ . At the highest level, all these paths depart from their sources and then visit the level-1 squares one-by-one, in a snake-like fashion, as in Figure 2(a). Consider now some level-1 square  $Q$ , and assume that its level-1 color is  $c_1(I)$ , where  $I \in \mathcal{I}_1$  is some level-1 interval of  $R^*$ . Then only the paths  $P \in \mathcal{P}$  that originate at the vertices of  $I$  will enter the

square  $Q$ ; the remaining paths will exploit the spacing between the level-1 squares in order to bypass it; the spacing between the level-1 squares is sufficient to allow this. Once we have defined this global routing, we need to specify how the routing is carried out inside each square. We employ the same procedure recursively. Consider some level-1 square  $Q$ , and let  $\mathcal{P}' \subseteq \mathcal{P}$  be the set of all paths that visit  $Q$ . Assume further that the level-1 color of  $Q$  is  $c_1(I)$ . Since we are only allowed to have at most  $d_2/16$  demand pairs in  $\mathcal{M}'$  whose level-1 color is  $c_1(I)$ ,  $|\mathcal{P}'| \leq d_2/16$ . Let  $\mathcal{Q}' \subseteq \mathcal{Q}_2$  be the set of all level-2 squares contained in  $Q$ . The paths in  $\mathcal{P}'$  will visit the squares of  $\mathcal{Q}'$  one-by-one in a snake-like fashion (but this part of the routing is performed inside  $Q$ ). As before, for every level-2 square  $Q' \subseteq Q$ , if the level-2 color of  $Q'$  is  $c_2(I')$ , then only those paths of  $\mathcal{P}'$  that originate at the vertices of  $I'$  will enter  $Q'$ ; the remaining paths will use the spacing between the level-2 squares to bypass  $Q'$ . Since  $|\mathcal{P}'| \leq d_2/16$ , and all level-2 squares are at distance at least  $d_2$  from each other, there is a sufficient spacing to allow this routing. We continue this process recursively, until, at the last level of the recursion, we route at most one path per color, to its destination vertex.

In order to complete the proof of the theorem, we need to show that there exists a good ensemble  $(\mathcal{J}, f, \mathcal{M}')$  of size  $|\mathcal{M}'| \geq |\text{OPT}|/2^{O(\sqrt{\log n} \cdot \log \log n)}$ , and that we can find such an ensemble efficiently.

## 2.5 The Existence of the Ensemble

The key notion that we use in order to show that a large good ensemble  $(\mathcal{J}, f, \mathcal{M}')$  exists is that of a *shadow property*. Suppose  $Q$  is some  $(d \times d)$  sub-grid of  $G$ , and let  $\hat{\mathcal{M}} \subseteq \mathcal{M}$  be some subset of the demand pairs. Among all demand pairs  $(s, t) \in \hat{\mathcal{M}}$  with  $t \in Q$ , let  $(s_1, t_1)$  be the one with  $s_1$  appearing earliest on the first row  $R^*$  of  $G$ , and let  $(s_2, t_2)$  be the one with  $s_2$  appearing latest on  $R^*$ . The *shadow of  $Q$  with respect to  $\hat{\mathcal{M}}$*  is the sub-path of  $R^*$  between  $s_1$  and  $s_2$ . Let  $N_{\hat{\mathcal{M}}}(Q)$  be the number of all demand pairs  $(s, t) \in \hat{\mathcal{M}}$  with  $s$  lying in the shadow of  $Q$  (that is,  $s$  lies between  $s_1$  and  $s_2$  on  $R^*$ ). We say that  $\hat{\mathcal{M}}$  has the *shadow property with respect to  $Q$*  iff  $N_{\hat{\mathcal{M}}}(Q) \leq d$ . We say that  $\hat{\mathcal{M}}$  has the *shadow property with respect to the hierarchical system  $\tilde{\mathcal{H}} = (\mathcal{Q}_1, \dots, \mathcal{Q}_\rho)$  of squares*, iff  $\hat{\mathcal{M}}$  has the shadow property with respect to every square in  $\bigcup_{h=1}^\rho \mathcal{Q}_h$ . Let  $\mathcal{P}^*$  be the optimal solution to the instance  $(G, \tilde{\mathcal{M}})$  of NDP, where  $\tilde{\mathcal{M}}$  only includes the demand pairs that belong to  $\tilde{\mathcal{H}}$ . Let  $\mathcal{M}^* \subseteq \tilde{\mathcal{M}}$  be the set of the demand pairs routed by  $\mathcal{P}^*$ . For every demand pair  $(s, t) \in \mathcal{M}^*$ , let  $P(s, t) \in \mathcal{P}^*$  be the path routing this demand pair. Intuitively, it feels like  $\mathcal{M}^*$  should have the shadow property. Indeed, let  $Q \in \bigcup_{h=1}^\rho \mathcal{Q}_h$  be some square of size  $(d_h \times d_h)$ , and let  $(s_1, t_1), (s_2, t_2) \in \mathcal{M}^*$  be defined for  $Q$  as before, so that the shadow of  $Q$  with respect to  $\mathcal{M}^*$  is the sub-path of  $R^*$  between  $s_1$  and  $s_2$ . Let  $P$  be any path of length at most  $2d_h$  connecting  $t_1$  to  $t_2$  in  $Q$ , and let  $\gamma$  be the closed curve consisting of the union of  $P(s_1, t_1)$ ,  $P$ ,  $P(s_2, t_2)$ , and the shadow of  $Q$ . Consider the disc  $D$  whose boundary is  $\gamma$ . The intuition is that, if  $(s, t) \in \mathcal{M}^*$  is a demand pair whose source lies in the shadow of  $Q$ , and destination lies outside of  $D$ , then  $P(s, t)$  must cross the path  $P$ , as it needs to escape the disc  $D$ . Since path  $P$  is relatively short, only a small number of such demand pairs may exist. The main difficulty with this argument is that we may have a large number of demand pairs  $(s, t)$ , whose source lies in the shadow of  $Q$ , and the destination lies in the disc  $D$ . Intuitively, this can only happen if  $P(s_1, t_1)$  and  $P(s_2, t_2)$  “capture” a large area of the grid. We show that, in a sense, this cannot happen too often, and that there is a subset  $\mathcal{M}^{**} \subseteq \mathcal{M}^*$  of at least  $|\mathcal{M}^*|/2^{O(\sqrt{\log n} \cdot \log \log n)}$  demand pairs, such that  $\mathcal{M}^{**}$  has the shadow property with respect to  $\tilde{\mathcal{H}}$ .

Finally, we show that there exists a good ensemble  $(\mathcal{J}, f, \mathcal{M}')$  with  $|\mathcal{M}'| \geq |\mathcal{M}^{**}|/2^{O(\sqrt{\log n} \cdot \log \log n)}$ . We construct the ensemble over the course of  $\rho$  iterations, starting with  $\mathcal{M}' = \mathcal{M}^{**}$ . In the  $h$ th iteration we construct the set  $\mathcal{I}_h$  of the level- $h$  intervals of  $R^*$ , assign level- $h$  colors to all level- $h$  squares of  $\tilde{\mathcal{H}}$ , and discard some demand pairs from  $\mathcal{M}'$ . Recall that  $\eta = 2^{\lceil \sqrt{\log n} \rceil}$ . In the first iteration, we let  $\mathcal{I}_1$  be a partition of the row  $R^*$  into intervals, each of which contains roughly  $\frac{d_1}{16\eta} = \frac{d_2}{16} \leq \frac{|\mathcal{M}^*|}{\eta}$  vertices

of  $S(\mathcal{M}')$ . Assume that these intervals are  $I_1, \dots, I_r$ , and that they appear in this left-to-right order on  $R^*$ . We call all intervals  $I_j$  where  $j$  is odd *interesting intervals*, and the remaining intervals  $I_j$  *uninteresting intervals*. We discard from  $\mathcal{M}'$  all demand pairs  $(s, t)$ , where  $s$  lies on an uninteresting interval. Consider now some level-1 square  $Q$ , and let  $\mathcal{M}(Q) \subseteq \mathcal{M}'$  be the set of all demand pairs whose destinations lie in  $Q$ . Since the original set  $\mathcal{M}^{**}$  of demand pairs had the shadow property with respect to  $Q$ , it is easy to verify that all source vertices of the demand pairs in  $\mathcal{M}(Q)$  must belong to a single interesting interval of  $\mathcal{I}_1$ . Let  $I$  be that interval. Then we color the square  $Q$  with the level-1 color  $c_1(I)$  corresponding to the interval  $I$ . This completes the first iteration. Notice that for each level-1 color  $c_1(I)$ , at most  $d_2/16$  demand pairs  $(s, t) \in \mathcal{M}'$  have  $s \in I$ . In the following iteration, we similarly partition every interesting level-1 interval into level-2 intervals that contain roughly  $d_3/16 \leq |\mathcal{M}^*|/\eta^2$  source vertices of  $\mathcal{M}'$  each, and then define a coloring of all level-2 squares similarly, while suitably updating the set  $\mathcal{M}'$  of the demand pairs. We continue this process for  $\rho$  iterations, eventually obtaining a good ensemble  $(\mathfrak{J}, f, \mathcal{M}')$ . Since we only discard a constant fraction of the demand pairs of  $\mathcal{M}'$  in every iteration, at the end,  $|\mathcal{M}'| \geq |\mathcal{M}^{**}|/2^{O(\rho)} = |\mathcal{M}^{**}|/2^{O(\sqrt{\log n})} \geq |\mathcal{M}^*|/2^{O(\sqrt{\log n} \cdot \log \log n)}$ .

## 2.6 Finding the Good Ensemble

In our final step, our goal is to find a good ensemble  $(\mathfrak{J}, f, \mathcal{M}')$  maximizing  $|\mathcal{M}'|$ . We show an efficient randomized  $2^{O(\sqrt{\log n} \cdot \log \log n)}$ -approximation algorithm for this problem. First, we show that, at the cost of losing a small factor in the approximation ratio, we can restrict our attention to a small collection  $\mathfrak{J}_1, \mathfrak{J}_2, \dots, \mathfrak{J}_z$  of hierarchical partitions of  $R^*$  into intervals, and that it is enough to obtain a  $2^{O(\sqrt{\log n} \cdot \log \log n)}$ -approximate solution for the problem of finding the largest ensemble  $(\mathfrak{J}_j, f, \mathcal{M}')$  for each such partition  $\mathfrak{J}_j$  separately.

We then fix one such hierarchical partition  $\mathfrak{J}_j$ , and design an LP-relaxation for the problem of computing a coloring  $f$  of  $\tilde{\mathcal{H}}$  and a collection  $\mathcal{M}'$  of demand pairs, such that  $(\mathfrak{J}_j, f, \mathcal{M}')$  is a good ensemble, while maximizing  $|\mathcal{M}'|$ . Finally, we design an efficient randomized LP-rounding  $2^{O(\sqrt{\log n} \cdot \log \log n)}$ -approximation algorithm for the problem.

## 2.7 Completing the Proof of Theorem 1.1

So far we have assumed that all source vertices lie on the top boundary of the grid, and all destination vertices are at a distance at least  $\Omega(\text{OPT})$  from the grid boundary. Let  $\mathcal{A}$  be the randomized efficient  $2^{O(\sqrt{\log n} \cdot \log \log n)}$ -approximation algorithm for this special case. We now extend it to the general **Restricted NDP-Grid** problem. For every destination vertex  $t$ , we identify the closest vertex  $\tilde{t}$  that lies on the grid boundary. Using standard grouping techniques, and at the cost of losing an additional  $O(\log n)$  factor in the approximation ratio, we can assume that all source vertices lie on the top boundary of the grid, all vertices in  $\{\tilde{t} \mid t \in T(\mathcal{M})\}$  lie on a single boundary edge of the grid (assume for simplicity that it is the bottom boundary), and that there is some integer  $d$ , such that for every destination vertex  $t \in T(\mathcal{M})$ ,  $d \leq d(t, \tilde{t}) < 2d$ . We show that we can define a collection  $\mathcal{Z} = \{Z_1, \dots, Z_r\}$  of disjoint square sub-grids of  $G$ , and a collection  $\mathcal{I} = \{I_1, \dots, I_r\}$  of disjoint sub-intervals of  $R^*$ , such that the bottom boundary of each sub-grid  $Z_i$  is contained in the bottom boundary of  $G$ , the top boundary of  $Z_i$  is within distance at least  $\text{OPT}$  from  $R^*$ ,  $Z_1, \dots, Z_r$  appear in this left-to-right order in  $G$ , and  $I_1, \dots, I_r$  appear in this left-to-right order on  $R^*$ . For each  $1 \leq j \leq r$ , we let  $\mathcal{M}_j$  denote the set of all demand pairs with the sources lying on  $I_j$  and the destinations lying in  $Z_j$ . For each  $1 \leq j \leq r$ , we then obtain a new instance  $(G, \mathcal{M}_j)$  of the NDP problem. We show that there exist a collection  $\mathcal{Z}$  of squares and a collection  $\mathcal{I}$  of intervals, such that the value of the optimal solution to each instance  $(G, \mathcal{M}_j)$ , that we denote by  $\text{OPT}_j$ , is at most  $d$ , while  $\sum_{j=1}^r \text{OPT}_j \geq \text{OPT}/2^{O(\sqrt{\log n} \cdot \log \log n)}$ . Moreover, it is not hard to show that, if we can compute, for each  $1 \leq j \leq r$ , a routing of some subset  $\mathcal{M}'_j \subseteq \mathcal{M}_j$

of demand pairs in  $G$ , then we can also route all demand pairs in  $\bigcup_{j=1}^r \mathcal{M}'_j$  simultaneously in  $G$ .

There are two problems with this approach. First, we do not know the set  $\mathcal{Z}$  of sub-grids of  $G$  and the set  $\mathcal{I}$  of intervals of  $R^*$ . Second, it is not clear how to solve each resulting problem  $(G, \mathcal{M}_j)$ . To address the latter problem, we define a simple mapping of all source vertices in  $S(\mathcal{M}_j)$  to the top boundary of grid  $Z_j$ , obtaining an instance of **Restricted NDP-Grid**, where all source vertices lie on the top boundary of the grid  $Z_j$ , and all destination vertices lie at a distance at least  $\text{OPT}_j \leq d$  from its boundary. We can then use algorithm  $\mathcal{A}$  in order to solve this problem efficiently. It is easy to see that, if we can route some subset  $\mathcal{M}'_j$  of the demand pairs via node-disjoint paths in  $Z_j$ , then we can extend this routing to the corresponding set of original demand pairs, whose sources lie on  $R^*$ .

Finally, we employ dynamic programming in order to find the set  $\mathcal{Z}$  of sub-grids of  $G$  and the set  $\mathcal{I}$  of intervals of  $I$ . For each such potential sub-grid  $Z$  and interval  $I$ , we use algorithm  $\mathcal{A}$  in order to find a routing of a large set of demand pairs of the corresponding instance defined inside  $Z$ , and then exploit the resulting solution values for each such pair  $(I, Z)$  in a simple dynamic program, that allows us to compute the set  $\mathcal{Z}$  of sub-grids of  $G$ , the set  $\mathcal{I}$  of intervals of  $I$ , and the final routing.

### 3 Preliminaries

All logarithms in this paper are to the base of 2.

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  $V(G^{h, \ell}) = \{v(i, j) \mid 1 \leq i \leq h, 1 \leq j \leq \ell\}$ , 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  $E^V = \{(v_{i,j}, v_{i+1,j}) \mid 1 \leq i < h, 1 \leq j \leq \ell\}$  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 two vertices  $u = v(i, j)$  and  $u' = v(i', j')$  of the grid, the shortest-path distance between them in  $G$  is denoted by  $d(u, u')$ . Given two vertex subsets  $X, Y \subseteq V(G^{h, \ell})$ , the distance between them is  $d(X, Y) = \min_{u \in X, u' \in Y} \{d(u, u')\}$ . Given a vertex  $v = v(i, j)$  of the grid, we denote by  $\text{row}(v)$  and  $\text{col}(v)$  the row and the column, respectively, that contain  $v$ . The *boundary of the grid* is  $\Gamma(G^{h, \ell}) = R_1 \cup R_h \cup W_1 \cup W_\ell$ . We sometimes refer to  $R_1$  and  $R_h$  as the top and the bottom boundaries of the grid respectively, and to  $W_1$  and  $W_\ell$  as the left and the right boundaries of the grid. We say that  $G^{h, \ell}$  is a *square grid* iff  $h = \ell$ .

Given a set  $\mathcal{R}$  of consecutive rows of a grid  $G = G^{\ell, h}$  and a set  $\mathcal{W}$  of consecutive columns of  $G$ , we let  $\Upsilon_G(\mathcal{R}, \mathcal{W})$  be the subgraph of  $G$  induced by the set  $\{v(j, j') \mid R_j \in \mathcal{R}, W_{j'} \in \mathcal{W}\}$  of vertices. We say that  $\Upsilon = \Upsilon_G(\mathcal{R}, \mathcal{W})$  is the *sub-grid of  $G$  spanned by the set  $\mathcal{R}$  of rows and the set  $\mathcal{W}$  of columns*. A sub-graph  $G' \subseteq G$  is called a *sub-grid of  $G$*  iff there is a set  $\mathcal{R}$  of consecutive rows and a set  $\mathcal{W}$  of consecutive columns of  $G$ , such that  $G' = \Upsilon_G(\mathcal{R}, \mathcal{W})$ . If additionally  $|\mathcal{R}| = |\mathcal{W}| = d$  for some integer  $d$ , then we say that  $G'$  is a *square of size  $(d \times d)$* .

In the NDP-Grid problem, the input is a grid  $G = G^{\ell, \ell}$  and a set  $\mathcal{M} = \{(s_1, t_1), \dots, (s_k, t_k)\}$  of pairs of its vertices, called demand pairs. We refer to the vertices in set  $S(\mathcal{M}) = \{s_1, \dots, s_k\}$  as source vertices, to the vertices in set  $T(\mathcal{M}) = \{t_1, \dots, t_k\}$  as destination vertices, and to the vertices of  $S(\mathcal{M}) \cup T(\mathcal{M})$  as terminals. The solution to the problem is a set  $\mathcal{P}$  of node-disjoint paths in  $G$ , where each path  $P \in \mathcal{P}$  connects some demand pair in  $\mathcal{M}$ . The goal is to maximize the number of the demand pairs routed, that is,  $|\mathcal{P}|$ . In this paper we consider a special case of NDP-Grid, that we call **Restricted NDP-Grid**, where all source vertices appear on the boundary of the grid. We denote by  $n = \ell^2$  the

number of vertices in the input graph  $G$ .

Given a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of the demand pairs, we denote by  $S(\mathcal{M}')$  and  $T(\mathcal{M}')$  the sets of all source and all destination vertices participating in the pairs in  $\mathcal{M}'$ , respectively.

## 4 The Algorithm

Throughout, we assume that we know the value  $\text{OPT}$  of the optimal solution; in order to do so, we simply guess the value  $\text{OPT}$  (that is, we go over all such possible values), and run our approximation algorithm for each such guessed value. It is enough to show that the algorithm returns a factor- $2^{O(\sqrt{\log n \cdot \log \log n})}$ -approximate solution whenever the value  $\text{OPT}$  has been guessed correctly. We use a parameter  $\eta = 2^{\lceil \sqrt{\log n} \rceil}$ . We first further restrict the problem and consider a special case where the destination vertices appear far from the grid boundary. We show an efficient randomized algorithm for approximately solving this special case of the problem, by proving the following theorem, which is one of our main technical contributions.

**Theorem 4.1** *There is an efficient randomized algorithm  $\mathcal{A}$ , that, given an instance  $(G, \mathcal{M})$  of Restricted NDP-Grid and an integer  $\text{OPT} > 0$ , such that the value of the optimal solution to  $(G, \mathcal{M})$  is at least  $\text{OPT}$ , and every destination vertex  $t \in T(\mathcal{M})$  lies at a distance at least  $\text{OPT}/\eta$  from  $\Gamma(G)$ , returns a solution that routes at least  $\text{OPT}/2^{O(\sqrt{\log n \cdot \log \log n})}$  demand pairs, with high probability.*

In Section 9 we complete the proof of Theorem 1.1 using algorithm  $\mathcal{A}$  from Theorem 4.1 as a subroutine, by reducing the given instance of Restricted NDP-Grid to a number of instances of the form required by Theorem 4.1, and then applying algorithm  $\mathcal{A}$  to each of them. In this section, we focus on the proof of Theorem 4.1.

Recall that our input grid  $G$  has dimensions  $(\ell \times \ell)$ . Throughout, for an integer  $r > 0$ , we denote  $[r] = \{1, \dots, r\}$ . Notice that by losing a factor 4 in the approximation guarantee, we can assume w.l.o.g. that all sources lie on the top row of  $G$ , and that  $\text{OPT} \leq \ell$ .

**Parameters.** The following parameters are used throughout the algorithm. Let  $c^* \geq 2$  be a large constant, whose value will be set later. Recall that  $\eta = 2^{\lceil \sqrt{\log n} \rceil}$ . Let  $\rho$  be the largest integer, for which  $\eta^{\rho+2} \leq \text{OPT}/2^{c^* \sqrt{\log n \log \log n}}$ . Intuitively, we will route no more than  $\eta^{\rho+2}$  demand pairs, and it is helpful that this number is significantly smaller than  $\text{OPT}$ . The parameter  $\rho$  will be used as the number of recursive levels in the hierarchical system of squares, and in the algorithm overall. Clearly,  $\rho \leq \sqrt{\log n}$ , and  $\text{OPT}/(\eta \cdot 2^{c^* \sqrt{\log n \log \log n}}) < \eta^{\rho+2} \leq \text{OPT}/2^{c^* \sqrt{\log n \log \log n}}$ . In particular,  $\eta^{\rho+2} < \ell$ , where  $\ell$  is the size of the side of the grid  $G$ .

It would be convenient for us if we could assume that  $\ell$  is an integral multiple of  $\eta^{\rho+2}$ . Since this is not true in general, below we define a sub-grid  $\tilde{G}$  of  $G$  of dimensions  $(\ell' \times \ell')$ , where  $\ell'$  is a large enough integral multiple of  $\eta^{\rho+2}$ . We show that  $\tilde{G}$  can be defined so that, in a sense, we can restrict our attention to the sub-problem induced by  $\tilde{G}$ , without paying much in the approximation factor.

To this end, we let  $\ell'$  be the largest integral multiple of  $\eta^{\rho+2}$  with  $\ell' < \ell$ , so  $\ell - \ell' \leq \eta^{\rho+2} \leq \text{OPT}/\eta$ . Let  $G'$  be the sub-grid of  $G$  spanned by its first  $\ell'$  columns and all its rows, and let  $G''$  be the sub-grid of  $G$  spanned by its last  $\ell'$  columns and all its rows. Finally, let  $\mathcal{M}'$  and  $\mathcal{M}''$  be the subsets of the demand pairs contained in  $G'$  and  $G''$ , respectively. Notice that, since we have assumed that all destination vertices lie at a distance at least  $\text{OPT}/\eta$  from  $\Gamma(G)$ , all vertices of  $T(\mathcal{M})$  are contained in both  $G'$  and  $G''$ . Consider two instances of the NDP problem, both of which are defined on the original graph  $G$ ; the first one uses the set  $\mathcal{M}'$  of the demand pairs, and the second one uses the set  $\mathcal{M}''$  of demand

pairs. Clearly, one of these instances has a solution of value at least  $\text{OPT}/2$ , as for each demand pair  $(s, t) \in \mathcal{M}$ , both  $s$  and  $t$  must belong to either  $\mathcal{M}'$  or to  $\mathcal{M}''$  (or both). We assume without loss of generality that the problem induced by the set  $\mathcal{M}'$  of demand pairs has solution of value at least  $\text{OPT}/2$ . In particular, we assume that all source vertices lie on the first row of  $G'$ , that we denote by  $R^*$  from now on. Notice however that we are only guaranteed that a routing of at least  $\text{OPT}/2$  demand pairs of  $\mathcal{M}'$  exists in the original graph  $G$ . For convenience, abusing the notation, from now on we will denote  $\mathcal{M}'$  by  $\mathcal{M}$  and  $\text{OPT}/2$  by  $\text{OPT}$ .

For each  $1 \leq h \leq \rho$ , we let  $d_h = \eta^{\rho-h+3}$ , so that  $d_1 = \eta^{\rho+2}$  and  $\text{OPT}/(\eta \cdot 2^{c^* \sqrt{\log n} \log \log n}) < d_1 \leq \text{OPT}/2^{c^* \sqrt{\log n} \log \log n}$ ;  $d_\rho = \eta^3$ ; and  $d_h = d_{h-1}/\eta$  for all  $h > 1$ . Throughout, we assume that  $n$  is large enough, so that, for example,  $\log n > 2^{27}$ .

One of the main concepts that we use is that of a hierarchical decomposition of the grid into squares, that we define next. Recall that  $G'$  is the sub-grid of  $G$  spanned by its first  $\ell'$  columns and all its rows. We let  $\tilde{G} \subseteq G'$  be the  $(\ell' \times \ell')$ -sub-grid of  $G'$ , that is spanned by all its columns and its  $\ell'$  bottommost rows. Notice that since  $\ell' < \ell$ , the top row  $R^*$  of  $G'$ , where the source vertices reside, is disjoint from  $\tilde{G}$ .

## 4.1 Hierarchical Systems of Squares

A subset  $I \subseteq [\ell']$  of consecutive integers is called an *interval*. We say that two intervals  $I, I'$  are *disjoint* iff  $I \cap I' = \emptyset$ , and we say that they are *d-separated* iff for every pair of integers  $i \in I, j \in I'$ ,  $|i - j| > d$ . A collection  $\mathcal{I}$  of intervals of  $[\ell']$  is called *d-canonical*, iff for each interval  $I \in \mathcal{I}$ ,  $|I| = d$ , and every pair of intervals in  $\mathcal{I}$  is *d-separated*.

Consider now the  $(\ell' \times \ell')$ -grid  $\tilde{G}$ . Given two intervals  $I, I'$  of  $[\ell']$ , we denote by  $Q(I, I')$  the sub-graph of  $\tilde{G}$  induced by all vertices in  $\{v(i, j) \mid i \in I, j \in I'\}$ . Given two sets  $\mathcal{I}, \mathcal{I}'$  of intervals of  $[\ell']$ , we let  $Q(\mathcal{I}, \mathcal{I}')$  be the corresponding set of sub-graphs of  $G'$ , that is,  $Q(\mathcal{I}, \mathcal{I}') = \{Q(I, I') \mid I \in \mathcal{I}, I' \in \mathcal{I}'\}$ .

**Definition.** Let  $\mathcal{Q}$  be any collection of sub-graphs of  $\tilde{G}$ , and let  $d$  be an integer. We say that  $\mathcal{Q}$  is a *d-canonical family of squares* iff there are two *d-canonical sets*  $\mathcal{I}, \mathcal{I}'$  of intervals, such that  $\mathcal{Q} = Q(\mathcal{I}, \mathcal{I}')$ . In particular, each sub-graph in  $\mathcal{Q}$  is a square of size  $(d \times d)$ , and for every pair  $Q, Q' \in \mathcal{Q}$  of distinct squares, for every pair  $v \in V(Q), v' \in V(Q')$  of vertices,  $d(v, v') \geq d$ .

We are now ready to define a hierarchical system of squares. We use the integral parameters  $\eta$  and  $\rho$  defined above, as well as the parameters  $\{d_h\}_{1 \leq h \leq \rho}$ .

**Definition.** A *hierarchical system of squares* is a sequence  $\tilde{\mathcal{H}} = (\mathcal{Q}_1, \mathcal{Q}_2, \dots, \mathcal{Q}_\rho)$  of sets of squares, such that:

- for all  $1 \leq h \leq \rho$ ,  $\mathcal{Q}_h$  is a  $d_h$ -canonical family of squares; and
- for all  $1 < h \leq \rho$ , for every square  $Q \in \mathcal{Q}_h$ , there is a square  $Q' \in \mathcal{Q}_{h-1}$ , such that  $Q \subseteq Q'$ .

Given a hierarchical system  $\tilde{\mathcal{H}} = (\mathcal{Q}_1, \mathcal{Q}_2, \dots, \mathcal{Q}_\rho)$ , of squares we denote by  $V(\tilde{\mathcal{H}})$  the set of all vertices lying in the squares of  $\mathcal{Q}_\rho$ , that is,  $V(\tilde{\mathcal{H}}) = \bigcup_{Q \in \mathcal{Q}_\rho} V(Q)$ . We say that the vertices of  $V(\tilde{\mathcal{H}})$  *belong to*  $\tilde{\mathcal{H}}$ . We need the following simple claim, whose proof appears in the Appendix.

**Claim 4.2** *There is an efficient algorithm, that constructs  $4^\rho$  hierarchical systems  $\tilde{\mathcal{H}}_1, \dots, \tilde{\mathcal{H}}_{4^\rho}$  of squares of  $\tilde{G}$ , such that every vertex of  $\tilde{G}$  belongs to exactly one system.*

Given a hierarchical system  $\tilde{\mathcal{H}} = (\mathcal{Q}_1, \dots, \mathcal{Q}_\rho)$  of squares and an integer  $1 \leq h \leq \rho$ , we call the squares in  $\mathcal{Q}_h$  *level- $h$  squares*. If  $h > 1$ , then for each level- $h$  square  $Q$ , we call the unique level- $(h-1)$  square  $Q'$  with  $Q \subseteq Q'$  the *parent-square* of  $Q$ , and we say that  $Q$  is a *child-square* of  $Q'$ .

From now on, we fix the collection  $\mathcal{F} = \{\tilde{\mathcal{H}}_1, \dots, \tilde{\mathcal{H}}_{4^\rho}\}$  of the hierarchical systems of squares of  $\tilde{G}$  that is given by Claim 4.2. Consider an optimal solution  $\mathcal{P}^*$  to the NDP instance  $(G, \mathcal{M})$ , and let  $\mathcal{M}^* \subseteq \mathcal{M}$  be the set of the demand pairs routed by this solution. For each  $1 \leq i \leq 4^\rho$ , we let  $\mathcal{M}_i^* \subseteq \mathcal{M}^*$  be the set of all demand pairs whose destinations belong to the system  $\tilde{\mathcal{H}}_i$ . Clearly, there is an index  $i^*$ , such that  $|\mathcal{M}_{i^*}^*| \geq |\mathcal{M}^*|/4^\rho$ . Even though we do not know the index  $i^*$ , we can run our approximation algorithm for each possible value of  $i^*$ . It is enough to show that our algorithm finds a  $2^{O(\sqrt{\log n \cdot \log \log n})}$ -approximate solution to instance  $(G, \mathcal{M})$  of NDP if the index  $i^*$  is guessed correctly. Therefore, we assume from now on that the index  $i^*$  is known to us. For convenience, we denote  $\tilde{\mathcal{H}}_{i^*}$  by  $\tilde{\mathcal{H}}$  from now on. We also denote  $\mathcal{M}_{i^*}^*$  by  $\mathcal{M}^{**}$ . Let  $\tilde{\mathcal{M}} \subseteq \mathcal{M}$  denote the set of all demand pairs whose destination vertices lie in the system  $\tilde{\mathcal{H}}$ . Then we are guaranteed that there is a solution to NDP instance  $(G, \tilde{\mathcal{M}})$  of value at least  $\text{OPT}/4^\rho \geq \text{OPT}/2^{2\sqrt{\log n}}$ .

## 4.2 The Shadow Property

**Definition.** Let  $Q \subseteq \tilde{G}$  be a square sub-grid of  $\tilde{G}$  of size  $(d \times d)$ , and let  $\hat{\mathcal{M}} \subseteq \mathcal{M}$  be any subset of the demand pairs, such that all vertices in  $S(\hat{\mathcal{M}})$  are distinct. The shadow of  $Q$  with respect to  $\hat{\mathcal{M}}$ ,  $J_{\hat{\mathcal{M}}}(Q)$ , is the shortest sub-path of row  $R^*$  of  $G'$  that contains the source vertices of all demand pairs  $(s, t) \in \hat{\mathcal{M}}$  with  $t \in Q$ . If no demand pairs in  $\hat{\mathcal{M}}$  have destinations in  $Q$ , then  $J_{\hat{\mathcal{M}}}(Q) = \emptyset$ . The length of the shadow,  $L_{\hat{\mathcal{M}}}(Q)$ , is the total number of vertices  $s \in S(\hat{\mathcal{M}})$ , that belong to  $J_{\hat{\mathcal{M}}}(Q)$ . Given a parameter  $\beta > 0$ , we say that  $Q$  has the  $\beta$ -shadow property with respect to  $\hat{\mathcal{M}}$ , iff  $L_{\hat{\mathcal{M}}}(Q) \leq \beta \cdot d$ .

We use the above definition in both the regimes where  $\beta \leq 1$  and  $\beta > 1$ . We need the following simple observation, whose proof is deferred to the Appendix.

**Observation 4.3** Let  $\mathcal{Q}$  be any collection of disjoint squares in  $\tilde{G}$ , let  $0 < \beta_1 < \beta_2$  be parameters, and let  $\hat{\mathcal{M}}$  be any set of demand pairs, such that all vertices in  $S(\hat{\mathcal{M}})$  are distinct, and every square in  $\mathcal{Q}$  has the  $\beta_2$ -shadow property with respect to  $\hat{\mathcal{M}}$ . Then there is an efficient algorithm to find a subset  $\hat{\mathcal{M}}' \subseteq \hat{\mathcal{M}}$  of at least  $\left\lfloor \frac{\beta_1 |\hat{\mathcal{M}}|}{4\beta_2} \right\rfloor$  demand pairs, such that every square in  $\mathcal{Q}$  has the  $\beta_1$ -shadow property with respect to  $\hat{\mathcal{M}}'$ .

The proof of the following theorem appears in Section 5.

**Theorem 4.4** Let  $\mathcal{Q}$  be a  $d$ -canonical set of squares, and let  $\hat{\mathcal{M}} \subseteq \mathcal{M}$  be a subset of demand pairs, such that:

- there is a set  $\mathcal{P}$  of node-disjoint paths in  $G$  routing all demand pairs in  $\hat{\mathcal{M}}$ ; and
- for each demand pair  $(s, t) \in \hat{\mathcal{M}}$ , there is a square  $Q \in \mathcal{Q}$  with  $t \in Q$ .

Then there is a subset  $\hat{\mathcal{M}}' \subseteq \hat{\mathcal{M}}$  of at least  $|\hat{\mathcal{M}}|/\log^6 n$  demand pairs, such that every square  $Q \in \mathcal{Q}$  has the 1-shadow property with respect to  $\hat{\mathcal{M}}'$ .

Recall that  $\mathcal{F}$  is the family of hierarchical systems of squares of  $\tilde{G}$  given by Claim 4.2,  $\tilde{\mathcal{H}} \in \mathcal{F}$  is the hierarchical system that we have guessed, and  $\tilde{\mathcal{M}} \subseteq \mathcal{M}$  is the set of all demand pairs whose

destinations belong to  $\tilde{\mathcal{H}}$ . Recall also that we have assumed that the value of the optimal solution for instance  $(G, \tilde{\mathcal{M}})$  of NDP-Grid is at least  $\text{OPT}/2^{2\sqrt{\log n}}$ , and we have denoted by  $\mathcal{M}^{**}$  the set of the demand pairs routed by such a solution.

**Corollary 4.5** *Denote  $\tilde{\mathcal{H}} = (\mathcal{Q}_1, \dots, \mathcal{Q}_\rho)$ . Then there is a set  $\tilde{\mathcal{M}}^* \subseteq \tilde{\mathcal{M}}$  of at least  $\frac{\text{OPT}}{2^{8\sqrt{\log n} \cdot \log \log n}}$  demand pairs, such that: (i) all demand pairs in  $\tilde{\mathcal{M}}^*$  can be simultaneously routed via node-disjoint paths in  $G$ ; and (ii) for each  $1 \leq h \leq \rho$ , all squares in set  $\mathcal{Q}_h$  have the  $1/\eta^2$ -shadow property with respect to  $\tilde{\mathcal{M}}^*$ .*

**Proof:** Notice that, since all demand pairs in  $\mathcal{M}^{**}$  are routable by node-disjoint paths, all vertices of  $S(\mathcal{M}^{**})$  are distinct.

We perform  $\rho$  iterations. The input to the  $j$ th iteration is a subset  $\mathcal{M}_{j-1} \subseteq \mathcal{M}^{**}$  of demand pairs, where the input to the first iteration is  $\mathcal{M}_0 = \mathcal{M}^{**}$ . In order to execute the  $j$ th iteration, we apply Theorem 4.4 to the set  $\mathcal{Q}_j$  of squares and the set  $\mathcal{M}_{j-1}$  of demand pairs, to obtain a subset  $\mathcal{M}_j \subseteq \mathcal{M}_{j-1}$  of demand pairs of cardinality at least  $|\mathcal{M}_{j-1}|/\log^6 n$ , such that all squares in  $\mathcal{Q}_j$  have the 1-shadow property with respect to  $\mathcal{M}_j$ . Consider the final set  $\mathcal{M}_\rho$  of demand pairs. Clearly,

$$|\mathcal{M}_\rho| \geq \frac{|\mathcal{M}_0|}{\log^{6\rho} n} \geq \frac{\text{OPT}}{2^{2\sqrt{\log n}} \cdot 2^{6\sqrt{\log n} \cdot \log \log n}} \geq \frac{\text{OPT}}{2^{7\sqrt{\log n} \cdot \log \log n}},$$

and for every  $1 \leq h \leq \rho$ , every square in  $\mathcal{Q}_h$  has the 1-shadow property with respect to  $\mathcal{M}_\rho$ . Applying Observation 4.3 to  $\mathcal{M}_\rho$  and the set  $\mathcal{Q} = \bigcup_{h=1}^\rho \mathcal{Q}_h$ , we obtain a subset  $\tilde{\mathcal{M}}^* \subseteq \mathcal{M}_\rho$  of at least  $\left\lfloor \frac{|\mathcal{M}_\rho|}{4\eta^2} \right\rfloor \geq \frac{|\text{OPT}|}{8\eta^2 \cdot 2^{7\sqrt{\log n} \cdot \log \log n}} \geq \frac{|\text{OPT}|}{2^{8\sqrt{\log n} \cdot \log \log n}}$  demand pairs, that has the required properties.  $\square$

### 4.3 Hierarchical Decomposition of Row $R^*$

Recall that  $R^*$  is the first row of the grid  $G'$ . We start with intuition. Recall that  $\tilde{\mathcal{M}}^*$  is the set of the demand pairs from Corollary 4.5. In general, we would like to define a hierarchical partition  $(\mathcal{J}_1, \mathcal{J}_2, \dots, \mathcal{J}_\rho)$  of  $R^*$ , that has the following property: for each level  $1 \leq h \leq \rho$ , every interval  $I \in \mathcal{J}_h$  contains either 0, or roughly  $d_h$  source vertices of the demand pairs in  $\tilde{\mathcal{M}}^*$ . Being able to find such a partition is crucial to our algorithm. If we knew the set  $\tilde{\mathcal{M}}^*$  of demand pairs, finding such a partition would have been trivial. But unfortunately, set  $\tilde{\mathcal{M}}^*$  depends on the optimal solution and is not known to us. We show instead that, at the cost of losing an additional  $2^{O(\sqrt{\log n} \log \log n)}$ -factor in the approximation ratio, we can define a small collection of such hierarchical partitions of  $R^*$ , one of which is guaranteed to have the above property. In each such hierarchical partition that we construct, for each level  $1 \leq h \leq \rho$ , all intervals in  $\mathcal{J}_h$  will have the same length, that we denote by  $\ell_h$ . The lengths  $\ell_h$  are not known to us a-priori, and differ from partition to partition, but we will be able to guess them from a small set of possibilities.

Suppose we are given a sequence  $L = (\ell_1, \ell_2, \dots, \ell_r)$  of integers, with  $\ell_1 > \ell_2 > \dots > \ell_r$ , where  $r > 0$  is some integer, and for each  $1 \leq h \leq r$ ,  $\ell_h$  is an integral power of  $\eta$ . A hierarchical  $L$ -decomposition of the first row  $R^*$  of the grid  $G'$  is defined as follows. First, we partition  $R^*$  into a collection  $\mathcal{J}_1$  of intervals, each of which contains exactly  $\ell_1$  consecutive vertices (recall that  $R^*$  contains  $\ell'$  vertices, and  $\ell'$  is an integral multiple of  $d_1 = \eta^{\rho+2}$ , which in turn is an integral power of  $\eta$ ). We call the intervals in  $\mathcal{J}_1$  *level-1 intervals*. Assume now that we are given, for some  $1 \leq h < r$ , a partition  $\mathcal{J}_h$  of  $R^*$  into level- $h$  intervals. We now define a partition  $\mathcal{J}_{h+1}$  of  $R^*$  into level- $(h+1)$  intervals as follows. Start with  $\mathcal{J}_{h+1} = \emptyset$ . For each level- $h$  interval  $I \in \mathcal{J}_h$ , partition  $I$  into consecutive intervals containing exactly  $\ell_{h+1}$  vertices each. Add all resulting intervals to  $\mathcal{J}_{h+1}$ . Note that for a fixed

sequence  $L = (\ell_1, \ell_2, \dots, \ell_r)$  of integers, the corresponding hierarchical  $L$ -decomposition  $(\mathcal{J}_1, \dots, \mathcal{J}_r)$  of  $R^*$  is unique.

Assume now that we are also given a collection  $\mathcal{M}' \subseteq \mathcal{M}$  of demand pairs. We say that  $\mathcal{M}'$  is *compatible* with the sequence  $L = (\ell_1, \dots, \ell_\rho)$  of integers, iff for each  $1 \leq h \leq \rho$ , for every interval  $I \in \mathcal{J}_h$ , either  $I \cap S(\mathcal{M}') = \emptyset$ , or:

$$\frac{d_h}{16\eta} \leq |I \cap S(\mathcal{M}')| \leq \frac{d_h}{4}.$$

The following theorem uses the hierarchical system  $\tilde{\mathcal{H}} \in \mathcal{F}$  of squares we have defined above and the set  $\tilde{\mathcal{M}}^* \subseteq \mathcal{M}$  of demand pairs from Corollary 4.5. Recall that all demand pairs in  $\tilde{\mathcal{M}}^*$  have their destinations in  $\tilde{\mathcal{H}}$ , and they can all be routed via node-disjoint pairs in  $G$ . Moreover, all squares that lie in the system  $\tilde{\mathcal{H}}$  have the  $1/\eta^2$ -shadow property with respect to  $\tilde{\mathcal{M}}^*$ .

**Theorem 4.6** *There is a sequence  $L = (\ell_1, \ell_2, \dots, \ell_\rho)$  of integers, with  $\ell_1 > \ell_2 > \dots > \ell_\rho$ , where each  $\ell_i$  is an integral power of  $\eta$ , and a subset  $\tilde{\mathcal{M}}^{**} \subseteq \tilde{\mathcal{M}}^*$  of demand pairs, such that:*

- $|\tilde{\mathcal{M}}^{**}| \geq \text{OPT}/2^{O(\sqrt{\log n} \log \log n)}$ ;
- the set  $\tilde{\mathcal{M}}^{**}$  of demand pairs is compatible with the sequence  $L$ .

**Proof:** Let  $\rho'$  be the largest integer such that  $\ell' \geq \eta^{\rho'}$ , and note that  $\rho' \leq \sqrt{\log n}$ . We next gradually modify the set  $\tilde{\mathcal{M}}^*$  of demand pairs, by starting with  $\mathcal{M}_{\rho'} = \tilde{\mathcal{M}}^*$ , and then gradually obtaining smaller and smaller subsets  $\mathcal{M}_{\rho'-1}, \mathcal{M}_{\rho'-2}, \dots, \mathcal{M}_0$  of  $\mathcal{M}_{\rho'}$ . Along the way, we will also define a sequence  $\gamma_{\rho'}, \dots, \gamma_1$  of integers that will be helpful for us later.

Consider first the sequence  $L' = (\ell'_1, \ell'_2, \dots, \ell'_{\rho'})$  of integers, where for  $1 \leq h \leq \rho'$ ,  $\ell'_h = \eta^{\rho'-h}$ . Let  $(\mathcal{J}'_1, \mathcal{J}'_2, \dots, \mathcal{J}'_{\rho'})$  be the corresponding hierarchical  $L'$ -decomposition of  $R^*$ . We start with  $\mathcal{M}_{\rho'} = \tilde{\mathcal{M}}^*$ , and then perform  $\rho'$  iterations. An input to iteration  $j$  is some subset  $\mathcal{M}_{\rho'-j+1}$  of  $\tilde{\mathcal{M}}^*$ . In order to execute iteration  $j$ , we consider the set  $\mathcal{J}_{\rho'-j+1}$  of all level- $(\rho' - j + 1)$  intervals of  $R^*$  in the  $L'$ -hierarchical decomposition of  $R^*$ . We partition these intervals into  $\lceil \log n \rceil$  groups: an interval  $I \in \mathcal{J}_{\rho'-j+1}$  belongs to group  $U_z$  iff:

$$2^z < |I \cap S(\mathcal{M}_{\rho'-j+1})| \leq 2^{z+1}.$$

We say that a demand pair  $(s, t)$  belongs to group  $U_z$  iff the interval  $I \in \mathcal{J}_{\rho'-j+1}$  containing  $s$  belongs to group  $U_z$ . Then there is an index  $z$ , such that  $U_z$  contains at least a  $(1/\lceil \log n \rceil)$ -fraction of the demand pairs of  $\mathcal{M}_{\rho'-j+1}$ . We then set  $\gamma_{\rho'-j+1} = 2^z$ , and we let  $\mathcal{M}_{\rho'-j}$  contain the set of all demand pairs of  $\mathcal{M}_{\rho'-j+1}$  that belong to  $U_z$ . Let  $\mathcal{M}_0$  be the set of demand pairs obtained after the last iteration of the algorithm. We then set  $\tilde{\mathcal{M}}^{**} = \mathcal{M}_0$ . It is easy to verify that:

$$|\tilde{\mathcal{M}}^{**}| \geq \frac{|\tilde{\mathcal{M}}^*|}{\lceil \log n \rceil^{\rho'}} \geq \frac{\text{OPT}}{2^{8\sqrt{\log n} \cdot \log \log n} \cdot \lceil \log n \rceil^{\sqrt{\log n}}} \geq \frac{\text{OPT}}{2^{10\sqrt{\log n} \cdot \log \log n}}.$$

Recall that we have defined integers  $\gamma_1, \gamma_2, \dots, \gamma_{\rho'}$ , and for each  $1 \leq h \leq \rho'$ , for every level- $h$  interval  $I_h \in \mathcal{J}'_h$ , either  $I_h$  contains no vertices of  $S(\tilde{\mathcal{M}}^{**})$ , or:

$$\gamma_h < |I \cap S(\tilde{\mathcal{M}}^{**})| \leq 2\gamma_h.$$

Since for each  $1 < h \leq \rho'$ , a level- $h$  interval is a sub-interval of a level- $(h-1)$  interval, it is easy to verify that  $\gamma_1 \geq \gamma_2 \geq \dots \geq \gamma_{\rho'}$ . Moreover, we claim that for each  $1 < h \leq \rho'$ ,  $\gamma_h \geq \gamma_{h-1}/(2\eta)$ . This is since a level- $(h-1)$  interval  $I$  is partitioned into exactly  $\eta$  level- $h$  intervals. If  $\gamma_h < \gamma_{h-1}/(2\eta)$ , then each such interval contributes fewer than  $\gamma_{h-1}/\eta$  source vertices to  $I$ , and so  $I$  contains fewer than  $\gamma_h$  vertices of  $S(\tilde{\mathcal{M}})$ , a contradiction. It is also easy to verify that  $\gamma_{\rho'} = 1$ , while, since  $|\mathcal{I}'_1| \leq \eta^2$ ,  $\gamma_1 \geq |\tilde{\mathcal{M}}^*|/(\eta^2 \lceil \log n \rceil) \geq \text{OPT}/(\lceil \log n \rceil \cdot 2^{9\sqrt{\log n} \cdot \log \log n})$ . We set the parameter  $c^*$ , that was used in the definition of the parameter  $\rho$ , so that  $\gamma_1 \geq d_1$  holds. For example,  $c^* \geq 11$  is sufficient.

We are now ready to define the sequence  $L = (\ell_1, \dots, \ell_\rho)$  of integers. For each  $1 \leq j \leq \rho$ , we let  $\ell_j$  be the largest integer  $\ell'_j \in L'$ , for which  $\gamma_{\ell'_j} \leq d_j/8$ . From the above discussion, such an integer exists for all  $1 \leq j \leq \rho$ , and we have that  $d_j/(16\eta) \leq \gamma_{\ell_j} \leq d_j/8$ . In particular for each interval  $I \in \mathcal{J}_h$ , if  $I \cap S(\tilde{\mathcal{M}}^{**}) \neq \emptyset$ , then  $d_j/(16\eta) \leq |I \cap S(\tilde{\mathcal{M}}^{**})| \leq d_j/4$ .

It is now easy to verify that  $\tilde{\mathcal{M}}^{**}$  is compatible with the sequence  $L$ , since the sets of intervals of the  $L$ -hierarchical partition  $(\mathcal{J}_1, \mathcal{J}_2, \dots, \mathcal{J}_\rho)$  all belong to the collection  $(\mathcal{J}'_1, \mathcal{J}'_2, \dots, \mathcal{J}'_{\rho'})$  corresponding to the  $L'$ -hierarchical partition, as  $L \subseteq L'$ .  $\square$

For our algorithm, we need to assume that we know the sequence  $L$  of integers given by Theorem 4.6. As sequence  $L$  consists of at most  $\sqrt{\log n}$  integers, each of which is an integral power of  $\eta$ , there are at most  $\sqrt{\log n}$  possible choices for each integer of  $L$ , and at most  $(\sqrt{\log n})^{\sqrt{\log n}}$  possible choices for the sequence  $L$ . Therefore, we can go over all such choices, and apply our algorithm to each of them separately. It is now sufficient to show that our algorithm succeeds in producing the right output when the sequence  $L$  is guessed correctly — that is, it has the properties guaranteed by Theorem 4.6. We assume from now on that the correct sequence  $L = (\ell_1, \dots, \ell_\rho)$  is given, and we denote by  $(\mathcal{J}_1, \dots, \mathcal{J}_\rho)$  the corresponding  $L$ -hierarchical decomposition of  $R^*$ , that can be computed efficiently given  $L$ . Intervals lying in set  $\mathcal{J}_h$  are called *level- $h$  intervals*. Recall that we have also assumed that we are given a lower bound  $\text{OPT}$  on the value of the optimal solution, and that we have guessed the hierarchical system  $\tilde{\mathcal{H}} \in \mathcal{F}$  of squares correctly. We denote  $\tilde{\mathcal{H}} = (\mathcal{Q}_1, \dots, \mathcal{Q}_\rho)$ . Recall also that  $\tilde{\mathcal{M}} \subseteq \mathcal{M}$  is the set of all demand pairs whose destinations lie in  $\tilde{\mathcal{H}}$ .

#### 4.4 Hierarchical System of Colors

For every level  $1 \leq h \leq \rho$ , for every level- $h$  interval  $I \in \mathcal{J}_h$ , we introduce a distinct level- $h$  color  $c_h(I)$ , and we color all vertices of  $I$  with this color. We let  $\chi_h$  be the set of all level- $h$  colors, so  $\chi_h = \{c_h(I) \mid I \in \mathcal{J}_h\}$ . Every vertex of  $R^*$  is now assigned  $\rho$  colors - one color for every level. For convenience, we define a set  $\chi_0$  of level-0 colors, that contains a unique color  $c_0$ . This naturally defines a hierarchical structure of the colors. Consider any level  $0 \leq h < \rho$ , and some level- $h$  interval  $I \in \mathcal{J}_h$ , with the corresponding color  $c_h(I)$ . For every level- $(h+1)$  interval  $I'$  contained in  $I$ , we say that the corresponding color  $c_{h+1}(I')$  is the child-color of  $c_h(I)$ . This also naturally defines a descendant-ancestor relation between colors: for  $1 \leq h \leq h' \leq \rho$ , we say that  $c_h(I)$  is an ancestor-color of  $c_{h'}(I')$  iff  $I' \subseteq I$ . In particular, each color is an ancestor-color of itself. If a color  $c$  is an ancestor-color of color  $c'$ , then we say that  $c'$  is a descendant-color of  $c$ . We let  $\tilde{\chi}(c) \subseteq \chi$  denote the set of all descendant-colors of  $c$ , and by  $\tilde{\chi}_h(c)$  the set of all descendants of  $c$  that belong to level  $h$ .

Given the hierarchical system of colors  $\mathcal{C} = (\chi_0, \chi_1, \dots, \chi_\rho)$  as above, a coloring of  $\tilde{\mathcal{H}}$  by  $\mathcal{C}$  is an assignment of colors to squares,  $f : \bigcup_{i=1}^{\rho} \mathcal{Q}_i \rightarrow \bigcup_{i=1}^{\rho} \chi_i$ , such that:

- For each  $1 \leq h \leq \rho$ , every level- $h$  square  $Q \in \mathcal{Q}_h$  is assigned a single level- $h$  color  $c \in \chi_h$ ; and
- For all  $1 < h \leq \rho$ , if  $Q \in \mathcal{Q}_h$  is a level- $h$  square, that is assigned a level- $h$  color  $c$ , and its parent-square  $Q' \in \mathcal{Q}_{h-1}$  is assigned a level- $(h-1)$  color  $c'$ , then  $c$  is a child-color of  $c'$ .

Let  $U$  denote the set of all vertices of  $\tilde{G}$  that belong to the hierarchical system  $\tilde{\mathcal{H}}$  of squares, that is,  $U = \bigcup_{Q \in \mathcal{Q}_\rho} V(Q)$ . Recall that  $R^* \cap U = \emptyset$ . Consider any valid assignment of colors, and let  $v \in U$  be some vertex. For  $1 \leq h \leq \rho$ , let  $Q_h$  be the level- $h$  square containing  $v$ , and let  $c_h$  be the color assigned to  $Q_h$ . Then we say that the level- $h$  color of  $v$  is  $c_h$ . We also say that the level-0 color of  $v$  is  $c_0$ . Therefore, every vertex  $v \in U(\mathcal{H}) \cup R^*$  is associated with a  $(\rho + 1)$ -tuple of colors  $(c_0(v), \dots, c_\rho(v))$ , where  $c_h(v)$  is its level- $h$  color (recall that we have already assigned colors to the vertices of  $R^*$ ). Moreover, for  $0 \leq h < \rho$ ,  $c_{h+1}(v)$  is a child-color of  $c_h(v)$ .

Assume now that we are given some coloring  $f$  of  $\tilde{\mathcal{H}}$  by  $\mathcal{C}$ , and some subset  $\tilde{\mathcal{M}}' \subseteq \tilde{\mathcal{M}}$  of demand pairs. We say that  $\tilde{\mathcal{M}}'$  is a *perfect set of demand pairs* for  $f$ , iff:

- All sources and all destinations of the pairs in  $\tilde{\mathcal{M}}'$  are distinct;
- For each demand pair  $(s, t) \in \tilde{\mathcal{M}}'$ , both  $s$  and  $t$  are assigned the same level- $\rho$  color (notice that this means that for every level  $1 \leq h \leq \rho$ , the level- $h$  color assigned to  $s$  and  $t$  is the same); and
- For every level  $1 \leq h \leq \rho$ , for every level- $h$  color  $c_h$ , the number of demand pairs  $(s, t) \in \tilde{\mathcal{M}}'$ , for which  $s$  and  $t$  are assigned the color  $c_h$  is at most  $d_h$ .

The rest of the argument consists of three parts. First, we prove that there exists a coloring  $f$  of  $\tilde{\mathcal{H}}$ , and a perfect set  $\tilde{\mathcal{M}}' \subseteq \tilde{\mathcal{M}}$  of demand pairs for  $f$ , so that  $|\tilde{\mathcal{M}}'|$  is quite large compared to  $\text{OPT}$ ; set  $\tilde{\mathcal{M}}'$  is obtained by appropriately modifying the set  $\tilde{\mathcal{M}}^{**}$  of demand pairs given by Theorem 4.6. Next, we show that, given any coloring  $f$  of  $\tilde{\mathcal{H}}$  by  $\mathcal{C}$  and any perfect set  $\tilde{\mathcal{M}}'$  of demand pairs for  $f$ , we can efficiently route a large fraction of the demand pairs in  $\tilde{\mathcal{M}}'$  in graph  $G$ . Finally, we show an efficient algorithm for (approximately) computing the largest set  $\tilde{\mathcal{M}}' \subseteq \tilde{\mathcal{M}}$  of demand pairs, together with a coloring  $f$  of  $\tilde{\mathcal{H}}$  by  $\mathcal{C}$ , such that  $\tilde{\mathcal{M}}'$  is perfect for  $f$ . The following three theorems summarize these three steps, and their proofs appear in the following sections.

**Theorem 4.7** *There is a coloring  $f$  of  $\tilde{\mathcal{H}}$  by  $\mathcal{C}$ , and a set  $\tilde{\mathcal{M}}' \subseteq \tilde{\mathcal{M}}$  of demand pairs, such that  $\tilde{\mathcal{M}}'$  is perfect for  $f$ , and  $|\tilde{\mathcal{M}}'| \geq \text{OPT}/2^{O(\sqrt{\log n} \cdot \log \log n)}$ .*

**Theorem 4.8** *There is an efficient algorithm, that, given a coloring  $f$  of  $\tilde{\mathcal{H}}$  by  $\mathcal{C}$ , and a set  $\tilde{\mathcal{M}}' \subseteq \tilde{\mathcal{M}}$  of demand pairs with  $|\tilde{\mathcal{M}}'| \leq d_1$ , such that  $\tilde{\mathcal{M}}'$  is perfect for  $f$ , routes at least  $|\tilde{\mathcal{M}}'|/2^{O(\sqrt{\log n})}$  demand pairs of  $\tilde{\mathcal{M}}'$  in  $G$ .*

**Theorem 4.9** *There is an efficient randomized algorithm, that, given:*

- a hierarchical system  $\tilde{\mathcal{H}} = (\mathcal{Q}_1, \dots, \mathcal{Q}_\rho) \in \mathcal{F}$  of squares, together with a subset  $\tilde{\mathcal{M}} \subseteq \mathcal{M}$  of demand pairs whose destinations belong to  $\tilde{\mathcal{H}}$ ; and
- a sequence  $L = (\ell_1, \dots, \ell_\rho)$  of integers, with  $\ell_1 \geq \ell_2 \geq \dots \geq \ell_\rho$ , such that each integer  $\ell_h$  is an integral power of  $\eta$ , together with the corresponding  $L$ -hierarchical decomposition of  $R^*$  and the corresponding hierarchical system  $\mathcal{C}$  of colors,

*computes a coloring  $f$  of  $\tilde{\mathcal{H}}$  by  $\mathcal{C}$  and a set  $\tilde{\mathcal{M}}' \subseteq \tilde{\mathcal{M}}$  of demand pairs that is perfect for  $f$ , such that with high probability,  $|\tilde{\mathcal{M}}'| \geq \text{OPT}'/2^{O(\sqrt{\log n})}$ , where  $\text{OPT}'$  is the cardinality of the largest set  $\tilde{\mathcal{M}}'' \subseteq \tilde{\mathcal{M}}$  of demand pairs, such that there is some coloring  $f'$  of  $\tilde{\mathcal{H}}$  by  $\mathcal{C}$ , for which  $\tilde{\mathcal{M}}''$  is a perfect set.*

It is now easy to complete the proof of Theorem 4.1. We assume that our algorithm is given a lower bound  $\text{OPT}$  on the value of the optimal solution. It starts by constructing the family  $\mathcal{F}$  of

hierarchical systems of squares using Claim 4.2. It then guesses one of the systems  $\tilde{\mathcal{H}} \in \mathcal{F}$ , and a sequence  $L = (\ell_1, \dots, \ell_\rho)$  of integers as described above. We can now efficiently compute the  $L$ -hierarchical decomposition  $(\mathcal{J}_1, \dots, \mathcal{J}_\rho)$  of row  $R^*$  of  $G'$ , and the corresponding hierarchical system  $\mathcal{C}$  of colors. We apply Theorem 4.9 in order to compute a coloring  $f$  of  $\tilde{\mathcal{H}}$  by  $\mathcal{C}$  and the corresponding set  $\tilde{\mathcal{M}}' \subseteq \tilde{\mathcal{M}}$  of demand pairs that is perfect for  $f$ . From Theorems 4.7 and 4.9, if  $\tilde{\mathcal{H}}$  and  $L$  were guessed correctly, then with high probability,  $|\tilde{\mathcal{M}}'| \geq \text{OPT}/2^{O(\sqrt{\log n \cdot \log \log n})}$ . If  $|\tilde{\mathcal{M}}'| > d_1$ , then we discard demand pairs from  $\tilde{\mathcal{M}}'$  until  $|\tilde{\mathcal{M}}'| = d_1$  holds. As  $d_1 = \eta^{\rho+2} \geq \text{OPT}/2^{O(\sqrt{\log n \cdot \log \log n})}$ , we still have that  $|\tilde{\mathcal{M}}'| \geq \text{OPT}/2^{O(\sqrt{\log n \cdot \log \log n})}$ . Finally, we use Theorem 4.8 in order to route at least  $|\tilde{\mathcal{M}}'|/2^{O(\sqrt{\log n})} \geq \text{OPT}/2^{O(\sqrt{\log n \cdot \log \log n})}$  demand pairs of  $\tilde{\mathcal{M}}'$  in  $G$ . We now turn to prove Theorem 4.4 and Theorems 4.7 — 4.9.

## 5 The Shadow Property

This section is dedicated to proving Theorem 4.4.

For every demand pair  $(s, t) \in \hat{\mathcal{M}}$ , we let  $P(s, t)$  denote the path routing this pair in  $\mathcal{P}$ . We assume that  $G$  is embedded into the plane in the natural manner, and we construct a number of discs in the plane. Whenever we refer to a disc, we mean a simple disc, whose boundary is a simple cycle. Let  $D^0$  denote the disc whose outer boundary is the boundary of the grid  $G$ .

**Definition.** We say that a disc  $D$  in the plane is canonical, if either (i)  $D = D^0$ , or (ii) we can partition the boundary  $\sigma(D)$  of  $D$  into four contiguous segments,  $\sigma_1(D), \sigma_2(D), \sigma_3(D), \sigma_4(D)$ , such that  $\sigma_1(D) \subseteq R^*$ ;  $\sigma_3(D)$  is contained in the boundary of some square  $Q \in \mathcal{Q}$ ;  $\sigma_2(D)$  is contained in some path  $P(s, t) \in \mathcal{P}$ , and  $\sigma_4(D)$  is contained in some path  $P(s', t') \in \mathcal{P}$  (where possibly  $(s, t) = (s', t')$ ).

**Definition.** We say that a collection  $\mathcal{D}$  of canonical discs is nested iff for every pair  $D, D' \in \mathcal{D}$ , either one of the discs is contained in the other, or they are disjoint. Moreover, if the boundaries of two distinct discs  $D, D' \in \mathcal{D}$  intersect, then  $\sigma(D) \cap \sigma(D') \subseteq R^*$  must hold.

Assume now that we are given a nested collection  $\mathcal{D}$  of canonical discs. Consider some disc  $D \in \mathcal{D}$ , and let  $\mathcal{D}' \subseteq \mathcal{D}$  be the set of all discs  $D'$  with  $D' \subseteq D$ . We define the region of the plane associated with  $D$  as  $R_{\mathcal{D}}(D) = D \setminus (\bigcup_{D' \in \mathcal{D}'} D')$ .

The idea is to compute a collection  $\mathcal{D} = \{D^0, D^1, \dots, D^z\}$  of nested canonical discs, and for every disc  $D^i$  a set  $\mathcal{M}_i \subseteq \hat{\mathcal{M}}$  of demand pairs, such that all demand pairs in  $\mathcal{M}_i$  are routed inside the region  $R_{\mathcal{D}}(D^i)$  by  $\mathcal{P}$ . Clearly, for  $i \neq i'$ ,  $\mathcal{M}_i \cap \mathcal{M}_{i'} = \emptyset$  will hold. We denote by  $\overline{\mathcal{M}} = \hat{\mathcal{M}} \setminus \bigcup_i \mathcal{M}_i$  the set of the demand pairs we discard, and we will ensure that  $|\overline{\mathcal{M}}|$  is small enough compared to  $|\hat{\mathcal{M}}|$ .

Given a set  $\overline{\mathcal{M}}$  of discarded demand pairs, we say that a square  $Q \in \mathcal{Q}$  is *active* iff at least one demand pair in  $\hat{\mathcal{M}} \setminus \overline{\mathcal{M}}$  has a destination vertex lying in  $Q$ . We say that the set  $\mathcal{D}$  of discs *respects* the active squares, iff for every active square  $Q \in \mathcal{Q}$ , there is a unique disc  $D^i \in \mathcal{D}$  with  $Q \subseteq R_{\mathcal{D}}(D^i)$ . We will ensure that the set  $\mathcal{D}$  of discs we construct respects all active squares.

**Heavy demand pairs.** Consider any demand pair  $(s, t) \in \hat{\mathcal{M}}$  and the path  $P(s, t)$  routing  $(s, t)$  in  $\mathcal{P}$ . We say that  $(s, t)$  *covers* the demand pair  $(s', t') \in \hat{\mathcal{M}}$  iff  $P(s, t)$  intersects the unique square  $Q \in \mathcal{Q}$  with  $t' \in Q$ . Notice that  $(s', t')$  may be covered by at most  $4d$  demand pairs. This is since  $t'$  belongs to exactly one square of  $\mathcal{Q}$ ; for every pair  $(s, t)$  covering  $(s', t')$ , path  $P(s, t)$  must contain a vertex from the boundary of that square; and the length of the boundary of each such square is at most  $4d$ . We say that demand pair  $(s, t)$  is *heavy* if it covers more than  $2^9 d \log n$  other pairs in  $\hat{\mathcal{M}}$ .

**Observation 5.1** *The number of heavy demand pairs is at most  $|\hat{\mathcal{M}}|/(2^7 \log n)$ .*

**Proof:** Assume otherwise. Then the total number of demand pairs covered by heavy pairs (counting multiplicities) is more than  $4|\hat{\mathcal{M}}|d$ . But every demand pair can be covered by at most  $4d$  demand pairs, a contradiction.  $\square$

**Partitioning algorithm.** Initially, we let  $\overline{\mathcal{M}}$  contain all heavy demand pairs, so  $|\overline{\mathcal{M}}| \leq |\hat{\mathcal{M}}|/(128 \log n)$ . We start with  $\mathcal{D} = \{D^0\}$ , and let the corresponding set of the demand pairs be  $\mathcal{M}_0 = \hat{\mathcal{M}} \setminus \overline{\mathcal{M}}$ . We then iterate. Let  $D^i \in \mathcal{D}$  be any disc and let  $\mathcal{M}_i$  be the corresponding set of demand pairs. Consider any pair  $(s, t), (s', t') \in \mathcal{M}_i$  of demand pairs. Let  $N_{\mathcal{M}_i}(t, t')$  be the number of the demand pairs  $(s'', t'') \in \mathcal{M}_i$  whose source lies between  $s$  and  $s'$  on  $R^*$ . Let  $N'_{\mathcal{M}_i}(t, t')$  denote the number of the remaining demand pairs in  $\mathcal{M}_i$  – those pairs whose sources do not lie between  $s$  and  $s'$  on  $R^*$ .

**Definition.** *We say that  $(s, t), (s', t')$  is a candidate pair for  $\mathcal{M}_i$  iff:*

- *there is a square  $Q \in \mathcal{Q}$  with  $t, t' \in Q$ ; and*
- *both  $N_{\mathcal{M}_i}(t, t'), N'_{\mathcal{M}_i}(t, t') > 2^{20}d \log^3 n$ .*

If a candidate pair  $(s, t), (s', t')$  exists in  $\mathcal{M}_i$ , then we do the following. Let  $Q \in \mathcal{Q}$  be the square that contains both  $t$  and  $t'$ . We define a new canonical disc  $D' \subseteq D^i$ . Let  $v$  be the first vertex on path  $P(s, t)$  that lies in  $Q$  (we view the path as directed from  $s$  to  $t$ ), and let  $P'(s, t)$  be the sub-path of  $P(s, t)$  between  $s$  and  $v$ . Similarly, let  $v'$  be the first vertex on path  $P(s', t')$  that lies in  $Q$ , and let  $P'(s', t')$  be the sub-path of  $P(s', t')$  between  $s'$  and  $v'$ . We let  $\sigma_3(D')$  be a path connecting  $v$  to  $v'$ , such that  $\sigma_3(D')$  is contained in the boundary of  $Q$ , and the length of  $\sigma_3(D')$  is at most  $2d$ . Let  $\sigma_2(D') = P'(s, t)$ ,  $\sigma_4(D') = P'(s', t')$ , and let  $\sigma_1(D')$  be the subpath of  $R^*$  between  $s$  and  $s'$ . Finally, let  $\sigma(D')$  be the concatenation of these four curves, and let  $D'$  be the disc whose boundary is  $\sigma(D')$ . Let  $\overline{\mathcal{M}}_i \subseteq \mathcal{M}_i$  contain all demand pairs  $(s'', t'')$ , such that either (i)  $P(s'', t'')$  crosses  $\sigma(D')$ ; or (ii)  $(s, t)$  covers  $(s'', t'')$ ; or (iii)  $(s', t')$  covers  $(s'', t'')$ . Recall that there are at most  $2d$  pairs of the first type (as all their paths must cross  $\sigma_3(D')$ ), and at most  $2^9 d \log n$  pairs of each of the remaining types (as we have discarded all heavy demand pairs). Altogether,  $|\overline{\mathcal{M}}_i| \leq 2d + 2^{10}d \log n$ . Notice that  $\overline{\mathcal{M}}_i$  contains the pairs  $(s, t)$  and  $(s', t')$ , and also all pairs whose destinations lie in the squares of  $\mathcal{Q}$  that  $\sigma(D')$  visits, including the pairs whose destinations lie in  $Q$ .

We create two new subsets of the demand pairs:  $\mathcal{M}' \subseteq \mathcal{M}_i \setminus \overline{\mathcal{M}}_i$  contains all demand pairs whose corresponding path in  $\mathcal{P}$  is contained in the new disc  $D'$ , and we update  $\mathcal{M}_i$  to contain all remaining demand pairs of the original set  $\mathcal{M}_i$  (excluding the pairs in  $\mathcal{M}'$  and  $\overline{\mathcal{M}}_i$ ). Therefore, for every demand pair in the new set  $\mathcal{M}_i$ , the path routing the pair in  $\mathcal{P}$  is contained in  $R_{\mathcal{D}}(D^i) \setminus D'$ . We then add  $D'$  to  $\mathcal{D}$ , with  $\mathcal{M}'$  as the corresponding set of the demand pairs, and we add all demand pairs in  $\overline{\mathcal{M}}_i$  to  $\overline{\mathcal{M}}$ . It is easy to verify that the resulting set of discs remains nested and canonical. From the above discussion, it also respects the currently active squares.

We have discarded at most  $2d + 2^{10}d \log n \leq 2^{11}d \log n$  demand pairs in the current iteration, but both  $\mathcal{M}'$  and the new set  $\mathcal{M}_i$  contain at least  $2^{20}d \log^3 n - 2^{10}d \log n - 2d \geq 2^{19}d \log^3 n$  demand pairs (after excluding the discarded demand pairs).

We continue this partitioning procedure, as long as there is some disc  $D^i \in \mathcal{D}$ , whose corresponding set  $\mathcal{M}_i$  of demand pairs contains a candidate pair  $(s, t), (s', t')$ . Let  $\mathcal{D} = \{D^0, \dots, D^z\}$  be the final set of discs, and for each  $i$ , let  $\mathcal{M}_i$  be the corresponding set of the demand pairs.

**Claim 5.2** *At the end of the algorithm,  $|\overline{\mathcal{M}}| < |\hat{\mathcal{M}}|/(64 \log n)$ .*

**Proof:** Recall that at the beginning of the algorithm,  $|\overline{\mathcal{M}}| \leq |\hat{\mathcal{M}}|/(128 \log n)$  held. Over the course of the algorithm, for every demand pair  $(s, t) \in \hat{\mathcal{M}} \setminus \overline{\mathcal{M}}$ , we define a budget of  $(s, t)$  as follows. Assume that  $(s, t)$  belongs to some set  $\mathcal{M}_i$ , corresponding to some disc  $D^i \in \mathcal{D}$ . We set the budget of  $(s, t)$  to be  $B(s, t) = \frac{\log(|\mathcal{M}_i|)}{128 \log^2 n}$ . Notice that this budget may change as the algorithm progresses. At the beginning of the algorithm, the total budget of all demand pairs is at most  $\frac{|\hat{\mathcal{M}}|}{128 \log n}$ , while  $|\overline{\mathcal{M}}| \leq \frac{|\hat{\mathcal{M}}|}{128 \log n}$ . It is now enough to prove that throughout the algorithm, the total budget of all demand pairs plus the cardinality of  $\overline{\mathcal{M}}$  do not increase.

Consider some iteration of the algorithm, where we processed some set  $\mathcal{M}_i$  of demand pairs as above, adding a new set  $\overline{\mathcal{M}}_i$  of demand pairs to  $\overline{\mathcal{M}}$ , and creating two new sets of demand pairs:  $\mathcal{M}'$  and  $\mathcal{M}'_i$  (that replaces  $\mathcal{M}_i$ ). Recall that we have guaranteed that  $|\overline{\mathcal{M}}_i| \leq 2^{11} d \log n$ , while  $|\mathcal{M}'|, |\mathcal{M}'_i| \geq 2^{19} d \log^3 n$ . Assume without loss of generality that  $|\mathcal{M}'| \leq |\mathcal{M}'_i|$ . Then the budget of every demand pair in  $\mathcal{M}'$  decreases by at least  $\frac{1}{128 \log^2 n}$  (as  $|\mathcal{M}'| \leq |\mathcal{M}_i|/2$ ), and the total decrease in the budgets is therefore at least  $\frac{|\mathcal{M}'|}{128 \log^2 n} \geq \frac{2^{19} d \log^3 n}{128 \log^2 n} \geq 2^{12} d \log n$ , while  $|\overline{\mathcal{M}}|$  increases by at most  $2^{11} d \log n$ . Therefore, the total budget of all demand pairs plus  $|\overline{\mathcal{M}}|$  do not increase.  $\square$

From the above discussion, the final set  $\mathcal{D}$  of discs is canonical, nested, and respects all remaining active squares.

**Boundary Pairs.** Consider some disc  $D^i \in \mathcal{D}$ , and assume first that  $|\mathcal{M}_i| \geq 2^{21} d \log^3 n$ . Let  $a_i, b_i$  be the two endpoints of  $\sigma_1(D^i)$ , and let  $\mathcal{Z}_0^i \subseteq \mathcal{M}_i$  contain  $2^{20} d \log^3 n$  demand pairs whose sources lie closest to  $a_i$  and  $2^{20} d \log^3 n$  demand pairs whose sources lie closest to  $b_i$ . If  $|\mathcal{M}_i| < 2^{21} d \log^3 n$ , then we set  $\mathcal{Z}_0^i = \mathcal{M}_i$ . In either case,  $|\mathcal{Z}_0^i| \leq 2^{21} d \log^3 n$ . We call the demand pairs in  $\mathcal{Z}_0^i$  *boundary pairs*. If  $|\mathcal{M}_i| > 2^{23} d \log^4 n$ , then we discard the boundary pairs from  $\mathcal{M}_i$ , and add them to  $\overline{\mathcal{M}}$ . After this procedure, we are still guaranteed that  $|\overline{\mathcal{M}}| \leq \frac{|\hat{\mathcal{M}}|}{64 \log n} + \frac{|\hat{\mathcal{M}}|}{4 \log n} \leq \frac{|\hat{\mathcal{M}}|}{2 \log n}$ . So we assume from now on that whenever the boundary pairs are present in  $\mathcal{M}_i$ ,  $|\mathcal{M}_i| \leq 2^{23} d \log^4 n$ .

**Forest of Discs.** It would be convenient for us to organize the discs in  $\mathcal{D}$  into a directed forest  $F$ , in a natural way. In each arborescence of the forest, the edges will be directed towards the root. The set of vertices of  $F$  is  $v_0, v_1, \dots, v_z$ , where  $v_i$  represents disc  $D^i$ . There is a directed edge  $(v_i, v_{i'})$  iff  $D^i \subsetneq D^{i'}$ , and there is no disc  $D^x \in \mathcal{D}$  with  $D^i \subsetneq D^x \subsetneq D^{i'}$ . Since the discs in  $\mathcal{D}$  are nested, it is immediate to verify that  $F$  is indeed a directed forest. For every arborescence  $\tau$  of  $F$ , the root of  $\tau$  is the vertex  $v_i$  corresponding to a disc  $D^i$  that has the largest area. For every vertex  $v_j \in V(F)$ , we let its weight  $w_j$  be  $|\mathcal{M}_j|$ . We denote by  $\mathcal{M}' = \bigcup_{i=0}^z \mathcal{M}_i$ , so  $|\mathcal{M}'| \geq |\hat{\mathcal{M}}|/2$ .

We use the following well-known result. For completeness, its proof appears in Appendix.

**Claim 5.3** *There is an efficient algorithm, that, given a directed forest  $F$  with  $n$  vertices, computes a partition  $\mathcal{Y} = \{Y_1, \dots, Y_{\lceil \log n \rceil}\}$  of  $V(F)$  into subsets, such that for each  $1 \leq j \leq \lceil \log n \rceil$ ,  $F[Y_j]$  is a collection of disjoint directed paths, that we denote by  $\mathcal{P}_j$ . Moreover, for all  $v, v' \in Y_j$ , if there is a directed path from  $v$  to  $v'$  in  $F$ , then they both lie on the same path in  $\mathcal{P}_j$ .*

Let  $\{Y_1, \dots, Y_{\lceil \log z \rceil}\}$  be the partition of  $V(F)$  given by Claim 5.3. Then there is some index  $1 \leq x \leq \lceil \log z \rceil$ , such that the total weight of all vertices in  $Y_x$  is at least  $|\mathcal{M}'|/\lceil \log z \rceil \geq |\hat{\mathcal{M}}|/(2 \log n)$ . Let  $\mathcal{D}' \subseteq \mathcal{D}$  be the set of all discs  $D^i$ , whose corresponding vertex  $v_i \in Y_x$ , and let  $\mathcal{M}'' \subseteq \mathcal{M}'$  be the set of all demand pairs that belong to all such discs, that is,  $\mathcal{M}'' = \bigcup_{D^i \in \mathcal{D}'} \mathcal{M}_i$ . Consider now some vertex  $v_i \in Y_x$ . Notice that at most one child-vertex  $v_{i'}$  of  $v_i$  in the forest  $F$  may belong to  $Y_x$ . If such a vertex  $v_{i'}$  does not exist, then all demand pairs in  $\mathcal{M}_i$  are called *right pairs*. Otherwise, for each demand pair

$(s, t) \in \mathcal{M}_i$ , if  $s$  appears to the left of  $\sigma_1(\mathcal{M}_{i'})$  on  $R^*$ , then we call it a *left pair*, and otherwise we call it a right pair. Therefore, all demand pairs in  $\mathcal{M}''$  are now partitioned into right pairs and left pairs. We assume without loss of generality that at least half of the demand pairs in  $\mathcal{M}''$  are left pairs; the other case is symmetric. Let  $\mathcal{M}''' \subseteq \mathcal{M}''$  be the set of all left pairs, so  $|\mathcal{M}'''| \geq |\mathcal{M}''|/2 \geq |\hat{\mathcal{M}}|/(4 \log n)$ .

**Claim 5.4** *The squares in  $\mathcal{Q}$  have the  $r$ -shadow property with respect to  $\mathcal{M}'''$ , for  $r = 2^{23} \log^4 n$ .*

Notice that, assuming the claim is correct, from Observation 4.3, there is a subset  $\hat{\mathcal{M}}' \subseteq \mathcal{M}'''$  of at least  $\frac{|\mathcal{M}'''}{2^{25} \log^4 n} \geq \frac{|\hat{\mathcal{M}}|}{2^{27} \log^5 n} \geq \frac{|\hat{\mathcal{M}}|}{\log^6 n}$  demand pairs, such that all squares in  $\mathcal{Q}$  have the 1-shadow property with respect to  $\hat{\mathcal{M}}'$  (we have used the assumption that  $n$  is large enough, so  $\log n \geq 2^{27}$ ). It now remains to prove Claim 5.4.

**Proof of Claim 5.4.** Assume otherwise, and let  $Q \in \mathcal{Q}$  be some square that does not have the  $r$ -shadow property with respect to  $\mathcal{M}'''$ . Since  $Q$  must be an active square, there is some disc  $D^i \in \mathcal{D}'$ , such that the square  $Q$  belongs to  $R_{\mathcal{D}}(D^i)$ . Consider the shadow  $J_{\tilde{\mathcal{M}}'''}(Q)$  of  $Q$ , and let  $s, s'$  be its left and right endpoints respectively. Then there are two demand pairs  $(s, t), (s', t') \in \mathcal{M}'''$ , with  $t, t' \in Q$ , such that at least  $rd$  vertices of  $S(\mathcal{M}''')$  lie between  $s$  and  $s'$  on  $R^*$ . Let  $S'$  be the set of all these source vertices, and let  $\mathcal{N} \subseteq \mathcal{M}'''$  be their corresponding demand pairs. We claim that  $\mathcal{N} \subseteq \mathcal{M}_i$ . Indeed, assume that there is some demand pair  $(s^*, t^*) \in \mathcal{N}$  that does not belong to  $\mathcal{M}_i$ . Then it must belong to some set  $\mathcal{M}_j$ , where  $v_j$  is a descendant of  $v_i$  in the forest  $F$ . This can only happen if  $v_j \in Y_x$ . Moreover, the segment  $\sigma_1(D^j)$  must be contained in the shadow  $L_{\tilde{\mathcal{M}}'''}(Q)$ . Let  $v_{j'}$  be the child vertex of  $v_i$  that lies in  $Y_x$  (such a vertex must exist due to the vertex  $v_j \in Y_x$ ). Then  $v_j$  is also the descendant of  $v_{j'}$ , and  $D^j \subseteq D^{j'} \subseteq D^i$ . In particular, the segment  $\sigma_1(D^{j'})$  must also be contained in the shadow  $L_{\tilde{\mathcal{M}}'''}(Q)$ . But then  $(s, t)$  is a left pair, while  $(s', t')$  is a right pair, which is impossible, since we have discarded all the right pairs. We conclude that  $\mathcal{N} \subseteq \mathcal{M}_i$ .

Since  $|\mathcal{M}_i| \geq |\mathcal{N}| > rd = 2^{23} d \log^4 n$ , we have deleted the boundary pairs from the set  $\mathcal{M}_i$ , and so in particular,  $(s, t)$  and  $(s', t')$  are not boundary pairs. It follows that  $(s, t)$  and  $(s', t')$  was a valid candidate pair, and we should have continued our partitioning algorithm.  $\square$

## 6 The Existence of Coloring

In this section we prove Theorem 4.7. Recall that we assume that we are given a lower bound  $\text{OPT}$  on the value of the optimal solution; a hierarchical system  $\tilde{\mathcal{H}} = (\mathcal{Q}_1, \dots, \mathcal{Q}_\rho) \in \mathcal{F}$ , and a sequence  $L = (\ell_1, \dots, \ell_\rho)$  of integers as in Theorem 4.6. As before, we denote by  $(\mathcal{J}_1, \dots, \mathcal{J}_\rho)$  the  $L$ -hierarchical decomposition of  $R^*$ . From Theorem 4.6, there is a set  $\tilde{\mathcal{M}}^{**}$  of demand pairs, with  $|\tilde{\mathcal{M}}^{**}| \geq \text{OPT}/2^{O(\sqrt{\log n} \log \log n)}$ , such that for all demand pairs in  $\tilde{\mathcal{M}}^{**}$ , their destinations belong to  $\tilde{\mathcal{H}}$ , so in particular  $\tilde{\mathcal{M}}^{**} \subseteq \tilde{\mathcal{M}}$ . Moreover, all demand pairs in  $\tilde{\mathcal{M}}^{**}$  can be routed via node-disjoint paths in  $G$ , all squares in  $\bigcup_{h=1}^\rho \mathcal{Q}_h$  have the  $1/\eta^2$ -shadow property with respect to  $\tilde{\mathcal{M}}^{**}$ , and the set  $\tilde{\mathcal{M}}^{**}$  of demand pairs is compatible with the sequence  $L$ . Recall that the latter means that for each  $1 \leq h \leq \rho$ , for every level- $h$  interval  $I \in \mathcal{J}_h$  of  $R^*$ , either  $I \cap S(\tilde{\mathcal{M}}^{**}) = \emptyset$ , or  $d_h/(16\eta) \leq |S(\tilde{\mathcal{M}}^{**}) \cap I| \leq d_h/4$ .

For each  $1 \leq h \leq \rho$ , for every level- $h$  square  $Q \in \mathcal{Q}_h$ , let  $\hat{\mathcal{M}}(Q) \subseteq \tilde{\mathcal{M}}^{**}$  denote the set of all demand pairs whose destinations lie in  $Q$ . We use the following simple observation.

**Observation 6.1** *For each  $1 \leq h \leq \rho$ , for every level- $h$  square  $Q \in \mathcal{Q}_h$ , the source vertices of the demand pairs in  $\hat{\mathcal{M}}(Q)$  belong to at most two level- $h$  intervals in  $\mathcal{J}_h$ .*

**Proof:** Assume otherwise, and let  $I', I'', I''' \in \mathcal{J}_h$  be three distinct level- $h$  intervals that each contain at last one source vertex of  $\hat{\mathcal{M}}(Q)$ . Assume that  $I', I'', I'''$  lie on  $R^*$  in this left-to-right order. Then

$I''$  is contained in the shadow  $L_{\hat{\mathcal{M}}}(Q)$ . From the  $1/\eta^2$ -shadow property of  $Q$ , the length of this shadow must be at most  $d_h/\eta^2$ . But from our construction, the length of the shadow is at least  $|S(\tilde{\mathcal{M}}^{**}) \cap I''| \geq d_h/(16\eta) > d_h/\eta^2$ , a contradiction.  $\square$

We start with the set  $\hat{\mathcal{M}}_0 = \tilde{\mathcal{M}}^{**}$  of demand pairs, and perform  $\rho$  iterations. In iteration  $h$ , the input is some subset  $\hat{\mathcal{M}}_{h-1} \subseteq \hat{\mathcal{M}}_0$  of demand pairs, and an assignment of level- $(h-1)$  colors to all level- $(h-1)$  squares, that has the following property: for each demand pair  $(s, t) \in \hat{\mathcal{M}}_{h-1}$ , if  $Q \in \mathcal{Q}_{h-1}$  is the level- $(h-1)$  square containing  $t$ , and  $I \in \mathcal{J}_{h-1}$  is the level- $(h-1)$  interval containing  $s$ , then  $Q$  and  $I$  are assigned the same level- $(h-1)$  color, or in other words,  $Q$  is assigned the level- $(h-1)$  color  $c_{h-1}(I)$ . Our starting point is the set  $\hat{\mathcal{M}}_0$  of demand pairs. We assume that the level-0 color of all vertices is  $c_0$ , so, letting  $\tilde{G}$  be a unique level-0 square, the invariant holds.

Consider now some iteration  $h$ , together with the corresponding set  $\hat{\mathcal{M}}_{h-1} \subseteq \tilde{\mathcal{M}}^{**}$  of demand pairs. Consider some level- $h$  square  $Q \in \mathcal{Q}_h$ , and assume that its parent-square is assigned the level- $(h-1)$  color  $c_{h-1}$ . Let  $\hat{\mathcal{M}}'(Q) \subseteq \hat{\mathcal{M}}_{h-1}$  be the set of all demand pairs, whose destinations lie in  $Q$ . From the above observation, there are at most two level- $h$  intervals of  $R^*$  (that we denote by  $I$  and  $I'$ , where possibly  $I = I'$ ) that contain the source vertices of  $\hat{\mathcal{M}}'(Q)$ . Assume w.l.o.g. that  $I$  contains at least half of these source vertices. From our invariant, all vertices of the parent-interval of  $I$  are also assigned the level- $(h-1)$  color  $c_{h-1}$  - same as the color assigned to the parent-square of  $Q$ . We then assign to  $Q$  the level- $h$  color  $c_h(I)$  - that is, the color associated with the interval  $I$ . We also discard from  $\hat{\mathcal{M}}'(Q)$  all demand pairs except those whose sources are contained in  $I$ , obtaining a new set  $\hat{\mathcal{M}}''(Q) \subseteq \hat{\mathcal{M}}'(Q)$  of demand pairs, containing at least  $|\hat{\mathcal{M}}''(Q)|/2$  demand pairs. We then set  $\hat{\mathcal{M}}_h = \bigcup_{Q \in \mathcal{Q}_h} \hat{\mathcal{M}}''(Q)$ .

Let  $\tilde{\mathcal{M}}' = \hat{\mathcal{M}}_\rho$  be the set of demand pairs obtained after the last iteration. Then it is immediate to verify that:

$$|\tilde{\mathcal{M}}'| \geq \frac{|\tilde{\mathcal{M}}^{**}|}{2^\rho} \geq \frac{|\tilde{\mathcal{M}}^{**}|}{2^{\sqrt{\log n}}} \geq \frac{\text{OPT}}{2^{O(\sqrt{\log n} \log \log n)}}.$$

Clearly, for each demand pair  $(s, t) \in \tilde{\mathcal{M}}'$ , both  $s$  and  $t$  are assigned the same level- $\rho$  color. Finally, fix some level  $1 \leq h \leq \rho$ , and some level- $h$  color  $c_h(I)$ , that is associated with some level- $h$  interval  $I$ . It remains to verify that at most  $d_h$  demand pairs are assigned color  $c_h(I)$ . This is clearly true, since all such demand pairs have their source vertices lying on  $I$ , and, since  $\tilde{\mathcal{M}}^{**}$  is compatible with  $L$  from Theorem 4.6, we are guaranteed that the number of such demand pairs is at most  $d_h/4$ .

## 7 Finding the Routing

The goal of this section is to prove Theorem 4.8. Recall that we are given as input a hierarchical system  $\tilde{\mathcal{H}} = (\mathcal{Q}_1, \dots, \mathcal{Q}_\rho) \in \mathcal{F}$  of squares and a hierarchical L-partition  $(\mathcal{J}_1, \dots, \mathcal{J}_\rho)$  of the row  $R^*$  of  $G'$ . We are also given a coloring  $f$  of  $\tilde{\mathcal{H}}$  by the corresponding hierarchical system  $\mathcal{C}$  of colors, and a set  $\tilde{\mathcal{M}}' \subseteq \tilde{\mathcal{M}}$  of at most  $d_1$  demand pairs that are perfect for  $f$ . In particular, we are guaranteed that for each demand pair  $(s, t) \in \tilde{\mathcal{M}}'$ , both  $s$  and  $t$  are assigned the same level- $\rho$  color, and, moreover, for every  $1 \leq h \leq \rho$ , for every level- $h$  color  $c_h$ , at most  $d_h$  demand pairs in  $\tilde{\mathcal{M}}'$  are assigned color  $c_h$ . Recall that all destination vertices appear within distance at least  $\text{OPT}/\eta > 4d_1$  from the grid boundary. For consistency, we define a single level-0 square  $Q_0$ , as follows: let  $\mathcal{R}$  contain all but the  $d_1$  top and the  $d_1$  bottom rows of the grid  $G$ , and let  $\mathcal{W}$  contain all but the first  $d_1$  and the last  $d_1$  columns of the grid  $G$ . Then  $Q_0$  is the sub-grid of  $G$  spanned by the rows in  $\mathcal{R}$  and the columns in  $\mathcal{W}$ . We discard from  $\mathcal{Q}_1$  all level-1 squares that are not contained in  $Q_0$ : such squares do not contain any destination vertices. We also discard from all other sets  $\mathcal{Q}_h$  all descendant-squares of the level-1

squares that were discarded. Abusing the notation, we denote the resulting hierarchical system of squares by  $\tilde{\mathcal{H}}$ . For convenience, we define a unique level-0 color  $c_0$ , that serves as a parent of every level-1 color, and we assign this color to  $Q_0$ .

Assume w.l.o.g. that  $\tilde{\mathcal{M}}' = \{(s_1, t_1), \dots, (s_z, t_z)\}$ , where  $s_1, \dots, s_z$  appear on  $R^*$  in this left-to-right order. We define a set  $\hat{\mathcal{M}} = \{(s_i, t_i) \mid i \equiv 1 \pmod{2\eta^3}\}$ , of  $\Omega(|\tilde{\mathcal{M}}'|/\eta^3)$  demand pairs. In the remainder of this section, we show an efficient algorithm for routing all demand pairs in  $\hat{\mathcal{M}}$ . Consider a level- $h$  color  $c_h$ , where  $1 \leq h \leq \rho$ . Let  $N(c_h)$  be the number of demand pairs  $(s, t) \in \hat{\mathcal{M}}$ , such that  $s$  and  $t$  have level- $h$  color  $c_h$ . Since  $d_h = \eta^{\rho-h+3}$ ,  $N(c_h) \leq \lceil d_h/(2\eta^3) \rceil \leq \eta^{\rho-h}$ , and  $|\hat{\mathcal{M}}| \leq d_1/(2\eta^3)$ . Notice that for every level- $\rho$  color  $c_\rho$ ,  $N(c_\rho) \leq 1$  holds.

We say that a level- $h$  square  $Q \in \mathcal{Q}_h$  is *empty* iff no demand pair in  $\hat{\mathcal{M}}$  has a destination in  $Q$ . As before, for each square  $Q$ , we let  $\Gamma(Q)$  denote its boundary. For each level  $1 \leq h \leq \rho$  and for each level- $h$  square  $Q_h \in \mathcal{Q}_h$ , let  $Q_h^+$  be an extended square, obtained by adding a margin of  $d_h/\eta$  on all sides of  $Q_h$ ; we also let  $Q_0^+$  be obtained from  $Q_0$  by adding a margin of  $d_1/\eta$  around it. Notice that the distance between  $\Gamma(Q_0^+)$  and  $\Gamma(G)$  remains at least  $d_1/2$ . Notice also that for each  $1 \leq h \leq \rho$ , for each pair  $Q, \hat{Q} \in \mathcal{Q}_h$  of level  $h$  squares,  $d(Q^+, \hat{Q}^+) \geq d_h/2$ .

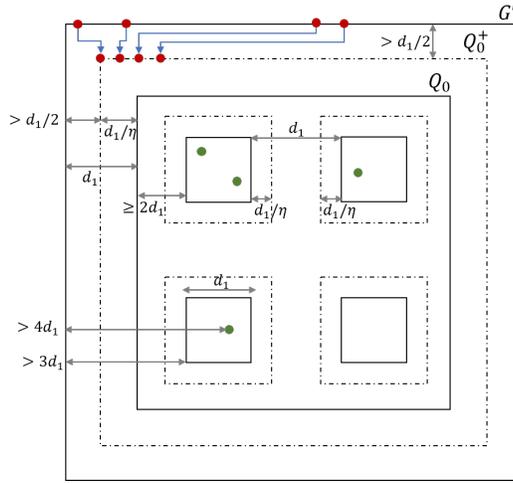


Figure 4: Square  $Q_0$  inside  $G$ . Original squares  $Q$  of  $\mathcal{H}$  are represented by solid lines and the corresponding augmented squares  $Q^+$  are represented by dashed lines. The sources  $s_j$  of  $\hat{\mathcal{M}}$  are represented by red nodes on the top boundary of  $G'$ , and their copies  $s'_j$  are on the top row of  $Q_0^+$ . The paths of  $\mathcal{P}_0$  are shown in blue.

Let  $I'_0$  be the set of the  $3 \cdot N(c_0)$  leftmost vertices on the top row of  $Q_0^+$ . We map the vertices of  $S(\hat{\mathcal{M}})$  to the vertices of  $I'_0$  as follows: if  $s$  is the  $j$ th vertex from the left in  $S(\hat{\mathcal{M}})$ , then it is mapped to the  $(3j)$ th leftmost vertex of  $I'_0$ , that we denote by  $s'$ . For every level  $1 \leq h \leq \rho$ , for every level- $h$  interval  $I_h \in \mathcal{I}_h$  of  $R^*$ , we can now naturally define a corresponding interval  $I'_h$  of  $I'_0$  as the smallest interval containing all vertices in  $\{s'_j \mid s_j \in I_h\}$ . We then let  $\mathcal{I}'_h = \{I'_h \mid I_h \in \mathcal{I}_h\}$ , so that  $(\mathcal{I}'_1, \dots, \mathcal{I}'_\rho)$  can be viewed as a hierarchical partition of  $I'_0$ . The parent-child relationship between the new intervals is defined exactly as before. Let  $\mathcal{I}'_0 = \{I'_0\}$ .

The following observation is immediate (see Figure 4):

**Observation 7.1** *There is an efficient algorithm to compute a set  $\mathcal{P}_0$  of node-disjoint paths, connecting every source  $s_j \in S(\hat{\mathcal{M}})$  to its corresponding copy  $s'_j \in I'_0$ , such that the paths in  $\mathcal{P}_0$  are internally disjoint from  $Q_0^+$ .*

In order to define the routing of all demand pairs in  $\hat{\mathcal{M}}$ , we use special a structure called a *snake*, which we define next. Recall that, given a set  $\mathcal{R}$  of consecutive rows of  $G$  and a set  $\mathcal{W}$  of consecutive columns of  $G$ , we denoted by  $\Upsilon(\mathcal{R}, \mathcal{W})$  the subgraph of  $G$  spanned by the rows in  $\mathcal{R}$  and the columns in  $\mathcal{W}$ ; we refer to such a graph as a *corridor*.

Let  $\Upsilon = \Upsilon(\mathcal{R}, \mathcal{W})$  be any corridor. Let  $R'$  and  $R''$  be the first and the last row of  $\mathcal{R}$  respectively, and let  $W'$  and  $W''$  be the first and the last column of  $\mathcal{W}$  respectively. Each of the four paths  $\Upsilon \cap R'$ ,  $\Upsilon \cap R''$ ,  $\Upsilon \cap W'$  and  $\Upsilon \cap W''$  is called a *boundary edge* of  $\Upsilon$ , and their union is called the *boundary* of  $\Upsilon$ . The width of the corridor  $\Upsilon$  is  $w(\Upsilon) = \min\{|\mathcal{R}|, |\mathcal{W}|\}$ . We say that two corridors  $\Upsilon, \Upsilon'$  are *internally disjoint*, iff every vertex  $v \in \Upsilon \cap \Upsilon'$  belongs to the boundaries of both corridors. We say that two internally disjoint corridors  $\Upsilon, \Upsilon'$  are *neighbors* iff  $\Upsilon \cap \Upsilon' \neq \emptyset$ .

We are now ready to define snakes. A snake  $\mathcal{Y}$  of length  $z$  is a sequence  $(\Upsilon_1, \Upsilon_2, \dots, \Upsilon_z)$  of  $z$  corridors that are pairwise internally disjoint, such that for all  $1 \leq z', z'' \leq z$ ,  $\Upsilon_{z'}$  is a neighbor of  $\Upsilon_{z''}$  iff  $|z' - z''| = 1$ . The width of the snake is defined to be the minimum of  $\min_{1 \leq z' < z'' \leq z} \{|\Upsilon_{z'} \cap \Upsilon_{z''}|\}$ , and  $\min_{1 \leq z' \leq z} \{w(\Upsilon_{z'})\}$ . We say that a node  $u \in \mathcal{Y}$  iff  $u \in \bigcup_{z'=1}^z \Upsilon_{z'}$ . Notice that given a snake  $\mathcal{Y}$ , there is a unique simple cycle  $\sigma(\mathcal{Y})$  in  $G$  such that all the nodes  $u \in \mathcal{Y}$  lie on or inside  $\sigma(\mathcal{Y})$  and all other nodes of  $G$  lie outside it. We call  $\sigma(\mathcal{Y})$  the *boundary* of  $\mathcal{Y}$ . We say that a pair of snakes  $\mathcal{Y}_1$  and  $\mathcal{Y}_2$  are *internally disjoint* iff there is a boundary edge  $\Sigma_1$  of some corridor of  $\mathcal{Y}_1$  and a boundary edge  $\Sigma_2$  of some corridor of  $\mathcal{Y}_2$  such that for every vertex  $u$  with  $u \in \mathcal{Y}_1$  and  $u \in \mathcal{Y}_2$ ,  $u$  must belong to  $V(\Sigma_1) \cap V(\Sigma_2)$ . We use the following simple claim, whose proof can be found, e.g. in [CKN17b].

**Claim 7.2** *Let  $\mathcal{Y} = (\Upsilon_1, \dots, \Upsilon_z)$  be a snake of width  $w$ , and let  $A$  and  $A'$  be two sets of vertices with  $|A| = |A'| \leq w - 2$ , such that the vertices of  $A$  lie on a single boundary edge of  $\Upsilon_1$  and the vertices of  $A'$  lie on a single boundary edge of  $\Upsilon_z$ . There is an efficient algorithm, that, given the snake  $\mathcal{Y}$ , and the sets  $A$  and  $A'$  of vertices as above, computes a set  $\mathcal{Q}$  of node-disjoint paths contained in  $\bigcup_{z'=1}^z \Upsilon_{z'}$ , that connect every vertex of  $A$  to a distinct vertex of  $A'$ .*

**Definition.** *Let  $0 \leq h \leq \rho$  be a level, let  $c_h \in \chi_h$  be a level- $h$  color, and let  $\mathcal{Y}$  be a snake. We say that a snake  $\mathcal{Y}$  is a valid level- $h$  snake for color  $c_h$  iff:*

- $\mathcal{Y} \subseteq Q_0^+$ ;
- $\mathcal{Y}$  has width at least  $3 \cdot N(c_h)$ ;
- $\mathcal{Y}$  contains the interval  $I' \in \mathcal{I}'_h$ , corresponding to the interval  $I \in \mathcal{I}_h$  that was used in defining color  $c_h$ , as a part of the top boundary of its first corridor; and
- for each non-empty level- $h$  square  $Q \in \mathcal{Q}_h$ , whose level- $h$  color is  $c_h$ ,  $Q^+$  is a corridor of  $\mathcal{Y}$ .

The following lemma is central to proving Theorem 4.8.

**Lemma 7.3** *There is an efficient algorithm, that constructs, for each level  $1 \leq h \leq \rho$ , for each level- $h$  color  $c_h$ , a valid level- $h$  snake  $\mathcal{Y}(c_h)$  for this color, such that all snakes of the same level are disjoint.*

We provide the proof of the lemma below, after completing the proof of Theorem 4.8 using it. Fix some level- $\rho$  color  $c_\rho$ . Recall that  $\hat{\mathcal{M}}$  contains at most one demand pair  $(s, t)$ , such that  $s$  and  $t$  are assigned the level- $\rho$  color  $c_\rho$ . Recall that  $s' \in I'_0$  is the vertex on the top boundary of  $Q_0^+$  to which  $s$  is mapped. We route the pair  $(s', t)$  inside the valid level- $\rho$  snake  $\mathcal{Y} = \mathcal{Y}(c_\rho)$  given by Lemma 7.3. This snake is guaranteed to contain both the level- $\rho$  interval  $I' \in \mathcal{I}'(\rho)$  to which  $s'$  belongs, and the non-empty level- $\rho$  square  $Q_\rho$  that contains  $t$  (as the level- $\rho$  color of  $Q_\rho$  must be  $c_\rho$ ). We select an arbitrary path  $P$  connecting  $s'$  to  $t$  inside  $\mathcal{Y}$ . In order to complete the routing, we combine these paths

with the set  $\mathcal{P}_0$  of paths computed in Observation 7.1. As all level- $\rho$  snakes are disjoint, we obtain a valid routing of at least  $|\hat{\mathcal{M}}| \geq |\tilde{\mathcal{M}}|/\eta^3 = |\tilde{\mathcal{M}}|/2^{O(\sqrt{\log n})}$  demand pairs. It now remains to prove Lemma 7.3.

**Proof of Lemma 7.3.** The proof is by induction on  $h$ . For level 0, we construct a single level-0 snake  $\mathcal{Y}(c_0)$ , consisting of a single corridor  $Q_0^+$ . Clearly, this is a valid level-0 snake for color  $c_0$ . Assume now that the claim holds for some level  $0 \leq h < \rho$ . We prove that it holds for level  $(h + 1)$ .

Fix a level- $h$  interval  $I_h \in \mathcal{I}_h$ , and its corresponding level- $h$  color  $c_h$ . Let  $\mathcal{Y} = \mathcal{Y}(c_h)$  be the valid level- $h$  snake for  $c_h$ , given by the induction hypothesis. Recall that for each non-empty level- $h$  square  $Q$  of color  $c_h$ ,  $Q^+$  is a corridor of  $\mathcal{Y}$ . The following claim completes the induction step and Lemma 7.3 follows.

**Claim 7.4** *There is an efficient algorithm, that, given a valid level- $h$  snake  $\mathcal{Y}(c_h)$  for a level- $h$  color  $c_h$ , constructs, for every child-color  $c_{h+1}$  of  $c_h$ , a valid level- $(h + 1)$  snake  $\mathcal{Y}(c_{h+1})$ , such that all resulting level- $(h + 1)$  snakes are disjoint from each other and are contained in  $\mathcal{Y}(c_h)$ .*

**Proof:** We will construct a number of *sub-snakes*<sup>3</sup> for each color  $c_{h+1}$ , and concatenate them together to construct a snake  $\mathcal{Y}(c_{h+1})$  with claimed properties. To this end, we will construct sub-snakes of two types: Sub-snakes of type 1 will be contained in  $\mathcal{Y}$ , but they will be internally disjoint from all non-empty level- $h$  squares of color  $c_h$ . Each sub-snake of type 2 will be contained inside some non-empty level- $h$  square of color  $c_h$ . In order to coordinate between these sub-snakes, we use *interfaces* on the top and bottom boundaries of such squares, that we define next.

For each non-empty level- $h$  square  $Q$  of color  $c_h$ , let  $I_Q$  be the shortest sub-path of the top boundary of  $Q^+$  that contains the  $3 \cdot N(c_h)$  leftmost vertices of the top boundary of  $Q^+$ . Similarly, let  $J_Q$  be the shortest sub-path of the bottom boundary of  $Q^+$  that contains the  $3 \cdot N(c_h)$  rightmost vertices of the bottom boundary of  $Q^+$ . Notice that  $I_Q$  and  $J_Q$  are well-defined since  $3 \cdot N(c_h) \leq 3 \cdot \lceil d_h/2\eta^3 \rceil \leq d_h$ .

Let  $c_h^1, c_h^2, \dots, c_h^z$  be the child-colors of  $c_h$ , indexed in the left-to-right order of their corresponding intervals on  $R^*$ . For a non-empty level- $h$  square  $Q$ , let  $\{I_Q^i \subseteq I_Q \mid i \in \{1, \dots, z\}\}$  be a collection of disjoint sub-paths of  $I_Q$  such that the following holds: (i) each sub-path  $I_Q^i$  contains  $3 \cdot N(c_h^i)$  vertices; and (ii) for all  $1 \leq i < i' \leq z$ ,  $I_Q^i$  appears to the left of  $I_Q^{i'}$  on the top boundary of  $Q^+$ . Similarly, let  $\{J_Q^i \subseteq J_Q \mid i \in \{1, \dots, z\}\}$  be an analogous collection of disjoint sub-paths of  $J_Q$ . Notice that there is a unique set of such sub-paths, since  $\sum_{i=1}^z 3 \cdot N(c_h^i) = 3 \cdot N(c_h)$ . The sub-paths  $I_Q^i$  and  $J_Q^i$  will act as interfaces between the sub-snakes corresponding to color  $c_h^i$ , and will be used to combine them together. Having described the interfaces, we now focus on constructing type-2 sub-snakes inside each non-empty square of level  $h$ . We will then construct type-1 sub-snakes, and finally combine them all together to obtain level- $(h + 1)$  snakes as claimed.

**Definition.** *Fix a non-empty level- $h$  square  $Q$  of color  $c_h$ , and let  $c_h^i$  be a child-color of  $c_h$ . Given a snake  $\hat{\mathcal{Y}}$ , we say that  $\hat{\mathcal{Y}}$  is a valid type-2 sub-snake of color  $c_h^i$  in square  $Q^+$  iff*

- $\hat{\mathcal{Y}} \subseteq Q^+$ ;
- $\hat{\mathcal{Y}}$  has width at least  $3 \cdot N(c_h^i)$ ;
- $\hat{\mathcal{Y}}$  contains  $I_Q^i$  as the top boundary of its first corridor and  $J_Q^i$  as the bottom boundary of its last corridor; and

---

<sup>3</sup>Each of the sub-snake that we construct is a snake as defined earlier. We use the term ‘sub-snake’ to distinguish them from the final snake  $\mathcal{Y}(c_{h+1})$ , that is obtained by concatenating a number of sub-snakes.

- for each non-empty level- $(h+1)$  child square  $\hat{Q}$  of  $Q$  of color  $c_h^i$ ,  $\hat{Q}^+$  is a corridor of  $\hat{\mathcal{Y}}$ .

**Claim 7.5** *There is an efficient algorithm, that, given a non-empty level- $h$  square  $Q$  of color  $c_h$  constructs, for every child-color  $c_h^i$  of  $c_h$ , a valid type-2 sub-snake  $\mathcal{Y}_Q(c_h^i)$  of color  $c_h^i$  in square  $Q^+$ , such that the resulting sub-snakes are pairwise disjoint.*

We prove the above claim after completing the proof of Claim 7.4, using the claim. We now turn to construct type-1 sub-snakes. We start with the following simple observation.

**Observation 7.6** *Let  $\mathcal{Y} = (\Upsilon_1, \Upsilon_2, \dots, \Upsilon_r)$  be a snake of width  $w$ . Let  $\Sigma_1$  be some boundary edge of  $\Upsilon_1$  that is contained in the boundary of  $\mathcal{Y}$ . Similarly, let  $\Sigma_r \neq \Sigma_1$  be some boundary edge of  $\Upsilon_r$  that is contained in the boundary of  $\mathcal{Y}$ . Let  $(A_1, A_2, \dots, A_j)$  be disjoint sub-paths on  $\Sigma_1$  and let  $(B_1, B_2, \dots, B_j)$  be disjoint sub-paths on  $\Sigma_r$  such that the sub-paths  $(A_1, \dots, A_j, B_j, B_{j-1}, \dots, B_1)$  appear in this clockwise order on the boundary of  $\mathcal{Y}$  and  $|\bigcup_{i=1}^j V(A_i)|, |\bigcup_{i=1}^j V(B_i)| \leq w$ . Then there is an efficient algorithm to construct pairwise disjoint snakes  $\mathcal{Y}_1, \dots, \mathcal{Y}_j$  such that for each  $1 \leq i \leq j$ , (i)  $\mathcal{Y}_i$  is contained in  $\mathcal{Y}$ ; (ii)  $A_i$  is contained as a part of the boundary of the first corridor of  $\mathcal{Y}_i$ ; (iii)  $B_i$  is contained as a part of the boundary of the last corridor of  $\mathcal{Y}_i$ ; and (iv) the width of the snake  $\mathcal{Y}_i$  is at least  $\min(|A_i|, |B_i|)$ .*

Recall that we are given a valid level- $h$  snake  $\mathcal{Y}$  for a level- $h$  color  $c_h$ . Also, recall that  $\mathcal{Y}$  is a sequence of corridors, containing the squares  $Q^+$  corresponding to non-empty level- $h$  color  $c_h$  squares  $Q$  as corridors. Let  $Q_1, \dots, Q_r$  be such squares, in the order of their appearance in  $\mathcal{Y}$ . From the above observation, it is now immediate to find pairwise disjoint sub-snakes  $\mathcal{Y}^1(c_h^i), \dots, \mathcal{Y}^r(c_h^i)$  for each child-color  $c_h^i$  of  $c_h$  such that:

- the top boundary of the first corridor of  $\mathcal{Y}^1(c_h^i)$  contains the interval  $I'(c_h^i) \in \mathcal{I}'_{h+1}$  corresponding to the interval  $I(c_h^i) \in \mathcal{I}_{h+1}$  that was used in defining color  $c_h^i$ ;
- for each  $2 \leq j \leq r$ , the top boundary of the first corridor of  $\mathcal{Y}^j(c_h^i)$  contains the interval  $J_{Q_{j-1}}^i$ ;
- for each  $1 \leq j \leq r$ , the bottom boundary of the last corridor of  $\mathcal{Y}^j(c_h^i)$  contains the interval  $I_{Q_j}^i$ ; and
- all the sub-snakes are pairwise disjoint, and are internally disjoint from the squares  $Q_1, \dots, Q_r$ .

For each child color  $c_h^i$  of  $c_h$ , we say that the sub-snakes  $\mathcal{Y}^1(c_h^i), \dots, \mathcal{Y}^r(c_h^i)$  are sub-snakes of type 1. Now we focus on combining the snakes of type 1 and type 2.

**Definition.** *Let  $\mathcal{Y}^1 = (\Upsilon_{\ell_1}^1, \dots, \Upsilon_{\ell_1}^1)$  and  $\mathcal{Y}^2 = (\Upsilon_1^2, \dots, \Upsilon_{\ell_2}^2)$  be a pair of snakes of width  $w_1$  and  $w_2$  respectively. We say that  $\mathcal{Y}^1$  and  $\mathcal{Y}^2$  are composable iff they are internally disjoint and there are boundary edges  $\Sigma^1$  of  $\Upsilon_{\ell_1}^1$  and  $\Sigma^2$  of  $\Upsilon_1^2$  such that  $|V(\Sigma^1) \cap V(\Sigma^2)| \geq \min(w_1, w_2)$ . We denote by  $\mathcal{Y}^1 \oplus \mathcal{Y}^2$  the sequence  $(\Upsilon_1^1, \dots, \Upsilon_{\ell_1}^1, \Upsilon_1^2, \dots, \Upsilon_{\ell_2}^2)$  of corridors.*

We begin with a simple observation that follows directly from the definition of snakes.

**Observation 7.7** *Let  $\mathcal{Y}^1$  and  $\mathcal{Y}^2$  be a pair of composable snakes, each of width at least  $w$ . Then  $\hat{\mathcal{Y}} = \mathcal{Y}^1 \oplus \mathcal{Y}^2$  is also a snake of width at least  $w$ .*

Notice that for each child-color  $c_h^i$  of  $c_h$  and for each  $1 \leq j \leq r$ , the pair of sub-snakes  $\mathcal{Y}^j(c_h^i)$  and  $\mathcal{Y}_{Q_j}(c_h^i)$  are composable. Similarly, for each  $1 \leq j < r$ , the pair of sub-snakes  $\mathcal{Y}_{Q_j}(c_h^i)$  and  $\mathcal{Y}^{j+1}(c_h^i)$  are composable. Claim 7.4 now follows by considering the snakes

$$\mathcal{Y}(c_h^i) = \mathcal{Y}^1(c_h^i) \oplus \mathcal{Y}_{Q_1}(c_h^i) \oplus \mathcal{Y}^2(c_h^i) \oplus \dots \oplus \mathcal{Y}^r(c_h^i) \oplus \mathcal{Y}_{Q_r}(c_h^i)$$

for all child-colors  $c_h^i$  of  $c_h$ . □

In the remainder of the section, we prove Claim 7.5.

**Proof of Claim 7.5.**

Recall that we are given a non-empty level- $h$  square  $Q$  of color  $c_h$ . Let  $\mathcal{D}$  be the set of all non-empty level- $(h+1)$  child squares of  $Q$ . Recall that given a pair of intervals  $I, I'$  of  $[\ell']$ , we denote by  $Q(I, I')$  the sub-graph induced by all vertices  $\{v(i, j) \mid i \in I, j \in I'\}$ . We say that a pair of squares  $Q_1, Q_2 \in \mathcal{D}$  are *vertically aligned* iff  $Q_1 = Q(I, I_1)$  and  $Q_2 = Q(I, I_2)$  for some intervals  $I, I_1, I_2$  of  $[\ell']$ . Let  $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_r$  be a partition of  $\mathcal{D}$  such that the following holds:

- for each  $1 \leq i \leq r$ , the squares in  $\mathcal{D}_i$  are vertically aligned; and
- for each  $1 \leq i < i' \leq r$ , all squares in  $\mathcal{D}_i$  appear to the left of those in  $\mathcal{D}_{i'}$ .

For each  $1 \leq i \leq r$ , we denote  $\mathcal{D}_i = (Q_i^1, \dots, Q_i^{k_i})$ , where for each  $1 \leq j < j' \leq k_i$ , the square  $Q_i^{j'}$  appears above  $Q_i^j$ .

Our goal, as before, is to construct a number of child-snakes<sup>4</sup> for each child-color  $c_h^i$  of  $c_h$ , which can be combined together to obtain the desired sub-snake  $\mathcal{Y}_Q(c_h^i)$ . To this end, we need interfaces, which we define next.

For each child-square  $\hat{Q} \in \mathcal{D}$  of  $Q$ , let  $R(\hat{Q}^+)$  and  $R'(\hat{Q}^+)$  be the rows containing the top and bottom boundaries of  $\hat{Q}^+$  respectively. Let  $\alpha(\hat{Q})$  be the sub-path containing all nodes of the top boundary of  $\hat{Q}^+$  along with  $3 \cdot N(c_h)$  nodes of  $R(\hat{Q}^+)$  lying immediately to the left and  $3 \cdot N(c_h)$  nodes of  $R'(\hat{Q}^+)$  lying immediately to the right of the top boundary of  $\hat{Q}^+$ . Similarly, define  $\beta(\hat{Q})$  analogously for the bottom boundary of  $\hat{Q}^+$ . We say that the vertices in paths  $\alpha(\hat{Q})$  and  $\beta(\hat{Q})$  are the *portals* of  $\hat{Q}$ . Let  $\{\alpha_i(\hat{Q}) \subseteq \alpha(\hat{Q}) \mid i \in \{1, \dots, z\}\}$  be a set of disjoint intervals in the natural left-to-right ordering with following two properties: (i) there are at least  $3 \cdot N(c_i)$  nodes in the interval  $\alpha_i(\hat{Q})$ ; and (ii) if the square  $\hat{Q}$  is assigned level- $(h+1)$  color  $c_h^i$ , then  $\alpha_i(\hat{Q})$  is the top boundary of  $\hat{Q}^+$ . We define the intervals  $\beta_i(\hat{Q}) \subseteq \beta(\hat{Q})$  corresponding to the bottom boundary of  $\hat{Q}^+$  for each  $1 \leq i \leq z$  in a similar fashion. The sub-paths  $\alpha_i(\hat{Q})$  and  $\beta_i(\hat{Q})$  will act as interfaces to connect child-snakes corresponding to color  $c_h^i$ . Now, we focus on constructing the child-snakes inside  $Q^+$ .

Notice that  $N_h(c_h) \leq \lceil d_h/2\eta^3 \rceil < d_h/\eta^3 \leq d_{h+1}/18$ , since  $\eta \geq 3$ . Recall that for each pair  $\hat{Q}, \tilde{Q} \in \mathcal{D}$  of distinct level- $(h+1)$  squares contained inside  $Q$ ,  $d(\hat{Q}^+, \tilde{Q}^+), d(\hat{Q}^+, \tau(Q^+)), d(\tilde{Q}^+, \tau(Q^+)) \geq d_{h+1}/2 > 9 \cdot N(c_h)$ . Hence, portals of any pair of distinct squares are separated by at least  $3 \cdot N(c_h)$  nodes from each other. Exploiting this spacing, it is now immediate to construct the child-snakes of the following five types (see Figure 5 for illustration).

- The set of child-snakes of type *A* contains, for each child-color  $c_h^i$  of  $c_h$ , a single snake  $\mathcal{Y}^A(c_h^i)$  of width  $3 \cdot N(c_h^i)$  with  $I_Q^i$  as the top boundary of its first corridor and  $\alpha(Q_1^1)$  as the bottom boundary of its last corridor.
- The set of child-snakes of type *B* contains, for each child-color  $c_h^i$  of  $c_h$ , a single snake  $\mathcal{Y}^B(c_h^i)$  of width  $3 \cdot N(c_h^i)$  with  $\beta_i(Q_r^{k_r})$  as the top boundary of its first corridor and  $J_Q^i$  as the bottom boundary of its last corridor.

---

<sup>4</sup>As in the case of sub-snakes, each of the child-snake that we construct, is a snake as defined earlier. We use the term ‘child-snake’ to distinguish them from the final combined sub-snake  $\mathcal{Y}_Q(c_h^i)$  that we construct.

- The set of child-snakes of type  $C$  contains, for each child-color  $c_h^i$  of  $c_h$ , for each  $1 \leq j \leq r$  and  $1 \leq j' < k_j$ , a snake  $\mathcal{Y}_j^{j'}(c_h^i)$  of width  $3 \cdot N(c_h^i)$  with  $\beta_i(Q_j^{j'})$  as the top boundary of its first corridor and  $\alpha(Q_j^{j'+1})$  as the bottom boundary of its last corridor.
- The child-snakes of type  $D$  consists of, for each child color  $c_h^i$  of  $c_h$  and  $1 \leq j < r$ , a snake  $\mathcal{Y}_j^{k_j}(c_h^i)$  of width  $3 \cdot N(c_h^i)$  with  $\beta_i(Q_j^{k_j})$  as the top boundary of its first corridor and  $\alpha_i(Q_j^1)$  as the bottom boundary of its last corridor.
- The child-snakes of type  $E$  consists of, for each child-color  $c_h^i$  of  $c_h$  and square  $Q_j^{j'} \in \mathcal{D}$ , a snake  $\hat{\mathcal{Y}}_j^{j'}(c_h^i)$  of width  $3 \cdot N(c_h^i)$  with  $\alpha_i(Q_j^{j'})$  as the top boundary of its first corridor and  $\beta_i(Q_j^{j'})$  as the bottom boundary of its last corridor. Moreover, if the level- $(h+1)$  color of  $Q_j^{j'}$  is  $c_h^i$ , then  $\hat{\mathcal{Y}}_j^{j'}(c_h^i)$  is a snake contains the corresponding extended square  $Q_j^{j'+}$  as a corridor.

Moreover, we also ensure that the child-snakes corresponding to different colors  $c_h^i$  are disjoint, and for each child color  $c_h^i$  the snakes are composable. Claim 7.5 now follows by considering the snakes

$$\mathcal{Y}_Q(c_h^i) = \mathcal{Y}^A(c_h^i) \oplus \hat{\mathcal{Y}}_1^1(c_h^i) \oplus \mathcal{Y}_1^1(c_h^i) \oplus \hat{\mathcal{Y}}_1^2(c_h^i) \oplus \mathcal{Y}_1^2(c_h^i) \oplus \dots \oplus \mathcal{Y}_r^{k_r-1}(c_h^i) \oplus \hat{\mathcal{Y}}_r^{k_r}(c_h^i) \oplus \mathcal{Y}^B(c_h^i)$$

for all child-colors  $c_h^i$  of  $c_h$ . □

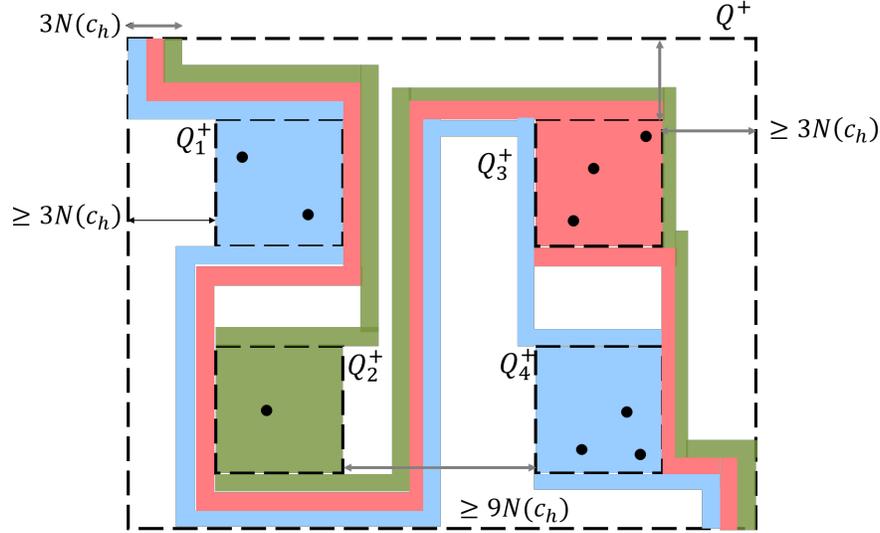


Figure 5: Squares  $Q_1, Q_2, Q_3, Q_4$  are children of level  $h$  square  $Q$  of color  $c_h$ . We label them in this way for simplicity. Here,  $Q_1$  and  $Q_2$ , as well as  $Q_3$  and  $Q_4$  are vertically aligned. The width of the snake responsible to route  $N(c_h)$  demand pairs, containing  $Q^+$  as a corridor, is  $3N(c_h)$ . The figure shows the corridors of snakes corresponding to the child-colors of  $c_h$  inside square  $Q^+$ .

## 8 Finding the Coloring of Squares and a Perfect Set of Demand Pairs

The goal of this section is to prove Theorem 4.9. It would be convenient for us to reformulate the problem slightly differently. Recall that we are given a hierarchical system  $\mathcal{C}$  of colors, and that for all source vertices  $s$  of the demand pairs in  $\tilde{\mathcal{M}}$ , their colorings are fixed: that is, each vertex  $s \in S(\tilde{\mathcal{M}})$  is

assigned, for each level  $1 \leq h \leq \rho$ , some level- $h$  color  $c_h$ . Consider now some demand pair  $(s, t) \in \tilde{\mathcal{M}}$ . For each level  $1 \leq h \leq \rho$ , if  $s$  is assigned the level- $h$  color  $c_h$ , then we will assign the same level- $h$  color  $c_h$  to  $t$ . Note that the same vertex of  $G$  may serve as a destination of several demand pairs; in such a case, we view this vertex as having different copies, each of which is colored according to the coloring of the corresponding source vertex.

For now, we will ignore the set  $S(\tilde{\mathcal{M}})$  of the source vertices of the demand pairs in  $\tilde{\mathcal{M}}$ , and we will focus on the multi-set  $U$  of the destination vertices in  $T(\tilde{\mathcal{M}})$ . Suppose we compute some coloring  $f$  of  $\tilde{\mathcal{H}}$  by  $\mathcal{C}$ . Consider some level- $\rho$  square  $Q$ , and some destination vertex  $t \in U \cap Q$ . Suppose square  $Q$  is assigned the level- $\rho$  color  $c_\rho$ . We say that  $t$  *agrees* with the coloring of  $Q$  iff  $t$  was also assigned the level- $\rho$  color  $c_\rho$ . We can now restate the problem slightly differently. We call the new problem Hierarchical Square Coloring (HSC). In this problem, we are given a hierarchical system  $\tilde{\mathcal{H}} = (\mathcal{Q}_1, \dots, \mathcal{Q}_\rho)$  of squares, a hierarchical system  $\mathcal{C}$  of colors, and a multi-set  $U$  of vertices that lie in  $\bigcup_{Q \in \mathcal{Q}_\rho} Q$ , together with an assignment of a level- $\rho$  color to each vertex in  $U$ . Notice that this assignment implicitly assigns, to each vertex  $v \in U$ , for each level  $1 \leq h < \rho$ , a unique level- $h$  color  $c_h$ , which is the level- $h$  ancestor of the level- $\rho$  color assigned to  $v$ . Our goal is to find a coloring  $f$  of  $\tilde{\mathcal{H}}$  by  $\mathcal{C}$ , and a subset  $U' \subseteq U$  of the vertices, such that:

- For each vertex  $t \in U'$ , if  $t$  belongs to a level- $\rho$  square  $Q$ , then the coloring of  $t$  agrees with the coloring of  $Q$ ; and
- For each level  $1 \leq h \leq \rho$ , for each level- $h$  color  $c_h$ , the total number of vertices in  $U'$  whose level- $h$  color is  $c_h$  is at most  $d_h$ .

The goal is to maximize  $|U'|$ . The main result of this section is a proof of the following theorem.

**Theorem 8.1** *There is an efficient randomized algorithm, that, given as input an instance  $(\tilde{\mathcal{H}}, \tilde{\mathcal{M}})$  of the HSC problem, with high probability returns an  $O(\log^4 n)$ -approximate solution to it.*

We first show that Theorem 4.9 follows from Theorem 8.1. Let  $U' \subseteq U$  be the (multi)-set of the destination vertices computed by the algorithm from Theorem 8.1. As every destination vertex in  $U'$  belongs to a unique demand pair in  $\tilde{\mathcal{M}}$ , set  $U'$  naturally defines a subset  $\tilde{\mathcal{M}}' \subseteq \tilde{\mathcal{M}}$  of demand pairs. It is immediate to verify that  $|\tilde{\mathcal{M}}'| \geq \Omega(\text{OPT}' / \log^4 n)$ . This is since the maximum-cardinality subset  $\tilde{\mathcal{M}}'' \subseteq \tilde{\mathcal{M}}$  of the demand pairs, for which a coloring  $f'$  of  $\tilde{\mathcal{H}}$  by  $\mathcal{C}$  exists, with  $\tilde{\mathcal{M}}''$  being perfect for  $f'$ , naturally defines a feasible solution to HSC, of value  $|\tilde{\mathcal{M}}''|$ . The set  $\tilde{\mathcal{M}}'$  of demand pairs has all properties of a perfect set, except that we are not guaranteed that the source and the destination vertices are all distinct. We rectify this as follows. Fix any level- $\rho$  color  $c_\rho$ . Recall that there are at most  $d_\rho = \eta^3 = 2^{O(\sqrt{\log n})}$  demand pairs  $(s, t) \in \tilde{\mathcal{M}}'$ , such that the level- $\rho$  color of  $s$  is  $c_\rho$ . We discard all such pairs but one from  $\tilde{\mathcal{M}}'$ . It is immediate to verify that  $|\tilde{\mathcal{M}}'| \geq \text{OPT}' / 2^{O(\sqrt{\log n})}$  still holds. We claim that all demand pairs in this final set  $\tilde{\mathcal{M}}'$  have distinct sources and destinations. Indeed, consider any two such distinct pairs  $(s, t), (s', t')$ . Assume first that  $s = s'$ . Let  $c_\rho$  be the level- $\rho$  color of  $s$ . Then  $\tilde{\mathcal{M}}'$  contains two demand pairs whose sources have the same level- $\rho$  color, a contradiction. Assume now that  $s \neq s'$  but  $t = t'$ . Then the level- $\rho$  colors of  $s$  and  $s'$  must be different, and so are the level- $\rho$  colors  $t$  and  $t'$ . However, both  $t$  and  $t'$  belong to the same level- $\rho$  square  $Q_\rho$ , and their color should agree with the level- $\rho$  color of  $Q_\rho$ , which is impossible.

The remainder of this section is dedicated to proving Theorem 8.1. We formulate an LP-relaxation of the problem, and then provide a randomized LP-rounding approximation algorithm for it.

For convenience, for every pair  $1 \leq h < h' \leq \rho$  of levels, for every level- $h$  square  $Q$ , we denote by  $\mathcal{D}_{h'}(Q)$  the set of all level- $h'$  squares that are contained in  $Q$ , and we think of them as level- $h'$

descendants of  $Q$ . Recall that for each level- $h$  color  $c_h$ , we have also denoted by  $\tilde{\chi}_{h'}(c)$  the set of all its descendant-colors that belong to level  $h'$ .

## 8.1 The Linear Program

For every level- $\rho$  square  $Q_\rho \in \mathcal{Q}_\rho$ , and every level- $\rho$  color  $c_\rho \in \chi_\rho$ , we let  $n(Q_\rho, c_\rho)$  denote the number of vertices of  $U \cap V(Q_\rho)$ , whose level- $\rho$  color is  $c_\rho$ .

Our linear program has three types of variables. Fix some level  $1 \leq h \leq \rho$  and a level- $h$  color  $c_h \in \chi_h$ . For every level- $h$  square  $Q_h \in \mathcal{Q}_h$ , we have a variable  $x(Q_h, c_h)$  indicating whether we choose color  $c_h$  for square  $Q_h$ . We also have a global variable  $Y_h(c_h)$ , counting the total number of all vertices in the final solution  $U'$ , whose level- $h$  color is  $c_h$ . Finally, for every level- $\rho$  square  $Q_\rho \in \mathcal{Q}_\rho$  and color  $c_\rho \in \chi_\rho$ , we have a variable  $y_\rho(Q_\rho, c_\rho)$ . Intuitively, if  $Q_\rho$  is assigned the color  $c_\rho$ , then  $y_\rho(Q_\rho, c_\rho)$  is the total number of vertices of  $U' \cap V(Q_\rho)$  in our solution  $U'$  (note that their level- $\rho$  color must be  $c_\rho$ ); otherwise,  $y_\rho(Q_\rho, c_\rho) = 0$ .

The objective function of the LP is:

$$\max \quad \sum_{c_\rho \in \chi_\rho} Y_\rho(c_\rho).$$

We now define the constraints of the LP. The first set of constraints requires that for each level- $\rho$  square  $Q_\rho$ , and each level- $\rho$  color  $c_\rho$ , if  $Q_\rho$  is assigned the color  $c_\rho$ , then  $y(Q_\rho, c_\rho)$  is bounded by  $n(Q_\rho, c_\rho)$ , and otherwise it is 0.

$$y(Q_\rho, c_\rho) \leq n(Q_\rho, c_\rho) \cdot x(Q_\rho, c_\rho) \quad \forall c_\rho \in \chi_\rho, \forall Q_\rho \in \mathcal{Q}_\rho \quad (1)$$

The next set of constraints states that for each level- $\rho$  color  $c_\rho$ , the total number of vertices of  $U'$  whose level- $\rho$  color is  $c_\rho$  is equal to the summation, over all level- $\rho$  squares  $Q_\rho$ , of the number of vertices of  $U' \cap V(Q)$ , whose level- $\rho$  color is  $c_\rho$ .

$$Y_\rho(c_\rho) = \sum_{Q_\rho \in \mathcal{Q}_\rho} y(Q_\rho, c_\rho) \quad \forall c_\rho \in \chi_\rho \quad (2)$$

The next set of constraints states that for every level  $h$  and every level- $h$  color  $c_h$ , the total number of vertices of  $U'$  whose level- $h$  color is  $c_h$  is equal to the total number of vertices of  $U'$  whose level- $\rho$  color is a descendant color of  $c_h$ .

$$Y_h(c_h) = \sum_{c_\rho \in \tilde{\chi}_\rho(c_h)} Y_\rho(c_\rho) \quad \forall 1 \leq h < \rho, \forall c_h \in \chi_h \quad (3)$$

The following set of constraints ensures that for each level  $h$ , and each level- $h$  color  $c_h$ , no more than  $d_h$  vertices of  $U'$  whose level- $h$  color is  $c_h$  are in the solution.

$$Y_h(c_h) \leq d_h \quad \forall 1 \leq h \leq \rho, \forall c_h \in \chi_h \quad (4)$$

The next set of constraints ensures that every level-1 square is assigned some level-1 color.

$$\sum_{c_1 \in \chi_1} x(Q_1, c_1) = 1 \quad \forall Q_1 \in \mathcal{Q}_1 \quad (5)$$

It is immediate to verify that all above constraints are satisfied by any valid integral solution to HSC, so the above LP is indeed a relaxation of HSC.

We now add a few additional constraints to coordinate between different levels, that will be crucial for our LP-rounding algorithm. Consider some level  $1 \leq h < \rho$ , and some level- $h$  square  $Q_h$ . Let  $Q_{h+1}$  be a child-square of  $Q_h$ . If  $Q_h$  is assigned the level- $h$  color  $c_h$ , then  $Q_{h+1}$  must be assigned some child-color of  $c_h$ . The following constraint expresses this requirement.

$$\sum_{c_{h+1} \in \tilde{\chi}_{h+1}(c_h)} x(Q_{h+1}, c_{h+1}) = x(Q_h, c_h) \quad \forall 1 \leq h < \rho; \quad \forall Q_h \in \mathcal{Q}_h; \quad \forall c_h \in \chi_h; \quad \forall Q_{h+1} \in \mathcal{D}_{h+1}(Q_h) \quad (6)$$

The next constraint requires that for every pair  $1 \leq h \leq h' \leq \rho$  of levels, for every level- $h$  square  $Q_h$  and every level- $h$  color  $c_h$ , if  $c_{h'}$  is a level- $h'$  descendant color of  $c_h$ , then the total number of vertices in  $U' \cap V(Q_h)$  that are assigned the level- $h'$  color  $c_{h'}$ , is at most  $d_{h'}$ , if square  $Q_h$  is assigned color  $c_h$ , and it is 0 otherwise (notice that we also require this for  $h = h'$ ):

$$\sum_{Q_\rho \in \mathcal{D}_\rho(Q_h)} \sum_{c_\rho \in \tilde{\chi}_\rho(c_{h'})} y(Q_\rho, c_\rho) \leq d_{h'} \cdot x(Q_h, c_h) \quad \forall 1 \leq h \leq h' \leq \rho; \quad \forall Q_h \in \mathcal{Q}_h; \quad \forall c_h \in \chi_h; \quad \forall c_{h'} \in \tilde{\chi}_{h'}(c_h) \quad (7)$$

(Note that indeed for an integral solution,  $\sum_{Q_\rho \in \mathcal{D}_\rho(Q_h)} \sum_{c_\rho \in \tilde{\chi}_\rho(c_{h'})} y(Q_\rho, c_\rho)$  is the total number of vertices in  $U' \cap V(Q_h)$  that are assigned the level- $h'$  color  $c_{h'}$ .)

Finally, we add non-negativity constraints:

$$x(Q_h, c_h) \geq 0 \quad \forall 1 \leq h \leq \rho; \quad \forall c_h \in \chi_h \quad (8)$$

$$y(Q_\rho, c_\rho) \geq 0 \quad \forall Q_\rho \in \mathcal{Q}_\rho; \quad \forall c_\rho \in \chi_\rho \quad (9)$$

This completes the description of the LP-relaxation. We now proceed to describe the LP-rounding algorithm.

## 8.2 LP-Rounding

For convenience, we denote  $|U| = n$ . Our LP-rounding algorithm proceeds in three stages. In the first stage we perform a simple randomized rounding, level-by-level. We will show that the expected number of vertices that we select to our solution  $U'$  is equal to the LP-solution value. However, it is possible that we select too many vertices of some color. In the second stage, we define, for every level  $h$ , and every level- $h$  color  $c_h$ , a bad event that too many vertices of color  $c_h$  are added to  $U'$ , and show that with high probability none of these bad events happen. If any of the bad events happen, we discard all vertices and return an empty solution. Notice that in this case, we will discard at most  $|U| = n$  vertices. Therefore, if we can prove that the probability of any bad event happening is less than  $1/(2n)$ , then the expected solution cost will remain close to the optimal one. The result of the second stage is a set  $U'$ , whose expected cardinality is close to the LP-solution cost, and this solution is almost feasible: namely, for every level  $h$  and every level- $h$  color  $c_h$ , the number of vertices in  $U'$  whose level- $h$  color is  $c_h$  is at most  $d_h \cdot \text{poly} \log(n)$ . In the third stage, we turn the solution  $U'$  into a feasible one, at the cost of losing a polylogarithmic factor on its cardinality.

For convenience, we let  $Q_0$  denote the whole grid  $\tilde{G}$ , and we define a unique level-0 color  $c_0$ ; every level-1 color  $c_1 \in \chi_1$  is a child-color of  $c_0$ . We assume that  $Q_0$  is assigned the level-0 color  $c_0$ .

## Stage 1: Randomized Rounding

We perform  $\rho$  iterations, where in iteration  $h$  we settle the colors of all level- $h$  squares, by suitably modifying the LP-solution. We will maintain the following invariant:

P1. For all  $1 \leq h \leq \rho$ , at the beginning of iteration  $h$ , the current LP-solution  $(x', y')$  has the following properties. For every level- $(h-1)$  square  $Q_{h-1}$ , if  $c_{h-1}$  is the level- $(h-1)$  color assigned to  $Q_{h-1}$ , then for every level  $h-1 < h' \leq \rho$ , for every level- $h'$  descendant-square  $Q_{h'} \in \mathcal{D}_{h'}(Q_{h-1})$ , and every level- $h'$  color  $c_{h'} \in \chi_{h'}$ :

- If  $c_{h'}$  is not a descendant-color of  $c_{h-1}$ , then  $x'(Q_{h'}, c_{h'}) = 0$  (and if  $h' = \rho$  then  $y'(Q_{h'}, c_{h'}) = 0$ );
- Otherwise,  $x'(Q_{h'}, c_{h'}) = x(Q_{h'}, c_{h'})/x(Q_{h-1}, c_{h-1})$ . Moreover, if  $h' = \rho$ , then  $y'(Q_{h'}, c_{h'}) = y(Q_{h'}, c_{h'})/x(Q_{h-1}, c_{h-1})$ .

(Here, the  $x$  and the  $y$ -variables are the ones from the original LP-solution).

In order to perform the first iteration, each level-1 square  $Q_1 \in \mathcal{Q}_1$  chooses one of the colors  $c_1 \in \chi_1$ , where each color  $c_1$  is chosen with probability  $x(Q_1, c_1)$  (recall that by Constraint (5),  $\sum_{c_1 \in \chi_1} x(Q_1, c_1) = 1$ ). Assume now that we have chosen color  $c_1$  for square  $Q_1$ . We update the LP-solution. For every level  $2 \leq h \leq \rho$ , every level- $h$  square  $Q_h \in \mathcal{D}_h(Q_1)$ , and every level- $h$  color  $c_h$ , we do the following: if  $c_h$  is not a descendant color of  $c_1$ , then we set  $x'(Q_h, c_h) = 0$  (if  $h = \rho$ , then we also set  $y'(Q_h, c_h) = 0$ ). Otherwise, we set  $x'(Q_h, c_h) = x(Q_h, c_h)/x(Q_1, c_1)$  (similarly, if  $h = \rho$ , then we also set  $y'(Q_h, c_h) = y(Q_h, c_h)/x(Q_1, c_1)$ ). We then continue to the next iteration.

The invariant clearly holds at the beginning of the second iteration. Assume now that it holds at the beginning of iteration  $h$ . The iteration is executed as follows. Consider some level- $h$  square  $Q_h$ . Suppose its parent square is  $Q_{h-1}$ , and it was assigned color  $c_{h-1}$ . Then Constraint (6), together with Invariant (P1) ensures that:

$$\sum_{c_h \in \tilde{\chi}_h(c_{h-1})} x'(Q_h, c_h) = 1.$$

Square  $Q_h$  chooses one of the child-colors  $c_h$  of  $c_{h-1}$ , where color  $c_h$  is chosen with probability  $x'(Q_h, c_h)$ . Assume that we have assigned color  $c_h$  to square  $Q_h$ . We now update the LP-solution. Consider some level  $h < h' \leq \rho$ , some level- $h'$  square  $Q_{h'} \in \mathcal{D}_{h'}(Q_h)$  that is a descendant of  $Q_h$ , and some level- $h'$  color  $c_{h'} \in \chi_{h'}$ . Assume first that if  $c_{h'}$  is not a descendant color of  $c_h$ . Then we set  $x''(Q_{h'}, c_{h'}) = 0$  (if  $h' = \rho$ , then we also set  $y''(Q_{h'}, c_{h'}) = 0$ ). Assume now that  $c_{h'}$  is a descendant color of  $c_h$ . Let  $Q_{h-1}$  be the parent-square of  $Q_h$ , and let  $c_{h-1}$  be the parent-color of  $c_h$ . Then we set:

$$x''(Q_{h'}, c_{h'}) = \frac{x'(Q_{h'}, c_{h'})}{x'(Q_h, c_h)} = \frac{x(Q_{h'}, c_{h'})/x(Q_{h-1}, c_{h-1})}{x(Q_h, c_h)/x(Q_{h-1}, c_{h-1})} = \frac{x(Q_{h'}, c_{h'})}{x(Q_h, c_h)}.$$

(We have used the fact that Invariant (P1) held at the beginning of the current iteration). Similarly, if  $h = \rho$ , then we set:

$$y''(Q_{h'}, c_{h'}) = \frac{y'(Q_{h'}, c_{h'})}{x'(Q_h, c_h)} = \frac{y(Q_{h'}, c_{h'})/x(Q_{h-1}, c_{h-1})}{x(Q_h, c_h)/x(Q_{h-1}, c_{h-1})} = \frac{y(Q_{h'}, c_{h'})}{x(Q_h, c_h)}.$$

We then replace solution  $(x', y')$  with the new solution  $(x'', y'')$ . This completes the description of iteration  $h$ . Observe that Invariant (P1) continues to hold at the end of the current iteration.

The final iteration  $\rho$  is executed as follows. For every level- $\rho$  square  $Q_\rho \in \mathcal{Q}_\rho$ , we select a level- $\rho$  color  $c_\rho$  for  $Q_\rho$  exactly as before, but we do not update the LP-solution. If  $\frac{y'(Q_\rho, c_\rho)}{x'(Q_\rho, c_\rho)} \geq 1$ , then we add  $\left\lceil \frac{y'(Q_\rho, c_\rho)}{x'(Q_\rho, c_\rho)} \right\rceil$  distinct vertices of  $Q_\rho \cap U$ , whose level- $\rho$  color is  $c_\rho$  to our solution  $U'$ . This quantity is guaranteed to be bounded by  $n(Q_\rho, c_\rho)$ , due to constraint (1) and Invariant (P1). Otherwise, with probability  $\frac{y'(Q_\rho, c_\rho)}{x'(Q_\rho, c_\rho)}$ , we add one vertex of  $Q_\rho \cap U$  whose level- $\rho$  color is  $c_\rho$  to our solution, and add 0 such vertices with the remaining probability. (Recall that  $\frac{y'(Q_\rho, c_\rho)}{x'(Q_\rho, c_\rho)} = \frac{y(Q_\rho, c_\rho)}{x(Q_\rho, c_\rho)}$  from Invariant (P1) — we use this fact later.) The following theorem concludes the analysis of Stage 1.

**Theorem 8.2** *The expected number of vertices added to  $U'$  at Stage 1 is at least  $\sum_{c_\rho \in \chi_\rho} Y_\rho(c_\rho)$ .*

**Proof:** Consider some color  $c_\rho \in \chi_\rho$ . Recall that  $Y_\rho(c_\rho) = \sum_{Q_\rho \in \mathcal{Q}_\rho} y(Q_\rho, c_\rho)$  due to Constraint (2). For each level- $\rho$  square  $Q_\rho$  we define a new random variable  $z(Q_\rho, c_\rho)$  to be the number of vertices of  $Q_\rho$ , whose level- $\rho$  color is  $c_\rho$ , that are added to the solution  $U'$  (this number may only be non-zero if  $Q_\rho$  is assigned the color  $c_\rho$ ). It is now enough to prove that  $\mathbf{E}[z(Q_\rho, c_\rho)] \geq y(Q_\rho, c_\rho)$ . From now on we fix a square  $Q_\rho$  and a color  $c_\rho$  and prove the above inequality for them. We use the following claim:

**Claim 8.3** *The probability that square  $Q_\rho$  is assigned color  $c_\rho$  is  $x(Q_\rho, c_\rho)$ .*

**Proof:** For all  $1 \leq h < \rho$ , let  $Q_h$  be the level- $h$  ancestor-square of  $Q_\rho$ , and let  $c_h$  the level- $h$  ancestor-color of  $c_\rho$ . For all  $1 \leq h \leq \rho$ , let  $\mathcal{E}_h$  be the event that square  $Q_h$  is assigned color  $c_h$ . Then:

$$\Pr[\mathcal{E}_1] = x(Q_1, c_1).$$

We now fix some  $1 < h \leq \rho$ , and analyze the probability of event  $\mathcal{E}_h$ , assuming that event  $\mathcal{E}_{h-1}$  has happened. Let  $x', y'$  denote the LP-solution at the beginning of iteration  $h$ . Then:

$$\Pr[\mathcal{E}_h \mid \mathcal{E}_{h-1}] = x'(Q_h, c_h) = \frac{x(Q_h, c_h)}{x(Q_{h-1}, c_{h-1})}.$$

Therefore:

$$\begin{aligned} \Pr[\mathcal{E}_\rho] &= \Pr[\mathcal{E}_\rho \mid \mathcal{E}_{\rho-1}] \cdot \Pr[\mathcal{E}_{\rho-1} \mid \mathcal{E}_{\rho-2}] \cdots \Pr[\mathcal{E}_2 \mid \mathcal{E}_1] \cdot \Pr[\mathcal{E}_1] \\ &= \frac{x(Q_\rho, c_\rho)}{x(Q_{\rho-1}, c_{\rho-1})} \cdot \frac{x(Q_{\rho-1}, c_{\rho-1})}{x(Q_{\rho-2}, c_{\rho-2})} \cdots \frac{x(Q_2, c_2)}{x(Q_1, c_1)} \cdot x(Q_1, c_1) \\ &= x(Q_\rho, c_\rho). \end{aligned}$$

□

We are now ready to complete the proof of Theorem 8.2. Assume first that  $\frac{y(Q_\rho, c_\rho)}{x(Q_\rho, c_\rho)} \geq 1$ . Then, if  $Q_\rho$  is assigned color  $c_\rho$  (which happens with probability  $x(Q_\rho, c_\rho)$ ), we select  $\left\lceil \frac{y(Q_\rho, c_\rho)}{x(Q_\rho, c_\rho)} \right\rceil$  vertices of  $Q_\rho$ , whose level- $\rho$  color is  $c_\rho$  to the solution. Therefore,  $\mathbf{E}[z(Q_\rho, c_\rho)] \geq y(Q_\rho, c_\rho)$ .

Assume now that  $\frac{y(Q_\rho, c_\rho)}{x(Q_\rho, c_\rho)} < 1$ . Then, if  $Q_\rho$  is assigned color  $c_\rho$  (which again happens with probability  $x(Q_\rho, c_\rho)$ ), we add one vertex of  $Q_\rho$ , whose level- $\rho$  color is  $c_\rho$ , with probability  $\frac{y(Q_\rho, c_\rho)}{x(Q_\rho, c_\rho)}$ . Clearly, in this case,  $\mathbf{E}[z(Q_\rho, c_\rho)] = y(Q_\rho, c_\rho)$ .  $\square$

## Stage 2: Ensuring that the Solution is Almost Feasible

In this stage, for each level  $1 \leq h \leq \rho$  and for each level- $h$  color  $c_h \in \chi_h$ , we inspect the number of vertices in the solution  $U'$  that we constructed at stage 1, whose level- $h$  color is  $c_h$ . If this number is greater than  $64d_h \log^3 n$ , then we say that color  $c_h$  has failed. In that case, we return an empty solution. We think of this as deleting all vertices from our solution, so overall at most  $n$  vertices are deleted. We will use the following standard Chernoff bound (see e.g. [DP09]).

**Theorem 8.4** *Let  $\{X_1, \dots, X_r\}$  be a collection of independent random variables taking values in  $[0, 1]$ , and let  $X = \sum_i X_i$ . Denote  $\mu = \mathbf{E}[X] = \sum_i \mathbf{E}[X_i]$ . Then for all  $0 < \epsilon < 1$ :*

$$\Pr[X > (1 + \epsilon)\mu] \leq e^{-\epsilon^2\mu/3},$$

and

$$\Pr[X < (1 - \epsilon)\mu] \leq e^{-\epsilon^2\mu/2}.$$

**Corollary 8.5** *Let  $X_1, \dots, X_r$  be a collection of independent random variables taking values in  $[0, d]$  for some integer  $d > 0$ , and let  $X = \sum_i X_i$ . Denote  $\mu = \mathbf{E}[X] = \sum_i \mathbf{E}[X_i]$ . Then for all  $0 < \epsilon < 1$ :*

$$\Pr[X > (1 + \epsilon)\mu] \leq e^{-\epsilon^2\mu/(3d)},$$

and

$$\Pr[X < (1 - \epsilon)\mu] \leq e^{-\epsilon^2\mu/(2d)}.$$

(The Corollary is obtained by replacing each variable  $X_i$  with a variable  $X'_i = X_i/d$  and then applying Theorem 8.4). Following is the main theorem for the analysis of Stage 2.

**Theorem 8.6** *Fix some level  $1 \leq h \leq \rho$  and some level- $h$  color  $c_h \in \chi_h$ . The probability that color  $c_h$  fails is at most  $1/(2n^3)$ .*

Note that there are at most  $n$  colors at each level, and  $\rho = O(\sqrt{\log n})$  levels. Therefore, using the union bound, we obtain the following corollary.

**Corollary 8.7** *The probability that any color fails is at most  $1/(2n)$ .*

We now turn to prove Theorem 8.6.

**Proof of Theorem 8.6.** We assume that the level  $h$  and the color  $c_h$  are fixed. For consistency, we define a fractional solution  $(x', y')$  obtained at the end of iteration  $\rho$ , as follows. Consider a level- $\rho$  square  $Q_\rho$ , and let  $c_\rho$  be the level- $\rho$  color that was assigned to it at the end of iteration  $\rho$ . Then we set  $x'(Q_\rho, c_\rho) = 1$ , and  $y'(Q_\rho, c_\rho) = y(Q_\rho, c_\rho)/x(Q_\rho, c_\rho)$ . For every other level- $\rho$  color  $c'_\rho$ , we set  $x'(Q_\rho, c'_\rho) = 0$ , and  $y'(Q_\rho, c'_\rho) = 0$ .

For simplicity, we denote  $\chi' = \tilde{\chi}_\rho(c_h)$  - the set of all descendant-colors of  $c_h$  at level  $\rho$ . For every level  $1 \leq h' \leq \rho$ , we define a random variable  $N_{h'}$ , whose intuitive meaning is the number of color- $c_h$  vertices in the fractional solution, that is obtained at the end of iteration  $h'$ .

Formally, let  $(x', y')$  be the solution obtained at the end of iteration  $h'$ . Then:

$$N_{h'} = \sum_{Q_\rho \in \mathcal{Q}_\rho} \sum_{c_\rho \in \mathcal{X}'} y'(Q_\rho, c_\rho).$$

For all  $1 \leq h' \leq \rho$ , we define a bad event  $\mathcal{E}_{h'}$  to be the event that  $N_{h'} > d_h \cdot 16(1 + 1/\log n)^{h'} \log^3 n$ . The following lemma is central to the analysis.

**Lemma 8.8**  $\Pr[\mathcal{E}_\rho] \leq \frac{1}{4n^3}$ .

If  $\mathcal{E}_\rho$  does not happen, then, if we denote by  $(x', y')$  the solution obtained at the end of iteration  $\rho$ :

$$N_\rho = \sum_{Q_\rho \in \mathcal{Q}_\rho} \sum_{c_\rho \in \mathcal{X}'} y'(Q_\rho, c_\rho) < 16d_h \left(1 + \frac{1}{\log n}\right)^\rho \log^3 n < 32d_h \log^3 n,$$

since  $\rho \leq \sqrt{\log n}$ .

Recall that for every square  $Q_\rho \in \mathcal{Q}_\rho$ , if  $y'(Q_\rho, c_\rho) \geq 1$ , then we add  $\lceil y'(Q_\rho, c_\rho) \rceil$  vertices of  $Q_\rho$ , whose level- $\rho$  color is  $c_\rho$  to the solution. Otherwise, with probability  $y'(Q_\rho, c_\rho)$ , we add only one such vertex to our solution. Overall, from Theorem 8.4, it is immediate to verify that, if event  $\mathcal{E}_\rho$  does not happen, then with probability at least  $(1 - 1/n^4)$ , the number of vertices in  $U'$ , whose level- $h$  color is  $c_h$  is at most  $2N_\rho \leq 64d_h \log^3 n$ . It now remains to prove Lemma 8.8.

**Proof of Lemma 8.8.** For convenience, we also define event  $\mathcal{E}_0$  that  $N_0 > d_h$ , which, from Constraints (3) and (4) happens with probability 0. Clearly,

$$\begin{aligned} \Pr[\mathcal{E}_\rho] &\leq \Pr[\mathcal{E}_\rho \mid \neg\mathcal{E}_{\rho-1}] \cdot \Pr[\neg\mathcal{E}_{\rho-1}] + \Pr[\mathcal{E}_{\rho-1}] \\ &\leq \Pr[\mathcal{E}_\rho \mid \neg\mathcal{E}_{\rho-1}] + \Pr[\mathcal{E}_{\rho-1} \mid \neg\mathcal{E}_{\rho-2}] \cdot \Pr[\neg\mathcal{E}_{\rho-2}] + \Pr[\mathcal{E}_{\rho-2}] \\ &\quad \vdots \\ &\leq \sum_{h'=1}^{\rho} \Pr[\mathcal{E}_{h'} \mid \neg\mathcal{E}_{h'-1}]. \end{aligned}$$

We now analyze  $\Pr[\mathcal{E}_{h'} \mid \neg\mathcal{E}_{h'-1}]$ . The following claim will finish the proof of the lemma.

**Claim 8.9** For all  $1 \leq h' \leq \rho$ ,  $\Pr[\mathcal{E}_{h'} \mid \neg\mathcal{E}_{h'-1}] \leq 1/(4n^4)$ .

**Proof:** Fix some  $1 \leq h' \leq \rho$ . Let  $(x', y')$  be the solution at the beginning of iteration  $h'$ , and let  $(x'', y'')$  be the solution at the end of iteration  $h'$ . We will repeatedly use Constraint (7) that we restate here; we slightly change the indexing to avoid confusion with the variables  $h$  and  $h'$  that we use here.

The constraint states that for all pairs  $1 \leq \tilde{h} \leq \tilde{h}' \leq \rho$  of levels, for every level- $\tilde{h}$  square  $Q_{\tilde{h}} \in \mathcal{Q}_{\tilde{h}}$  and for every pair of a level- $\tilde{h}$  color  $c_{\tilde{h}}$  and a level- $\tilde{h}'$  color  $c_{\tilde{h}'}$ , that is a descendant-color of  $c_{\tilde{h}}$ , the following inequality holds:

$$\sum_{Q_\rho \in \mathcal{D}_\rho(Q_{\tilde{h}})} \sum_{c_\rho \in \tilde{\mathcal{X}}_\rho(c_{\tilde{h}'})} y(Q_\rho, c_\rho) \leq d_{\tilde{h}'} \cdot x(Q_{\tilde{h}}, c_{\tilde{h}}). \quad (10)$$

Intuitively, if  $Q_{\tilde{h}}$  is assigned the color  $c_{\tilde{h}}$ , then at most  $d_{\tilde{h}}$  vertices of  $Q_{\tilde{h}}$  may be assigned the level- $\tilde{h}'$  color  $c_{\tilde{h}'}$ ; otherwise, none of them can. Notice that due to Invariant (P1), the above constraint remains valid with respect to the solution  $(x', y')$ , if  $\tilde{h} \geq h'$ .

Consider some level- $h'$  square  $Q_{h'} \in \mathcal{Q}_{h'}$ . We assume that its parent-square is  $Q_{h'-1}$ , and that it was assigned some color  $c_{h'-1}$ . Recall that  $(x', y')$  is the LP-solution at the beginning of iteration  $h'$ , and  $(x'', y'')$  is the LP-solution at the end of iteration  $h'$ . We define:

$$M'(Q_{h'}) = \sum_{Q_\rho \in \mathcal{D}_\rho(Q_{h'})} \sum_{c_\rho \in \mathcal{X}'} y'(Q_\rho, c_\rho),$$

and we define  $M''(Q_{h'})$  similarly with respect to  $y''$ . Intuitively,  $M'(Q_{h'})$  is the number of vertices of  $Q_{h'}$  that are (possibly fractionally) assigned the level- $h$  color  $c_h$  at the beginning of iteration  $h'$ , while  $M''(Q_{h'})$  reflects the same quantity at the end of iteration  $h'$ . The values of variables  $M'(Q_{h'})$  are fixed at the beginning of iteration  $h'$ , while the values  $M''(Q_{h'})$  are random variables. Then:

$$N_{h'-1} = \sum_{Q_{h'} \in \mathcal{Q}_{h'}} M'(Q_{h'}),$$

and since we assume that event  $\mathcal{E}_{h'-1}$  did not happen,  $N_{h'-1} \leq d_h \cdot 16(1 + 1/\log n)^{h'-1} \log^3 n$ . Similarly,

$$N_{h'} = \sum_{Q_{h'} \in \mathcal{Q}_{h'}} M''(Q_{h'}),$$

and it is enough to prove that the probability that  $N_{h'} > 16(1 + 1/\log n)^{h'} \log^3 n$  is less than  $1/(4n^4)$ .

We consider two cases, depending on whether  $h > h'$  holds.

**Case 1:**  $h > h'$ . Let  $c_{h'}$  be the unique ancestor-color of color  $c_h$ , that belongs to level  $h'$ . Using Constraint (10) with  $\tilde{h} = h'$  and  $\tilde{h}' = h$ , we get that of each such level- $h'$  square  $Q_{h'}$ :

$$M'(Q_{h'}) \leq d_h x'(Q_{h'}, c_{h'}).$$

If  $Q_{h'}$  is assigned the color  $c_{h'}$  (which happens with probability at most  $x'(Q_{h'}, c_{h'})$ ), then  $M''(Q_{h'}) = M'(Q_{h'})/x'(Q_{h'}, c_{h'}) \leq d_h$ ; otherwise  $M''(Q_{h'}) = 0$ . Therefore,  $\{M''(Q_{h'})\}_{Q_{h'} \in \mathcal{Q}_{h'}}$  is a collection of independent random variables taking values between 0 and  $d_h$ , and:

$$\mathbf{E}[N_{h'}] = \sum_{Q_{h'} \in \mathcal{Q}_{h'}} \mathbf{E}[M''(Q_{h'})] = \sum_{Q_{h'} \in \mathcal{Q}_{h'}} M'(Q_{h'}) = N_{h'-1} \leq d_h \cdot 16(1 + 1/\log n)^{h'-1} \log^3 n.$$

Let  $\mu = d_h \cdot 16(1 + 1/\log n)^{h'-1} \log^3 n$ . From Corollary 8.5:

$$\mathbf{Pr}[N_{h'} \geq (1 + 1/\log n)\mu] \leq e^{-\mu/(3d_h \log^2 n)} \leq e^{-5 \log n} \leq 1/(4n^4),$$

as required

**Case 2:**  $h \leq h'$ . Let  $Q_{h'} \in \mathcal{Q}_{h'}$  be any level- $h'$  square, and let  $c_{h'} \in \tilde{\chi}_{h'}(c_h)$  be any level- $h'$  descendant-color of  $c_h$ . Using Constraint (10) with  $\tilde{h} = \tilde{h}' = h'$ , we get that:

$$\sum_{Q_\rho \in \mathcal{D}_\rho(Q_{h'})} \sum_{c_\rho \in \tilde{\chi}_\rho(c_{h'})} y'(Q_\rho, c_\rho) \leq d_{h'} \cdot x'(Q_{h'}, c_{h'}).$$

Denote the left-hand-side by this inequality by  $M'(Q_{h'}, c_{h'})$ , so:

$$M'(Q_{h'}, c_{h'}) = \sum_{Q_\rho \in \mathcal{D}_\rho(Q_{h'})} \sum_{c_\rho \in \tilde{\chi}_\rho(c_{h'})} y'(Q_\rho, c_\rho),$$

and

$$M''(Q_{h'}) = \sum_{c_{h'} \in \tilde{\chi}_{h'}(c_h)} M'(Q_{h'}, c_{h'}).$$

For each color  $c_{h'} \in \tilde{\chi}_{h'}(c_h)$ , if  $Q_{h'}$  is assigned the color  $c_{h'}$  (which happens with probability  $x'(Q_{h'}, c_{h'})$ ), then  $M''(Q_{h'}) = M'(Q_{h'}, c_{h'})/x'(Q_{h'}, c_{h'}) \leq d_{h'} \leq d_h$ . If none of the colors in  $\tilde{\chi}_{h'}(c_h)$  are assigned to  $Q_{h'}$ , then  $M''(Q_{h'}) = 0$ . Therefore, the expectation of  $M''(Q_{h'})$  is:

$$\mathbf{E} [M''(Q_{h'})] = \sum_{c_{h'} \in \tilde{\chi}_{h'}(c_h)} M'(Q_{h'}, c_{h'}) = M'(Q_{h'}).$$

Overall, we conclude that variables in set  $\{M''(Q_{h'})\}_{Q_{h'} \in \mathcal{Q}_{h'}}$  are independent random variables, taking values in  $[0, d_h]$ , and that:

$$\mathbf{E} [N_{h'}] = \sum_{Q_{h'} \in \mathcal{Q}_{h'}} \mathbf{E} [M''(Q_{h'})] = \sum_{Q_{h'} \in \mathcal{Q}_{h'}} M'(Q_{h'}) = N_{h'-1} \leq d_h \cdot 16(1 + 1/\log n)^{h'-1} \log^3 n.$$

Let  $\mu = d_h \cdot 16(1 + 1/\log n)^{h'-1} \log^3 n$ . From Corollary 8.5:

$$\Pr [N_{h'} \geq (1 + 1/\log n)\mu] \leq e^{-\mu/(3d_h \log^2 n)} \leq e^{-5 \log n} \leq 1/(4n^4).$$

□

□

□

### Stage 3: Turning the Solution into a Feasible One

If any of the colors fail, then we return an empty solution. Assume now that no color fails. Let  $U' \subseteq U$  be the set of all vertices chosen by the current solution.

**Claim 8.10** *There is an efficient algorithm to compute a subset  $U'' \subseteq U'$  of vertices, with  $|U''| \geq \frac{|U'|}{256 \log^4 n}$ , such that  $U''$  is a feasible solution to the HSC problem.*

**Proof:** If, for every level  $h$ , for every level- $h$  color  $c_h$ , at most  $d_h$  vertices of  $U'$  are assigned the level- $h$  color  $c_h$ , then we return the set  $U'$  of vertices, which is a feasible solution to the HSC problem instance.

Otherwise, denote  $r = |U'|$ , and consider an ordering  $(u_1, u_2, \dots, u_r)$  of the vertices of  $U'$ , that has the following property: for every level  $1 \leq h \leq \rho$ , for every level- $h$  color  $c_h \in \chi_h$ , the vertices of  $U'$

whose level- $h$  color is  $c_h$  appear consecutively in this ordering. Due to the nested definition of the colors, it is easy to see that such an ordering exists (in particular, we can use the ordering induced by the ordering of the original source vertices on  $R^*$ ). We now let  $U''$  contain all vertices  $u_i \in U'$ , where  $i = 1 \pmod{\lceil 128 \log^4 n \rceil}$ . It is now immediate to verify that  $|U''| \geq \frac{|U'|}{256 \log^4 n}$ , and that it is a feasible solution to the HSC problem.  $\square$

Let  $\text{OPT}_{\text{LP}}$  denote the value of the optimal LP solution. Let  $N_1$  be the total number of vertices that belong to the solution at the end of phase 1, so  $\mathbf{E}[N_1] = \text{OPT}_{\text{LP}}$ . Let  $N_2$  be the total number of vertices that were deleted during the second phase. Then with probability at least  $(1 - 1/(2n))$ ,  $N_2 = 0$ , and with the remaining probability,  $N_2 \leq n$ . Therefore,  $\mathbf{E}[N_2] \leq 1/2$ . Finally, let  $N_3$  denote the cardinality of the final set  $U''$  of vertices that our algorithm returns. Then:

$$N_3 \geq \frac{N_1 - N_2}{256 \log^4 n},$$

and therefore,  $\mathbf{E}[N_3] \geq \frac{\text{OPT}_{\text{LP}}}{512 \log^4 n}$ . Since  $U''$  is a feasible solution to the HSC problem,  $N_3 \leq \text{OPT}_{\text{LP}}$ .

**Claim 8.11**  $\Pr \left[ N_3 \geq \frac{\text{OPT}_{\text{LP}}}{1024 \log^4 n} \right] \geq \frac{1}{1024 \log^4 n}$

**Proof:** Assume otherwise. Then:

$$\mathbf{E}[N_3] < \left( \frac{1}{1024 \log^4 n} \cdot \text{OPT}_{\text{LP}} \right) + \frac{\text{OPT}_{\text{LP}}}{1024 \log^4 n} < \left( \frac{\text{OPT}_{\text{LP}}}{512 \log^4 n} \right),$$

a contradiction.  $\square$

We run the above algorithm  $c \log^5 n$  times independently (for some large constant  $c$ ), and return a solution of largest cardinality across all runs. The probability that the solution value is less than  $\text{OPT}_{\text{LP}}/(1024 \log^4 n)$  is then bounded by  $1/\text{poly}(n)$ . This concludes the proof of Theorem 8.1.

## 9 Approximating Restricted NDP-Grid

In this section we complete the proof of Theorem 1.1. Suppose we are given an instance  $(G, \mathcal{M})$  of Restricted NDP-Grid. We assume that we know the value  $\text{OPT}$  of the optimal solution to this instance, by going over all such possible choices, and running the algorithm on each of them. Let  $\Gamma_1, \Gamma_2, \Gamma_3$  and  $\Gamma_4$  denote the four boundary edges of the grid  $G$ , and let  $\Gamma = \bigcup_{q=1}^4 \Gamma_q$  be its full boundary. For every destination vertex  $t \in T(\mathcal{M})$ , we denote by  $\tilde{t}$  the vertex of  $\Gamma$  minimizing the distance  $d(t, \tilde{t})$ , breaking ties arbitrarily. Recall that, given any subset  $\mathcal{M}' \subseteq \mathcal{M}$  of the demand pairs, we denoted by  $S(\mathcal{M}')$  and  $T(\mathcal{M}')$  the sets of all vertices that serve as the sources and the destinations of the demand pairs in  $\mathcal{M}'$ , respectively. We denote by  $\tilde{T}(\mathcal{M}') = \{\tilde{t} \mid t \in T(\mathcal{M}')\}$ .

Given two disjoint sub-paths  $\pi, \pi' \subseteq \Gamma$  (that we call intervals), and a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of demand pairs, we let  $\mathcal{M}'_{\pi, \pi'} \subseteq \mathcal{M}'$  contain all demand pairs  $(s, t)$  with  $s \in \pi$  and  $\tilde{t} \in \pi'$ .

### 9.1 Special Instances

It will be convenient for us to define special sub-instances of instance  $(G, \mathcal{M})$ , that have a specific structure. We start by defining interesting pairs of intervals.

**Definition.** Let  $I, I' \subseteq \Gamma$  be two disjoint intervals, and let  $d > 0$  be an integer. We say that  $(I, I')$  is a  $d$ -interesting pair of intervals, iff:

- interval  $I'$  is contained in a single boundary edge of  $G$ , and every vertex of  $I'$  is within distance at least  $16d$  from each of the remaining three boundary edges of  $G$ ;
- interval  $I$  is contained in a single boundary edge of  $G$  (possibly the same as  $I'$ );
- $d \leq |V(I')| \leq \sqrt{n}/2$ , and  $d(I, I') \geq 16d$ .

**Definition.** Suppose we are given an integer  $d > 0$ , and a pair  $(I, I')$  of disjoint intervals of  $\Gamma$ . Assume further that we are given a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of the demand pairs, and consider the sub-instance  $(G, \mathcal{M}')$  of the problem, defined over the same graph  $G$ , with the set  $\mathcal{M}'$  of demand pairs. We say that  $(G, \mathcal{M}')$  is a valid  $(I, I', d)$ -instance iff: (i)  $S(\mathcal{M}') \subseteq I$ ; (ii)  $\tilde{T}(\mathcal{M}') \subseteq I'$ ; and (iii) for each  $t \in T(\mathcal{M}')$ ,  $d \leq d(t, \tilde{t}) < 2d$ . If, additionally,  $(I, I')$  is a  $d$ -interesting pair of intervals, then we say that  $(G, \mathcal{M}')$  is a perfect  $(I, I', d)$ -instance.

The main idea of the algorithm is to compute a partition of the set  $\mathcal{M}$  of the demand pairs into subsets  $\mathcal{M}_1, \dots, \mathcal{M}_z$ , where each set  $\mathcal{M}_i$  is a perfect  $(I, I', d)$ -instance for some pair  $(I, I')$  of intervals. We will then solve the problem defined by each such sub-instance separately, by first casting it as a special case in which the destinations lie far from the grid boundary, and then employing Theorem 4.1. We now define the modified instances, to which Theorem 4.1 will eventually be applied.

## 9.2 Modified Instances

Assume that we are given an integer  $d > 0$ , and a  $d$ -interesting pair  $(I, I')$  of intervals, together with a perfect  $(I, I', d)$ -sub-instance  $(G, \mathcal{M}')$  of  $(G, \mathcal{M})$ . We define a corresponding *modified instance*  $(G', \mathcal{M}'')$ . The underlying graph  $G'$  will be an appropriately chosen sub-grid of  $G$ , and each demand pair in the new set  $\mathcal{M}''$  will correspond to a unique demand pair in  $\mathcal{M}'$ .

In order to define the grid  $G'$ , we assume without loss of generality that  $I'$  is a sub-path of the bottom boundary edge of the grid. Let  $\mathcal{W}''$  be the subset of the columns of  $G$  that intersect  $I'$ , and let  $\mathcal{R}''$  be the set of  $4d$  bottommost rows of  $G$ . Let  $G'' \subseteq G$  be the sub-grid of  $G$  spanned by the set  $\mathcal{W}''$  of columns and the set  $\mathcal{R}''$  of rows. In order to obtain the final graph  $G'$ , we add to  $\mathcal{W}''$   $4d$  columns lying immediately to its left and  $4d$  consecutive columns lying immediately to its right, denoting the resulting set of columns by  $\mathcal{W}'$ . We also add to  $\mathcal{R}''$  a set of  $|\mathcal{W}'| - 4d$  rows lying immediately above  $\mathcal{R}''$ , obtaining a set  $\mathcal{R}'$  of rows of  $G$ . Graph  $G'$  is the sub-grid of  $G$  spanned by the columns of  $\mathcal{W}'$  and the rows of  $\mathcal{R}'$ . Notice that  $G'$  is a square grid.

The set  $\mathcal{M}''$  of demand pairs is constructed as follows. For each demand pair  $(s, t) \in \mathcal{M}'$ , we add a demand pair  $(s', t')$  to  $\mathcal{M}''$ . Vertex  $t'$  is mapped to the same location as vertex  $t$  in  $G'$ . Vertex  $s'$  is mapped to one of the vertices on the top row of  $G'$ , using the following procedure. Let  $X$  denote the set of  $|S(\mathcal{M}')|$  leftmost vertices on the top row of  $G'$ . Let  $v$  be the bottom left corner of the grid  $G'$ . The traversal of  $\Gamma(G)$  in the clock-wise direction, starting from  $v$ , defines an ordering  $\pi$  of the vertices in  $S(\mathcal{M}')$ . Similarly, a traversal of  $\Gamma(G')$  in the clock-wise direction, starting from  $v$ , defines an ordering  $\pi'$  of the vertices in  $X$ . Consider some vertex  $s \in S(\mathcal{M}')$ , and assume that it is the  $i$ th vertex of  $S(\mathcal{M}')$  according to the ordering  $\pi$ . We map it to the  $i$ th vertex of  $X$ , according to the ordering  $\pi'$ . This finishes the definition of the modified instance  $(G', \mathcal{M}'')$  corresponding to the valid  $(I, I', d)$  sub-instance  $(G, \mathcal{M}')$  of  $(G, \mathcal{M})$ . We need the following simple claim, whose proof is straightforward and is deferred to the Appendix.

**Claim 9.1** *Let  $d > 0$  be an integer, let  $(I, I')$  be a  $d$ -interesting pair of intervals, and let  $(G, \mathcal{M}')$  be a perfect  $(I, I', d)$  sub-instance of  $(G, \mathcal{M})$ . Let  $\text{OPT}'$  be the value of the optimal solution to instance  $(G, \mathcal{M}')$ . Then the value of the optimal solution to the corresponding modified instance  $(G', \mathcal{M}'')$  is at least  $\min\{\text{OPT}', d\}$ .*

Let  $d > 0$  be an integer, let  $(I, I')$  be a  $d$ -interesting pair of intervals, and let  $(G, \mathcal{M}')$  be a valid  $(I, I', d)$  sub-instance of  $(G, \mathcal{M})$ . Let  $\text{OPT}'$  be the value of the optimal solution to this instance. Consider the corresponding modified instance  $(G', \mathcal{M}'')$  of  $(G, \mathcal{M}')$ . Then  $(G', \mathcal{M}'')$ , together with the parameter  $\text{OPT} = \min\{d, \text{OPT}'\}$  is a valid input to Theorem 4.1. Even though we do not know the value of this parameter, we can try all  $d$  possibilities for it, and apply Theorem 4.1 to each one of them, selecting the best resulting solution. We are then guaranteed to compute, with high probability, a solution to instance  $(G', \mathcal{M}'')$ , whose value is at least  $\frac{\min\{d, \text{OPT}'\}}{2^{O(\sqrt{\log n} \log \log n)}}$ . We let  $\mathcal{A}(G', \mathcal{M}'')$  denote the resulting solution, and we let  $|\mathcal{A}(G', \mathcal{M}'')|$  denote its value.

### 9.3 Main Partitioning Theorem

The following theorem is central to our proof.

**Theorem 9.2** *Suppose we are given two disjoint intervals  $\pi, \pi'$  of  $\Gamma$ , each of which is contained in a single boundary edge of  $G$ , an integer  $d > 0$ , and a valid  $(\pi, \pi', d)$ -instance  $(G, \mathcal{M}')$ . Assume further that  $|\mathcal{M}'| \geq 1024d$ , and all demand pairs in  $\mathcal{M}'$  can be simultaneously routed via node-disjoint paths. Let  $z = \left\lceil \frac{|\mathcal{M}'|}{160d} \right\rceil - 1$ . Then there is a collection  $\Sigma = \{\sigma_1, \dots, \sigma_z\}$  of disjoint sub-intervals of  $\pi$ , and a collection  $\Sigma' = \{\sigma'_1, \dots, \sigma'_z\}$  of disjoint sub-intervals of  $\pi'$ , such that:*

- *the intervals  $\sigma_1, \sigma_2, \dots, \sigma_z, \sigma'_z, \dots, \sigma'_2, \sigma'_1$  appear on  $\Gamma$  in this circular order;*
- *for every pair of intervals  $\sigma \in \Sigma \cup \Sigma'$  and  $\sigma' \in \Sigma'$ , at least  $16d$  vertices of  $\Gamma$  separate the two intervals;*
- *every interval  $\sigma' \in \Sigma'$  contains at least  $16d$  and at most  $\sqrt{n}/2$  vertices, and lies within distance at least  $16d$  from every boundary edge of  $G$ , except for that containing  $\pi'$ ; and*
- *for all  $1 \leq i \leq z$ , the value of the optimal solution of the NDP instance  $(G, \mathcal{M}'_{\sigma_i, \sigma'_i})$  is at least  $d$ .*

Notice that in particular, for each  $1 \leq i \leq z$ ,  $(\sigma_i, \sigma'_i)$  is a  $d$ -interesting pair of intervals. Consider the set  $\tilde{\mathcal{M}} \subseteq \mathcal{M}$  of all demand pairs, such that for each  $(s, t) \in \tilde{\mathcal{M}}$ ,  $s \in \pi$ ,  $\tilde{t} \in \pi'$ , and  $d \leq d(t, \tilde{t}) < 2d$ , and apply Theorem 9.2 to it. Consider the resulting sets  $\Sigma = \{\sigma_1, \dots, \sigma_z\}$  and  $\Sigma' = \{\sigma'_1, \dots, \sigma'_z\}$  of segments. Then for all  $1 \leq i \leq z$ ,  $(G, \tilde{\mathcal{M}}_{\sigma_i, \sigma'_i})$  is a perfect  $(\sigma_i, \sigma'_i, d)$ -instance. Let  $(G_i, \tilde{\mathcal{M}}'_{\sigma_i, \sigma'_i})$  denote the corresponding modified instance for  $(G, \tilde{\mathcal{M}}_{\sigma_i, \sigma'_i})$ . Then all resulting sub-graphs  $G_1, \dots, G_z$  of  $G$  are disjoint. Moreover, since the length of each interval  $\sigma'_i$  is bounded by  $\sqrt{n}/2$ , if we assume w.l.o.g. that  $\pi'$  is contained in the bottom boundary of the grid  $G$ , then, assuming that  $d < \sqrt{n}/8$ , the sub-graphs  $G_1, \dots, G_z$  must be disjoint from the top  $\sqrt{n}/2 - 4d$  rows of  $G$ . For each  $1 \leq i \leq z$ , with high probability,  $|\mathcal{A}(G_i, \tilde{\mathcal{M}}'_{\sigma_i, \sigma'_i})| \geq d/2^{O(\sqrt{\log n} \log \log n)}$ , and so  $\sum_{i=1}^z |\mathcal{A}(G_i, \tilde{\mathcal{M}}'_{\sigma_i, \sigma'_i})| \geq |\tilde{\mathcal{M}}|/2^{O(\sqrt{\log n} \log \log n)}$  with high probability. We now turn to prove Theorem 9.2.

**Proof of Theorem 9.2.** Let  $\mathcal{P}^*$  be the set of node-disjoint paths routing the demand pairs in  $\mathcal{M}'$ . For each demand pair  $(s, t) \in \mathcal{M}'$ , let  $P(s, t) \in \mathcal{P}^*$  be the path routing it. For convenience, we assume that  $\pi'$  is contained in the bottom boundary of  $G$ .

For every interval  $\sigma' \subseteq \pi'$ , let  $\mathcal{M}'(\sigma') \subseteq \mathcal{M}'$  be the set of all demand pairs  $(s, t)$  with  $\tilde{t} \in \sigma'$ . Note that it is possible that for  $t, t' \in T(\mathcal{M}')$  with  $t \neq t'$ ,  $\tilde{t} = \tilde{t}'$ . Given a vertex  $v \in \sigma'$ , we let  $I(v) \subseteq \pi'$  be the shortest interval, whose leftmost vertex is  $v$ , such that  $I(v)$  contains at least  $16d$  vertices, and  $|\mathcal{M}'(I(v))| \geq 16d$ . If no such interval exists, then  $I(v)$  is undefined.

**Observation 9.3** *For each vertex  $v \in \pi'$ , if  $I(v)$  is defined, then  $|\mathcal{M}'(I(v))| \leq 20d$ .*

**Proof:** Let  $v'$  be the rightmost vertex of  $I(v)$ , and let  $I' = I(v) \setminus \{v'\}$ . Since we have chosen  $I(v)$  over the shorter interval  $I'$ , either  $|\mathcal{M}'(I')| < 16d$ , or  $|V(I')| < 16d$ . Assume first that the former is true. Since all demand pairs in  $\mathcal{M}'$  are routable via node-disjoint paths, all vertices in  $S(\mathcal{M}')$  are distinct. As for each terminal  $t \in T(\mathcal{M}')$ ,  $d(t, \tilde{t}) \leq 2d$ , there may be at most  $2d$  demand pairs  $(s, t) \in \mathcal{M}'$  with  $\tilde{t} = v'$ . But  $\mathcal{M}'(I')$  contains all demand pairs of  $\mathcal{M}'(I(v))$  except for those pairs  $(s, t)$  with  $\tilde{t} = v'$ . Therefore,  $|\mathcal{M}'(I(v))| \leq |\mathcal{M}'(I')| + 2d \leq 18d$ .

Assume now that  $|V(I')| < 16d$ . Let  $G' \subseteq G$  be the sub-grid of  $G$ , whose bottom boundary is  $I(v)$ , and whose height is  $2d$ . Then the width of  $G'$  is  $16d$ , and its top, left, and right boundary has total length  $20d$ . All paths routing the demand pairs in  $\mathcal{M}'(I(v))$  have to cross the boundary of  $G'$ , and so  $|\mathcal{M}'(I(v))| \leq 20d$ .  $\square$

Let  $\tilde{T} = \{\tilde{t} \mid t \in T(\mathcal{M}')\}$ . We start by defining a sequence  $(\mu_1, \dots, \mu_{2z+2})$  of  $2z + 2$  disjoint intervals of  $\pi'$ , as follows. Let  $v_1$  be the first vertex of  $\pi'$ . We then set  $\mu_1 = I(v_1)$ . Assume now that we have defined intervals  $\mu_1, \dots, \mu_i$ . Let  $v_{i+1}$  be the vertex lying immediately to the right of the right endpoint of  $\mu_i$ . We then set  $\mu_{i+1} = I(v_{i+1})$ . Notice that from Observation 9.3, for each  $i$ ,  $16d \leq |\mathcal{M}'(\mu_i)| \leq 20d$ . Since  $z = \left\lceil \frac{|\mathcal{M}'|}{160d} \right\rceil - 1$ , all intervals  $\mu_1, \dots, \mu_{2z+2}$  are well-defined. For convenience, we denote, for each  $1 \leq i \leq 2z$ ,  $\mathcal{M}'(\mu_i)$  by  $\mathcal{M}_i$ . For every vertex  $t \in T(\mathcal{M}')$ , let  $Q_t$  be the shortest path connecting  $t$  to  $\tilde{t}$  in  $G$ , so  $Q_t$  is contained in the column of  $G$  which contains  $t$ . We also let  $U(\mu_i) = \bigcup_{(s,t) \in \mathcal{M}_i} V(Q_t)$ .

For all  $1 \leq i, j \leq 2z$ , we say that intervals  $\mu_i$  and  $\mu_j$  are *neighbors* iff  $|i - j| \leq 1$ . Consider now some interval  $\mu_i$ , with  $1 \leq i \leq 2z$  and some demand pair  $(s, t) \in \mathcal{M}_i$ . We say that  $(s, t)$  is a *bad demand pair* iff path  $P(s, t)$  contains a vertex of  $U(\mu_j)$  for some interval  $\mu_j$  that is not a neighbor of  $\mu_i$ . Otherwise, we say that  $(s, t)$  is a *good demand pair*.

**Claim 9.4** *For each  $1 \leq i \leq 2z$ , at most  $8d$  demand pairs of  $\mathcal{M}_i$  are bad.*

We prove the claim below, after we complete the proof of Theorem 9.2 using it. Note that at most one interval  $\mu_j$  may contain more than  $\sqrt{n}/2$  vertices; we assume without loss of generality that, if such an interval  $\mu_j$  exists, then  $j$  is an odd integer. The set  $\Sigma' = \{\sigma'_1, \dots, \sigma'_z\}$  of intervals of  $\pi'$  is defined as follows: for each  $1 \leq i \leq z$ , we let  $\sigma'_i = \mu_{2i}$ . Notice that every pair of such intervals is separated by at least  $16d$  vertices, since each odd-indexed interval  $\mu_j$  must contain at least  $16d$  vertices by definition. Since we discard the first and the last intervals  $\mu_j$ , the distance from each resulting interval  $\sigma'_i$  to interval  $\pi$ , and to each of the remaining three boundaries of  $G$  is also at least  $16d$ .

In order to define the set  $\Sigma = \{\sigma_1, \dots, \sigma_z\}$  of intervals, fix some  $1 \leq i \leq z$ . Let  $\mathcal{M}'_i \subseteq \mathcal{M}_{2i}$  be the set of all good demand pairs in  $\mathcal{M}_{2i}$ . We let  $\sigma_i$  be the smallest sub-interval of  $\pi$ , containing all vertices of  $S(\mathcal{M}'_i)$ .

It is immediate to verify that, for each  $1 \leq i \leq z$ , the value of the optimal solution of the NDP instance  $(G, \mathcal{M}'_{\sigma_i, \sigma'_i})$  is at least  $d$ , since we can route all demand pairs of  $\mathcal{M}'_i$  in  $G$ . From our construction, every interval  $\sigma'_i \in \Sigma'$  contains at least  $16d$  and at most  $\sqrt{n}/2$  vertices, and it is separated by at least  $16d$  vertices from all other intervals in  $\Sigma \cup \Sigma'$ , and from each of the remaining three boundaries of  $G$ . It remains to show that the intervals in  $\Sigma$  are all disjoint, and that  $\sigma_1, \sigma_2, \dots, \sigma_z, \sigma'_z, \dots, \sigma'_2, \sigma'_1$  appear on  $\Gamma$  in this circular order.

Fix some pair of indices  $1 \leq i < j \leq z$ . If we traverse  $\Gamma$  in counter-clock-wise order, starting from the first vertex of  $\sigma'_1$ , then we will first encounter all vertices of  $\sigma'_i$ , then all vertices of  $\sigma'_j$ , and finally all vertices of  $\sigma_i \cup \sigma_j$ . It is enough to show that we will encounter every vertex of  $\sigma_j$  before we encounter any vertex of  $\sigma_i$ . In particular, it is enough to show that for every pair  $(s, t) \in \mathcal{M}'_j$  and  $(s', t') \in \mathcal{M}'_i$  of demand pairs, we will encounter  $s$  before  $s'$  in our traversal.

Assume for contradiction that this is false for some  $(s, t) \in \mathcal{M}'_j$  and  $(s', t') \in \mathcal{M}'_i$ . Then  $\tilde{t}', \tilde{t}, s', s$  appear on  $\Gamma$  in this circular order (see Figure 6). Since both demand pairs  $(s, t)$  and  $(s', t')$  are good, path  $P(s, t)$  may not contain a vertex of  $Q_{t'}$ , and path  $P(s', t')$  may not contain a vertex of  $Q_t$  (see Figure 6). But then the union of the path  $P(s', t')$  and  $Q_{t'}$  separates  $s$  from  $t$  in  $G$ , a contradiction. It now remains to prove Claim 9.4.

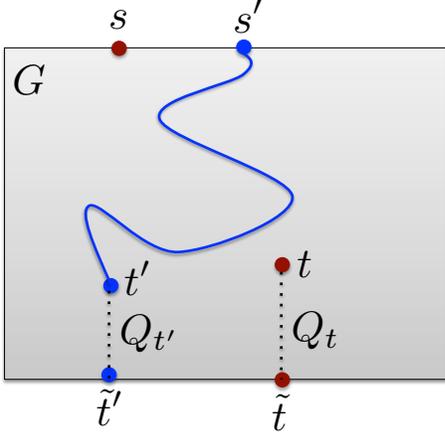


Figure 6: Ordering of the terminals of good demand pairs.

**Proof of Claim 9.4.** Fix some index  $1 \leq i \leq 2z$ , and let  $\mathcal{N}$  be the set of all bad demand pairs in  $\mathcal{M}_i$ . We further partition  $\mathcal{N}$  into two subsets: set  $\mathcal{N}'$  contains all demand pairs  $(s, t)$  whose corresponding path  $P(s, t)$  contains a vertex in some set  $U(\mu_j)$  for  $j > i + 1$ , and set  $\mathcal{N}''$  containing all remaining demand pairs (for each such demand pair  $(s, t)$ , the corresponding path  $P(s, t)$  must contain a vertex in some set  $U(\mu_j)$ , for some  $j < i - 1$ ). It is enough to show that  $|\mathcal{N}'|, |\mathcal{N}''| \leq 4d$ . We show that  $|\mathcal{N}'| \leq 4d$ ; the bound for  $\mathcal{N}''$  is proved similarly. Given a pair  $(s, t) \in \mathcal{N}'$ , and a pair  $(\hat{s}, \hat{t}) \in \mathcal{M}_{i+1}$ , we say that there is a *conflict* between  $(s, t)$  and  $(\hat{s}, \hat{t})$  iff  $P(s, t)$  contains a vertex of  $Q_{\hat{t}}$ .

We need the following simple observation.

**Observation 9.5** *Let  $(s, t)$  be any demand pair in  $\mathcal{N}'$ . Then there are at least  $12d$  demand pairs  $(\hat{s}, \hat{t}) \in \mathcal{M}_{i+1}$ , such that there is a conflict between  $(s, t)$  and  $(\hat{s}, \hat{t})$ .*

Assume first that the observation is correct, and assume for contradiction that  $|\mathcal{N}'| > 4d$ . Then there are at least  $48d^2$  pairs  $((s, t), (\hat{s}, \hat{t}))$  with  $(s, t) \in \mathcal{N}'$  and  $(\hat{s}, \hat{t}) \in \mathcal{M}_{i+1}$ , such that there is a conflict between  $(s, t)$  and  $(\hat{s}, \hat{t})$ . However,  $|\mathcal{M}_{i+1}| \leq 20d$ , and, since, for each pair  $(\hat{s}, \hat{t}) \in \mathcal{M}_{i+1}$ ,  $Q_{\hat{t}}$  contains at most  $2d$  vertices, there may be at most  $2d$  pairs  $(s, t) \in \mathcal{N}'$  that conflict with  $(\hat{s}, \hat{t})$ , a contradiction. We now proceed to prove the observation.

Let  $v$  be some vertex on path  $P(s, t)$ , that belongs to a set  $U(\mu_j)$ , for some  $j > i + 1$ . Then  $v$  lies on some path  $Q_{t'}$ , for some  $t' \in T(\mathcal{M}_j)$  (see Figure 7). We construct a (not necessarily simple) curve  $\gamma$ , by concatenating the path  $Q_t$ ; the sub-path of  $\pi'$  between  $\tilde{t}$  and  $\tilde{t}'$ ; the sub-path of  $Q_{t'}$  between  $\tilde{t}'$  and  $v$ ; and the sub-path of  $P(s, t)$  between  $v$  and  $t$  (see Figure 8). Consider now some demand pair  $(\hat{s}, \hat{t}) \in \mathcal{M}_{i+1}$ . If path  $P(s, t)$  does not contain a vertex of  $Q_{\hat{t}}$ , then curve  $\gamma$  separates  $\tilde{t}$  from  $\hat{s}$ , and so

path  $P(\hat{s}, \hat{t})$  has to cross the curve  $\gamma$ . This is only possible if path  $P(\hat{s}, \hat{t})$  contains a vertex of  $Q_t \cup Q_{t'}$ . But since both  $Q_t$  and  $Q_{t'}$  contain at most  $2d$  vertices each, at most  $4d$  paths of  $\mathcal{P}^*$  may intersect the curve  $\gamma$ . Therefore, there are at least  $16d - 4d = 12d$  demand pairs in  $\mathcal{M}_{i+1}$ , whose corresponding path in  $\mathcal{P}^*$  does not intersect the curve  $\gamma$ . From the above discussion, for each such pair  $(\hat{s}, \hat{t})$ , path  $P(s, t)$  must contain a vertex of  $Q_{\hat{t}}$ , and so  $(s, t)$  conflicts with  $(\hat{s}, \hat{t})$ . We conclude that  $(s, t)$  conflicts with at least  $12d$  pairs of  $\mathcal{M}_{i+1}$ .

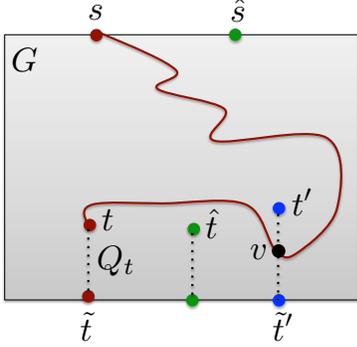


Figure 7: Illustration to the proof of Observation 9.5. Demand pairs  $(s, t)$  and  $(\hat{s}, \hat{t})$ . Path  $P(s, t)$  is shown in red.

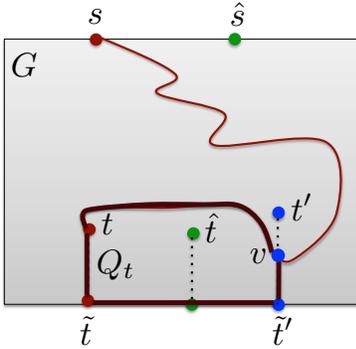


Figure 8: Illustration to the proof of Observation 9.5. The curve  $\gamma$  is shown in brown.

□

□

## 9.4 The Algorithm

If all vertices of  $T(\mathcal{M})$  lie on the boundary  $\Gamma$  of  $G$ , then we can solve the problem efficiently using standard dynamic programming. Therefore, we assume from now on, that for each  $t \in T(\mathcal{M})$ ,  $d(t, \tilde{t}) \geq 1$ .

Recall that  $\Gamma_1, \dots, \Gamma_4$  denote the four boundary edges of  $G$ . We partition the set  $\mathcal{M}$  of demand pairs into 16 subsets  $\mathcal{M}_{q,q'}$ , for  $1 \leq q, q' \leq 4$ : demand pair  $(s, t)$  belongs to set  $\mathcal{M}_{q,q'}$  iff  $s \in \Gamma_q$  and  $\tilde{t} \in \Gamma_{q'}$ . We then solve each of the resulting instances  $(G, \mathcal{M}_{q,q'})$  separately, and return the best of the resulting solutions. Since one of these instances is guaranteed to have a solution of value at least  $\text{OPT}/16$ , it is enough to show a factor- $2^{O(\sqrt{\log n \log \log n})}$ -approximation algorithm for each such instance separately. Therefore, from now on we focus on one such instance  $(G, \mathcal{M}_{q,q'})$ , and we assume w.l.o.g. that the value of the optimal solution in this instance is at least  $\text{OPT}/16$ .

We further partition the set  $\mathcal{M}_{q,q'}$  of demand pairs into the following subsets: we let  $\mathcal{M}_{q,q'}^0$  contain all demand pairs  $(s, t)$  with  $d(t, \tilde{t}) \geq \text{OPT}/\eta$ , and for  $0 < r < \log n$ , we let  $\mathcal{M}_{q,q'}^r$  contain all demand pairs  $(s, t)$  with  $\frac{\text{OPT}}{\eta \cdot 2^r} \leq d(t, \tilde{t}) < \frac{\text{OPT}}{\eta \cdot 2^{r-1}}$ . Clearly, there is an index  $0 < r^* < \log n$ , for which the value of the optimal solution of instance  $(G, \mathcal{M}_{q,q'}^{r^*})$  is at least  $\Omega(\text{OPT}/\log n)$ . Therefore, it is enough to compute a factor- $2^{O(\sqrt{\log n} \log \log n)}$ -approximation for each such instance separately. Notice that we can compute a factor- $2^{O(\sqrt{\log n} \log \log n)}$ -approximate solution for instance  $\mathcal{M}_{q,q'}^0$  using Theorem 4.1. Therefore, we now focus on an instance  $(G, \mathcal{M}_{q,q'}^{r^*})$ , for some fixed  $r^* > 0$ .

In order to simplify the notation, we denote by  $\mathcal{M}$  the corresponding set  $\mathcal{M}_{q,q'}^{r^*}$  of demand pairs; we assume that we are given an integer  $1 \leq d \leq \text{OPT}/2\eta$ , such that all destination vertices  $t \in T(\mathcal{M})$  have  $d \leq d(t, \tilde{t}) < 2d$ . We denote by  $\text{OPT}'$  the value of the optimal solution of the resulting instance  $(G, \mathcal{M})$ , and we assume that  $\Theta(\text{OPT}/\log n) \leq \text{OPT}' \leq \text{OPT}$ . Moreover, we can assume that  $\text{OPT}' > 2^{13}d$ , since otherwise we can directly apply Theorem 4.1 to instance  $(G, \mathcal{M})$  with parameter  $\text{OPT} = d$  to obtain an  $2^{O(\sqrt{\log n} \log \log n)}$ -approximation. We now consider three cases, depending on the location of the boundary edges  $\Gamma_q$  and  $\Gamma_{q'}$ .

**Case 1:  $\Gamma_q$  and  $\Gamma_{q'}$  are opposite boundary edges.** Without loss of generality, we assume that  $\Gamma_q$  is the top row of the grid  $G$ , and it contains the vertices of  $S(\mathcal{M})$ , while  $\Gamma_{q'}$  is the bottom row of the grid, and it contains the vertices  $\{\tilde{t} \mid t \in T(\mathcal{M})\}$ .

Let  $\mathcal{M}' \subseteq \mathcal{M}$  be the set of the demand pairs routed by the optimal solution. Let  $I$  denote the top boundary of the grid  $G$ , and let  $I'$  denote its bottom boundary. Theorem 9.2 guarantees the existence of a set  $\Sigma = \{\sigma_1, \dots, \sigma_z\}$  of disjoint sub-intervals of  $I$ , and a set  $\Sigma' = \{\sigma'_1, \dots, \sigma'_z\}$  of disjoint sub-intervals of  $I'$ , for  $z = \left\lfloor \frac{|\mathcal{M}'|}{160d} \right\rfloor - 1$ , such that the intervals  $\sigma_1, \sigma_2, \dots, \sigma_z, \sigma'_z, \dots, \sigma'_2, \sigma'_1$  appear on  $\Gamma$  in this circular order; every pair  $\sigma \in \Sigma \cup \Sigma'$ ,  $\sigma' \in \Sigma'$  of intervals is separated by at least  $16d$  vertices of  $\Gamma$ ; the length of each interval  $\sigma' \in \Sigma'$  is at least  $16d$  and at most  $\sqrt{n}/2$ , and for each  $1 \leq i \leq z$ , if we consider the perfect  $(\sigma_i, \sigma'_i, d)$ -instance  $(G, \tilde{\mathcal{M}}_i)$ , where  $\tilde{\mathcal{M}}_i = \mathcal{M}_{\sigma_i, \sigma'_i}$ , and denote by  $(G_i, \tilde{\mathcal{M}}'_i)$  the corresponding modified instance, then  $|\mathcal{A}(G_i, \tilde{\mathcal{M}}'_i)| \geq d/2^{O(\sqrt{\log n} \log \log n)}$ . Moreover, for every  $1 \leq i \leq z$ ,  $(\sigma_i, \sigma'_i)$  is a  $d$ -interesting pair of intervals. Notice, however, that, since the set  $\mathcal{M}'$  of demand pairs routed by the optimal solution is not known to us, we cannot compute the sets  $\Sigma$  and  $\Sigma'$  of intervals directly. We now show how to overcome this difficulty.

For every  $d$ -interesting pair  $(\sigma, \sigma')$  of intervals, with  $\sigma \subseteq I$ ,  $\sigma' \subseteq I'$ , consider the perfect  $(\sigma, \sigma', d)$ -instance  $(G, \mathcal{M}_{\sigma, \sigma'})$ , and the corresponding modified instance  $(G'_{\sigma, \sigma'}, \mathcal{M}'_{\sigma, \sigma'})$ . We can then compute an approximate solution  $\mathcal{A}(G'_{\sigma, \sigma'}, \mathcal{M}'_{\sigma, \sigma'})$  of value  $|\mathcal{A}(G'_{\sigma, \sigma'}, \mathcal{M}'_{\sigma, \sigma'})|$  to this instance, using the algorithm from Theorem 4.1, together with the parameter  $\text{OPT} = d$ . We can assume w.l.o.g. that, if  $\mathcal{M}'_{\sigma, \sigma'} \neq \emptyset$ , then  $|\mathcal{A}(G'_{\sigma, \sigma'}, \mathcal{M}'_{\sigma, \sigma'})| \geq 1$ . Let  $c$  be a constant, so that the approximation factor of Algorithm  $\mathcal{A}$  given in Theorem 4.1 is  $2^{c\sqrt{\log n} \log \log n}$ . Assume now that we have computed the values  $|\mathcal{A}(G'_{\sigma, \sigma'}, \mathcal{M}'_{\sigma, \sigma'})|$  for all  $d$ -interesting pairs  $\sigma \subseteq I$  and  $\sigma' \subseteq I'$  of intervals. We say that a pair  $(\sigma, \sigma')$  of interesting intervals is a *good pair* iff  $|\mathcal{A}(G'_{\sigma, \sigma'}, \mathcal{M}'_{\sigma, \sigma'})| \geq \left\lceil d/2^{c\sqrt{\log n} \log \log n} \right\rceil$ . It is now enough to compute a collection  $\Sigma = \{\sigma_1, \dots, \sigma_z\}$  of disjoint sub-intervals of  $I$ , and a collection  $\Sigma' = \{\sigma'_1, \dots, \sigma'_z\}$  of disjoint sub-intervals of  $I'$ , for  $z = \left\lfloor \frac{|\mathcal{M}'|}{160d} \right\rfloor - 1$ , such that the intervals  $\sigma_1, \sigma_2, \dots, \sigma_z, \sigma'_z, \dots, \sigma'_2, \sigma'_1$  appear on  $\Gamma$  in this circular order; for all  $i$ ,  $(\sigma_i, \sigma'_i)$  is a good pair of intervals, and  $\sigma'_i$  is at a distance at least  $16d$  from every interval in  $\Sigma \cup \Sigma'$ . This can be done by using simple dynamic programming.

Assume now that we have computed the collections  $\Sigma$  and  $\Sigma'$  of intervals as above. For each  $1 \leq i \leq z$ , let  $\mathcal{P}_i = \mathcal{A}(G'_{\sigma_i, \sigma'_i}, \mathcal{M}'_{\sigma_i, \sigma'_i})$  be the set of paths routed by the solution of value  $|\mathcal{A}(G'_{\sigma_i, \sigma'_i}, \mathcal{M}'_{\sigma_i, \sigma'_i})|$  that we have computed. Let  $\mathcal{N}'_i \subseteq \mathcal{M}'(\sigma_i, \sigma'_i)$  be the set of the demand pairs routed by this solution, and

let  $\mathcal{N}_i \subseteq \mathcal{M}(\sigma_i, \sigma'_i)$  be the set of the original demand pairs corresponding to the pairs in  $\mathcal{N}'_i$ . For convenience, we denote  $G_i = G'_{\sigma_i, \sigma'_i}$ . Finally, let  $\tilde{\mathcal{N}}_i$  denote the set of pairs  $(s, s')$ , where  $(s, t) \in \mathcal{N}_i$ , and  $(s', t') \in \mathcal{N}'_i$  is the corresponding demand pair in the modified instance.

From the above discussion,  $\sum_{i=1}^z |\mathcal{N}_i| \geq \text{OPT}'/2^{O(\sqrt{\log n} \cdot \log \log n)}$ . It is now enough to show that all demand pairs in set  $\bigcup_{i=1}^z \mathcal{N}_i$  can be routed in  $G$ . In order to do so, it is enough to show that all pairs in  $\bigcup_{i=1}^z \tilde{\mathcal{N}}_i$  can be routed in graph  $G$  via paths that are internally disjoint from  $\bigcup_{i=1}^z G_i$ , as we can exploit the paths in  $\bigcup_{i=1}^z \mathcal{P}_i$  in order to complete the routing. Recall that the graphs  $G_i$  are disjoint from the top  $\sqrt{n}/4$  rows of the grid, and the source vertices of the demand pairs in  $\bigcup_{i=1}^z \tilde{\mathcal{N}}_i$  appear in the same left-to-right order as their destination vertices. It is now immediate to complete the routing of the demand pairs in  $\bigcup_{i=1}^z \tilde{\mathcal{N}}_i$ .

**Case 2:  $\Gamma_q$  and  $\Gamma_{q'}$  are neighboring boundary edges.** Assume w.l.o.g. that  $\Gamma_q$  is the left boundary edge of  $G$  and  $\Gamma_{q'}$  is its bottom boundary edge. This case is dealt with very similarly to Case 1, with minor changes. As before, we assume that we are given a set  $\mathcal{M}$  of demand pairs and an integer  $d$ , such that for each pair  $(s, t) \in \mathcal{M}$ ,  $s \in \Gamma_q$ ,  $\tilde{t} \in \Gamma_{q'}$ , and  $d \leq d(t, \tilde{t}) < 2d$ . We assume that we know the value  $\text{OPT}'$  of the optimal solution to instance  $(G, \mathcal{M})$ , and we assume w.l.o.g. that  $\text{OPT}' \geq 2^{13}d$ . We let  $I$  be the left boundary edge of  $G$ , and we let  $I'$  be the bottom boundary edge of  $G$ , excluding its  $\lceil \text{OPT}'/16 \rceil$  leftmost vertices. Let  $\mathcal{M}' \subseteq \mathcal{M}$  be the subset of all demand pairs  $(s, t)$  with  $\tilde{t} \in I'$ . Let  $\text{OPT}''$  be the value of the optimal solution to instance  $(G, \mathcal{M}')$ .

**Observation 9.6**  $\text{OPT}'' \geq \text{OPT}'/2$ .

**Proof:** Let  $\mathcal{P}^*$  be the optimal solution to instance  $(G, \mathcal{M})$ , and let  $\mathcal{M}^*$  be the set of the demand pairs routed by it. Let  $\hat{\mathcal{M}} = \mathcal{M}^* \setminus \mathcal{M}'$ . Then it is enough to show that  $|\hat{\mathcal{M}}| \leq \text{OPT}'/2$ . Let  $R''$  be the bottom row of  $G$ , and let  $I'' = R'' \setminus I'$ . Then for each demand pair  $(s, t) \in \hat{\mathcal{M}}$ ,  $\tilde{t} \in I''$  must hold. Let  $Q$  be a sub-grid of  $G$ , whose bottom boundary is  $I''$ , and whose height is  $\lceil \text{OPT}'/16 \rceil > 2d$ . Then the boundary of  $Q$  has length at most  $\text{OPT}'/4 + 4$ , and every path routing a demand pair in  $\hat{\mathcal{M}}$  must cross the boundary of  $Q$ . Therefore,  $|\hat{\mathcal{M}}| \leq \text{OPT}'/2$ .  $\square$

The remainder of the algorithm is exactly the same as in Case 2. The only difference is in how the demand pairs in set  $\tilde{\mathcal{N}} = \bigcup_{i=1}^z \tilde{\mathcal{N}}_i$  are routed. This is done by utilizing the  $\text{OPT}'/16$  first columns of  $G$  and the  $\text{OPT}'/2$  top rows of  $G$ . Since we can assume that  $|\tilde{\mathcal{N}}| < \text{OPT}'/64$ , it is straightforward to find a suitable routing.

**Case 3:  $\Gamma_q = \Gamma_{q'}$ .** We assume w.l.o.g. that  $\Gamma_q$  is the bottom boundary of  $G$ . We assume that we are given a set  $\mathcal{M}$  of demand pairs and an integer  $d$ , such that for each pair  $(s, t) \in \mathcal{M}$ ,  $s, \tilde{t} \in \Gamma_q$ , and  $d \leq d(t, \tilde{t}) < 2d$ . We assume that we know the value  $\text{OPT}'$  of the optimal solution to instance  $(G, \mathcal{M})$ , and we assume w.l.o.g. that  $\text{OPT}' \geq 2^{13}d$ .

We partition the set  $\mathcal{M}$  of the demand pairs into three subsets: set  $\mathcal{M}^0$  contains all pairs  $(s, t)$  with  $|\text{col}(s) - \text{col}(t)| \leq 2d$ ; set  $\mathcal{M}^1$  contains all remaining pairs  $(s, t)$  with  $\text{col}(s) < \text{col}(t)$ ; and set  $\mathcal{M}^2$  contains all remaining demand pairs. We deal with the demand pairs in  $\mathcal{M}^0$  using the following claim.

**Claim 9.7** *There is an efficient randomized algorithm that computes a factor- $2^{O(\sqrt{\log n} \log \log n)}$  approximation to instance  $(G, \mathcal{M}^0)$  of NDP.*

**Proof:** Let  $\rho$  be a random integer between 0 and  $4d$ . We define a collection  $\mathcal{Q}' = \{Q_1, \dots, Q_r\}$  of square sub-grids of  $G$  as follows. Let  $Q_1$  be a sub-grid of size  $((4d + \rho) \times (4d + \rho))$ , containing the bottom left corner of  $G$  as the bottom left corner of  $Q_1$ . Assume now that we have defined  $Q_1, \dots, Q_i$ .

Let  $N$  be the number of columns of  $G$  that do not intersect the grids of  $Q_1, \dots, Q_i$ . If  $N > 8d$ , then we let  $Q_{i+1}$  be the  $(4d \times 4d)$  sub-grid of  $G$ , whose bottom left corner is the vertex serving as the bottom right corner of  $Q_i$ . Otherwise, we set  $r = i + 1$ , and we let  $Q_r$  be the square grid whose bottom left corner is the vertex serving as the bottom right corner of  $Q_i$ , and whose bottom right corner is the bottom right corner of the grid  $G$ . Notice that each grid  $Q_i$  has width and height  $4d$ , except for  $Q_1$  and  $Q_r$ , whose widths and heights may be between  $4d$  and  $8d$ .

For  $1 \leq i \leq r$ , let  $\mathcal{N}_i \subseteq \mathcal{M}^0$  be the set of all demand pairs  $(s, t)$  with  $s, t \in V(Q_i)$ . Let  $\mathcal{N} = \bigcup_{i=1}^r \mathcal{N}_i$ . Notice that each demand pair  $(s, t) \in \mathcal{M}^0$  belongs to  $\mathcal{N}$  with constant probability, and therefore, with high probability,  $\text{OPT}(G, \mathcal{N}) = \Omega(\text{OPT}(G, \mathcal{M}^0))$ .

We partition the set  $\mathcal{Q}'$  of the squares into four subsets,  $\mathcal{Q}^1, \dots, \mathcal{Q}^4$ , where for  $1 \leq i \leq 4$ ,  $\mathcal{Q}^i$  contains all squares  $Q_j$ , where  $j = i \pmod{4}$ . For  $1 \leq i \leq 4$ , let  $\mathcal{N}^i \subseteq \mathcal{N}$  be the subset of all demand pairs  $(s, t)$ , such that for some  $Q \in \mathcal{Q}^i$ ,  $s, t \in V(Q)$ . We solve each problem  $(G, \mathcal{N}^i)$  separately and return the best of the resulting four solutions.

We now fix  $1 \leq i \leq 4$ , and solve the problem  $(G, \mathcal{N}^i)$ . Consider some square  $Q_j \in \mathcal{Q}^i$ . Let  $Q_j^+$  be obtained by adding a margin of  $4d$  columns to the left and to the right of  $Q_j$ , and  $8d$  rows above  $Q_j$  (if  $j = 1$  or  $j = r$ , then we do not add columns on one of the sides of  $Q_j$ , and we add only  $4d$  rows above  $Q_j$ , so that  $Q_j^+$  is a square grid). Notice that from our construction, if  $Q_j, Q_{j'} \in \mathcal{Q}^i$ , then  $Q_j^+, Q_{j'}^+$  are disjoint. Using the same reasoning as in the proof of Claim 9.1,  $\text{OPT}(Q_j^+, \mathcal{N}_j) \geq \text{OPT}(G, \mathcal{N}_j)$ . Notice that  $\text{OPT}(G, \mathcal{N}_j) \leq 4d$ , as all source vertices of the demand pairs routed by the optimal solution must be distinct, and they lie on the bottom boundary of  $Q_j$ . It is then easy to see that  $(Q^+, \mathcal{N}_j)$  is a special case of Restricted NDP-Grid where all destination vertices lie at a distance at least  $\text{OPT}/4$  from the grid boundary, and we can find a  $2^{O(\sqrt{\log n} \log \log n)}$ -approximation for it using Theorem 4.1.  $\square$

If the optimal solution value to instance  $(G, \mathcal{M}^0)$  is at least  $\text{OPT}'/3$ , then we obtain a  $2^{O(\sqrt{\log n} \log \log n)}$ -approximation to this instance, and a  $2^{O(\sqrt{\log n} \log \log n)}$ -approximation overall, using Claim 9.7. Therefore, we assume from now on that any optimal solution must route at least  $2\text{OPT}'/3$  demand pairs from  $\mathcal{M}^1 \cup \mathcal{M}^2$ . We compute an approximate solution to each of the problems  $(G, \mathcal{M}^1)$  and  $(G, \mathcal{M}^2)$  separately, and take the better of the two solutions. Clearly, it is enough to design a factor- $2^{O(\sqrt{\log n} \log \log n)}$ -approximation to each of these two problems separately. We now focus on one of these problem, say  $(G, \mathcal{M}^1)$ , and for simplicity denote  $\mathcal{M}^1$  by  $\mathcal{M}$  and the value of the optimal solution to instance  $(G, \mathcal{M}^1)$  by  $\text{OPT}'$ .

Recall that  $\ell = \sqrt{n}$  is the length of the grid  $G$ . We can naturally associate with each demand pair  $(s, t) \in \mathcal{M}$  an interval  $I(s, t) \subseteq [\ell]$ : the left endpoint of  $I(s, t)$  is the index of  $\text{col}(s)$ , and its right endpoint is the index of  $\text{col}(t)$ .

Next, we partition the demand pairs in  $\mathcal{M}$  into  $h = \log n$  classes  $\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_h$ , as follows. Pair  $(s, t)$  belongs to class  $\mathcal{M}_i$  iff  $2^{i-1} \leq |I(s, t)| < 2^i$ . Note that for  $i \leq \log d$ ,  $\mathcal{M}_i = \emptyset$ , as all such pairs belong to  $\mathcal{M}^0$ . Thus, we obtain a collection of  $h$  instances  $(G, \mathcal{M}_i)$  of the problem. As before, we will compute a factor- $2^{O(\sqrt{\log n} \log \log n)}$ -approximation to each of these instances separately, and return the best of the resulting solutions. We now fix some index  $\log d \leq i \leq h$ , and design a factor- $2^{O(\sqrt{\log n} \log \log n)}$ -approximation algorithm to the corresponding instance  $(G, \mathcal{M}_i)$ .

Let  $\rho$  be an integer chosen uniformly at random from  $[0, 2^{i+3})$ . Let  $Z$  be the set of all integers  $1 \leq z \leq \ell$ , such that  $z = \rho + j \cdot 2^{i+3}$  for some integer  $j$ . Let  $\mathcal{M}' \subseteq \mathcal{M}$  contain all demand pairs  $(s, t)$ , such that there is some number  $z \in Z$  with  $\text{col}(s) < z < \text{col}(t)$ , and for all  $z' \in Z$ ,  $|\text{col}(s) - z'|, |\text{col}(t) - z'| \geq 2^{i-1}/4$ . Let  $I(s, t)$  be an interval whose left endpoint is  $\text{col}(s)$  and right endpoint is  $\text{col}(t)$ , and let  $I'(s, t) \subseteq I(s, t)$  be its sub-interval that excludes the first and the last  $2^{i-1}/4$  vertices of  $I(s, t)$ . Then it is easy to see that  $(s, t) \in \mathcal{M}'$  iff some integer in  $I'(s, t)$  belongs to  $Z$ , and that this happens with probability at least  $1/32$  over the choice of  $\rho$ . We now focus on solving the instance  $(G, \mathcal{M}')$  of the problem, and we

denote by  $\text{OPT}''$  the value of its optimal solution, so  $\mathbf{E}[\text{OPT}''] = \Omega(\text{OPT}')$ .

Let  $Z'$  contain all integers that lie halfway between consecutive pairs of integers of  $Z$ , that is:  $Z' = \{z + 2^{i+2} \mid z \in Z \text{ and } z + 2^{i+2} \leq \ell\}$ . We partition the grid  $G$  into sub-grids  $H_1, \dots, H_r$ , by deleting all columns  $W_z$ , where  $z \in Z'$ . Then for every demand pair  $(s, t) \in \mathcal{M}'$ , there is a unique sub-grid  $H_j$  containing both  $s$  and  $t$ . For each  $1 \leq j \leq r$ , let  $\hat{\mathcal{M}}_j \subseteq \mathcal{M}'$  be the set of all demand pairs  $(s, t)$  with both  $s$  and  $t$  contained in  $H_j$ . For each  $1 \leq j \leq r$ , we then consider the instance  $(G, \hat{\mathcal{M}}_j)$ . We can further partition  $H_j$  into two sub-grids  $H'_j$  and  $H''_j$  along the unique column  $W_z$  that is contained in  $H_j$ , and for which  $z \in Z$ . Notice that all source vertices lie on the bottom boundary of  $H'_j$  at distance at least  $2^{i-3}$  from the left and the right boundaries of  $H'_j$ , while all destination vertices lie in  $H''_j$ , at distance at least  $2^{i-3}$  from the left and the right boundaries of  $H''_j$ . Let  $I$  and  $I'$  denote the bottom boundaries of  $H'_j$  and  $H''_j$ , respectively.

Assume first that the value  $\text{OPT}_j$  of the optimal solution to instance  $(G, \hat{\mathcal{M}}_j)$  is at least  $1024d$ . Then we employ an algorithm similar to the one for Case 1: we use dynamic programming, together with the algorithm given by Theorem 4.1 to compute a collection  $\Sigma = \{\sigma_1, \dots, \sigma_z\}$  of disjoint sub-intervals of  $I$ , and a set  $\Sigma' = \{\sigma'_1, \dots, \sigma'_z\}$  of disjoint sub-intervals of  $I'$ , for  $z = \left\lceil \frac{|\mathcal{M}'|}{160d} \right\rceil - 1$ , such that the intervals  $\sigma_1, \sigma_2, \dots, \sigma_z, \sigma'_1, \sigma'_2, \dots, \sigma'_z$  appear on  $\Gamma$  in this circular order, and for all  $q$ ,  $(\sigma_q, \sigma'_q)$  is a good pair of intervals. For each  $1 \leq q \leq z$ , let  $(G'_{\sigma_q, \sigma'_q}, \hat{\mathcal{M}}'_{\sigma_q, \sigma'_q})$  denote the corresponding modified instance. The key is to notice that all such graphs  $G'_{\sigma_q, \sigma'_q}$  are mutually disjoint and are contained in  $H''_j$ . We construct the set  $\tilde{\mathcal{N}} = \bigcup_{q=1}^z \tilde{\mathcal{N}}_q$  of new demand pairs as before, and show that they can be routed in  $H_j$  via paths that are internally disjoint from  $\bigcup_{q=1}^z G'_{\sigma_q, \sigma'_q}$ . The routing is straightforward and exploits the columns of  $H'_j$  and the top  $\sqrt{n}/4$  rows of  $H''_j$ .

Finally, assume that the value of the optimal solution  $\text{OPT}_j$  to instance  $(G, \hat{\mathcal{M}}_j)$  is less than  $1024d$ . Note that  $\text{OPT}_j < 2^{i+2}$  must also hold, as interval  $I$  contains at most  $2^{i+2}$  vertices.

Let  $\hat{I}'$  be obtained from  $I'$  by discarding its first  $2^{i-3}$  and its last  $2^{i-3}$  vertices, and recall that  $2^i \geq d$ . Then  $(I, \hat{I}')$  is a  $d$ -interesting pair of intervals, and  $(G, \hat{\mathcal{M}}_j)$  is a perfect  $(I, \hat{I}', d)$  sub-instance of  $(G, \mathcal{M})$ . Let  $(\tilde{H}_j, \hat{\mathcal{M}}'_j)$  be the corresponding modified instance. Notice that  $\tilde{H}_j \subseteq H''_j$ . Then, from Claim 9.1, the value of the optimal solution to instance  $(\tilde{H}_j, \hat{\mathcal{M}}'_j)$  is at least  $\min\{\text{OPT}_j, d\} = \Omega(\text{OPT}_j)$ . We can use Theorem 4.1 in order to compute a  $2^{O(\sqrt{\log n} \log \log n)}$ -approximate solution  $\mathcal{P}_j$  to instance  $(\tilde{H}_j, \hat{\mathcal{M}}'_j)$ . From the above discussion  $|\mathcal{P}_j| \geq \frac{\text{OPT}_j}{2^{O(\sqrt{\log n} \log \log n)}}$ . It now remains to route the vertices of  $S(\hat{\mathcal{M}}_j)$  to the corresponding vertices of  $S(\hat{\mathcal{M}}'_j)$  via node-disjoint paths that are internally disjoint from  $\tilde{H}_j$ . This is done in exactly the same way as before.

## 10 Approximating NDP/EDP on Walls with Sources on the Boundary

In this section we prove Theorem 1.3, by extending our results for Restricted NDP-Grid to EDP and NDP on 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  $e_1^j, \dots, e_{h-1}^j$  be the edges of  $W_j$  indexed in their top-to-bottom order. Let  $E^*(G) \subseteq E(G)$  contain all edges  $e_z^j$ , where  $z \neq j \pmod{2}$ , and let  $\hat{G}$  be the graph obtained from  $G \setminus E^*(G)$ , by deleting all degree-1 vertices from it. Graph  $\hat{G}$  is called a *wall of length  $\ell/2$  and height  $h$*  (see Figure 9). 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

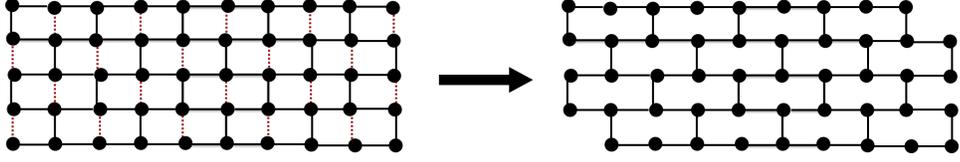


Figure 9: The red dotted edges are deleted to obtain wall  $\hat{G}$  from grid  $G$ .

denote them by  $R_1, \dots, R_h$  in this top-to-bottom order; notice that  $R_j$  is a sub-path of the  $j$ th row of  $G$  for all  $j$ . Graph  $\hat{G}$  contains a unique collection  $\mathcal{W}$  of  $\ell/2$  node-disjoint paths that connect vertices of  $R_1$  to vertices of  $R_h$  and are internally disjoint from  $R_1$  and  $R_h$ . We refer to the paths in  $\mathcal{W}$  as the *columns* of  $\hat{G}$ , and denote them by  $W_1, \dots, W_{\ell/2}$  in this left-to-right order. Paths  $W_1, W_{\ell/2}, R_1$  and  $R_h$  are called the left, right, top and bottom boundary edges of  $\hat{G}$ , respectively, and their union is the boundary of  $\hat{G}$ , that we denote by  $\Gamma(\hat{G})$ .

In both NDP-Wall and EDP-Wall problems, the input is a wall  $\hat{G}$  of length  $\sqrt{n}/2$  and height  $\sqrt{n}$  (we assume that  $\sqrt{n}/2$  is an integer), and a collection  $\mathcal{M} = \{(s_1, t_1), \dots, (s_k, t_k)\}$  of  $k$  demand pairs. The goal is to route the maximum number of the demand pairs via paths that are node-disjoint (for NDP-Wall) or edge-disjoint (for EDP-Wall). We will exploit the  $(\sqrt{n} \times \sqrt{n})$  grid  $G$ , and the fact that  $\hat{G} \subseteq G$  can be obtained from  $G$  as described above.

For any two nodes  $u, v$  of  $\hat{G}$ , we denote the shortest-path distance between them in  $\hat{G}$  by  $\hat{d}(u, v)$ . Distances between subsets of vertices, and between a vertex and a subset of vertices, are defined as before, using  $\hat{d}$ .

The following simple observation relates the values of the optimal solutions to the EDP-Wall and NDP-Wall problems, defined over the same wall  $\hat{G}$  and the same set  $\mathcal{M}$  of demand pairs.

**Observation 10.1** *Let  $\hat{G}$  be a wall of length  $\sqrt{n}/2$  and height  $\sqrt{n}$ , and let  $\mathcal{M} = \{(s_1, t_1), \dots, (s_k, t_k)\}$  be a set of  $k$  demand pairs. Let  $\text{OPT}_{\text{EDP}}$  be the value of the optimal solution to the EDP-Wall problem instance  $(\hat{G}, \mathcal{M})$ , and let  $\text{OPT}_{\text{NDP}}$  be the value of the optimal solution to the NDP-Wall problem instance  $(\hat{G}, \mathcal{M})$ . Then:  $\text{OPT}_{\text{NDP}} \leq \text{OPT}_{\text{EDP}} \leq O(\text{OPT}_{\text{NDP}})$ .*

**Proof:** Observe first that for any set  $\mathcal{P}$  of paths in  $\hat{G}$ , if the paths in  $\mathcal{P}$  are mutually node-disjoint, then they are also mutually edge-disjoint. Therefore,  $\text{OPT}_{\text{NDP}} \leq \text{OPT}_{\text{EDP}}$ .

To show the other direction, consider any set  $\mathcal{P}$  of paths, such that the paths in  $\mathcal{P}$  are edge-disjoint. It is enough to show that there is a subset  $\mathcal{P}' \subseteq \mathcal{P}$  of paths that are node-disjoint, and  $|\mathcal{P}'| \geq \Omega(|\mathcal{P}|)$ . Since the maximum vertex degree in  $\hat{G}$  is 3, the only way for two paths  $P, P' \in \mathcal{P}$  to share a vertex  $x$  is when  $x$  is an endpoint of at least one of the two paths.

We construct a directed graph  $H$ , whose vertex set is  $\{v_P \mid P \in \mathcal{P}\}$ , and there is an edge  $(v_P, v_{P'})$  iff an endpoint of  $P'$  belongs to  $P$ . It is immediate to verify that the maximum in-degree of any vertex in  $H$  is 4. Therefore, there is a set  $U \subseteq V(H)$  of  $\Omega(|V(H)|)$  vertices, such that no two vertices of  $U$  are connected with an edge. We let  $\mathcal{P}' = \{P \mid v_P \in U\}$ . Then the paths in  $\mathcal{P}'$  are node-disjoint, and  $|\mathcal{P}'| = \Omega(|\mathcal{P}|)$ .  $\square$

Since for any set  $\mathcal{P}$  of node-disjoint paths, the paths in  $\mathcal{P}$  are also mutually edge-disjoint, it is now enough to prove Theorem 1.3 for the NDP-Wall problem, with all source vertices lying on the wall boundary. We do so in the remainder of this section.

Our main approach is to consider the corresponding instance  $(G, \mathcal{M})$  of the NDP-Grid problem on the underlying grid  $G$ , and to exploit our algorithm for this problem. However, we cannot use this approach directly, since the boundary of the wall  $\hat{G}$  is not contained in the boundary of the grid  $G$ ,

and so some of the source vertices of  $S(\mathcal{M})$  may not lie on  $\Gamma(G)$ . We overcome this difficulty by mapping each vertex in  $S(\mathcal{M})$  to its closest vertex lying on  $\Gamma(G)$ . Formally, we define a new set  $\mathcal{M}'$  of demand pairs, containing, for each pair  $(s, t) \in \mathcal{M}$ , a new pair  $(s', t')$ , with  $t' = t$ , and  $s'$  defined as follows. If  $s \in \Gamma(G)$ , then  $s' = s$ ; otherwise, either  $s$  belongs to the first column of  $\hat{G}$ , in which case  $s'$  is defined to be the vertex of  $\Gamma(G)$  lying immediately to the left of  $s$ ; or  $s$  belongs to the last column of  $\hat{G}$ , in which case  $s'$  is defined to be the vertex of  $\Gamma(G)$  lying immediately to its right (see Figure 9). The following simple observation shows that the solution value does not change by much.

**Observation 10.2** *Let  $\text{OPT}$  be the value of the optimal solution to instance  $(\hat{G}, \mathcal{M})$  of NDP-Wall, and let  $\text{OPT}'$  be defined in the same way for instance  $(\hat{G}, \mathcal{M}')$ . Then  $\text{OPT}' = \Omega(\text{OPT})$ . Moreover, given any set  $\mathcal{P}'$  of node-disjoint paths routing a subset of the demand pairs in  $\mathcal{M}'$ , there is an efficient algorithm to compute a set  $\mathcal{P}$  of  $\Omega(|\mathcal{P}'|)$  node-disjoint paths routing a subsets of the demand pairs in  $\mathcal{M}$ .*

**Proof:** In order to prove the first assertion, let  $\mathcal{P}^*$  be the optimal solution to instance  $(\hat{G}, \mathcal{M})$ . We show that there is a set  $\mathcal{P}'$  of  $\Omega(|\mathcal{P}^*|)$  paths routing demand pairs in  $\mathcal{M}'$ . Let  $\mathcal{P}''$  be obtained from  $\mathcal{P}^*$  as follows. Consider any path  $P \in \mathcal{P}^*$ , and assume that it routes some pair  $(s, t) \in \mathcal{M}$ . If  $(s, t) \in \mathcal{M}'$ , then we add  $P$  to  $\mathcal{P}''$ ; otherwise, we extend  $P$  by adding the edge  $(s, s')$  to it, and add the resulting path to  $\mathcal{P}''$ . Consider the final set  $\mathcal{P}''$  of paths. While  $|\mathcal{P}''| = |\mathcal{P}^*|$ , it is possible that the paths in  $\mathcal{P}''$  are no longer node-disjoint. However, a pair  $P_1, P_2 \in \mathcal{P}''$  of paths may share a vertex  $x$  iff  $x$  is an endpoint of at least one of the two paths. We can now employ the same argument as in the proof of Observation 10.1 to obtain a subset  $\mathcal{P}' \subseteq \mathcal{P}''$  of node-disjoint paths, with  $|\mathcal{P}'| \geq \Omega(|\mathcal{P}''|)$ .

The proof of the second assertion is almost identical.  $\square$

From now on, it is sufficient to design a randomized  $2^{O(\sqrt{\log n \cdot \log \log n})}$ -approximation algorithm to the new instance  $(\hat{G}, \mathcal{M}')$  of NDP-Wall. Therefore, from now on we will assume that all source vertices lie on  $\Gamma(\hat{G}) \cap \Gamma(G)$ .

Our first step is, as before, to consider a special case of the problem, where the destination vertices lie far from the boundary of the wall. We prove the following analogue of Theorem 4.1.

**Theorem 10.3** *There is an efficient randomized algorithm  $\hat{\mathcal{A}}$ , that, given an instance  $(\hat{G}, \mathcal{M})$  of the NDP-Wall problem, with all sources lying on  $\Gamma(G) \cap \Gamma(\hat{G})$  (where  $G$  is the grid corresponding to  $\hat{G}$ ), and an integer  $\text{OPT}$ , such that the value of the optimal solution to instance  $(\hat{G}, \mathcal{M})$  of NDP-Wall is at least  $\text{OPT}$ , and every destination vertex lies at a distance at least  $\text{OPT}/\eta$  from  $\Gamma(\hat{G})$  in  $\hat{G}$ , returns a solution that routes at least  $\text{OPT}/2^{O(\sqrt{\log n \cdot \log \log n})}$  demand pairs in  $\hat{G}$ , with high probability.*

**Proof:** Consider the Restricted NDP-Grid instance defined over the grid  $G$ , with the set  $\mathcal{M}$  of the demand pairs. Since  $\hat{G} \subseteq G$ , the value of the optimal solution to this instance is at least  $\text{OPT}$ . We can apply Algorithm  $\mathcal{A}$  to instance  $(G, \mathcal{M})$  of Restricted NDP-Grid, to obtain a set  $\mathcal{P}$  of node-disjoint paths routing some subset  $\tilde{\mathcal{M}} \subseteq \mathcal{M}$  of demand pairs in graph  $G$ , such that with high probability,  $|\tilde{\mathcal{M}}| \geq \text{OPT}/2^{O(\sqrt{\log n \cdot \log \log n})}$ . We now show that all demand pairs in  $\tilde{\mathcal{M}}$  can also be routed in the wall  $\hat{G}$ . Indeed, in the analysis in Section 7, we have constructed disjoint level- $\rho$  snakes of width at least 3, such that each demand pair  $(s, t) \in \tilde{\mathcal{M}}$  is contained in a distinct snake, that we denote by  $\mathcal{Y}(s, t)$ . With a loss of an additive constant in the approximation guarantee, we can ensure that no snake passes through any corner of  $G$ .

**Observation 10.4** *Let  $\mathcal{Y} = (\Upsilon_1, \dots, \Upsilon_z)$  be any snake of width at least 3 in the grid  $G$ . Let  $H$  be the union of all graphs  $\Upsilon_j \cap \hat{G}$ , for  $1 \leq j \leq z$ . Then  $H$  is a connected sub-graph of  $\hat{G}$ .*

It is now immediate to obtain a routing of the demand pairs in  $\tilde{\mathcal{M}}$  via node-disjoint paths in  $\hat{G}$ . For every snake  $\mathcal{Y}(s, t)$ , let  $H(s, t)$  be the corresponding sub-graph of  $\hat{G}$ , given by the above observation. Then all graphs  $\{H(s, t)\}_{(s, t) \in \tilde{\mathcal{M}}}$  are mutually disjoint, and each such graph is connected. We simply connect  $s$  to  $t$  by any path contained in  $H(s, t)$ .  $\square$

Finally, we complete the proof of Theorem 1.3, by removing the assumption that the destination vertices lie far from the wall boundaries. As before, we assume that we are given an instance  $(\hat{G}, \mathcal{M})$  of NDP-Wall, such that, if  $G$  denotes the corresponding grid graph, then all source vertices in  $S(\mathcal{M})$  lie on  $\Gamma(G) \cap \Gamma(\hat{G})$ . As before, we use the algorithm from Section 9 on the underlying grid  $G$  and adapt the resulting solution to the wall  $\hat{G}$ .

Recall that for every destination vertex  $t \in T(\mathcal{M})$ , we denoted by  $\tilde{t}$  the vertex of  $\Gamma(G)$  minimizing  $d(t, \tilde{t})$ , breaking ties arbitrarily. For any subset  $\mathcal{M}' \subseteq \mathcal{M}$  of demand pairs, we denoted by  $\tilde{T}(\mathcal{M}') = \{\tilde{t} \mid t \in T(\mathcal{M}')\}$ . If all vertices of  $\mathcal{T}(\mathcal{M})$  lie on the boundary  $\Gamma(\hat{G})$  of wall  $\hat{G}$ , then we can efficiently obtain a constant-factor approximation for the problem using standard dynamic programming techniques. Thus, we assume from now on that  $\hat{d}(t, \tilde{t}) \geq 1$  for each  $t \in T(\mathcal{M})$ .

Recall that we showed that it suffices to consider the case where all sources  $S(\mathcal{M})$  lie on a single boundary edge  $\Gamma_q$  of  $G$ , and all vertices of  $\tilde{T}(\mathcal{M})$  lie on a single boundary edge  $\Gamma_{q'}$  of  $G$ , where possibly  $\Gamma_q = \Gamma_{q'}$ . We further saw that it suffices to consider only two cases: (i)  $d(t, \tilde{t}) \geq \text{OPT}/\eta$  for all destination nodes  $t$ ; or (ii) we are given an integer  $1 \leq d \leq \text{OPT}/2\eta$  such that  $d \leq d(t, \tilde{t}) < 2d$  for all destination nodes  $t$ , and  $\text{OPT} > 2^{13}d$ . In case (i), we can compute a factor  $2^{O(\sqrt{\log n \cdot \log \log n})}$ -approximate solution directly using Theorem 10.3. In the remainder of this section, we consider case (ii).

Let  $\mathcal{M}'$  be the set of the demand pairs that was chosen to be routed by our algorithm from Section 9. Then all vertices of  $\tilde{T}(\mathcal{M}')$  lie on the same boundary edge of the grid – we assume w.l.o.g. that it is the bottom boundary. Recall that we have defined a collection  $\{(G_1, \mathcal{M}_1), \dots, (G_z, \mathcal{M}_z)\}$  of modified sub-instances, such that all grids  $G_1, \dots, G_z$  are disjoint from each other and from the top  $\sqrt{n}/4$  rows of  $G$ , and in each of the instances  $(G_i, \mathcal{M}_i)$ , the destination vertices lie far enough from the boundary of the grid  $G_i$ , so that Theorem 4.1 could be applied to each such sub-instance separately. Let  $\mathcal{M}'_i \subseteq \mathcal{M}'$  be the subset of the demand pairs corresponding to the pairs in  $\mathcal{M}_i$ . For each  $1 \leq i \leq z$ , we have also defined a set  $\tilde{N}_i$  of pairs, connecting the original source vertices in  $S(\mathcal{M}'_i)$  to their corresponding source vertices in  $S(\mathcal{M}_i)$ . We have implicitly constructed, for each  $1 \leq i \leq z$ , a snake  $\mathcal{Y}_i$ , inside which the demand pairs in  $\tilde{N}_i$  are routed. All snakes  $\mathcal{Y}_1, \dots, \mathcal{Y}_z$  are disjoint, and have width at least  $3|\mathcal{N}_i|$  each. We can now translate this routing into a set of node-disjoint paths routing the demand pairs in  $\mathcal{M}'$  in the wall  $\hat{G}$ . The routing inside each sub-grid  $G_i$  is altered exactly like in the proof of Theorem 10.3; the routings of the sets  $\tilde{N}_i$  of demand pairs exploit the same snakes  $\mathcal{Y}_i$  inside the wall  $\hat{G}$ .

## 11 Approximation Algorithm for the Special Case with Sources Close to the Grid Boundary

The goal of this section is to prove Theorem 1.2. We assume that we are given an instance  $(G, \mathcal{M})$  of NDP-Grid and an integer  $\delta > 0$ , such that every source vertex is at a distance at most  $\delta$  from the grid boundary. Our goal is to design an efficient randomized factor- $(\delta \cdot 2^{O(\sqrt{\log n \cdot \log \log n})})$ -approximation algorithm for this special case of the problem. We can assume that  $\delta < n^{1/4}$ , as there is an  $\tilde{O}(n^{1/4})$ -approximation algorithm for NDP-Grid [CK15]. For every source vertex  $s \in S(\mathcal{M})$ , let  $\tilde{s}$  be the vertex lying closest to  $s$  on  $\Gamma(G)$ ; for every destination vertex  $t \in T(\mathcal{M})$ , let  $\tilde{t}$  be defined similarly. Recall that for each  $s \in S(\mathcal{M})$ ,  $d(s, \tilde{s}) \leq \delta$ . For every subset  $\mathcal{M}' \subseteq \mathcal{M}$  of the demand pairs, we denote

$\tilde{S}(\mathcal{M}') = \{\tilde{s} \mid s \in S(\mathcal{M}')\}$ , and similarly  $\tilde{T}(\mathcal{M}') = \{\tilde{t} \mid t \in T(\mathcal{M}')\}$ . Using the same arguments as in Section 9, at the cost of losing an  $O(\log n)$ -factor in the approximation ratio, we can assume that all vertices of  $\tilde{S}(\mathcal{M})$  are contained in a single boundary edge of the grid  $G$ , that we denote by  $\Gamma$ , and that all vertices of  $\tilde{T}(\mathcal{M})$  are contained in a single boundary edge of the grid  $G$ , that we denote by  $\Gamma'$ , where possibly  $\Gamma = \Gamma'$ . Moreover, we can assume that there is some integer  $d$ , such that for all  $t \in T(\mathcal{M})$ ,  $d \leq d(t, \tilde{t}) < 2d$ .

Let  $c$  be a sufficiently large constant. Assume first that  $d \leq \delta \cdot 2^{c\sqrt{\log n} \log \log n}$ . In this case, we can obtain an  $(\delta \cdot 2^{O(\sqrt{\log n} \log \log n)})$ -approximation using the following claim.

**Claim 11.1** *There is an efficient algorithm to compute a factor  $(\delta \cdot 2^{O(\sqrt{\log n} \log \log n)})$ -approximation for the special case where  $d \leq \delta \cdot 2^{c\sqrt{\log n} \log \log n}$ .*

**Proof:** Let  $\hat{\mathcal{M}} = \{(\tilde{s}, \tilde{t}) \mid (s, t) \in \mathcal{M}\}$  be a new set of demand pairs. Note that all vertices participating in the demand pairs in  $\hat{\mathcal{M}}$  lie on the boundary of  $G$ . Therefore, we can efficiently find an optimal solution to the NDP problem instance  $(G, \hat{\mathcal{M}})$  using standard dynamic programming. Next, we show that  $\text{OPT}(G, \hat{\mathcal{M}}) = \Omega(\text{OPT}(G, \mathcal{M}) / (\delta \cdot 2^{O(\sqrt{\log n} \log \log n)}))$ . For every vertex  $v \in S(\mathcal{M}) \cup T(\mathcal{M})$ , let  $U_v$  be the shortest path connecting  $v$  to  $\tilde{v}$ .

Let  $\mathcal{P}^*$  be the optimal solution to instance  $(G, \mathcal{M})$ , and let  $\mathcal{M}^* \subseteq \mathcal{M}$  be the set of the demand pairs routed by  $\mathcal{P}^*$ . For each demand pair  $(s, t) \in \mathcal{M}^*$ , let  $P(s, t) \in \mathcal{P}^*$  be the path routing it.

We construct a directed graph  $H$ , whose vertex set is  $V(H) = \{v(s, t) \mid (s, t) \in \mathcal{M}^*\}$ , and there is a directed edge from  $v(s, t)$  to  $v(s', t')$  iff the path  $P(s', t')$  routing the pair  $(s', t')$  in the optimal solution  $\mathcal{P}^*$  contains a vertex of  $U_s \cup U_t$ . It is easy to verify that the out-degree of every vertex of  $H$  is at most  $d + \delta$ , and so there is an independent set  $I$  in  $H$  containing at least  $\Omega(|V(H)| / (d + \delta)) = \Omega(\text{OPT}(G, \mathcal{M}) / (\delta \cdot 2^{O(\sqrt{\log n} \log \log n)}))$  vertices of  $I$ .

For each demand pair  $(s, t) \in \mathcal{M}^*$  with  $v(s, t) \in I$ , let  $P'(s, t)$  be the concatenation of  $U_s$ ,  $P(s, t)$  and  $U_t$ . Then  $\{P'(s, t) \mid v(s, t) \in I\}$  is a collection of  $\Omega(\text{OPT}(G, \mathcal{M}) / (\delta \cdot 2^{O(\sqrt{\log n} \log \log n)}))$  node-disjoint paths routing demand pairs in  $\hat{\mathcal{M}}$ . We conclude that  $\text{OPT}(G, \hat{\mathcal{M}}) = \Omega(\text{OPT}(G, \mathcal{M}) / (\delta \cdot 2^{O(\sqrt{\log n} \log \log n)}))$ .

Let  $\mathcal{P}'$  be a solution to instance  $(G, \hat{\mathcal{M}})$ , obtained by the constant-factor approximation algorithm, so  $|\mathcal{P}'| = \Omega(\text{OPT}(G, \hat{\mathcal{M}}) / (\delta \cdot 2^{O(\sqrt{\log n} \log \log n)}))$ , and let  $\hat{\mathcal{M}}' \subseteq \hat{\mathcal{M}}$  be the set of the demand pairs routed. Since all source and all destination of  $\hat{\mathcal{M}}'$  appear on the boundary of  $G$ ,  $|\hat{\mathcal{M}}'| = O(\sqrt{n})$ . Moreover, the demand pairs in  $\hat{\mathcal{M}}'$  must be non-crossing: that is, for every pair  $(s, t), (s', t') \in \hat{\mathcal{M}}'$ , the circular ordering of the corresponding terminals on the boundary of the grid is either  $(s, s', t', t)$  or  $(s, t, t', s')$ . Since we have assumed that  $\delta \leq n^{1/4}$  and  $d \leq \delta \cdot 2^{c\sqrt{\log n} \log \log n}$ , it is easy to verify that we can route  $|\Omega(\hat{\mathcal{M}}')|$  demand pairs from the original set  $\mathcal{M}$ , that correspond to the demand pairs in  $\hat{\mathcal{M}}'$ . □

From now on we assume that  $d > \delta \cdot 2^{c\sqrt{\log n} \log \log n}$ . Next, we define a new set  $\tilde{\mathcal{M}}$  of demand pairs, as follows:  $\tilde{\mathcal{M}} = \{(\tilde{s}, t) \mid (s, t) \in \mathcal{M}\}$ , so all source vertices of the demand pairs in  $\tilde{\mathcal{M}}$  lie on the boundary edge  $\Gamma$  of  $G$ . Let  $\text{OPT}'$  be the value of the optimal solution to problem  $(G, \tilde{\mathcal{M}})$ . The proof of the following claim is almost identical to the proof of Claim 11.1 and is omitted here.

**Claim 11.2**  $\text{OPT}' \geq \Omega(\text{OPT}(G, \mathcal{M}) / \delta)$ .

From now on we focus on instance  $(G, \tilde{\mathcal{M}})$  of Restricted NDP-Grid. Recall that for each source vertex  $s \in S(\mathcal{M})$ , we denoted by  $U_s$  the shortest path connecting  $s$  to  $\tilde{s}$ . We say that a path  $P$  routing a demand pair  $(\tilde{s}, t) \in \tilde{\mathcal{M}}$  is *canonical* iff  $U_s \subseteq P$ . We say that set  $\mathcal{P}$  of node-disjoint paths routing demand pairs in  $\tilde{\mathcal{M}}$  is canonical iff every path in  $\mathcal{P}$  is canonical. We will show that we can apply our  $2^{O(\sqrt{\log n} \log \log n)}$ -approximation algorithm to instance  $(G, \tilde{\mathcal{M}})$  to obtain a routing in which all paths

are canonical. For convenience, we assume w.l.o.g. that  $\Gamma$  is the top boundary edge of the grid. Therefore, a canonical path connecting a demand pair  $(\tilde{s}, t) \in \tilde{\mathcal{M}}$  must follow the column  $\text{col}(\tilde{s})$  of  $G$  until it reaches the vertex  $s$ .

Assume first that  $d > \text{OPT}'/\eta$ , where  $\eta$  is defined as before. Then we obtain a special case of the problem where the destination vertices are sufficiently far from the boundary of the grid, and can use the algorithm from Theorem 4.1 to find a  $2^{O(\sqrt{\log n} \cdot \log \log n)}$ -approximate solution to instance  $(G, \tilde{\mathcal{M}})$ . However, we slightly modify the routing to ensure that the paths in our solution are canonical. Recall that the routing itself is recursive. Let  $\mathcal{Q}_1$  be the set of all level-1 squares, and let  $\mathcal{Q}'_1 \subseteq \mathcal{Q}_1$  be the set of all non-empty level-1 squares. Recall that each square  $Q \in \mathcal{Q}_1$  has size  $(d_1 \times d_1)$ , where  $d_1 \leq \text{OPT}'/2^{c^* \sqrt{\log n} \log \log n}$ . For each square  $Q \in \mathcal{Q}'_1$ , we have created an extended square  $Q^+$ , by adding a margin of size  $d_1/\eta$  around  $Q$ . At the highest level of the recursion, the paths depart from the source vertices of the demand pairs we have chosen to route, and then visit the squares in  $\{Q^+ \mid Q \in \mathcal{Q}'_1\}$  one-by-one, in a snake-like fashion, entering and leaving each such square  $Q^+$  on a pre-selected set of vertices. Recall that we have assumed that  $d > \delta \cdot 2^{c \sqrt{\log n} \log \log n}$ , and  $d > \text{OPT}'/\eta$ . Then it is easy to see that the top  $d/2 \gg \delta$  rows of  $G$  are disjoint from the squares of  $\mathcal{Q}'_1$ , each of which must contain a destination vertex. As the algorithm routes at most  $d_1$  paths, this routing can be accomplished via canonical paths. The recursively defined routing inside the level-1 squares remains unchanged.

From now on, we assume that  $d < \text{OPT}'/\eta$ , and we now follow the algorithm from Section 9.

As before, we consider three sub-cases. Recall that  $\Gamma'$  is the boundary edge of  $G$  containing the vertices of  $\tilde{T}(\tilde{\mathcal{M}})$ . In the first sub-case,  $\Gamma'$  is the the bottom boundary of the grid. The algorithm constructs a number of square grids of height less than  $\ell/2$ , whose bottom boundaries are contained in the bottom boundary of the grid  $G$ . Each such sub-grid defines a modified instance, that is solved separately. Eventually, we need to connect the source vertices of the routed demand pairs to their counterparts in the modified instances. As the top  $\ell/2$  rows of  $G$  are disjoint from the modified instances, it is easy to see that this can be done via paths that are canonical.

The second sub-case is when  $\Gamma'$  is either the left or the right boundary edge of  $G$ ; assume w.l.o.g. that it is the left boundary edge. In this case, we have discarded the demand pairs whose destinations lie in the top  $\text{OPT}'/16$  rows of the grid  $G$ . As before, we then define a collection of modified instances, each of which is defined over a square sub-grid of  $G$  of width less than  $\ell/2$ . The left boundary of each sub-grid is contained in the left boundary of  $G$ , and each such sub-grid is disjoint from the top  $\text{OPT}'/16$  rows of  $G$ . As before, we need to connect the source vertices of the demand pairs routed in the modified sub-instances to their counterparts. In order to do so, we utilize the top  $\text{OPT}'/16$  rows and the rightmost  $\ell/2$  columns of  $G$ . This provides sufficient space to perform the routing via canonical paths, as  $\delta \ll d < \text{OPT}'/\eta$ .

The final sub-case is when  $\Gamma' = \Gamma$ . This case is in turn partitioned into two sub-cases. The first sub-case is when a large number of the demand pairs routed by the optimal solution belong to the set  $\mathcal{M}^0$ : the set of all demand pairs  $(s, t)$  with  $|\text{col}(s) - \text{col}(t)| \leq 2d$ . In this case, we defined a collection of disjoint sub-grids of  $G$ , of size  $(\Theta(d) \times \Theta(d))$ , reducing the problem to a number of disjoint sub-instances that are dealt with using Theorem 4.1. In each such sub-instance, we can ensure that the routing is canonical as before, since the value of the optimal solution in each sub-instance is  $O(d)$ , and  $\delta \ll d$ . In the second sub-case, we partition the grid  $G$  into a number of sub-instances, where each sub-instance is defined by a pair of consecutive vertical strips of  $G$ . One of the two strips contains all source vertices, and the other all destination vertices of the resulting sub-instance. We then define a modified instance inside the strip containing the destinations, and route the source vertices to the corresponding sources of the routed modified demand pairs. This routing utilizes the vertical strip containing the sources, and the bottom half of the vertical strip containing the destinations. As before,

it is easy to ensure that this routing is canonical.

## A Proofs Omitted from Section 4

### A.1 Proof of Claim 4.2

Before we define a hierarchical partition of  $\tilde{G}$  into squares, we need to define a hierarchical system of intervals.

**Definition.** *Given an integer  $1 \leq \rho' \leq \rho$ , a  $\rho'$ -hierarchical system of intervals is a sequence  $\mathcal{H} = (\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_{\rho'})$  of sets of intervals, such that:*

- *for all  $1 \leq h \leq \rho'$ ,  $\mathcal{I}_h$  is a  $d_h$ -canonical family of intervals; and*
- *for all  $1 < h \leq \rho'$ , for every interval  $I \in \mathcal{I}_h$ , there is an interval  $I' \in \mathcal{I}_{h-1}$ , such that  $I \subseteq I'$ .*

We let  $U(\mathcal{H}) = \bigcup_{I \in \mathcal{I}_{\rho'}} I$ , and we say that the integers in  $U(\mathcal{H})$  belong to the system  $\mathcal{H}$ .

We use the following simple observation.

**Observation A.1** *There is an efficient algorithm that constructs a collection  $\mathcal{H}_1, \dots, \mathcal{H}_{2^{\rho}}$  of  $2^{\rho}$   $\rho$ -hierarchical systems of intervals, such that every integer in  $[\ell']$  belongs to exactly one such system.*

**Proof:** It is enough to prove that there is an efficient algorithm, that, given an integer  $1 \leq \rho' \leq \rho$ , constructs a collection  $\mathcal{H}_1, \dots, \mathcal{H}_{2^{\rho'}}$  of  $2^{\rho'}$   $\rho'$ -hierarchical systems of intervals, such that every integer in  $[\ell']$  belongs to exactly one such system. The proof is by induction on  $\rho'$ . The base case is when  $\rho' = 1$ . We partition  $[\ell']$  into consecutive intervals, where every interval contains exactly  $d_1$  integers (recall that  $\ell'$  is an integral multiple of  $d_1 = \eta^{\rho+2}$ ). Let  $(I_1, I_2, \dots, I_r)$  be the resulting sequence of intervals, where we assume that the intervals appear in the sequence in their natural order. Let  $\mathcal{I}_1$  be the set of all odd-indexed intervals and  $\mathcal{I}'_1$  the set of all even-indexed intervals in the sequence. Clearly, each of  $\mathcal{I}_1$  and  $\mathcal{I}'_1$  is a  $d_1$ -canonical set of intervals. We define two 1-hierarchical systems, the first one containing only the set  $\mathcal{I}_1$ , and the second one containing only the set  $\mathcal{I}'_1$ . Note that every integer in  $[\ell']$  belongs to exactly one resulting system.

We assume now that the statement holds for all integers between 1 and  $(\rho' - 1)$ , for some  $\rho' > 1$ , and we prove it for  $\rho'$ . We assume that we are given a collection of  $2^{\rho'-1}$   $(\rho' - 1)$ -hierarchical systems of intervals, such that every integer in  $[\ell']$  belongs to exactly one system. Let  $\mathcal{H}$  be one such  $(\rho' - 1)$ -hierarchical system of intervals. We will construct two  $\rho'$ -systems,  $\mathcal{H}'$  and  $\mathcal{H}''$ , such that every integer that belongs to  $\mathcal{H}$  will belong to exactly one of the two systems. This is enough in order to complete the proof of the observation.

Assume that  $\mathcal{H} = (\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_{\rho'-1})$ . For simplicity, denote  $\mathcal{I}_{\rho'-1}$  by  $\mathcal{I}$ . We now construct two new sets  $\mathcal{I}', \mathcal{I}''$  of intervals, as follows. Start with  $\mathcal{I}' = \mathcal{I}'' = \emptyset$ , and process every interval  $I \in \mathcal{I}$  one-by-one. Consider some interval  $I \in \mathcal{I}$ . We partition  $I$  into consecutive intervals containing exactly  $d_{\rho'}$  integers each. Let  $\{I_1, \dots, I_r\}$  be the resulting partition, where we assume that the intervals are indexed in their natural order. We add to  $\mathcal{I}'$  all resulting odd-indexed intervals, and to  $\mathcal{I}''$  all resulting even-indexed intervals. Once every interval  $I \in \mathcal{I}$  is processed in this manner, we obtain our final sets  $\mathcal{I}', \mathcal{I}''$  of intervals. It is immediate to verify that each set  $\mathcal{I}', \mathcal{I}''$  is  $d_{\rho'}$ -canonical; that  $\bigcup_{I' \in \mathcal{I}' \cup \mathcal{I}''} I' = \bigcup_{I \in \mathcal{I}} I$ ; and that every integer of  $\bigcup_{I \in \mathcal{I}} I$  belongs to exactly one interval of  $\mathcal{I}' \cup \mathcal{I}''$ . We then set  $\mathcal{H}' = (\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_{\rho'-1}, \mathcal{I}')$ , and  $\mathcal{H}'' = (\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_{\rho'-1}, \mathcal{I}'')$ .  $\square$

From Observation A.1, we can construct  $2^\rho$  hierarchical  $\rho$ -systems  $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_{2^\rho}$  of intervals of  $[\ell']$ . For every pair  $1 \leq i, j \leq 2^\rho$  of integers, we construct a single hierarchical family  $\tilde{\mathcal{H}}_{i,j}$  of squares, such that  $V(\tilde{\mathcal{H}}_{i,j}) = \{v_{x,y} \mid x \in U(\mathcal{H}_i), y \in U(\mathcal{H}_j)\}$ . Since every integer in  $[\ell']$  belongs to exactly one set  $U(\mathcal{H}_z)$  for  $1 \leq z \leq 2^\rho$ , it is immediate to verify that every vertex of  $G'$  belongs to exactly one resulting hierarchical family  $\tilde{\mathcal{H}}_{i,j}$  of squares.

We now define the construction of the system  $\tilde{\mathcal{H}}_{i,j} = (\mathcal{Q}_1^{i,j}, \mathcal{Q}_2^{i,j}, \dots, \mathcal{Q}_\rho^{i,j})$ . Denote  $\mathcal{H}_i = (I_1, \dots, I_\rho)$  and  $\mathcal{H}_j = (I'_1, \dots, I'_\rho)$ . The construction is simple: for all  $1 \leq r \leq \rho$ , we let  $\mathcal{Q}_r^{i,j} = \mathcal{Q}(\mathcal{I}_r, \mathcal{I}'_r)$ . From the above discussion, since sets  $\mathcal{I}_r, \mathcal{I}'_r$  are  $d_r$ -canonical, so is set  $\mathcal{Q}_r^{i,j}$ . It is also easy to verify that the set of vertices contained in the squares of  $\mathcal{Q}_r^{i,j}$  is exactly  $\{v(x, y) \mid x \in \bigcup_{I \in \mathcal{I}_r} I, y \in \bigcup_{I' \in \mathcal{I}'_r} I'\}$ , and that  $\tilde{\mathcal{H}}_{i,j}$  is indeed a hierarchical system of squares of  $G$ .

## A.2 Proof of Observation 4.3

Assume that  $\hat{\mathcal{M}} = \{(s_1, t_1), \dots, (s_z, t_z)\}$ , where the source vertices  $s_1, \dots, s_z$  appear in this left-to-right order on  $R^*$ . We then let  $\hat{\mathcal{M}}' = \{(s_i, t_i) \mid i \equiv 1 \pmod{2 \lceil \beta_2 / \beta_1 \rceil}\}$ . Clearly,  $|\hat{\mathcal{M}}'| \geq \left\lfloor \frac{|\hat{\mathcal{M}}|}{2 \lceil \beta_2 / \beta_1 \rceil} \right\rfloor \geq \left\lfloor \frac{\beta_1 |\hat{\mathcal{M}}|}{4 \beta_2} \right\rfloor$ .

We claim that every square  $Q \in \mathcal{Q}$  has the  $\beta_2$ -shadow property with respect to  $\hat{\mathcal{M}}'$ . Indeed, let  $Q \in \mathcal{Q}$  be any such square, and assume that its dimensions are  $(d \times d)$ . Then  $J_{\hat{\mathcal{M}}'}(Q) \subseteq J_{\hat{\mathcal{M}}}(Q)$ , and  $J_{\hat{\mathcal{M}}}(Q)$  contained at most  $\beta_2 d$  source vertices of the demand pairs of  $\hat{\mathcal{M}}$ . From our construction of  $\hat{\mathcal{M}}'$ , it contains at most  $\left\lceil \frac{\beta_2 d}{2 \lceil \beta_2 / \beta_1 \rceil} \right\rceil \leq \beta_1 d$  demand pairs of  $\hat{\mathcal{M}}'$ .

## B Proof of Claim 5.3

**Proof:** We compute a partition  $\mathcal{Y}(\tau)$  for every tree  $\tau \in F$  separately. The partition is computed in iterations, where in the  $j$ th iteration we compute the set  $Y_j(\tau) \subseteq V(\tau)$  of vertices, together with the corresponding collection  $\mathcal{P}_j(\tau)$  of paths. For the first iteration, if  $\tau$  contains a single vertex  $v$ , then we add this vertex to  $Y_1(\tau)$  and terminate the algorithm. Otherwise, for every leaf  $v$  of  $\tau$ , let  $P(v)$  be the longest directed path of  $\tau$ , starting at  $v$ , that only contains degree-1 and degree-2 vertices, and does not contain the root of  $\tau$ . We then add the vertices of  $P(v)$  to  $Y_1(\tau)$ , and the path  $P(v)$  to  $\mathcal{P}_1(\tau)$ . Once we process all leaf vertices of  $\tau$ , the first iteration terminates. It is easy to see that all resulting vertices in  $Y_1(\tau)$  induce a collection  $\mathcal{P}_1(\tau)$  of disjoint paths in  $\tau$ , and moreover if  $v, v' \in Y_1(\tau)$ , and there is a path from  $v$  to  $v'$  in  $\tau$ , then  $v, v'$  lie on the same path in  $\mathcal{P}_1(\tau)$ . We then delete all vertices of  $Y_1(\tau)$  from  $\tau$ .

The subsequent iterations are executed similarly, except that the tree  $\tau$  becomes smaller, since we delete all vertices that have been added to the sets  $Y_j(\tau)$  from the tree.

It is now enough to show that this process terminates after  $\lceil \log n \rceil$  iterations. In order to do so, we can describe each iteration slightly differently. Before each iteration starts, we gradually contract every edge  $e$  of the current tree, such that at least one endpoint of  $e$  has degree 2 in the tree, and  $e$  is not incident on the root of  $\tau$ . We then obtain a tree in which every inner vertex (except possibly the root) has degree at least 3, and delete all leaves from this tree. The number of vertices remaining in the contracted tree after each such iteration therefore decreases by at least factor 2. It is easy to see that the number of iteration in this procedure is the same as the number of iterations in our algorithm, and is bounded by  $\lceil \log n \rceil$ . For each  $1 \leq j \leq \lceil \log n \rceil$ , we then set  $Y_j = \bigcup_{\tau \in F} Y_j(\tau)$ .  $\square$

## C Proof of Claim 9.1

Let  $\mathcal{P}^*$  be the optimal solution to instance  $(G, \mathcal{M})$ . If  $|\mathcal{P}^*| > d$ , then we discard paths from  $\mathcal{P}^*$  arbitrarily, until  $|\mathcal{P}^*| = d$  holds. Recall that  $G'$  is a sub-grid of  $G$  spanned by some subset  $\mathcal{W}'$  of its columns and some subset  $\mathcal{R}'$  of its rows. Recall also that we have defined a sub-grid  $G'' \subseteq G'$ , obtained from  $G'$  by deleting  $4d$  of its leftmost columns,  $4d$  of its rightmost columns, and  $|\mathcal{W}'| - 4d$  of its topmost rows.

Let  $\mathcal{M}^* \subseteq \mathcal{M}'$  be the set of the demand pairs routed by  $\mathcal{P}^*$ . For every path  $P \in \mathcal{P}^*$ , we define a collection  $\Sigma(P)$  of sub-paths of  $P$ , as follows. Assume that  $P$  routes some demand pair  $(s, t)$ . Let  $x_1, x_2, \dots, x_r$  be all vertices of  $\Gamma(G'')$  that appear on  $P$ , and assume that they appear on  $P$  in this order (where we view  $P$  as directed from  $s$  to  $t$ ). Denote  $x_0 = s$  and  $x_{r+1} = t$ . We then let  $\Sigma(P)$  contain, for each  $0 \leq i \leq r$ , the sub-path of  $P$  from  $x_i$  to  $x_{i+1}$ . We say that a segment  $\sigma \in \Sigma(P)$  is of type 1 if one of its endpoints is the source  $s$ ; we say that it is of type 2 if both its endpoints belong to  $\Gamma(G'')$ , and  $\sigma$  is internally disjoint from  $G''$ ; otherwise we say that it is of type 3. We let  $\Sigma_1$  contain all type-1 segments in all sets  $\Sigma(P)$  for  $P \in \mathcal{P}^*$ , and we define  $\Sigma_2$  and  $\Sigma_3$  similarly for all type-2 and type-3 segments.

Notice that all type-3 segments are contained in  $G''$ . Let  $\tilde{\mathcal{M}}_2$  be the set of all pairs  $(u, v)$  of vertices, such that some segment  $\sigma \in \Sigma_2$  connects  $u$  to  $v$ . We also define a set  $\tilde{\mathcal{M}}_1$  of pairs of vertices, corresponding to the segments of  $\Sigma_1$  as follows. Let  $\sigma \in \Sigma_1$  be any type-1 segment, and assume that its endpoints are  $s$  and  $v$ , with  $s \in S(\mathcal{M}')$  and  $v \in \Gamma(G'')$ . Let  $s' \in \Gamma(G')$  be the new source vertex to which  $s$  was mapped. Then we add  $(s', v)$  to  $\tilde{\mathcal{M}}_1$ . In order to complete the proof of the claim, it is now enough to prove the following observation.

**Observation C.1** *There is a set  $\mathcal{P}$  of node-disjoint paths that routes all pairs in  $\tilde{\mathcal{M}}_1 \cup \tilde{\mathcal{M}}_2$  in graph  $G'$ , so that the paths in  $\mathcal{P}$  are internally disjoint from  $G''$ .*

Indeed, combining the paths in  $\mathcal{P}$  with the segments in  $\Sigma_3$  provides a set of node-disjoint paths in graph  $G'$  that routes  $|\mathcal{M}^*|$  demand pairs of  $\mathcal{M}''$  – the demand pairs corresponding to the pairs in  $\mathcal{M}^*$ .

The proof of the above observation is straightforward; we only provide its sketch here. Let  $J$  be the sub-path of  $\Gamma(G'')$ , obtained by deleting all vertices lying on the bottom boundary of  $G''$  from it, excluding the two bottom corners of  $G''$ . For every pair  $(u, v) \in \tilde{\mathcal{M}}_2$ , we can think of the corresponding sub-path of  $J$  between  $u$  and  $v$  as an interval  $I(u, v)$ . It is immediate to verify that all intervals defined by the pairs in  $\tilde{\mathcal{M}}_2$  are nested: that is, if  $(u, v), (u', v') \in \tilde{\mathcal{M}}_2$ , then either intervals  $I(u, v)$  and  $I(u', v')$  are disjoint, or one of them is contained in the other.

Let  $\mathcal{I} = \{I(u, v) \mid (u, v) \in \tilde{\mathcal{M}}_2\}$  be the corresponding set of intervals. Notice that no interval in  $\mathcal{I}$  may contain a destination vertex  $v$  of any demand pair  $(s, v) \in \tilde{\mathcal{M}}_1$ . Let  $\mathcal{I}^0 \subseteq \mathcal{I}$  be the set of all intervals containing both the top left and the top right corners of  $G''$ . Let  $\mathcal{I}^1 \subseteq \mathcal{I} \setminus \mathcal{I}^0$  be the set of all intervals containing the top left corner of  $G''$ , and similarly, let  $\mathcal{I}^2 \subseteq \mathcal{I} \setminus \mathcal{I}^0$  be the set of all intervals containing the top right corner of  $G''$ . Observe that for all  $0 \leq i \leq 2$ , for every pair  $I, I' \in \mathcal{I}^i$  of intervals, one of the intervals is contained in the other. Let  $\mathcal{I}' = \mathcal{I} \setminus \left(\bigcup_{i=0}^2 \mathcal{I}^i\right)$  be the set of all remaining intervals.

We partition the intervals of  $\mathcal{I}'$  into levels, as follows. We say that  $I(u, v)$  is a level-1 interval, iff no other interval of  $\mathcal{I}'$  is contained in it. Let  $\mathcal{I}_1$  denote the set of all level-1 intervals. Assume now that we have defined the sets  $\mathcal{I}_1, \dots, \mathcal{I}_i$  of intervals of levels  $1, \dots, i$ , respectively. We say that an interval  $I \in \mathcal{I}'$  belongs to level  $(i + 1)$  iff it does not contain any interval from set  $\mathcal{I}' \setminus \left(\bigcup_{i'=1}^i \mathcal{I}_{i'}\right)$ .

We say that an interval  $I(u, v) \in \mathcal{I}'$  is a left interval, iff  $u$  and  $v$  belong to the left boundary edge of  $G''$ . Similarly, we say that  $I(u, v) \in \mathcal{I}'$  is a right interval, iff  $u$  and  $v$  belong to the right boundary edge

of  $G''$ . Otherwise, we say that it is a top interval. In that case, both  $u$  and  $v$  must belong to the top boundary of  $G''$ . It is easy to verify that, if  $I$  is a left or a right interval, then it must belong to levels  $1, \dots, 2d$ , as the height of the grid  $G''$  is  $4d$ . Moreover, if  $h$  is the largest level to which a left interval belongs, then  $h + |\mathcal{I}_0| + |\mathcal{I}_1| \leq 4d$ , and if  $h'$  is the largest level to which a right interval belongs, then  $h + |\mathcal{I}_0| + |\mathcal{I}_2| \leq 4d$ .

Let  $W'_1, W'_2, \dots, W'_{4d}$  denote the  $4d$  columns of  $G'$  that lie to the left of  $G''$ , and assume that they are indexed in the right-to-left order. For each level  $1 \leq r \leq 4d$ , for every level- $r$  interval  $I(u, v)$ , we route the demand pair  $(u, v)$  via the column  $W'_r$  in a straightforward manner: we connect  $u$  and  $v$  to  $W'_r$  by horizontal paths, and then complete the routing using the corresponding sub-path of  $W'_r$ .

The routing of the right intervals is performed similarly. The top intervals are routed similarly by exploiting the rows of  $G'$  that lie above  $G''$  – recall that there are  $|\mathcal{W}''| + 4d$  such rows, where  $|\mathcal{W}''|$  is the width of  $G''$ . We then add paths routing the pairs corresponding to the intervals in  $\mathcal{I}^1, \mathcal{I}^2$  and  $\mathcal{I}^0$  in a straightforward manner (see Figure 10).

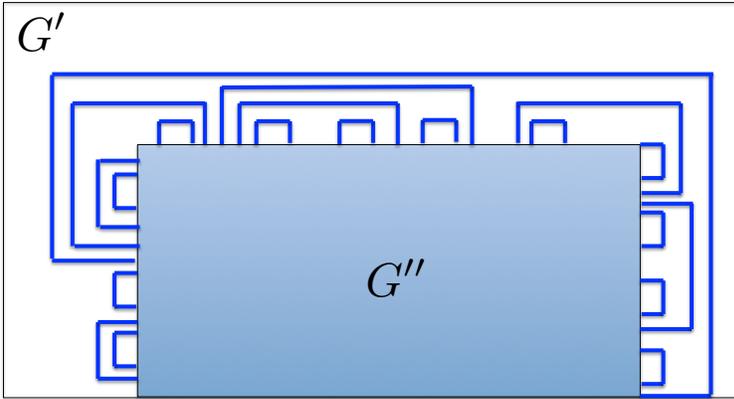


Figure 10: Routing in the modified graph.

Let  $H$  be the graph obtained from  $G'$ , after we delete all non-boundary vertices of  $G''$  from it, and all vertices that participate in the routing of the pairs in  $\tilde{\mathcal{M}}_2$  that we just defined. It is easy to verify that all demand pairs in  $\tilde{\mathcal{M}}_1$  can be routed in  $H$ . In order to do so, we set up a flow network, where we start from the graph  $H$ , and add two special vertices  $s$  and  $t$  to it. We connect  $s$  to every vertex in  $S(\tilde{\mathcal{M}}_1)$ , and we connect  $t$  to every vertex in  $T(\tilde{\mathcal{M}}_1)$ , setting the capacity of every vertex of  $H$  to be 1. It is easy to verify that there is an  $s$ - $t$  flow of value  $|\tilde{\mathcal{M}}_1|$  in this network (since every cut separating  $s$  from  $t$  must contain at least  $|\tilde{\mathcal{M}}_1|$  vertices). From the integrality of flow, and due to the way in which the mapping between the vertices of  $S(\mathcal{M}')$  and  $S(\mathcal{M}'')$  was defined, we can obtain an integral routing of all demand pairs in  $\tilde{\mathcal{M}}_1$  via node-disjoint paths in  $H$ .

## References

- [ACG<sup>+</sup>10] Matthew Andrews, Julia Chuzhoy, Venkatesan Guruswami, Sanjeev Khanna, Kunal Talwar, and Lisa Zhang. Inapproximability of edge-disjoint paths and low congestion routing on undirected graphs. *Combinatorica*, 30(5):485–520, 2010.
- [AKW00] Alok Aggarwal, Jon Kleinberg, and David P. Williamson. Node-disjoint paths on the mesh and a new trade-off in VLSI layout. *SIAM J. Comput.*, 29(4):1321–1333, February 2000.

- [And10] Matthew Andrews. Approximation algorithms for the edge-disjoint paths problem via Raecke decompositions. In *Proceedings of IEEE FOCS*, pages 277–286, 2010.
- [AR95] Yonatan Aumann and Yuval Rabani. Improved bounds for all optical routing. In *Proceedings of the sixth annual ACM-SIAM symposium on Discrete algorithms*, SODA '95, pages 567–576, Philadelphia, PA, USA, 1995. Society for Industrial and Applied Mathematics.
- [AZ06] Matthew Andrews and Lisa Zhang. Logarithmic hardness of the undirected edge-disjoint paths problem. *J. ACM*, 53(5):745–761, September 2006.
- [CC] Chandra Chekuri and Julia Chuzhoy. Half-integral all-or-nothing flow. Unpublished Manuscript.
- [CE13] Chandra Chekuri and Alina Ene. Poly-logarithmic approximation for maximum node disjoint paths with constant congestion. In *Proc. of ACM-SIAM SODA*, 2013.
- [Chu16] Julia Chuzhoy. Routing in undirected graphs with constant congestion. *SIAM J. Comput.*, 45(4):1490–1532, 2016.
- [CK15] Julia Chuzhoy and David H. K. Kim. On approximating node-disjoint paths in grids. In Naveen Garg, Klaus Jansen, Anup Rao, and José D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, AP-PROX/RANDOM 2015, August 24-26, 2015, Princeton, NJ, USA*, volume 40 of *LIPICs*, pages 187–211. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.
- [CKL16] Julia Chuzhoy, David H. K. Kim, and Shi Li. Improved approximation for node-disjoint paths in planar graphs. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2016, pages 556–569, New York, NY, USA, 2016. ACM.
- [CKN17a] Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat. Almost polynomial hardness of node-disjoint paths in grids. Unpublished Manuscript, 2017.
- [CKN17b] Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat. New hardness results for routing on disjoint paths. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 86–99. ACM, 2017.
- [CKS04] Chandra Chekuri, Sanjeev Khanna, and F Bruce Shepherd. Edge-disjoint paths in planar graphs. In *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on*, pages 71–80. IEEE, 2004.
- [CKS05] Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd. Multicommodity flow, well-linked terminals, and routing problems. In *Proc. of ACM STOC*, pages 183–192, 2005.
- [CKS06] Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd. An  $O(\sqrt{n})$  approximation and integrality gap for disjoint paths and unsplittable flow. *Theory of Computing*, 2(1):137–146, 2006.
- [CL16] 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.
- [CS78] M. Cutler and Y. Shiloach. Permutation layout. *Networks*, 8:253–278, 1978.

- [DP09] Devdatt Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, New York, NY, USA, 1st edition, 2009.
- [EIS76] Shimon Even, Alon Itai, and Adi Shamir. On the complexity of timetable and multicommodity flow problems. *SIAM J. Comput.*, 5(4):691–703, 1976.
- [FMS16] Krzysztof Fleszar, Matthias Mnich, and Joachim Spoerhase. New Algorithms for Maximum Disjoint Paths Based on Tree-Likeness. In Piotr Sankowski and Christos Zaroliagis, editors, *24th Annual European Symposium on Algorithms (ESA 2016)*, volume 57 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 42:1–42:17, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [Kar75] R. Karp. On the complexity of combinatorial problems. *Networks*, 5:45–68, 1975.
- [KK13] Ken-Ichi Kawarabayashi and Yusuke Kobayashi. An  $O(\log n)$ -approximation algorithm for the edge-disjoint paths problem in Eulerian planar graphs. *ACM Trans. Algorithms*, 9(2):16:1–16:13, March 2013.
- [Kle05] Jon Kleinberg. An approximation algorithm for the disjoint paths problem in even-degree planar graphs. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science, FOCS '05*, pages 627–636, Washington, DC, USA, 2005. IEEE Computer Society.
- [KS04] Stavros G. Kolliopoulos and Clifford Stein. Approximating disjoint-path problems using packing integer programs. *Mathematical Programming*, 99:63–87, 2004.
- [KT95] Jon M. Kleinberg and Éva Tardos. Disjoint paths in densely embedded graphs. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, pages 52–61, 1995.
- [KT98] Jon M. Kleinberg and Éva Tardos. Approximations for the disjoint paths problem in high-diameter planar networks. *J. Comput. Syst. Sci.*, 57(1):61–73, 1998.
- [KvL84] MR Kramer and Jan van Leeuwen. The complexity of wire-routing and finding minimum area layouts for arbitrary vlsi circuits. *Advances in computing research*, 2:129–146, 1984.
- [Lyn75] James F. Lynch. The equivalence of theorem proving and the interconnection problem. *SIGDA Newsl.*, 5(3):31–36, September 1975.
- [Räc02] Harald Räcke. Minimizing congestion in general networks. In *Proc. of IEEE FOCS*, pages 43–52, 2002.
- [RS88] Neil Robertson and Paul D. Seymour. Graph minors. VII. disjoint paths on a surface. *J. Comb. Theory, Ser. B*, 45(2):212–254, 1988.
- [RS90] N. Robertson and P. D. Seymour. Outline of a disjoint paths algorithm. In *Paths, Flows and VLSI-Layout*. Springer-Verlag, 1990.
- [RS95] Neil Robertson and Paul D Seymour. Graph minors. XIII. the disjoint paths problem. *Journal of Combinatorial Theory, Series B*, 63(1):65–110, 1995.
- [RT87] Prabhakar Raghavan and Clark D. Tompson. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7:365–374, December 1987.

- [RZ10] Satish Rao and Shuheng Zhou. Edge disjoint paths in moderately connected graphs. *SIAM J. Comput.*, 39(5):1856–1887, 2010.
- [SCS11] Loïc Seguin-Charbonneau and F. Bruce Shepherd. Maximum edge-disjoint paths in planar graphs with congestion 2. In *Proceedings of the 2011 IEEE 52Nd Annual Symposium on Foundations of Computer Science, FOCS '11*, pages 200–209, Washington, DC, USA, 2011. IEEE Computer Society.