

# 1 Improved Approximation for Node-Disjoint Paths 2 in Grids with Sources on the Boundary

3 **Julia Chuzhoy**<sup>1</sup>

4 Toyota Technological Institute at Chicago, 6045 S. Kenwood Ave., Chicago, Illinois 60637, USA  
5 cjulia@ttic.edu

6 **David H. K. Kim**<sup>2</sup>

7 Computer Science Department, University of Chicago, 1100 East 58th Street, Chicago, Illinois  
8 60637, USA  
9 hongk@cs.uchicago.edu

10 **Rachit Nimavat**<sup>3</sup>

11 Toyota Technological Institute at Chicago, 6045 S. Kenwood Ave., Chicago, Illinois 60637, USA  
12 nimavat@ttic.edu

## 13 — Abstract —

14 We study the classical Node-Disjoint Paths (NDP) problem: given an undirected  $n$ -vertex graph  
15  $G$ , together with a set  $\{(s_1, t_1), \dots, (s_k, t_k)\}$  of pairs of its vertices, called source-destination, or  
16 demand pairs, find a maximum-cardinality set  $\mathcal{P}$  of mutually node-disjoint paths that connect  
17 the demand pairs. The best current approximation for the problem is achieved by a simple greedy  
18  $O(\sqrt{n})$ -approximation algorithm. Until recently, the best negative result was an  $\Omega(\log^{1/2-\epsilon} n)$ -  
19 hardness of approximation, for any fixed  $\epsilon$ , under standard complexity assumptions.

20 A special case of the problem, where the underlying graph is a grid, has been studied extensively.  
21 The best current approximation algorithm for this special case achieves an  $\tilde{O}(n^{1/4})$ -approximation  
22 factor. On the negative side, a recent result by the authors shows that NDP is hard to approximate  
23 to within factor  $2^{\Omega(\sqrt{\log n})}$ , even if the underlying graph is a subgraph of a grid, and all source  
24 vertices lie on the grid boundary. In a very recent follow-up work, the authors further show that  
25 NDP in grid graphs is hard to approximate to within factor  $\Omega(2^{\log^{1-\epsilon} n})$  for any constant  $\epsilon$  under  
26 standard complexity assumptions, and to within factor  $n^{\Omega(1/(\log \log n)^2)}$  under randomized ETH.

27 In this paper we study the NDP problem in grid graphs, where all source vertices  $\{s_1, \dots, s_k\}$   
28 appear on the grid boundary. Our main result is an efficient randomized  $2^{O(\sqrt{\log n \cdot \log \log n})}$ -  
29 approximation algorithm for this problem. Our result in a sense complements the  $2^{\Omega(\sqrt{\log n})}$ -  
30 hardness of approximation for sub-graphs of grids with sources lying on the grid boundary, and  
31 should be contrasted with the above-mentioned almost polynomial hardness of approximation of  
32 NDP in grid graphs (where the sources and the destinations may lie anywhere in the grid).

33 Much of the work on approximation algorithms for NDP relies on the multicommodity flow  
34 relaxation of the problem, which is known to have an  $\Omega(\sqrt{n})$  integrality gap, even in grid graphs,  
35 with all source and destination vertices lying on the grid boundary. Our work departs from this  
36 paradigm, and uses a (completely different) linear program only to select the pairs to be routed,  
37 while the routing itself is computed by other methods.

<sup>1</sup> Supported in part by NSF grants CCF-1318242 and CCF-1616584.

<sup>2</sup> Supported in part by NSF grants CCF-1318242 and CCF-1616584.

<sup>3</sup> Supported in part by NSF grant CCF-1318242.



38 **2012 ACM Subject Classification** Theory of computation → Routing and network design prob-  
39 lems

40 **Keywords and phrases** Node-disjoint paths, approximation algorithms, routing and layout

41 **Digital Object Identifier** 10.4230/LIPIcs.ICALP.2018.<article-no>

## 42 **1** Introduction

43 We study the classical Node-Disjoint Paths (NDP) problem, where the input consists of  
44 an undirected  $n$ -vertex graph  $G$  and a collection  $\mathcal{M} = \{(s_1, t_1), \dots, (s_k, t_k)\}$  of pairs of its  
45 vertices, called *source-destination* or *demand* pairs. We say that a path  $P$  *routes* a demand  
46 pair  $(s_i, t_i)$  iff the endpoints of  $P$  are  $s_i$  and  $t_i$ . The goal is to compute a maximum-  
47 cardinality set  $\mathcal{P}$  of node-disjoint paths, where each path  $P \in \mathcal{P}$  routes a distinct demand  
48 pair in  $\mathcal{M}$ . We denote by NDP-Planar the special case of the problem when the underlying  
49 graph  $G$  is planar, and by NDP-Grid the special case where  $G$  is a square grid<sup>4</sup>. We refer to  
50 the vertices in set  $S = \{s_1, \dots, s_k\}$  as *source vertices*; to the vertices in set  $T = \{t_1, \dots, t_k\}$   
51 as *destination vertices*, and to the vertices in set  $S \cup T$  as *terminals*.

52 NDP is a fundamental graph routing problem that has been studied extensively in both  
53 graph theory and theoretical computer science communities. Robertson and Seymour [31, 33]  
54 explored the problem in their Graph Minor series, providing an efficient algorithm for NDP  
55 when the number  $k$  of the demand pairs is bounded by a constant. But when  $k$  is a part  
56 of input, the problem becomes NP-hard [20, 18], even in planar graphs [27], and even in  
57 grid graphs [26]. The best current approximation factor of  $O(\sqrt{n})$  for NDP is achieved by a  
58 simple greedy algorithm [25]. Until recently, this was also the best approximation algorithm  
59 for NDP-Planar and NDP-Grid. A natural way to design approximation algorithms for NDP  
60 is via the multicommodity flow relaxation: instead of connecting each routed demand pair  
61 with a path, send maximum possible amount of (possibly fractional) flow between them.  
62 The optimal solution to this relaxation can be computed via a standard linear program.  
63 The  $O(\sqrt{n})$ -approximation algorithm of [25] can be cast as an LP-rounding algorithm of  
64 this relaxation. Unfortunately, it is well-known that the integrality gap of this relaxation is  
65  $\Omega(\sqrt{n})$ , even when the underlying graph is a grid, with all terminals lying on its boundary. In  
66 a recent work, Chuzhoy and Kim [12] designed an  $\tilde{O}(n^{1/4})$ -approximation for NDP-Grid, thus  
67 bypassing this integrality gap barrier. Their main observation is that, if all terminals lie close  
68 to the grid boundary (say within distance  $O(n^{1/4})$ ), then a simple dynamic programming-  
69 based algorithm yields an  $O(n^{1/4})$ -approximation. On the other hand, if, for every demand  
70 pair, either the source or the destination lies at a distance at least  $\Omega(n^{1/4})$  from the grid  
71 boundary, then the integrality gap of the multicommodity flow relaxation improves, and  
72 one can obtain an  $\tilde{O}(n^{1/4})$ -approximation via LP-rounding. A natural question is whether  
73 the integrality gap improves even further, if all terminals lie further away from the grid  
74 boundary. Unfortunately, the authors show in [12] that the integrality gap remains at  
75 least  $\Omega(n^{1/8})$ , even if all terminals lie within distance  $\Omega(\sqrt{n})$  from the grid boundary. The  
76  $\tilde{O}(n^{1/4})$ -approximation algorithm for NDP-Grid was later extended and generalized to an

---

<sup>4</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.

77  $\tilde{O}(n^{9/19})$ -approximation algorithm for NDP-Planar [13].

78 On the negative side, until recently, only an  $\Omega(\log^{1/2-\epsilon} n)$ -hardness of approximation was  
 79 known for the general version of NDP, for any constant  $\epsilon$ , unless  $\text{NP} \subseteq \text{ZPTIME}(n^{\text{poly} \log n})$  [4,  
 80 3], and only APX-hardness was known for NDP-Planar and NDP-Grid [12]. In a recent  
 81 work [15], the authors have shown that NDP is hard to approximate to within a  $2^{\Omega(\sqrt{\log n})}$   
 82 factor unless  $\text{NP} \subseteq \text{DTIME}(n^{O(\log n)})$ , even if the underlying graph is a planar graph with  
 83 maximum vertex degree at most 3, and all source vertices lie on the boundary of a single  
 84 face. The result holds even when the input graph  $G$  is a vertex-induced subgraph of a  
 85 grid, with all sources lying on the grid boundary. In a very recent work [14], the au-  
 86 thors show that NDP-Grid is  $2^{\Omega(\log^{1-\epsilon} n)}$ -hard to approximate for any constant  $\epsilon$  assuming  
 87  $\text{NP} \not\subseteq \text{BPTIME}(n^{\text{poly} \log n})$ , and moreover, assuming randomized ETH, the hardness of ap-  
 88 proximation factor becomes  $n^{\Omega(1/(\log \log n)^2)}$ . We note that the instances constructed in these  
 89 latter hardness proofs require all terminals to lie far from the grid boundary.

90 In this paper we explore NDP-Grid. This important special case of NDP was initially motiv-  
 91 ated by applications in VLSI design, and has received a lot of attention since the 1960's. We  
 92 focus on a restricted version of NDP-Grid, that we call Restricted NDP-Grid: here, in addition  
 93 to the graph  $G$  being a square grid, we also require that all source vertices  $\{s_1, \dots, s_k\}$  lie  
 94 on the grid boundary. We do not make any assumptions about the locations of the des-  
 95 tination vertices, that may appear anywhere in the grid. The best current approximation  
 96 algorithm for Restricted NDP-Grid is the same as that for the general NDP-Grid, and achieves  
 97 a  $\tilde{O}(n^{1/4})$ -approximation [12]. Our main result is summarized in the following theorem.

98 **► Theorem 1.** *There is an efficient randomized  $2^{O(\sqrt{\log n \cdot \log \log n})}$ -approximation algorithm  
 99 for Restricted NDP-Grid.*

100 This result in a sense complements the  $2^{\Omega(\sqrt{\log n})}$ -hardness of approximation of NDP on  
 101 sub-graphs of grids with all sources lying on the grid boundary of [15]<sup>5</sup>, and should be  
 102 contrasted with the recent almost polynomial hardness of approximation of [14] for NDP-  
 103 Grid mentioned above. Our algorithm departs from previous work on NDP in that it does  
 104 not use the multicommodity flow relaxation. Instead, we define sufficient conditions that  
 105 allow us to route a subset  $\mathcal{M}'$  of demand pairs via disjoint paths, and show that there  
 106 exists a subset of demand pairs satisfying these conditions, whose cardinality is at least  
 107  $\text{OPT}/2^{O(\sqrt{\log n \cdot \log \log n})}$ , where OPT is the value of the optimal solution. It is then enough  
 108 to compute a maximum-cardinality subset of the demand pairs satisfying these conditions.  
 109 We write an LP-relaxation for this problem and design a  $2^{O(\sqrt{\log n \cdot \log \log n})}$ -approximation  
 110 LP-rounding algorithm for it. We emphasize that the linear program is only used to select  
 111 the demand pairs to be routed, and not to compute the routing itself.

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

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

<sup>5</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.


117 We note that for instances of NDP-Grid where both the sources and the destinations are  
 118 within distance at most  $\delta$  from the grid boundary, it is easy to obtain an efficient  $O(\delta)$ -  
 119 approximation algorithm (see, e.g. [12]).

120 A problem closely related to NDP is the Edge-Disjoint Paths (EDP) problem. It is defined  
 121 similarly, except that now the paths chosen to route the demand pairs may share vertices,  
 122 and are only required to be edge-disjoint. The approximability status of EDP is very similar  
 123 to that of NDP: there is an  $O(\sqrt{n})$ -approximation algorithm [10], and an  $\Omega(\log^{1/2-\epsilon} n)$ -  
 124 hardness of approximation for any constant  $\epsilon$ , unless  $\text{NP} \subseteq \text{ZPTIME}(n^{\text{poly} \log n})$  [4, 3]. As  
 125 in the NDP problem, we can use the standard multicommodity flow LP-relaxation of the  
 126 problem, in order to obtain the  $O(\sqrt{n})$ -approximation algorithm, and the integrality gap of  
 127 the LP-relaxation is  $\Omega(\sqrt{n})$  even in planar graphs. Recently, Fleszar et al. [19] designed an  
 128  $O(\sqrt{r} \cdot \log(kr))$ -approximation algorithm for EDP, where  $r$  is the feedback vertex set number  
 129 of the input graph  $G = (V, E)$  — the smallest number of vertices that need to be deleted  
 130 from  $G$  in order to turn it into a forest.

131 Several special cases of EDP have better approximation algorithms: an  $O(\log^2 n)$ -approximation  
 132 is known for even-degree planar graphs [9, 8, 22], and an  $O(\log n)$ -approximation is known  
 133 for nearly-Eulerian uniformly high-diameter planar graphs, and nearly-Eulerian densely em-  
 134 bedded graphs, including grid graphs [5, 24, 23]. Furthermore, an  $O(\log n)$ -approximation  
 135 algorithm is known for EDP on 4-edge-connected planar, and Eulerian planar graphs [21]. It  
 136 appears that the restriction of the graph  $G$  to be Eulerian, or near-Eulerian, makes the EDP  
 137 problem on planar graphs significantly simpler, and in particular improves the integrality  
 138 gap of the standard multicommodity flow LP-relaxation.

139 The analogue of the grid graph for the EDP problem is the wall graph (see Figure 1): the  
 140 integrality gap of the multicommodity flow relaxation for EDP on wall graphs is  $\Omega(\sqrt{n})$ . The  
 141  $\tilde{O}(n^{1/4})$ -approximation algorithm of [12] for NDP-Grid extends to EDP on wall graphs, and  
 142 the  $2^{\Omega(\sqrt{\log n})}$ -hardness of approximation of [15] for NDP-Planar also extends to EDP on sub-  
 143 graphs of walls, with all sources lying on the top boundary of the wall. The recent hardness  
 144 result of [14] for NDP-Grid also extends to an  $2^{\Omega(\log^{1-\epsilon} n)}$ -hardness of EDP on wall graphs,  
 145 assuming  $\text{NP} \not\subseteq \text{BPTIME}(n^{\text{poly} \log n})$ , and to  $n^{\Omega(1/(\log \log n)^2)}$ -hardness assuming randomized  
 146 ETH. We extend our results to EDP and NDP on wall graphs:

147 ► **Theorem 3.** *There is an efficient randomized  $2^{O(\sqrt{\log n \cdot \log \log n})}$ -approximation algorithm  
 148 for EDP and for NDP on wall graphs, when all source vertices lie on the wall boundary.*



■ **Figure 1** A wall graph.

149 **Other related work.**

150 Cutler and Shiloach [17] studied an even more restricted version of NDP-Grid, where all  
 151 source vertices lie on the top row  $R^*$  of the grid, and all destination vertices lie on a single

152 row  $R'$  of the grid, far enough from its top and bottom boundaries. They considered three  
153 different settings of this special case. In the packed-packed setting, all sources appear consec-  
154 utively on  $R^*$ , and all destinations appear consecutively on  $R'$  (but both sets may appear in  
155 an arbitrary order). They show a necessary and a sufficient condition for all demand pairs to  
156 be routable via node-disjoint paths in this setting. The second setting is the packed-spaced  
157 setting. Here, the sources again appear consecutively on  $R^*$ , but all destinations are at a  
158 distance at least  $d$  from each other. For this setting, the authors show that if  $d \geq k$ , then  
159 all demand pairs can be routed. We note that [12] extended their algorithm to a more  
160 general setting, where the destination vertices may appear anywhere in the grid, as long as  
161 the distance between any pair of the destination vertices, and any destination vertex and  
162 the boundary of the grid, is at least  $\Omega(k)$ . Robertson and Seymour [32] provided sufficient  
163 conditions for the existence of node-disjoint routing of a given set of demand pairs in the  
164 more general setting of graphs drawn on surfaces, and they designed an algorithm whose  
165 running time is  $\text{poly}(n) \cdot f(k)$  for finding the routing, where  $f(k)$  is at least exponential in  $k$ .  
166 Their result implies the existence of the routing in grids, when the destination vertices are  
167 sufficiently far from each other and from the grid boundaries, but it does not provide an effi-  
168 cient algorithm to compute such a routing. The third setting studied by Cutler and Shiloach  
169 is the spaced-spaced setting, where the distances between every pair of source vertices, and  
170 every pair of destination vertices are at least  $d$ . The authors note that they could not come  
171 up with a better algorithm for this setting, than the one provided for the packed-spaced  
172 case. Aggarwal, Kleinberg, and Williamson [1] considered a special case of NDP-Grid, where  
173 the set of the demand pairs is a permutation: that is, every vertex of the grid participates  
174 in exactly one demand pair. They show that  $\Omega(\sqrt{n}/\log n)$  demand pairs are routable in  
175 this case via node-disjoint paths. They further show that if all terminals are at a distance  
176 at least  $d$  from each other, then at least  $\Omega(\sqrt{nd}/\log n)$  pairs are routable.

177 A variation of the NPD and EDP problems, where small congestion is allowed, has been a  
178 subject of extensive study, starting with the classical paper of Raghavan and Thompson [29]  
179 that introduced the randomized rounding technique. We say that a set  $\mathcal{P}$  of paths causes  
180 congestion  $c$ , if at most  $c$  paths share the same vertex or the same edge, for the NDP and  
181 the EDP settings respectively. A recent line of work [9, 28, 2, 30, 11, 16, 7, 6] has lead to an  
182  $O(\text{poly log } k)$ -approximation for both NDP and EDP problems with congestion 2. For planar  
183 graphs, a constant-factor approximation with congestion 2 is known [34].

## 184 **Organization.**

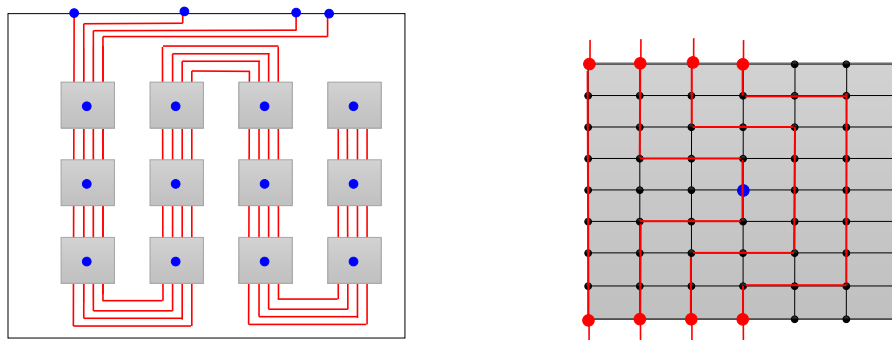
185 The majority of this extended abstract is dedicated to a detailed but informal overview of  
186 the proofs of Theorem 1 and Theorem 2. The formal proofs, as well as the extension to EDP  
187 and NDP on wall graphs, are deferred to the full version of the paper.

## 188 **2 High-Level Overview of the Algorithm**

189 The goal of this section is to provide an informal high-level overview of the main result of the  
190 paper – the proof of Theorem 1. With this goal in mind, the values of various parameters are  
191 given imprecisely in this section, in a way that best conveys the intuition. The full version  
192 of the paper contains a formal description of the algorithm and the precise settings of all  
193 parameters.

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

199 We assume that our input graph  $G$  is the  $(\ell \times \ell)$ -grid, and we denote by  $n = \ell^2$  the  
 200 number of its vertices. We further assume that the set of the demand pairs is  $\mathcal{M} =$   
 201  $\{(s_1, t_1), \dots, (s_k, t_k)\}$ , with the vertices in set  $S = \{s_1, \dots, s_k\}$  called source vertices; the  
 202 vertices in set  $T = \{t_1, \dots, t_k\}$  called destination vertices; and the vertices in  $S \cup T$  called  
 203 terminals. Let  $\text{OPT}$  denote the value of the optimal solution to the NDP instance  $(G, \mathcal{M})$ .  
 204 We assume that the vertices of  $S$  lie on the top boundary of the grid, that we denote by  $R^*$ ,  
 205 and the vertices of  $T$  lie sufficiently far from the grid boundary – say, at a distance at least  
 206  $\text{OPT}$  from it. For a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of the demand pairs, we denote by  $S(\mathcal{M}')$  and  $T(\mathcal{M}')$   
 207 the sets of the source and the destination vertices of the demand pairs in  $\mathcal{M}'$ , respectively.  
 208 As our starting point, we consider a simple observation of Chuzhoy and Kim [12], that gener-  
 209 alizes the results of Cutler and Shiloach [17]. Suppose we are given an instance of NDP-Grid  
 210 with  $k$  demand pairs, where the sources lie on the top boundary of the grid, and the desti-  
 211 nation vertices may appear anywhere in the grid, but the distance between every pair of the  
 212 destination vertices, and every destination vertex and the boundary of the grid, is at least  
 213  $(8k + 8)$  – we call such instances *spaced-out instances*. In this case, all demand pairs in  $\mathcal{M}$   
 214 can be efficiently routed via node-disjoint paths, as follows. Consider, for every destination  
 215 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  
 216 center of  $B_i$ . We construct a set  $\mathcal{P}$  of  $k$  node-disjoint paths, that originate at the vertices of  
 217  $S$ , and traverse the sub-grids  $B_i$  one-by-one in a snake-like fashion (see a schematic view on  
 218 Figure 2a). We call this part of the routing *global routing*. The *local routing* needs to specify  
 219 how the paths in  $\mathcal{P}$  traverse each box  $B_i$ . This is done in a straightforward manner, while  
 220 ensuring that the unique path originating at vertex  $s_i$  visits the vertex  $t_i$  (see Figure 2b).  
 221 By suitably truncating the final set  $\mathcal{P}$  of paths, we obtain a routing of all demand pairs in  
 222  $\mathcal{M}$  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.

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

230 For simplicity, let us first describe the algorithm with just two recursive levels. Suppose we  
 231 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  
 232 demand pairs that we would like to route. Denote  $|\mathcal{M}'| = k'$ , and assume that we are given  
 233 a collection  $\mathcal{Q}$  of square sub-grids of  $G$ , of size  $(4k' \times 4k')$  each (that we call *squares*), such  
 234 that every pair  $Q, Q' \in \mathcal{Q}$  of distinct squares is at a distance at least  $4k'$  from each other.  
 235 Assume further that each such sub-grid  $Q \in \mathcal{Q}$  is assigned a color  $\chi(Q) \in \{c_1, \dots, c_z\}$ , such  
 236 that, if  $Q$  is assigned the color  $c_j$ , then all demand pairs  $(s, t) \in \mathcal{M}'$  whose destination  $t$  lies  
 237 in  $Q$  have their source  $s \in I_j$  (so intuitively, each color  $c_j$  represents an interval  $I_j$ ). Let  
 238  $\mathcal{M}'_j \subseteq \mathcal{M}'$  be the set of all demand pairs  $(s, t) \in \mathcal{M}'$  with  $s \in I_j$ . We would like to ensure  
 239 that  $|\mathcal{M}'_j|$  is roughly  $k'/z$ , and that all destination vertices of  $T(\mathcal{M}'_j)$  are at a distance at  
 240 least  $|\mathcal{M}'_j|$  from each other. We claim that if we could find the collection  $\{I_1, \dots, I_z\}$  of the  
 241 intervals of the first row, a collection  $\mathcal{Q}$  of sub-grids of  $G$ , a coloring  $\chi : \mathcal{Q} \rightarrow \{c_1, \dots, c_z\}$ ,  
 242 and a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of the demand pairs with these properties, then we would be able to  
 243 route all demand pairs in  $\mathcal{M}'$ .

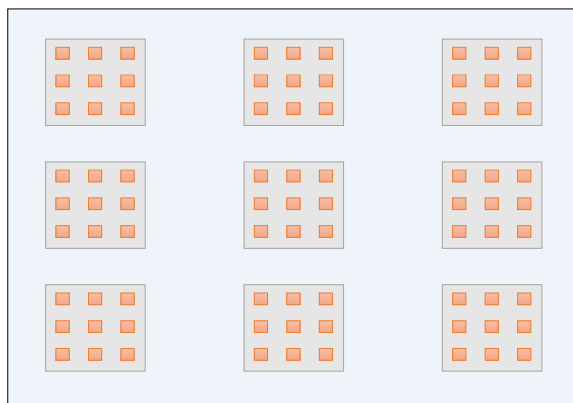
244 In order to do so, for each square  $Q \in \mathcal{Q}$ , we construct an augmented square  $Q^+$ , by adding  
 245 a margin of  $k'$  rows and columns around  $Q$ . Our goal is to construct a collection  $\mathcal{P}$  of node-  
 246 disjoint paths routing the demand pairs in  $\mathcal{M}'$ . We start by constructing a global routing,  
 247 where all paths in  $\mathcal{P}$  originate from the vertices of  $S(\mathcal{M}')$  and then visit the squares in  
 248  $\{Q^+ \mid Q \in \mathcal{Q}\}$  in a snake-like fashion, just like we did for the spaced-out instances described  
 249 above (see Figure 2a). Consider now some square  $Q \in \mathcal{Q}$  and the corresponding augmented  
 250 square  $Q^+$ . Assume that  $\chi(Q) = c_j$ , and let  $\mathcal{P}_j \subseteq \mathcal{P}$  be the set of paths originating at the  
 251 source vertices that lie in  $I_j$ . While traversing the square  $Q^+$ , we ensure that only the paths  
 252 in  $\mathcal{P}_j$  enter the square  $Q$ ; the remaining paths use the margins on the left and on the right of  
 253  $Q$  in order to traverse  $Q^+$ . This can be done because the sources of the paths in  $\mathcal{P}_j$  appear  
 254 consecutively on  $R^*$ , relatively to the sources of all paths in  $\mathcal{P}$ . In order to complete the  
 255 local routing inside the square  $Q$ , observe that the destination vertices appear far enough  
 256 from each other, and so we can employ the simple algorithm for spaced-out instances inside  
 257  $Q$ .

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

270 conditions is called a *good ensemble*. We show that a good ensemble with a large enough  
 271 set  $\mathcal{M}'$  of demand pairs exists, and then design an approximation algorithm for computing  
 272 a good ensemble maximizing  $|\mathcal{M}'|$ . We now describe each of these ingredients in turn.

## 273 2.1 A Hierarchical System of Squares.

274 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  
 275 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  
 276 *level- $h$  squares*); every such square  $Q \in \mathcal{Q}_h$  has size  $(d_h \times d_h)$ , and every pair of distinct  
 277 squares  $Q, Q' \in \mathcal{Q}_h$  are within distance at least  $d_h$  from each other (see Figure 3). We  
 278 require that for each  $1 < h \leq \rho$ , for every square  $Q \in \mathcal{Q}_h$ , there is a unique square  $Q' \in \mathcal{Q}_{h-1}$   
 279 (called the *parent-square* of  $Q$ ) that contains  $Q$ . We say that a demand pair  $(s, t)$  *belongs*  
 280 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  
 281 simple efficient algorithm to construct  $2^{O(\sqrt{\log n})}$  such hierarchical systems of squares, so  
 282 that every demand pair belongs to at least one of them. Each such system  $\tilde{\mathcal{H}}$  of squares  
 283 induces an instance of NDP — the instance is defined over the same graph  $G$ , and the set  
 284  $\tilde{\mathcal{M}} \subseteq \mathcal{M}$  of demand pairs that belong to the system  $\tilde{\mathcal{H}}$ . It is then enough to obtain a factor  
 285  $2^{O(\sqrt{\log n \cdot \log \log n})}$ -approximation algorithm for each resulting instance  $(G, \tilde{\mathcal{M}})$  separately.  
 286 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  
 287 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  
 288 on designing a  $2^{O(\sqrt{\log n \cdot \log \log n})}$ -approximation algorithm for instance  $(G, \tilde{\mathcal{M}})$ .



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

## 289 2.2 A Hierarchical Partition of the Top Grid Boundary.

290 Recall that  $R^*$  denotes the first row of the grid. A hierarchical partition  $\mathcal{J}$  of  $R^*$  is a sequence  
 291  $\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_\rho$  of sets of sub-paths of  $R^*$ , such that for each  $1 \leq h \leq \rho$ , the paths in  $\mathcal{I}_h$  (that  
 292 we refer to as *level- $h$  intervals*) partition the vertices of  $R^*$ . We also require that for all  
 293  $1 < h \leq \rho$ , every level- $h$  interval  $I \in \mathcal{I}_h$  is contained in a unique level- $(h - 1)$  interval  
 294  $I' \in \mathcal{I}_{h-1}$ , that we refer to as the *parent-interval* of  $I$ . For every level  $1 \leq h \leq \rho$ , we define a  
 295 collection  $\chi_h$  of colors, containing one color  $c_h(I)$  for each level- $h$  interval  $I \in \mathcal{I}_h$ . If  $I' \in \mathcal{I}_h$   
 296 is a parent-interval of  $I \in \mathcal{I}_{h+1}$ , then we say that color  $c_h(I')$  is a *parent-color* of  $c_{h+1}(I)$ .



## 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  $\mathfrak{I}$ .

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  $(\mathfrak{I}, f, \mathcal{M}')$  is a *good ensemble* iff  $\mathfrak{I}$  is a hierarchical partition of  $R^*$  into intervals;  $f$  is a valid coloring of the squares in  $\tilde{\mathcal{H}}$  with respect to  $\mathfrak{I}$ ; 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  $(\mathfrak{I}, f, \mathcal{M}')$ , 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 2a. 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

341 squares are at distance at least  $d_2$  from each other, there is a sufficient spacing to allow this  
 342 routing. We continue this process recursively, until, at the last level of the recursion, we  
 343 route at most one path per color, to its destination vertex.

344 In order to complete the proof of the theorem, we need to show that there exists a good  
 345 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  
 346 ensemble efficiently.

## 347 2.5 The Existence of the Ensemble.

348 The key notion that we use in order to show that a large good ensemble  $(\mathcal{J}, f, \mathcal{M}')$  exists  
 349 is that of a *shadow property*. Suppose  $Q$  is some  $(d \times d)$  sub-grid of  $G$ , and let  $\hat{\mathcal{M}} \subseteq \mathcal{M}$   
 350 be some subset of the demand pairs. Among all demand pairs  $(s, t) \in \hat{\mathcal{M}}$  with  $t \in Q$ , let  
 351  $(s_1, t_1)$  be the one with  $s_1$  appearing earliest on the first row  $R^*$  of  $G$ , and let  $(s_2, t_2)$  be the  
 352 one with  $s_2$  appearing latest on  $R^*$ . The *shadow of  $Q$  with respect to  $\hat{\mathcal{M}}$*  is the sub-path of  
 353  $R^*$  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  $s$   
 354 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  
 355 *shadow property with respect to  $Q$*  iff  $N_{\hat{\mathcal{M}}}(Q) \leq d$ . We say that  $\hat{\mathcal{M}}$  has the *shadow property*  
 356 *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  
 357 property with respect to every square in  $\bigcup_{h=1}^\rho \mathcal{Q}_h$ . Let  $\mathcal{P}^*$  be the optimal solution to the  
 358 instance  $(G, \tilde{\mathcal{M}})$  of NDP, where  $\tilde{\mathcal{M}}$  only includes the demand pairs that belong to  $\tilde{\mathcal{H}}$ . Let  
 359  $\mathcal{M}^* \subseteq \tilde{\mathcal{M}}$  be the set of the demand pairs routed by  $\mathcal{P}^*$ . For every demand pair  $(s, t) \in \mathcal{M}^*$ ,  
 360 let  $P(s, t) \in \mathcal{P}^*$  be the path routing this demand pair. Intuitively, it feels like  $\mathcal{M}^*$  should  
 361 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  
 362 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  
 363 to  $\mathcal{M}^*$  is the sub-path of  $R^*$  between  $s_1$  and  $s_2$ . Let  $P$  be any path of length at most  $2d_h$   
 364 connecting  $t_1$  to  $t_2$  in  $Q$ , and let  $\gamma$  be the closed curve consisting of the union of  $P(s_1, t_1)$ ,  $P$ ,  
 365  $P(s_2, t_2)$ , and the shadow of  $Q$ . Consider the disc  $D$  whose boundary is  $\gamma$ . The intuition is  
 366 that, if  $(s, t) \in \mathcal{M}^*$  is a demand pair whose source lies in the shadow of  $Q$ , and destination  
 367 lies outside of  $D$ , then  $P(s, t)$  must cross the path  $P$ , as it needs to escape the disc  $D$ . Since  
 368 path  $P$  is relatively short, only a small number of such demand pairs may exist. The main  
 369 difficulty with this argument is that we may have a large number of demand pairs  $(s, t)$ ,  
 370 whose source lies in the shadow of  $Q$ , and the destination lies in the disc  $D$ . Intuitively, this  
 371 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,  
 372 in a sense, this cannot happen too often, and that there is a subset  $\mathcal{M}^{**} \subseteq \mathcal{M}^*$  of at least  
 373  $|\mathcal{M}^*|/2^{O(\sqrt{\log n \cdot \log \log n})}$  demand pairs, such that  $\mathcal{M}^{**}$  has the shadow property with respect  
 374 to  $\tilde{\mathcal{H}}$ .

375 Finally, we show that there exists a good ensemble  $(\mathcal{J}, f, \mathcal{M}')$  with  $|\mathcal{M}'| \geq |\mathcal{M}^{**}|/2^{O(\sqrt{\log n \cdot \log \log n})}$ .

376 We construct the ensemble over the course of  $\rho$  iterations, starting with  $\mathcal{M}' = \mathcal{M}^{**}$ . In the  
 377  $h$ th iteration we construct the set  $\mathcal{I}_h$  of the level- $h$  intervals of  $R^*$ , assign level- $h$  colors to all  
 378 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  
 379 the first iteration, we let  $\mathcal{I}_1$  be a partition of the row  $R^*$  into intervals, each of which contains  
 380 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  
 381 that they appear in this left-to-right order on  $R^*$ . We call all intervals  $I_j$  where  $j$  is odd  
 382 *interesting intervals*, and the remaining intervals  $I_j$  *uninteresting intervals*. We discard from  
 383  $\mathcal{M}'$  all demand pairs  $(s, t)$ , where  $s$  lies on an uninteresting interval. Consider now some  
 384 level-1 square  $Q$ , and let  $\mathcal{M}(Q) \subseteq \mathcal{M}'$  be the set of all demand pairs whose destination lies  
 385 in  $Q$ . Since the original set  $\mathcal{M}^{**}$  of demand pairs had the shadow property with respect

386 to  $Q$ , it is easy to verify that all source vertices of the demand pairs in  $\mathcal{M}(Q)$  must belong  
 387 to a single interesting interval of  $\mathcal{I}_1$ . Let  $I$  be that interval. Then we color the square  $Q$   
 388 with the level-1 color  $c_1(I)$  corresponding to the interval  $I$ . This completes the first iteration.  
 389 Notice that for each level-1 color  $c_1(I)$ , at most  $d_2/16$  demand pairs  $(s, t) \in \mathcal{M}'$  have  
 390  $s \in I$ . In the following iteration, we similarly partition every interesting level-1 interval  
 391 into level-2 intervals that contain roughly  $d_3/16 \leq |\mathcal{M}'|/\eta^2$  source vertices of  $\mathcal{M}'$  each, and  
 392 then define a coloring of all level-2 squares similarly, while suitably updating the set  $\mathcal{M}'$  of  
 393 the demand pairs. We continue this process for  $\rho$  iterations, eventually obtaining a good  
 394 ensemble  $(\mathcal{J}, f, \mathcal{M}')$ . Since we only discard a constant fraction of the demand pairs of  $\mathcal{M}'$  in  
 395 every iteration, at the end,  $|\mathcal{M}'| \geq |\mathcal{M}^{**}|/2^\rho = |\mathcal{M}^{**}|/2^{O(\sqrt{\log n})} \geq |\mathcal{M}^*|/2^{O(\sqrt{\log n \cdot \log \log n})}$ .

## 396 2.6 Finding the Good Ensemble.

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

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

## 407 2.7 Completing the Proof of Theorem 1.

408 So far we have assumed that all source vertices lie on the top boundary of the grid, and  
 409 all destination vertices are at a distance at least  $\Omega(\text{OPT})$  from the grid boundary. Let  $\mathcal{A}$   
 410 be the randomized efficient  $2^{O(\sqrt{\log n \cdot \log \log n})}$ -approximation algorithm for this special case.  
 411 We now extend it to the general Restricted NDP-Grid problem. For every destination vertex  
 412  $t$ , we identify the closest vertex  $\tilde{t}$  that lies on the grid boundary. Using standard grouping  
 413 techniques, at the cost of losing an additional  $O(\log n)$  factor in the approximation ratio,  
 414 we can assume that all source vertices lie on the top boundary of the grid, all vertices in  
 415  $\{\tilde{t} \mid t \in T(\mathcal{M})\}$  lie on a single boundary edge of the grid (assume for simplicity that it is the  
 416 bottom boundary), and that there is some integer  $d$ , such that for every destination vertex  
 417  $t \in T(\mathcal{M})$ ,  $d \leq d(t, \tilde{t}) < 2d$ . We show that we can define a collection  $\mathcal{Z} = \{Z_1, \dots, Z_r\}$  of  
 418 disjoint square sub-grids of  $G$ , and a collection  $\mathcal{I} = \{I_1, \dots, I_r\}$  of disjoint sub-intervals of  
 419  $R^*$ , such that the bottom boundary of each sub-grid  $Z_i$  is contained in the bottom boundary  
 420 of  $G$ , the top boundary of  $Z_i$  is within distance at least  $\text{OPT}$  from  $R^*$ ,  $Z_1, \dots, Z_r$  appear in  
 421 this left-to-right order in  $G$ , and  $I_1, \dots, I_r$  appear in this left-to-right order on  $R^*$ . For each  
 422  $1 \leq j \leq r$ , we let  $\mathcal{M}_j$  denote the set of all demand pairs with the sources lying on  $I_j$  and  
 423 the destinations lying in  $Z_j$ . For each  $1 \leq j \leq r$ , we then obtain a new instance  $(G, \mathcal{M}_j)$   
 424 of the NDP problem. We show that there exist a collection  $\mathcal{Z}$  of squares and a collection  
 425  $\mathcal{I}$  of intervals, such that the value of the optimal solution to each instance  $(G, \mathcal{M}_j)$ , that  
 426 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  
 427 is not hard to show that, if we can compute, for each  $1 \leq j \leq r$ , a routing of some subset

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

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

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

### 445 **3** **Approximation Algorithm for the Special Case with Sources Close** 446 **to the Grid Boundary**

447 In this section we provide a sketch of the proof of Theorem 2. We assume that we are given  
 448 an instance  $(G, \mathcal{M})$  of NDP-Grid and an integer  $\delta > 0$ , such that every source vertex is at  
 449 a distance at most  $\delta$  from the grid boundary. Our goal is to design an efficient random-  
 450 ized factor- $(\delta \cdot 2^{O(\sqrt{\log n \cdot \log \log n})})$ -approximation algorithm for this special case. For every  
 451 terminal  $v \in S(\mathcal{M}) \cup T(\mathcal{M})$ , let  $\tilde{v}$  be the vertex lying closest to  $v$  on the boundary of the  
 452 grid  $G$ . Using standard grouping techniques, at the cost of losing an  $O(\log n)$ -factor in the  
 453 approximation ratio, we can assume that there is some integer  $d$ , such that for all  $t \in T(\mathcal{M})$ ,  
 454  $d \leq d(t, \tilde{t}) < 2d$ .

455 Assume first that  $d \leq \delta \cdot 2^{O(\sqrt{\log n \cdot \log \log n})}$ . Let  $\hat{\mathcal{M}} = \{(\tilde{s}, \tilde{t}) \mid (s, t) \in \mathcal{M}\}$  be a new set of  
 456 demand pairs, so that all vertices participating in these demand pairs lie on the boundary  
 457 of  $G$ . We can efficiently find an optimal solution to the NDP problem instance  $(G, \hat{\mathcal{M}})$  using  
 458 standard dynamic programming. We then show that  $\text{OPT}(G, \hat{\mathcal{M}}) = \Omega(\text{OPT}(G, \mathcal{M}) / (\delta \cdot$   
 459  $2^{O(\sqrt{\log n \cdot \log \log n})}))$ , obtaining an  $(\delta \cdot 2^{O(\sqrt{\log n \cdot \log \log n})})$ -approximation algorithm.

460 From now on we assume that  $d > \delta \cdot 2^{\Omega(\sqrt{\log n \cdot \log \log n})}$ . Next, we define a new set  $\tilde{\mathcal{M}}$  of  
 461 demand pairs:  $\tilde{\mathcal{M}} = \{(\tilde{s}, t) \mid (s, t) \in \mathcal{M}\}$ , so all source vertices of the demand pairs in  $\tilde{\mathcal{M}}$  lie  
 462 on the boundary of  $G$ , obtaining an instance of Restricted NDP-Grid. Let  $\text{OPT}'$  be the value  
 463 of the optimal solution to problem  $(G, \tilde{\mathcal{M}})$ . We show that  $\text{OPT}' \geq \Omega(\text{OPT}(G, \mathcal{M}) / \delta)$ .

464 We then focus on instance  $(G, \tilde{\mathcal{M}})$  of Restricted NDP-Grid. We say that a path  $P$  routing a  
 465 demand pair  $(\tilde{s}, t) \in \tilde{\mathcal{M}}$  is *canonical* iff it contains the original source  $s$ . The crux of the proof  
 466 is to show that we can modify the routing produced by the  $2^{O(\sqrt{\log n \cdot \log \log n})}$ -approximation  
 467 algorithm to instance  $(G, \tilde{\mathcal{M}})$ , so that in the resulting routing all paths are canonical. In  
 468 order to do so, we utilize the fact that the destination vertices lie much further from the  
 469 grid boundaries than the source vertices. This creates sufficient margins around the grid  
 470 boundaries that allow us to modify the routing to turn it a canonical one.

471 ——— **References** ———

- 472 **1** Alok Aggarwal, Jon Kleinberg, and David P. Williamson. Node-disjoint paths on the  
473 mesh and a new trade-off in VLSI layout. *SIAM J. Comput.*, 29(4):1321–1333, Feb-  
474 ruary 2000. URL: <http://dx.doi.org/10.1137/S0097539796312733>, doi:10.1137/  
475 S0097539796312733.
- 476 **2** Matthew Andrews. Approximation algorithms for the edge-disjoint paths problem via  
477 Raecke decompositions. In *Proceedings of IEEE FOCS*, pages 277–286, 2010. URL: [http://](http://dx.doi.org/10.1109/FOCS.2010.33)  
478 [dx.doi.org/10.1109/FOCS.2010.33](http://dx.doi.org/10.1109/FOCS.2010.33), doi:<http://dx.doi.org/10.1109/FOCS.2010.33>.
- 479 **3** Matthew Andrews, Julia Chuzhoy, Venkatesan Guruswami, Sanjeev Khanna, Kunal Tal-  
480 war, and Lisa Zhang. Inapproximability of edge-disjoint paths and low congestion routing  
481 on undirected graphs. *Combinatorica*, 30(5):485–520, 2010.
- 482 **4** Matthew Andrews and Lisa Zhang. Logarithmic hardness of the undirected edge-disjoint  
483 paths problem. *J. ACM*, 53(5):745–761, September 2006. URL: [http://doi.acm.org/10.](http://doi.acm.org/10.1145/1183907.1183910)  
484 [1145/1183907.1183910](http://doi.acm.org/10.1145/1183907.1183910), doi:10.1145/1183907.1183910.
- 485 **5** Yonatan Aumann and Yuval Rabani. Improved bounds for all optical routing. In *Proceed-*  
486 *ings of the sixth annual ACM-SIAM symposium on Discrete algorithms*, SODA '95, pages  
487 567–576, Philadelphia, PA, USA, 1995. Society for Industrial and Applied Mathematics.  
488 URL: <http://dl.acm.org/citation.cfm?id=313651.313820>.
- 489 **6** Chandra Chekuri and Julia Chuzhoy. Half-integral all-or-nothing flow. Unpublished Ma-  
490 nuscript.
- 491 **7** Chandra Chekuri and Alina Ene. Poly-logarithmic approximation for maximum node dis-  
492 joint paths with constant congestion. In *Proc. of ACM-SIAM SODA*, 2013.
- 493 **8** Chandra Chekuri, Sanjeev Khanna, and F Bruce Shepherd. Edge-disjoint paths in planar  
494 graphs. In *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Sym-*  
495 *posium on*, pages 71–80. IEEE, 2004.
- 496 **9** Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd. Multicommodity flow, well-  
497 linked terminals, and routing problems. In *Proc. of ACM STOC*, pages 183–192, 2005.
- 498 **10** Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd. An  $O(\sqrt{n})$  approximation and  
499 integrality gap for disjoint paths and unsplittable flow. *Theory of Computing*, 2(1):137–146,  
500 2006.
- 501 **11** Julia Chuzhoy. Routing in undirected graphs with constant congestion. *SIAM J. Com-*  
502 *put.*, 45(4):1490–1532, 2016. URL: <https://doi.org/10.1137/130910464>, doi:10.1137/  
503 130910464.
- 504 **12** Julia Chuzhoy and David H. K. Kim. On approximating node-disjoint paths in  
505 grids. In Naveen Garg, Klaus Jansen, Anup Rao, and José D. P. Rolim, editors,  
506 *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Tech-*  
507 *niques, APPROX/RANDOM 2015, August 24–26, 2015, Princeton, NJ, USA*, volume 40  
508 of *LIPICs*, pages 187–211. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.  
509 URL: <http://dx.doi.org/10.4230/LIPICs.APPROX-RANDOM.2015.187>, doi:10.4230/  
510 LIPICs.APPROX-RANDOM.2015.187.
- 511 **13** Julia Chuzhoy, David H. K. Kim, and Shi Li. Improved approximation for node-disjoint  
512 paths in planar graphs. In *Proceedings of the 48th Annual ACM SIGACT Symposium on*  
513 *Theory of Computing*, STOC 2016, pages 556–569, New York, NY, USA, 2016. ACM. URL:  
514 <http://doi.acm.org/10.1145/2897518.2897538>, doi:10.1145/2897518.2897538.
- 515 **14** Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat. Almost polynomial hardness of  
516 node-disjoint paths in grids. Unpublished Manuscript, 2017.
- 517 **15** Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat. New hardness results for routing on  
518 disjoint paths. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings*  
519 *of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017*,

- 520 *Montreal, QC, Canada, June 19-23, 2017*, pages 86–99. ACM, 2017. URL: <http://doi.acm.org/10.1145/3055399.3055411>, doi:10.1145/3055399.3055411.
- 521
- 522 **16** Julia Chuzhoy and Shi Li. A polylogarithmic approximation algorithm for edge-disjoint  
523 paths with congestion 2. *J. ACM*, 63(5):45:1–45:51, 2016. URL: <http://dl.acm.org/citation.cfm?id=2893472>.
- 524
- 525 **17** M. Cutler and Y. Shiloach. Permutation layout. *Networks*, 8:253–278, 1978. doi:10.1002/  
526 net.3230080308.
- 527 **18** Shimon Even, Alon Itai, and Adi Shamir. On the complexity of timetable and multicom-  
528 modity flow problems. *SIAM J. Comput.*, 5(4):691–703, 1976. URL: <http://dx.doi.org/10.1137/0205048>, doi:10.1137/0205048.
- 529
- 530 **19** Krzysztof Fleszar, Matthias Mnich, and Joachim Spoerhase. New Algorithms for Maximum  
531 Disjoint Paths Based on Tree-Likeness. In Piotr Sankowski and Christos Zaroliagis, edit-  
532 ors, *24th Annual European Symposium on Algorithms (ESA 2016)*, volume 57 of *Leibniz*  
533 *International Proceedings in Informatics (LIPIcs)*, pages 42:1–42:17, Dagstuhl, Germany,  
534 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: <http://drops.dagstuhl.de/opus/volltexte/2016/6354>, doi:10.4230/LIPIcs.ESA.2016.42.
- 535
- 536 **20** R. Karp. On the complexity of combinatorial problems. *Networks*, 5:45–68, 1975.
- 537 **21** Ken-Ichi Kawarabayashi and Yusuke Kobayashi. An  $O(\log n)$ -approximation algorithm  
538 for the edge-disjoint paths problem in Eulerian planar graphs. *ACM Trans. Algorithms*,  
539 9(2):16:1–16:13, March 2013. URL: <http://doi.acm.org/10.1145/2438645.2438648>,  
540 doi:10.1145/2438645.2438648.
- 541 **22** Jon Kleinberg. An approximation algorithm for the disjoint paths problem in even-degree  
542 planar graphs. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Com-*  
543 *puter Science, FOCS '05*, pages 627–636, Washington, DC, USA, 2005. IEEE Computer So-  
544 ciety. URL: <http://dx.doi.org/10.1109/SFCS.2005.18>, doi:10.1109/SFCS.2005.18.
- 545 **23** Jon M. Kleinberg and Éva Tardos. Disjoint paths in densely embedded graphs. In *Pro-*  
546 *ceedings of the 36th Annual Symposium on Foundations of Computer Science*, pages 52–61,  
547 1995.
- 548 **24** Jon M. Kleinberg and Éva Tardos. Approximations for the disjoint paths problem in high-  
549 diameter planar networks. *J. Comput. Syst. Sci.*, 57(1):61–73, 1998.
- 550 **25** Stavros G. Kolliopoulos and Clifford Stein. Approximating disjoint-path problems using  
551 packing integer programs. *Mathematical Programming*, 99:63–87, 2004. doi:10.1007/  
552 s10107-002-0370-6.
- 553 **26** MR Kramer and Jan van Leeuwen. The complexity of wire-routing and finding minimum  
554 area layouts for arbitrary vlsi circuits. *Advances in computing research*, 2:129–146, 1984.
- 555 **27** James F. Lynch. The equivalence of theorem proving and the interconnection prob-  
556 lem. *SIGDA Newsl.*, 5(3):31–36, September 1975. URL: <http://doi.acm.org/10.1145/1061425.1061430>,  
557 doi:10.1145/1061425.1061430.
- 558 **28** Harald Räcke. Minimizing congestion in general networks. In *Proc. of IEEE FOCS*, pages  
559 43–52, 2002.
- 560 **29** Prabhakar Raghavan and Clark D. Tompson. Randomized rounding: a technique for  
561 provably good algorithms and algorithmic proofs. *Combinatorica*, 7:365–374, Decem-  
562 ber 1987. URL: <http://portal.acm.org/citation.cfm?id=45291.45296>, doi:10.1007/  
563 BF02579324.
- 564 **30** Satish Rao and Shuheng Zhou. Edge disjoint paths in moderately connected graphs. *SIAM*  
565 *J. Comput.*, 39(5):1856–1887, 2010.
- 566 **31** N. Robertson and P. D. Seymour. Outline of a disjoint paths algorithm. In *Paths, Flows*  
567 *and VLSI-Layout*. Springer-Verlag, 1990.

- 568 **32** Neil Robertson and Paul D. Seymour. Graph minors. VII. disjoint paths on a surface.  
569 *J. Comb. Theory, Ser. B*, 45(2):212–254, 1988. URL: [http://dx.doi.org/10.1016/](http://dx.doi.org/10.1016/0095-8956(88)90070-6)  
570 [0095-8956\(88\)90070-6](http://dx.doi.org/10.1016/0095-8956(88)90070-6), doi:10.1016/0095-8956(88)90070-6.
- 571 **33** Neil Robertson and Paul D Seymour. Graph minors. XIII. the disjoint paths problem.  
572 *Journal of Combinatorial Theory, Series B*, 63(1):65–110, 1995.
- 573 **34** Loïc Seguin-Charbonneau and F. Bruce Shepherd. Maximum edge-disjoint paths in planar  
574 graphs with congestion 2. In *Proceedings of the 2011 IEEE 52Nd Annual Symposium*  
575 *on Foundations of Computer Science*, FOCS '11, pages 200–209, Washington, DC, USA,  
576 2011. IEEE Computer Society. URL: <http://dx.doi.org/10.1109/FOCS.2011.30>, doi:  
577 [10.1109/FOCS.2011.30](http://dx.doi.org/10.1109/FOCS.2011.30).