

# A Polylogarithmic Approximation Algorithm for Edge-Disjoint Paths with Congestion 2

Julia Chuzhoy\*      Shi Li†

December 20, 2016

## Abstract

In the Edge-Disjoint Paths with Congestion problem (EDPwC), we are given an undirected  $n$ -vertex graph  $G$ , a collection  $\mathcal{M} = \{(s_1, t_1), \dots, (s_k, t_k)\}$  of pairs of vertices called demand pairs and an integer  $c$ . The goal is to connect the maximum possible number of the demand pairs by paths, so that the maximum edge congestion - the number of paths sharing any edge - is bounded by  $c$ . When the maximum allowed congestion is  $c = 1$ , this is the classical Edge-Disjoint Paths problem (EDP).

The best current approximation algorithm for EDP achieves an  $O(\sqrt{n})$ -approximation, by rounding the standard multi-commodity flow relaxation of the problem. This matches the  $\Omega(\sqrt{n})$  lower bound on the integrality gap of this relaxation. We show an  $O(\text{poly log } k)$ -approximation algorithm for EDPwC with congestion  $c = 2$ , by rounding the same multi-commodity flow relaxation. This gives the best possible congestion for a sub-polynomial approximation of EDPwC via this relaxation. Our results are also close to optimal in terms of the number of pairs routed, since EDPwC is known to be hard to approximate to within a factor of  $\tilde{\Omega}((\log n)^{1/(c+1)})$  for any constant congestion  $c$ . Prior to our work, the best approximation factor for EDPwC with congestion 2 was  $\tilde{O}(n^{3/7})$ , and the best algorithm achieving a polylogarithmic approximation required congestion 14.

---

\*Toyota Technological Institute at Chicago, Chicago, IL 60637. Email: [cjulia@ttic.edu](mailto:cjulia@ttic.edu). Supported in part by NSF CAREER grant CCF-0844872, Sloan Research Fellowship and NSF grant CCF-1318242.

†Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY 14260. Email: [shil@buffalo.edu](mailto:shil@buffalo.edu). Part of the work was done while the author was affiliated with the Toyota Technological Institute at Chicago.

## 1 Introduction

One of the central and most extensively studied graph routing problems is the Edge-Disjoint Paths problem (EDP). In this problem, we are given an undirected  $n$ -vertex graph  $G = (V, E)$ , and a collection  $\mathcal{M} = \{(s_1, t_1), \dots, (s_k, t_k)\}$  of  $k$  pairs of vertices, called source-sink pairs, or demand pairs. The goal is to find a collection  $\mathcal{P}$  of edge-disjoint paths, connecting the maximum possible number of the demand pairs.

Robertson and Seymour [RS90] have shown that EDP can be solved efficiently, when the number  $k$  of the demand pairs is bounded by a constant. However, for general values of  $k$ , it is NP-hard to even decide whether all pairs can be simultaneously routed via edge-disjoint paths [Kar72, EIS76]. A standard approach to designing approximation algorithms for EDP and other routing problems, is to first compute a multi-commodity flow relaxation, where instead of connecting the demand pairs with paths, we are only required to send the maximum amount of multi-commodity flow between the demand pairs, with at most one flow unit sent between every pair. Such a fractional solution can be computed efficiently by using the standard multi-commodity flow LP-relaxation, and it can then be rounded to obtain an integral solution. Indeed, the best current approximation algorithm for the EDP problem, due to Chekuri, Khanna and Shepherd [CKS06], achieves an  $O(\sqrt{n})$ -approximation using this approach. Unfortunately, a simple example by Garg, Vazirani and Yannakakis [GVY97] (see also Section D in the Appendix), shows that the integrality gap of the multi-commodity flow relaxation can be as large as  $\Omega(\sqrt{n})$ , thus implying that the algorithm of [CKS06] is essentially the best possible for EDP, when using this approach. This integrality gap appears to be a major barrier to obtaining better approximation algorithms for EDP. With the current best hardness of approximation factor standing on  $\Omega(\log^{1/2-\epsilon} n)$  for any constant  $\epsilon$  (unless NP is contained in  $\text{ZPTIME}(n^{\text{poly log } n})$  [AZ06, ACG<sup>+</sup>10]), the approximability of EDP remains one of the central open problems in the area of routing.

A natural question is whether we can obtain better approximation algorithms by slightly relaxing the disjointness requirement, and allowing the paths to share edges. We say that a set  $\mathcal{P}$  of paths is an  $\alpha$ -approximate solution with congestion  $c$ , iff the paths in  $\mathcal{P}$  connect at least  $\text{OPT}/\alpha$  of the demand pairs, while every edge of  $G$  appears on at most  $c$  paths in  $\mathcal{P}$ . Here,  $\text{OPT}$  is the value of the optimal solution to EDP, where no congestion is allowed. This relaxation of the EDP problem is called EDP with congestion (EDPwC). The EDPwC problem is a natural framework to study the tradeoff between the number of pairs routed and the congestion, and it is useful in scenarios where we can afford a small congestion on edges.

The classical randomized rounding technique of Raghavan and Thompson [RT87] gives a constant-factor approximation for EDPwC, when the congestion  $c$  is  $\Omega(\log n / \log \log n)$ . More generally, for any congestion value  $c$ , factor  $O(n^{1/c})$ -approximation algorithms are known for EDPwC [AR06, BS00, KS04]. Recently, Andrews [And10] has shown a randomized  $O(\text{poly log } n)$ -approximation algorithm with congestion  $c = O(\text{poly log log } n)$ , and Chuzhoy [Chu12] has shown a randomized  $O(\text{poly log } k)$ -approximation algorithm with congestion 14. For the congestion value  $c = 2$ , Kawarabayashi and Kobayashi [KK11] have recently shown an  $\tilde{O}(n^{3/7})$ -approximation algorithm, thus improving the best previously known  $O(\sqrt{n})$ -approximation for  $c = 2$  [AR06, BS00, KS04]. We note that all the above mentioned algorithms rely on the standard multi-commodity flow LP relaxation of the problem. It is easy to see that the values of the optimal solution of this LP relaxation for the EDP problem, when no congestion is allowed, and for the EDPwC problem, when congestion  $c$  is allowed, are within a factor  $c$  from each other. Therefore, the statements of these results remain valid even when the approximation factor is computed with respect to the optimal solution to the EDPwC problem.

In this paper, we show a randomized  $O(\text{poly log } k)$ -approximation algorithm for EDPwC with congestion 2. Given an instance  $(G, \mathcal{M})$  of the EDP problem, our algorithm with high probability routes at least  $\Omega(\text{OPT} / \text{poly log } k)$  pairs with congestion 2, where  $\text{OPT}$  is the maximum number of pairs that

can be routed with no congestion. Our algorithm also achieves an  $O(\text{poly log } k)$ -approximation when compared with the optimal solution to EDPwC with congestion 2. Like all the algorithms for EDP and EDPwC mentioned above, our algorithm also performs a rounding of the standard multi-commodity flow relaxation for EDP. Therefore, our result shows that when congestion 2 is allowed, the integrality gap of this relaxation improves from  $\Omega(\sqrt{n})$  to polylogarithmic. Our result is essentially optimal with respect to this relaxation, both for the congestion and the number of pairs routed, in the following sense. As observed above, if we are interested in obtaining a sub-polynomial approximation for EDP via the multi-commodity flow relaxation, then the best congestion we can hope for is 2. On the other hand, Andrews et al. [ACG<sup>+</sup>10] have shown that the integrality gap of the multi-commodity flow relaxation for EDPwC is  $\Omega\left(\left(\frac{\log n}{(\log \log n)^2}\right)^{1/(c+1)}\right)$  for any constant congestion  $c$ . In particular, the integrality gap for congestion 2 is polylogarithmic, although the degree of the logarithm is much lower than the degree we obtain in our approximation algorithm. Andrews et al. [ACG<sup>+</sup>10] have also shown that for any constant  $\epsilon$ , for any  $1 \leq c \leq O\left(\frac{\log \log n}{\log \log \log n}\right)$ , there is no  $O\left((\log n)^{\frac{1-\epsilon}{c+1}}\right)$ -approximation algorithm for EDPwC with congestion  $c$ , unless  $\text{NP} \subseteq \text{ZPTIME}(n^{\text{poly log } n})$ . In particular, this gives an  $\Omega\left(\log^{(1-\epsilon)/3} n\right)$ -hardness of approximation for EDPwC with congestion 2.

While the approximability status of the EDP problem remains open, our results show a fundamental difference between routing with congestion 1 and routing with congestion 2 or higher. Suppose we are given a solution  $\mathcal{P}$  to the EDP problem that connects  $D$  of the demand pairs with congestion  $c$ , and we are interested in obtaining another solution with a lower congestion. Our results provide an efficient randomized algorithm to find a solution connecting  $\Omega(D/(c \text{ poly log } k))$  of the demand pairs with congestion 2. That is, we can lower the congestion to 2 with only a factor  $(c \text{ poly log } k)$  loss in the number of the demand pairs routed<sup>1</sup>. However, if we are interested in routing with no congestion, then we may have to lose an  $\Omega(\sqrt{n})$ -factor in the number of pairs routed. For example, in the integrality gap construction of Garg, Vazirani and Yannakakis [GVY97], there is a solution that routes  $k = \Theta(\sqrt{n})$  pairs with congestion 2, but if we require a routing with congestion 1, then at most one pair can be routed (see section D of the Appendix).

We note that better approximation algorithms are known for some special cases of the EDP problem. Rao and Zhou [RZ10] have shown that if the value of the global minimum cut in the input graph  $G$  is  $\Omega(\log^5 n)$ , then there is an efficient randomized  $O(\text{poly log } n)$ -approximation algorithm for EDP. Interestingly, this algorithm is also based on the multi-commodity flow relaxation. The EDP problem is also known to have polylogarithmic approximation algorithms on bounded-degree expander graphs [LR99, BFU94, BFSU94, KR96, Fri00], and constant-factor approximation algorithms on trees [GVY97, CMS07], grids and grid-like graphs [AR95, AGLR94, KT98, KT95]. Routing problems have also been extensively studied on planar graphs. Chekuri, Khanna and Shepherd [CKS04a, CKS09] have shown a poly-logarithmic approximation algorithm for EDPwC with congestion 2 and a constant approximation algorithm with congestion 4 on planar graphs. Both results have recently been improved by Seguin-Charbonneau and Shepherd [SCS11], who showed a constant factor approximation algorithm with congestion 2. When no congestion is allowed, Kleinberg [Kle05] has shown an  $O(\log^2 n)$ -approximation for Eulerian planar graphs. Kawarabayashi and Kobayashi [KK13] have recently improved this result to an  $O(\log n)$ -approximation, for both Eulerian and 4-connected planar graphs. However, improving the  $O(\sqrt{n})$ -approximation algorithm for EDP on general planar graphs still remains elusive. Chekuri and Ene [CE13] have recently generalized the result of [Chu12] to obtain an approximation algorithm for the closely related Node-Disjoint Paths problem, where the paths in the solution are required to be node-disjoint, instead of edge-disjoint. They show a  $\text{poly log}(k)$ -

---

<sup>1</sup>We can view the set  $\mathcal{P}$  of paths as a fractional solution to the EDP problem instance, where  $1/c$  flow units are sent along each path. This gives a fractional solution of value  $D/c$ . We can then use our algorithm to route  $\Omega(D/(c \text{ poly log } k))$  of the demand pairs with congestion 2, by rounding this fractional solution.

approximation with constant congestion for this problem.

A problem closely related to EDP is **Congestion Minimization**, where the goal is to route all the demand pairs, while minimizing the maximum congestion. The classical randomized rounding technique of Raghavan and Thompson [RT87], when applied to the multi-commodity flow relaxation of the problem, achieves an  $O(\log n / \log \log n)$ -approximation. This is the best currently known approximation algorithm for the problem. The best current hardness of approximation ratio, due to Andrews and Zhang [AZ07], is  $\Omega\left(\frac{\log \log n}{\log \log \log n}\right)$ , under the assumption that  $\text{NP} \not\subseteq \text{ZPTIME}(n^{\text{poly} \log n})$ .

**Our results.** Our main result is summarized in the following theorem.

**Theorem 1.1** *There is an efficient randomized algorithm, that, given a graph  $G$ , and a collection  $\mathcal{M}$  of  $k$  source-sink pairs, with high probability finds a routing of  $\Omega(\text{OPT}/(\text{poly} \log k))$  of the pairs in  $\mathcal{M}$  with congestion at most 2, where  $\text{OPT}$  is the maximum number of pairs that can be routed with congestion 2.*

We remark that our algorithm in fact routes  $\Omega(\text{OPT}_{\text{LP}}/(\text{poly} \log k))$  demand pairs, where  $\text{OPT}_{\text{LP}}$  is the value of the optimal solution to the standard multicommodity LP-relaxation of the EDPwC problem with congestion 2.

We now provide an overview of our techniques and compare them to [Chu12]. We denote by  $\mathcal{T}$  the set of vertices participating in the demand pairs, and we call them *terminals*. As in previous work [CKS05, RZ10, And10, Chu12], we use the notion of well-linkedness. Given a graph  $G = (V, E)$  and a subset  $\mathcal{T}$  of its vertices called terminals, we say that  $\mathcal{T}$  is  $\alpha$ -well-linked in  $G$ , iff for any partition  $(A, B)$  of  $V$ ,  $|E(A, B)| \geq \alpha \cdot \min\{|A \cap \mathcal{T}|, |B \cap \mathcal{T}|\}$ . Chekuri, Khanna and Shepherd [CKS04b, CKS05] have shown an efficient algorithm, that, given any EDP instance  $(G, \mathcal{M})$ , partitions it into a number of sub-instances  $(G_1, \mathcal{M}_1), \dots, (G_\ell, \mathcal{M}_\ell)$ , such that, on the one hand, for each  $i$ , the set of terminals participating in  $\mathcal{M}_i$  is 1-well-linked in  $G_i$ , and on the other hand, the sum of the values of the optimal fractional solutions in all these instances is  $\Omega(\text{OPT}/\log^2 k)$ . Therefore, it is enough to find a polylogarithmic approximation with congestion 2 in each such sub-instance separately. From now on we assume that we are given an instance  $(G, \mathcal{M})$ , where the set  $\mathcal{T}$  of terminals is 1-well-linked in  $G$ .

Chekuri, Khanna and Shepherd [CKS04b, CKS05, CKS04a] have suggested the following high-level approach to solving EDP instances  $(G, \mathcal{M})$ , where the terminals are well-linked in  $G$ . They start by defining a graph called a *crossbar*: a graph  $H$  with a subset  $Y \subseteq V(H)$  of vertices is called a crossbar for  $Y$  with congestion  $c$ , iff any matching over the vertices of  $Y$  can be routed with congestion at most  $c$  in  $H$ . They then note that if we could show an algorithm that finds a crossbar  $(H, Y)$  in graph  $G$ , with  $|Y| = k/\text{poly} \log k$ , and constant congestion, then we can obtain a polylogarithmic approximation to EDPwC with constant congestion. An algorithm for constructing such a crossbar with constant congestion follows from the work of [Chu12].

We follow this approach, and define a structure that we call a *good crossbar*, which gives slightly stronger properties than the general crossbar defined above. For a subset  $S \subseteq V$  of vertices, let  $\text{out}(S)$  denote the set of edges with one endpoint in  $S$  and one endpoint in  $V \setminus S$ . Informally, we say that a vertex set  $S$  has the  $\alpha$ -bandwidth property with respect to a subset  $\Gamma \subseteq \text{out}(S)$  of edges, iff the set  $\Gamma$  of terminals is  $\alpha$ -well-linked in graph  $G[S]$ . Formally, we require that for any partition  $(A, B)$  of  $S$ ,  $|E(A, B)| \geq \alpha \cdot \min\{|\text{out}(A) \cap \Gamma|, |\text{out}(B) \cap \Gamma|\}$ .

A good crossbar consists of three parts. The first part is a *good family of vertex subsets*  $\mathcal{F} = \{S_1, \dots, S_\gamma\}$ , where  $\gamma = O(\log^2 n)$ . The sets  $S_1, \dots, S_\gamma$  are all mutually disjoint, and only contain non-terminal vertices. Each set  $S_j$  is associated with a subset  $\Gamma_j^* \subseteq \text{out}(S_j)$  of  $k^* = k/\text{poly} \log k$  edges, such that  $S_j$  has the 1-bandwidth property with respect to  $\Gamma_j^*$ . The second part of the good crossbar is a collection  $\mathcal{M}^* \subseteq \mathcal{M}$  of  $k^*/2$  demand pairs, and the third part is a collection  $\tau = T_1, \dots, T_{k^*}$  of

$k^*$  trees. Let  $\mathcal{T}^*$  denote the set of all terminals participating in the demand pairs in  $\mathcal{M}^*$ . Then for each  $1 \leq i \leq k^*$ , the tree  $T_i$  contains a distinct terminal  $t_i \in \mathcal{T}^*$ , and a distinct edge  $e_{i,j} \in \Gamma_j^*$ , for all  $1 \leq j \leq \gamma$ . In other words, for each set  $S_j \in \mathcal{F}$ , we can denote  $\Gamma_j^* = \{e_{1,j}, \dots, e_{k^*,j}\}$ , where  $e_{i,j} \in E(T_i)$  for all  $1 \leq i \leq k^*$ . (See Figure 1.)




Figure 1: A good crossbar

Given a good crossbar  $(\mathcal{F}, \mathcal{M}^*, \tau)$ , we say that the congestion of this crossbar is  $c$  iff every edge of  $G$  appears in at most  $c$  trees in  $\tau$ , and every edge of  $G$  with both endpoints in the same set  $S_j \in \mathcal{F}$  appears in at most  $c - 1$  such trees. Chuzhoy [Chu12] has implicitly defined a good crossbar, and has shown that, given a good crossbar that causes congestion  $c$ , there is an efficient randomized algorithm to route  $\Omega(k^*/\text{poly log } k)$  demand pairs with congestion at most  $c$  in graph  $G$ . This algorithm uses the cut-matching game of Khandekar, Rao and Vazirani [KRV09] to embed an expander into  $G$ , and then finds a routing in this expander using the algorithm of Rao and Zhou [RZ10]. She has also shown an efficient algorithm that either finds a routing of  $O(k/\text{poly log } k)$  demand pairs with congestion 14, or constructs a good crossbar with congestion 14, thus obtaining an  $O(\text{poly log } k)$ -approximation with congestion 14 for EDPwC.

We follow a similar approach here, except that we construct a good crossbar with congestion 2. Combining this with the result of [Chu12], we obtain an  $O(\text{poly log } k)$ -approximation to EDPwC with congestion 2. We now provide more details about the construction of the good crossbar of [Chu12], and the barriers that we need to overcome in order to reduce the congestion to 2. The algorithm of [Chu12] consists of three steps. In the first step, a good family  $\mathcal{F} = \{S_1, \dots, S_\gamma\}$  of vertex subsets is constructed, and for each  $1 \leq j \leq \gamma$ , a subset  $\Gamma_j^* \subseteq \text{out}(S_j)$  is selected. Additionally, for each set  $S_j \in \mathcal{F}$ , there is a set  $\mathcal{P}_j$  of  $k^*$  edge-disjoint paths in graph  $G$ , connecting every edge in  $\Gamma_j^*$  to a distinct terminal in  $\mathcal{T}$ . In the second step, we construct a family  $\tau' = \{T'_1, \dots, T'_{k^*}\}$  of trees, where every tree  $T'_i \in \tau'$  contains a distinct edge  $e_{i,j} \in \Gamma_j^*$  for each  $1 \leq j \leq \gamma$ . In the third step, we select a subset  $\mathcal{M}^* \subseteq \mathcal{M}$  of  $k^*/2$  pairs, and let  $\mathcal{T}^*$  be the set of terminals participating in the pairs in  $\mathcal{M}^*$ . We then find a collection of paths that cause a small constant congestion in  $G$ , connecting each terminal in  $\mathcal{T}^*$  to a distinct tree in  $\tau'$ . The union of these new paths with the trees in  $\tau'$  gives the final collection  $\tau$  of trees. In order to construct the trees in  $\tau'$ , consider the graph  $G'$  obtained from  $G$  by contracting every set  $S_j \in \mathcal{F}$  into a super-node  $v_j$ . Then the problem of finding the set  $\tau'$  of trees is similar to the problem of packing Steiner trees in  $G'$ . Moreover, the collections  $\{\mathcal{P}_j\}_{j=1}^\gamma$  of paths can be interpreted as a fractional solution of value  $\Omega(k/\text{poly log } k)$  to this problem. We can then round this solution to obtain the family  $\tau'$  of trees, using known algorithms for packing Steiner trees. The fact that we need to deal with the clusters  $S_j$  instead of the super-nodes  $v_j$  makes the problem more challenging technically, but due to the fact that the sets  $S_j$  have the bandwidth property with respect to the sets  $\Gamma_j^*$  of edges, the problem is still tractable. The bandwidth property of the sets  $S_j \in \mathcal{F}$  is exploited in

this algorithm, by routing several matchings between the edges of  $\Gamma_j^*$  across each such set  $S_j$ , in order to simulate the super-node  $v_j$ .

There are several steps contributing to the accumulation of congestion in this construction. We mention the two main barriers to reducing the congestion to 2 here. The first problem is that, given the good family  $\mathcal{F}$  of vertex subsets, it is possible that for some set  $S_j \in \mathcal{F}$  the paths in  $\mathcal{P}_j$  may contain edges lying in other clusters  $G[S_{j'}]$  where  $S_{j'} \in \mathcal{F}$ . Moreover, these paths do not necessarily enter and leave the set  $S_{j'}$  through the edges in  $\mathcal{T}_{j'}^*$ . The construction of the trees in  $\tau'$  uses each cluster  $G[S_{j'}]$  with  $S_{j'} \in \mathcal{F}$  to route several matchings between the edges of  $\Gamma_{j'}^*$ . If the edges of  $G[S_{j'}]$  additionally participate in the paths in  $\{\mathcal{P}_j\}_{j=1}^\gamma$ , then this leads to accumulation of congestion. Indeed, an edge of graph  $G$  may appear in up to 8 trees of  $\tau'$  in the construction of [Chu12]. Ideally, we would like to construct a good family  $\mathcal{F}$  of vertex subsets, where for each set  $S_j \in \mathcal{F}$ , the paths in  $\mathcal{P}_j$  do not contain any vertices lying in other subsets  $S_{j'} \in \mathcal{F}$ . However, it is not clear whether it is possible to achieve such strong properties.

The second difficulty is that step 2 and step 3 are executed separately, each contributing to the total congestion. It appears that one has to incur a congestion of at least 2 when constructing the set  $\tau'$  of trees, using current techniques. We then need to additionally select a subset  $\mathcal{M}^* \subseteq \mathcal{M}$  of the demand pairs, and to route the terminals participating in  $\mathcal{M}^*$  to the trees in  $\tau'$ , thus increasing the congestion beyond 2.

In order to find a good family  $\mathcal{F}$  of vertex subsets, the algorithm of [Chu12] performs a number of iterations. In each iteration we start with what is called a legal contracted graph  $G'$ . This graph is associated with a partition  $\mathcal{C}$  of  $V(G)$  into clusters, such that each terminal  $t \in T$  belongs to a separate cluster (that is,  $\{t\} \in \mathcal{C}$ ), each set  $C \in \mathcal{C}$  has the bandwidth property with respect to  $\text{out}(C)$ , and  $|\text{out}(C)|$  is small. Graph  $G'$  is obtained from  $G$  by contracting every cluster  $C \in \mathcal{C}$  into a super-node. We call such a clustering  $\mathcal{C}$  a *good clustering*. At the beginning of the algorithm,  $\mathcal{C}$  contains each vertex of  $G$  in a separate cluster, and  $G' = G$ . In every iteration, given a legal contracted graph  $G'$ , the algorithm either computes a good family  $\mathcal{F}$  of vertex subsets, or produces a new legal contracted graph  $G''$ , containing strictly fewer edges than  $G'$ . This guarantees that after  $n$  iterations, the algorithm produces a good family  $\mathcal{F}$  of vertex subsets.

In order to overcome the two problems mentioned above and avoid accumulating the congestion, we combine all three steps of the algorithm together. We define a potential function  $\varphi$  over partitions  $\mathcal{C}$  of  $V(G)$ , where  $\varphi(\mathcal{C})$  roughly measures the number of edges in the graph obtained from  $G$  by contracting every cluster in  $\mathcal{C}$  into a super-node. The potential function  $\varphi$  has additional useful properties, that allow us to perform a number of standard operations (such as the bandwidth decomposition) on the clusters of  $\mathcal{C}$ , without increasing the potential value.

As before, our algorithm consists of a number of iterations (that we call phases). In each such phase, we start with some legal contracted graph  $G'$  and a corresponding good clustering  $\mathcal{C}'$ . We then either construct a good crossbar, or produce a new good clustering  $\mathcal{C}''$  with  $\varphi(\mathcal{C}'') < \varphi(\mathcal{C}')$ , together with the corresponding new legal contracted graph  $G''$ . Each phase is executed as follows. We start with a collection of  $\gamma$  different clusterings  $\mathcal{C}_1, \dots, \mathcal{C}_\gamma$  of the vertices of  $G$ , where for each  $i$ ,  $\mathcal{C}_i$  is a partition of  $V(G)$  into clusters, and  $\varphi(\mathcal{C}_i) < \varphi(\mathcal{C}')$  (but  $\mathcal{C}_i$  is not necessarily a good clustering). We then perform a number of iterations. In each iteration, we select a family  $\mathcal{F} = \{S_1, \dots, S_\gamma\}$  of disjoint subsets of non-terminal vertices, that we treat as a potential good family vertex subsets, and we try to find a subset  $\mathcal{M}^* \subseteq \mathcal{M}$  of  $k^*$  demand pairs and a family  $\tau$  of trees, to complete the construction of a good crossbar. If we do not succeed in constructing a good crossbar in the current iteration, then we use the family  $\mathcal{F}$  of vertex subsets to refine one of the current clusterings  $\mathcal{C}_i$ , such that the potential of the new clustering goes down by a significant amount. This ensures that after polynomially-many iterations, we will succeed in either constructing a good crossbar, or a good clustering  $\mathcal{C}^*$  with  $\varphi(\mathcal{C}^*) < \varphi(\mathcal{C}')$ .

This combination of all three steps of the algorithm of [Chu12] appears necessary to overcome the two barriers described above. For example, it is possible that the family  $\mathcal{F}$  of vertex subsets is a good family, but we are still unable to extend it to a good crossbar (for instance due to the paths in  $\mathcal{P}_j$  using the edges of  $G[S_{j'}]$ , as described above). Still, we will be able to make progress in such cases by refining one of the current clusterings  $\mathcal{C}_i$ . Similarly, we construct the trees and connect the terminals participating in pairs in  $\mathcal{M}^*$  to them simultaneously, to avoid accumulating congestion. Again, whenever we are unable to do so, we will be able to refine the one of the current clusterings  $\mathcal{C}_i$ .

**Subsequent Work** Following our work, Chekuri and Chuzhoy [CC14] have extended and generalized our construction of a good crossbar to handle node-disjoint routing. They used this construction to obtain the first polynomial bounds for the Excluded Grid Theorem of Robertson and Seymour [RS86]. This bound was recently improved in [Chu16]. We note that the algorithmic framework and the technical tools introduced in this paper, including the potential function for clusterings and legal contracted graphs, were used in both these results.

**Organization** We start with preliminaries in Section 2 and provide an overview of our algorithm in Section 3. We develop the machinery to analyze vertex clusterings in Section 4, and complete the algorithm description in Sections 5 and 6. For convenience, a list of parameters is provided in Section E, and the known lower bounds on the integrality gap of the multi-commodity flow relaxation for EDP are provided in Section D of the Appendix.

## 2 Preliminaries

We assume that we are given an undirected  $n$ -vertex graph  $G = (V, E)$ , and a set  $\mathcal{M} = \{(s_1, t_1), \dots, (s_k, t_k)\}$  of  $k$  source-sink pairs, that we also call demand pairs. We denote by  $\mathcal{T}$  the set of vertices that participate in pairs in  $\mathcal{M}$ , and we call them *terminals*. Let  $\text{OPT}$  denote the maximum number of demand pairs that can be simultaneously routed via edge-disjoint paths. Given any subset  $\mathcal{M}' \subseteq \mathcal{M}$  of the demand pairs, we denote by  $\mathcal{T}(\mathcal{M}') \subseteq \mathcal{T}$  the subset of terminals participating in the pairs in  $\mathcal{M}'$ .

Using a standard transformation, we can assume without loss of generality that every terminal  $t \in \mathcal{T}$  participates in exactly one source-sink pair: otherwise, for each demand pair in which  $t$  participates, we can add a new terminal to the graph, that replaces  $t$  in the demand pair, and connect this new terminal to  $t$ . Similarly, using standard transformations, we can assume that the degree of every terminal is exactly 1, and the degree of every non-terminal vertex is at most 4. In order to achieve the latter property, we replace every vertex  $v$  whose degree  $d_v > 4$  with the  $(d_v \times d_v)$ -grid, and connect the edges incident to  $v$  to the vertices of the first row of the grid. It is easy to verify that these transformations do not affect the solution value for any congestion parameter. From now on we assume that the degree of every terminal is 1, the degree of every non-terminal vertex is at most 4, and every terminal participates in exactly one demand pair.

For any subset  $S \subseteq V$  of vertices, we denote by  $\text{out}_G(S) = E_G(S, V \setminus S)$ , and by  $E_G(S)$  the subset of edges with both endpoints in  $S$ , omitting the subscript  $G$  when clear from context. Given a subset  $S$  of vertices of  $G$ , we sometimes denote  $\bar{S} = V(G) \setminus S$ , and we sometimes refer to the partition  $(S, \bar{S})$  of  $V(G)$  as a cut. Throughout the paper, we say that a random event succeeds with high probability, if the probability of its success is at least  $(1 - 1/\text{poly}(n))$ . All logarithms are to the base of 2.

Let  $\mathcal{P}$  be any collection of paths in graph  $G$ . We say that paths in  $\mathcal{P}$  cause congestion  $\eta$  in  $G$ , iff for every edge  $e \in E$ , at most  $\eta$  paths in  $\mathcal{P}$  contain  $e$ . Given any pair  $(E_1, E_2)$  of subsets of edges, we denote by  $F : E_1 \rightsquigarrow_\eta E_2$  the flow where every edge in  $E_1$  sends one flow unit to the edges in  $E_2$ , and the congestion due to the flow  $F$  is at most  $\eta$ . If additionally every edge in  $E_2$  receives at most one flow unit, then we denote this flow by  $F : E_1 \stackrel{1:1}{\rightsquigarrow}_\eta E_2$ . Similarly, given a set  $\mathcal{P}$  of paths connecting the edges of  $E_1$  to the edges of  $E_2$ , we denote  $\mathcal{P} : E_1 \rightsquigarrow_\eta E_2$  iff  $\mathcal{P} = \{P_e \mid e \in E_1\}$ , where  $e$  is the first

edge on  $P_e$ , and the total congestion caused by the paths in  $\mathcal{P}$  is at most  $\eta$ . If every edge of  $E_2$  has at most one path terminating at it, then we denote  $\mathcal{P} : E_1 \overset{1:1}{\rightsquigarrow}_\eta E_2$ . We use a similar notation for flows and path sets connecting subsets of vertices to each other, or a subset vertices with a subset of edges.

Given a subset  $S$  of vertices and two subsets  $E_1, E_2 \subseteq \text{out}(S)$  of edges, we say that the flow  $F : E_1 \rightsquigarrow_\eta E_2$  is *contained in*  $S$  iff every flow-path is completely contained in  $G[S]$ , except for its first and last edges, that belong to  $\text{out}(S)$ . Similarly, we say that a set  $\mathcal{P} : E_1 \rightsquigarrow_\eta E_2$  of paths is contained in  $S$  iff all inner edges on every path in  $\mathcal{P}$  belong to  $G[S]$ .

**Sparsest Cut and the Flow-Cut Gap** Suppose we are given a graph  $G = (V, E)$ , and a subset  $\mathcal{T} \subseteq V$  of  $k$  terminals. Given a partition  $(S, \bar{S})$  of  $V$ , with  $S \cap \mathcal{T}, \bar{S} \cap \mathcal{T} \neq \emptyset$ , the *sparsity* of the partition is  $\Phi(S) = \frac{|E(S, \bar{S})|}{\min\{|S \cap \mathcal{T}|, |\bar{S} \cap \mathcal{T}|\}}$ , and the value of the sparsest cut in  $G$  is defined to be:  $\Phi(G) = \min_{\substack{S \subseteq V \\ S \cap \mathcal{T}, \bar{S} \cap \mathcal{T} \neq \emptyset}} \{\Phi(S)\}$ . The goal of the sparsest cut problem is, given an input graph  $G$  and a set  $\mathcal{T}$  of terminals, to find a cut of minimum sparsity. Arora, Rao and Vazirani [ARV09] have shown an  $O(\sqrt{\log k})$ -approximation algorithm for the sparsest cut problem. We denote this algorithm by  $\mathcal{A}_{\text{ARV}}$ , and its approximation factor by  $\alpha_{\text{ARV}}(k) = O(\sqrt{\log k})$ .

**Well-Linkedness and Bandwidth Property** We define the the notion of well-linkedness, that has been used extensively in algorithms for network routing, e.g. in [Räc02, CKS04b, CKS05, RZ10, And10, Chu12].

**Definition 1.** We say that a set  $\mathcal{T}$  of vertices is  $\alpha$ -well-linked in graph  $G$ , iff for every partition  $(A, B)$  of  $V(G)$ ,  $|E_G(A, B)| \geq \alpha \cdot \min\{|A \cap \mathcal{T}|, |B \cap \mathcal{T}|\}$ .

Notice that if  $\mathcal{T}$  is  $\alpha$ -well-linked in  $G$ , then for any pair  $(\mathcal{T}_1, \mathcal{T}_2)$  of subsets of  $\mathcal{T}$  with  $|\mathcal{T}_1| = |\mathcal{T}_2|$ , we can efficiently find a collection  $\mathcal{P}$  of paths,  $\mathcal{P} : \mathcal{T}_1 \overset{1:1}{\rightsquigarrow}_{[1/\alpha]} \mathcal{T}_2$ . This follows from the max-flow min-cut theorem and the integrality of flow (a formal proof can be found in the full version of [CC14]).

Next, we define the *bandwidth property* - a notion closely related to well-linkedness, and similar to the notion of bandwidth property defined in [Räc02] (this notion was also defined in [Chu12]). Intuitively, we say that a cluster  $S$  has the  $\alpha$ -bandwidth property, iff the boundary  $\text{out}(S)$  of the cluster is  $\alpha$ -well-linked in  $G[S]$ . The formal definition appears below.

**Definition 2.** Given a graph  $G$ , a subset  $S$  of its vertices, and a parameter  $\alpha > 0$ , we say that  $S$  has the  $\alpha$ -bandwidth property, iff for every partition  $(A, B)$  of  $S$ :

$$|E(A, B)| \geq \alpha \cdot \min\{|\text{out}(A) \cap \text{out}(S)|, |\text{out}(B) \cap \text{out}(S)|\}.$$

Given a subset  $\Gamma \subseteq \text{out}(S)$  of edges, we say that  $S$  has the  $\alpha$ -bandwidth property with respect to  $\Gamma$  iff for any partition  $(A, B)$  of  $S$ ,  $|E(A, B)| \geq \alpha \cdot \min\{|\text{out}(A) \cap \Gamma|, |\text{out}(B) \cap \Gamma|\}$ .

We also need a more general notion of bandwidth, introduced in [Chu12]. Intuitively, it handles subsets  $S$  of vertices, where  $|\text{out}(S)|$  may be large, but we only need to route small amounts of flow across  $S$ .

**Definition 3.** Let  $S$  be a subset of vertices of a graph  $G$ . For an integer  $k > 0$  and for a parameter  $0 < \alpha < 1$ , we say that set  $S$  has the  $(k, \alpha)$ -bandwidth property iff for every subset  $\Gamma \subseteq \text{out}(S)$  with  $|\Gamma| \leq k$ , set  $S$  has the  $\alpha$ -bandwidth property with respect to  $\Gamma$ .

Note that if  $|\text{out}(S)| \leq k$ , then set  $S$  has the  $(k, \alpha)$ -bandwidth property iff it has the  $\alpha$ -bandwidth property, and the two definitions of bandwidth property become equivalent.



Assume now that  $S$  does not have the  $(k, \alpha)$ -bandwidth property. Then there must be a partition  $(X, Y)$  of  $S$ , such that  $|E(X, Y)| < \alpha \cdot \min\{|\text{out}(X) \cap \text{out}(S)|, |\text{out}(Y) \cap \text{out}(S)|, k/2\}$ . We say that  $(X, Y)$  is a  $(k, \alpha)$ -violating cut for  $S$ .

Given a graph  $G$ , a subset  $S \subseteq V(G)$  of vertices, and a subset  $\Gamma \subseteq \text{out}(S)$  of edges, we define an instance  $\text{SC}(G, S, \Gamma)$  of the sparsest cut problem as follows. First, we sub-divide every edge  $e \in \Gamma$  by a new vertex  $t_e$ , and we let  $\mathcal{T}(\Gamma) = \{t_e \mid e \in \Gamma\}$ . Let  $G_S(\Gamma)$  be the sub-graph of the resulting graph, induced by  $S \cup \mathcal{T}(\Gamma)$ . The instance  $\text{SC}(G, S, \Gamma)$  of the sparsest cut problem is defined over the graph  $G_S(\Gamma)$ , where the vertices of  $\mathcal{T}(\Gamma)$  serve as terminals. Observe that for all  $0 < \alpha < 1$ , the following are equivalent:

- the value of the sparsest cut in  $\text{SC}(G, S, \Gamma)$  is at least  $\alpha$ ;
- set  $S$  has the  $\alpha$ -bandwidth property with respect to  $\Gamma$  in graph  $G$ ; and
- set  $\mathcal{T}(\Gamma)$  is  $\alpha$ -well-linked in graph  $G_S(\Gamma)$ .

If  $\Gamma = \text{out}_G(S)$ , then we will denote the corresponding instance of the sparsest cut problem by  $\text{SC}(G, S)$ .

**The Grouping Technique** The grouping technique was first introduced by Chekuri, Khanna and Shepherd [CKS04b], and has since been widely used in algorithms for network routing [CKS05, RZ10, And10, Chu12], to boost the network connectivity and well-linkedness parameters. We start with a simple standard grouping procedure, summarized in the following theorem; its proof can be found e.g. in [CKS04b].

**Theorem 2.1** *Suppose we are given a connected graph  $G = (V, E)$ , with weights  $w(v)$  on vertices  $v \in V$ , and a parameter  $p$ , such that for each  $v \in V$ ,  $0 \leq w(v) \leq p$ , and  $\sum_{v \in V} w(v) \geq p$ . Then we can efficiently find a partition  $\mathcal{G}$  of  $V$ , and for each group  $U \in \mathcal{G}$ , find a tree  $T_U \subseteq G$  containing all vertices of  $U$ , such that the trees  $\{T_U\}_{U \in \mathcal{G}}$  are edge-disjoint, and for each  $U \in \mathcal{G}$ ,  $p \leq \sum_{v \in U} w(v) \leq 3p$ .*

We also use a more sophisticated grouping technique, due to Chekuri, Khanna and Shepherd [CKS04b]. The next two theorems summarize the specific settings in which the technique is used. Since these settings are slightly different from the setting of [CKS04b], we provide their proofs for completeness in Section A in the Appendix. The proofs closely follow the arguments of [CKS04b].

**Theorem 2.2** *Suppose we are given a connected graph  $G = (V, E)$  with maximum vertex degree at most 4, and  $r$  subsets  $\mathcal{T}_1, \dots, \mathcal{T}_r$  of vertices called terminals (we do not require that they are disjoint). Assume further that the set  $\bigcup_{j=1}^r \mathcal{T}_j$  of terminals is  $\alpha$ -well-linked in  $G$ , for some  $0 < \alpha < 1$ ,  $|\mathcal{T}_1| = k_1$ , and that for each  $1 < j \leq r$ ,  $|\mathcal{T}_j| = k_2$ , where  $k_1, k_2 \geq \Omega(r^2 \log r / \alpha)$ . Then there is an efficient randomized algorithm that with high probability computes, for each  $1 \leq j \leq r$ , a subset  $\mathcal{T}'_j \subseteq \mathcal{T}_j$  of terminals, such that all sets  $\{\mathcal{T}'_j\}_{j=1}^r$  are mutually disjoint,  $|\mathcal{T}'_j| = \Omega(\alpha/r^2) \cdot |\mathcal{T}_j|$  for all  $1 \leq j \leq r$ , and  $\bigcup_{j=1}^r \mathcal{T}'_j$  is 1-well-linked in  $G$ .*

**Theorem 2.3** *Suppose we are given a connected graph  $G = (V, E)$  with maximum vertex degree at most 4, two subsets  $\mathcal{T}_1, \mathcal{T}_2$  of vertices called terminals, and a perfect matching  $\mathcal{M}$  over the terminals of  $\mathcal{T}_1$ . Assume further that  $|\mathcal{T}_1| = k_1$ ,  $|\mathcal{T}_2| = k_2$ , where  $k_1, k_2 \geq 100/\alpha$ , and the set  $\mathcal{T}_1 \cup \mathcal{T}_2$  of terminals is  $\alpha$ -well-linked in  $G$ , for some  $0 < \alpha < 1$ . Then there is an efficient algorithm, that either computes a routing of a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of  $\Omega(\alpha k_1)$  pairs on edge-disjoint paths, or returns two **disjoint** subsets  $\mathcal{T}'_1 \subseteq \mathcal{T}_1$ ,  $\mathcal{T}'_2 \subseteq \mathcal{T}_2$ , so that  $|\mathcal{T}'_1| = \Omega(\alpha k_1)$ ,  $|\mathcal{T}'_2| = \Omega(\alpha k_2)$ , set  $\mathcal{T}'_1 \cup \mathcal{T}'_2$  is 1-well-linked in  $G$ , and for every pair  $(s, t) \in \mathcal{M}$ , either both  $s$  and  $t$  belong to  $\mathcal{T}'_1$ , or neither of them does.*

**Expanders and the Cut-Matching Game.** We say that a (multi)-graph  $G = (V, E)$  is an  $\alpha$ -expander, iff  $\min_{\substack{S \subseteq V: \\ 0 < |S| \leq |V|/2}} \left\{ \frac{|E(S, \bar{S})|}{|S|} \right\} \geq \alpha$ .

We use the cut-matching game of Khandekar, Rao and Vazirani [KRV09]. In this game, we are given a set  $V$  of  $N$  vertices, where  $N$  is even, and two players: a cut player, whose goal is to construct an expander  $X$  over the set  $V$  of vertices, and a matching player, whose goal is to delay its construction. The game is played in iterations. We start with the graph  $X$  containing the set  $V$  of vertices, and no edges. In each iteration  $j$ , the cut player computes a bi-partition  $(A_j, B_j)$  of  $V$  into two equal-sized sets, and the matching player returns some perfect matching  $M_j$  between the two sets. The edges of  $M_j$  are then added to  $X$ . Khandekar, Rao and Vazirani have shown that there is a strategy for the cut player, guaranteeing that after  $O(\log^2 N)$  iterations we obtain a  $\frac{1}{2}$ -expander with high probability. Subsequently, Orecchia et al. [OSVV08] have shown the following improved bound:

**Theorem 2.4 ([OSVV08])** *There is an efficient randomized algorithm for the cut player, such that, no matter how the matching player plays, after  $\gamma_{\text{CMG}}(N) = O(\log^2 N)$  iterations, graph  $X$  is an  $\alpha_{\text{CMG}}(N) = \Omega(\log N)$ -expander, with constant probability.*

### 3 Algorithm Overview

Throughout the paper, we denote by  $\gamma_{\text{CMG}} = \gamma_{\text{CMG}}(k) = O(\log^2 k)$  the parameter from Theorem 2.4. The central combinatorial object in our algorithm is what we call a *good crossbar*, that we define below. (See Figure 1 for illustration.)

**Definition 4.** Given a graph  $G = (V, E)$ , a set  $\mathcal{M}$  of  $k$  source-sink pairs, and a parameter  $k^* = k / \text{poly log } k$ , a good crossbar consists of the following three components:

1. A family  $\mathcal{S}^* = \{S_1^*, \dots, S_{\gamma_{\text{CMG}}}^*\}$  of disjoint subsets of **non-terminal** vertices. Each set  $S_j^* \in \mathcal{S}^*$  is associated with a subset  $\Gamma_j^* \subseteq \text{out}(S_j^*)$  of  $2k^*$  edges, and  $S_j^*$  has the 1-bandwidth property with respect to  $\Gamma_j^*$ .
2. A subset  $\mathcal{M}^* \subseteq \mathcal{M}$  of  $k^*$  demand pairs. Let  $\mathcal{T}^* = \mathcal{T}(\mathcal{M}^*)$  be the corresponding set of terminals.
3. A collection  $\tau^* = \{T_1, \dots, T_{2k^*}\}$  of trees in graph  $G$ . Each tree  $T_i \in \tau^*$ , contains a distinct terminal  $t_i \in \mathcal{T}^*$ , and for each  $1 \leq j \leq \gamma_{\text{CMG}}$ , tree  $T_i$  contains a distinct edge  $e_{i,j} \in \Gamma_j^*$ . In other words,  $\mathcal{T}^* = \{t_1, \dots, t_{2k^*}\}$ , where  $t_i \in T_i$  for each  $1 \leq i \leq 2k^*$ ; and for each  $1 \leq j \leq \gamma_{\text{CMG}}$ ,  $\Gamma_j^* = \{e_{1,j}, \dots, e_{2k^*,j}\}$ , where  $e_{i,j} \in T_i$  for each  $1 \leq i \leq 2k^*$ .

We say that the congestion of the good crossbar is  $c$  iff every edge of  $G$  participates in at most  $c$  trees in  $\tau^*$ , while every edge in  $\bigcup_{j=1}^{\gamma_{\text{CMG}}} E(S_j^*)$  belongs to at most  $c - 1$  such trees.

Our main result is that if the set  $\mathcal{T}$  of terminals is 1-well-linked in  $G$ , then there is an efficient randomized algorithm, that either finds a good congestion-2 crossbar, or routes a subset of  $k / \text{poly log } k$  demand pairs with congestion at most 2 in  $G$ . We summarize this result in the following theorem.

**Theorem 3.1** *Assume that we are given an undirected graph  $G = (V, E)$  with vertex degrees at most 4, and a set  $\mathcal{M}$  of  $k$  demand pairs, defined over a set  $\mathcal{T}$  of terminals. Assume further that the degree of every terminal is 1, every terminal participates in exactly one pair in  $\mathcal{M}$ , and  $\mathcal{T}$  is 1-well-linked in  $G$ . Then there is an efficient randomized algorithm, that with high probability outputs one of the following: (i) Either a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of  $k / \text{poly log } k$  demand pairs and the routing of the pairs in  $\mathcal{M}'$  with congestion at most 2 in  $G$ ; or (ii) a good congestion-2 crossbar  $(\mathcal{S}^*, \mathcal{M}^*, \tau^*)$ .*

The proof of Theorem 1.1 follows from the proof of Theorem 3.1 using techniques from previous work, and it appears in Section B of the Appendix. We note that the algorithm of [Chu12] also proceeded by (implicitly) constructing a good congestion-14 crossbar. Our main challenge is to reduce the congestion of the good crossbar to 2. From now on we focus on proving Theorem 3.1.

As in [Chu12], throughout the algorithm, we maintain a partition  $\mathcal{C}$  of the vertices of  $G$  into clusters, and a legal contracted graph of  $G$ , obtained by contracting each cluster  $C \in \mathcal{C}$  into a super-node  $v_C$ . The algorithm performs a number of phases, where in every phase we either find the desired routing of a subset of demand pairs in graph  $G$ , or compute a good congestion-2 crossbar, or find a new partition  $\mathcal{C}'$  of the vertices of  $G$  into clusters, whose corresponding contracted graph is smaller than the current contracted graph (with respect to the potential function that we define). In the next section we define several types of clusterings the algorithm uses, the notion of the legal contracted graph, and several operations on a given clustering, that are used throughout the algorithm.

## 4 Vertex Clusterings and Legal Contracted Graphs

Our algorithm uses the following parameters. Let  $\gamma_{\text{CMG}} = \gamma_{\text{CMG}}(k) = O(\log^2 k)$  be the parameter for the number of iterations in the cut-matching game from Theorem 2.4. We let  $\gamma = 2^{24} \gamma_{\text{CMG}}^4$ . We will also use the following two parameters for the bandwidth property:  $\alpha = \frac{1}{2^{11} \gamma \log k} = \Omega\left(\frac{1}{\log^9 k}\right)$ , used to perform the bandwidth decomposition, and  $\alpha_{\text{BW}} = \frac{\alpha}{\alpha_{\text{ARV}}(k)} = \Omega\left(\frac{1}{\log^{9.5} k}\right)$  - the parameter that we achieve for the bandwidth property. Finally, we use a parameter  $k_1 = \frac{k}{192 \gamma^3 \log \gamma} = \frac{k}{\text{poly log } k}$ , and we assume that the parameter  $k$  is large enough, so  $k_1 > 4/\alpha$  (otherwise we can simply route one of the demand pairs to obtain an  $O(\text{poly log } k)$ -approximation). We say that a cluster  $C \subseteq V(G)$  is *large* if  $|\text{out}(C)| \geq k_1$ , and we say that it is *small* otherwise.

**Definition 5.** Given a partition  $\mathcal{C}$  of the vertices of  $V(G)$  into clusters, we say that  $\mathcal{C}$  is an *acceptable clustering* of  $G$  iff:

- Every terminal  $t \in \mathcal{T}$  is in a separate cluster, that is,  $\{t\} \in \mathcal{C}$ ;
- Each small cluster  $C \in \mathcal{C}$  has the  $\alpha_{\text{BW}}$ -bandwidth property; and
- For each large cluster  $C \in \mathcal{C}$ ,  $G[C]$  is connected.

An acceptable clustering that contains no large clusters is called a *good clustering*.

**Definition 6.** Given a good clustering  $\mathcal{C}$  of  $G$ , a graph  $H_{\mathcal{C}}$  is a legal contracted graph of  $G$  associated with  $\mathcal{C}$ , iff we can obtain  $H_{\mathcal{C}}$  from  $G$  by contracting every  $C \in \mathcal{C}$  into a super-node  $v_C$ . We remove all self-loops, but we do not remove parallel edges. (Note that the terminals are not contracted since each terminal has its own cluster).

The following claim was proved in [Chu12]. We include its proof here for completeness, since we have changed some parameters.

**Claim 4.1** *If  $G'$  is a legal contracted graph for  $G$ , then  $G' \setminus \mathcal{T}$  contains at least  $k/3$  edges.*

**Proof:** For each terminal  $t \in \mathcal{T}$ , let  $e_t$  be the unique edge adjacent to  $t$  in  $G'$ , and let  $u_t$  be the other endpoint of  $e_t$ . We partition the terminals in  $\mathcal{T}$  into groups, where two terminals  $t, t'$  belong to the same group iff  $u_t = u_{t'}$ . Let  $\mathcal{G}$  be the resulting partition of the terminals. Since the degree of

every vertex in  $G'$  is at most  $k_1$ , each group  $U \in \mathcal{G}$  contains at most  $k_1$  terminals. Next, we find two disjoint subsets  $X, Y$  of  $\mathcal{T}$  such that  $|X| = |Y| = \lceil k/3 \rceil$ , and for each group  $U \in \mathcal{G}$ , either  $U \cap X = \emptyset$ , or  $U \cap Y = \emptyset$ . We can find such a partition by greedily processing each group  $U \in \mathcal{G}$ , and adding all terminals of  $U$  to one of the subsets  $X$  or  $Y$ , that currently contains fewer terminals. Finally, we remove terminals from set  $X$  until  $|X| = \lceil k/3 \rceil$ , and we do the same for  $Y$ . Since the terminals are 1-well-linked in graph  $G'$ , it is possible to route  $\lceil k/3 \rceil$  flow units from the terminals in  $X$  to the terminals in  $Y$ , with congestion at most 1. Since no group  $U$  is split between the two sets  $X$  and  $Y$ , each flow-path must contain at least one edge of  $G' \setminus \mathcal{T}$ . Therefore, the number of edges in  $G' \setminus \mathcal{T}$  is at least  $k/3$ .  $\square$

Given **any** partition  $\mathcal{C}$  of the vertices of  $G$ , we define a potential  $\varphi(\mathcal{C})$  for this clustering. The idea is that  $\varphi(\mathcal{C})$  will serve as a tight bound on the number of edges connecting different clusters in  $\mathcal{C}$ . At the same time, the potential function is designed in such a way, that we can perform a number of operations on the current clustering (such as, for example, bandwidth decompositions of the clusters), without increasing the potential.

At a high level, our algorithm maintains a good clustering  $\mathcal{C}$  of  $V(G)$ , where at the beginning, every vertex belongs to a separate cluster. The algorithm consists of a number of phases, where in each phase we start with some good clustering  $\mathcal{C}$ , and either route  $k/\text{poly log } k$  of the demand pairs with congestion at most 2 in  $G$ , or find a good congestion-2 crossbar in  $G$ , or produce another good clustering  $\mathcal{C}'$  with  $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - 1$ . After a polynomial number of phases, we will therefore either obtain a routing of  $k/\text{polylog}(k)$  demand pairs with congestion at most 2, or find a good congestion-2 crossbar.

Each phase of the algorithm is executed as follows. We start with the current good clustering  $\mathcal{C}$ , and construct  $\gamma$  acceptable clustering  $\mathcal{C}_1, \dots, \mathcal{C}_\gamma$ , where for each  $1 \leq i \leq \gamma$ ,  $\varphi(\mathcal{C}_i) < \varphi(\mathcal{C}) - 1$ . We then perform a number of iterations, where in each iteration we either route  $k/\text{poly log } k$  demand pairs in  $G$ , or find a good congestion-2 crossbar, or revisit one of the current acceptable clusterings  $\mathcal{C}_i$ , and perform an operation that produces another acceptable clustering, whose potential is strictly smaller than that of  $\mathcal{C}_i$ . In the end, we will either successfully route  $k/\text{poly log } k$  demand pairs in  $G$ , or construct a good congestion-2 crossbar in  $G$ , or we will find a good clustering  $\mathcal{C}_i$  with  $\varphi(\mathcal{C}_i) < \varphi(\mathcal{C}) - 1$ . We now proceed to define the potential function, and the operations we perform on clusterings.

Suppose we are given **any** partition  $\mathcal{C}$  of the vertices of  $G$ . For any integer  $h$ , we define a potential  $\varphi(h)$ , as follows. For  $h < k_1$ ,  $\varphi(h) = 4\alpha \log h$ . In order to define  $\varphi(h)$  for  $h \geq k_1$ , we consider the sequence  $\{n_0, n_1, \dots\}$  of numbers, where  $n_i = (\frac{3}{2})^i k_1$ . The potentials for these numbers are  $\varphi(n_0) = \varphi(k_1) = 4\alpha \log k_1 + 4\alpha$ , and for  $i > 0$ ,  $\varphi(n_i) = 4\frac{\alpha k_1}{n_i} + \varphi(n_{i-1})$ . Notice that for all  $i$ ,  $\varphi(n_i) \leq 12\alpha + 4\alpha \log k_1 \leq 8\alpha \log k_1 \leq \frac{1}{2^{8\gamma}}$ .

We now partition all integers  $h > k_1$  into sets  $L_1, L_2, \dots$ , where set  $L_i$  contains all integers  $h$  with  $n_{i-1} \leq h < n_i$ . For  $h \in L_i$ , we define  $\varphi(h) = \varphi(n_{i-1})$ . This finishes the definition of potentials of integers. Clearly, for all  $h$ ,  $\varphi(h) \leq \frac{1}{2^{8\gamma}}$ .

Assume now that we are given some edge  $e \in E$ . If both endpoints of  $e$  belong to the same cluster of  $\mathcal{C}$ , then we set its potential  $\varphi(e) = 0$ . Otherwise, if  $e = (u, v)$ , and  $u \in C$  with  $|\text{out}(C)| = h$ , while  $v \in C'$  with  $|\text{out}(C')| = h'$ , then we set  $\varphi(e) = 1 + \varphi(h) + \varphi(h')$ . We think of  $\varphi(h)$  as the contribution of  $u$ , and  $\varphi(h')$  the contribution of  $v$  to  $\varphi(e)$ . Notice that  $\varphi(e) \leq 1.1$ . Finally, we set  $\varphi(\mathcal{C}) = \sum_{e \in E} \varphi(e)$ .

Suppose we are given any partition  $\mathcal{C}$  of  $V(G)$ . We show that we can perform a bandwidth decomposition of small clusters in  $\mathcal{C}$ , without increasing the potential.

**Theorem 4.1** *Let  $\mathcal{C}$  be any partition of  $V(G)$ , and let  $C \in \mathcal{C}$  be any small cluster, such that  $G[C]$  is connected. Then there is an efficient algorithm that finds a partition  $\mathcal{W}$  of  $C$  into small clusters, such that each cluster  $R \in \mathcal{W}$  has the  $\alpha_{\text{BW}}$ -bandwidth property, and additionally, if  $\mathcal{C}'$  is a partition*

obtained from  $\mathcal{C}$  by removing  $C$  and adding the clusters of  $\mathcal{W}$  to it, then  $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C})$ .

**Proof:** We perform a standard bandwidth decomposition of  $C$  (sometimes called well-linked decomposition). We maintain a partition  $\mathcal{W}$  of  $C$ , where at the beginning,  $\mathcal{W} = \{C\}$ . We then perform a number of iterations.

In each iteration, we select a cluster  $S \in \mathcal{W}$ , and consider the corresponding instance  $\text{SC}(G, S)$  of the sparsest cut problem. We run the algorithm  $\mathcal{A}_{\text{ARV}}$  on  $\text{SC}(G, S)$ . If the sparsity of the cut produced by the algorithm is less than  $\alpha$ , then we obtain a partition  $(X, Y)$  of  $S$ , with  $|E(X, Y)| < \alpha \cdot \min\{|\text{out}(S) \cap \text{out}(X)|, |\text{out}(S) \cap \text{out}(Y)|\}$ . In this case, we remove  $S$  from  $\mathcal{W}$ , and add  $X$  and  $Y$  to  $\mathcal{W}$  instead. The algorithm ends when for every cluster  $S \in \mathcal{W}$ , algorithm  $\mathcal{A}_{\text{ARV}}$  returns a partition of sparsity at least  $\alpha$ . We are then guaranteed that every cluster in  $\mathcal{W}$  has the  $\alpha/\alpha_{\text{ARV}}(k) = \alpha_{\text{BW}}$ -bandwidth property, and it is easy to verify that all these clusters are small.

It now only remains to show that the potential does not increase. Each iteration of the algorithm is associated with a partition of the vertices of  $G$ , obtained from  $\mathcal{C}$  by removing  $C$  and adding all clusters of the current partition  $\mathcal{W}$  of  $C$  to it. It is enough to show that if  $\mathcal{C}'$  is the current partition of  $V(G)$ , and  $\mathcal{C}''$  is the partition obtained after one iteration, where a set  $S \in \mathcal{W}$  was replaced by two sets  $X$  and  $Y$ , then  $\varphi(\mathcal{C}'') \leq \varphi(\mathcal{C}')$ .

Assume without loss of generality that  $|\text{out}(X)| \leq |\text{out}(Y)|$ , so  $|\text{out}(X)| \leq 2|\text{out}(S)|/3$ . Let  $h = |\text{out}(S)|$ ,  $h_1 = |\text{out}(X)|$ ,  $h_2 = |\text{out}(Y)|$ , and recall that  $h, h_1, h_2 < k_1$ . The changes to the potential are the following:

- The potential of the edges in  $\text{out}(Y) \cap \text{out}(S)$  only goes down.
- The potential of every edge in  $\text{out}(X) \cap \text{out}(S)$  goes down by  $\varphi(h) - \varphi(h_1) = 4\alpha \log h - 4\alpha \log h_1 = 4\alpha \log \frac{h}{h_1} \geq 4\alpha \log 1.5 \geq 2.3\alpha$ , since  $h_1 \leq 2h/3$ . So the total decrease in the potential of the edges in  $\text{out}(X) \cap \text{out}(S)$  is at least  $2.3\alpha \cdot |\text{out}(X) \cap \text{out}(S)|$ .
- The edges in  $E(X, Y)$  did not contribute to the potential initially, and now contribute at most  $1 + \varphi(h_1) + \varphi(h_2) \leq 2$  each. Notice that  $|E(X, Y)| \leq \alpha \cdot |\text{out}(X) \cap \text{out}(S)|$ , and so they contribute at most  $2\alpha \cdot |\text{out}(X) \cap \text{out}(S)|$  in total.

Clearly, the overall potential goes down. □

Assume that we are given an acceptable clustering  $\mathcal{C}$  of  $G$ . We now define two operations on  $G$ , each of which produces a new acceptable clustering of  $G$ , whose potential is strictly smaller than  $\varphi(\mathcal{C})$ .

**Action 1: Partitioning a large cluster.** Suppose we are given an acceptable clustering  $\mathcal{C}$  of  $G$ , a large cluster  $C \in \mathcal{C}$ , and a  $(k_1, \alpha)$ -violating partition  $(X, Y)$  of  $C$ . In order to perform this operation, we first replace  $C$  with  $X$  and  $Y$  in  $\mathcal{C}$ . If, additionally, any of the clusters,  $X$  or  $Y$ , is small, then we perform a bandwidth decomposition of that cluster using Theorem 4.1, and update  $\mathcal{C}$  with the resulting partition. Clearly, the final partitioning  $\mathcal{C}'$  is an acceptable clustering. We denote this operation by  $\text{PARTITION}(C, X, Y)$ .

**Claim 4.2** *Let  $\mathcal{C}'$  be the outcome of operation  $\text{PARTITION}(C, X, Y)$ . Then  $\varphi(\mathcal{C}') < \varphi(\mathcal{C}) - 1/n$ .*

**Proof:** Let  $\mathcal{C}''$  be the clustering obtained from  $\mathcal{C}$ , by replacing  $C$  with  $X$  and  $Y$ . From Theorem 4.1, it is enough to prove that  $\varphi(\mathcal{C}'') < \varphi(\mathcal{C}) - 1/n$ .

Assume without loss of generality that  $|\text{out}(X)| \leq |\text{out}(Y)|$ . Let  $h = |\text{out}(C)|$ ,  $h_1 = |\text{out}(X)|$ ,  $h_2 = |\text{out}(Y)|$ , so  $h_1 < 2h/3$ . Assume that  $h \in L_i$ . Then either  $h_1 \in L_{i'}$  for  $i' \leq i - 1$ , or  $h_1 < k_1$ . The

potential of the edges in  $\text{out}(Y) \cap \text{out}(C)$  does not increase. The only other changes in the potential are the following: the potential of each edge in  $\text{out}(X) \cap \text{out}(C)$  decreases by  $\varphi(h) - \varphi(h_1)$ , and the potential of every edge in  $E(X, Y)$  increases from 0 to at most 1.1. We consider two cases.

First, if  $h_1 < k_1$ , then  $\varphi(h) \geq 4\alpha + \varphi(h_1)$ . So the potential of each edge in  $\text{out}(X) \cap \text{out}(C)$  decreases by at least  $4\alpha$ , and the overall decrease in potential due to these edges is at least  $4\alpha |\text{out}(X) \cap \text{out}(C)|$ . The total increase in potential due to the edges in  $E(X, Y)$  is bounded by  $1.1|E(X, Y)| \leq 1.1\alpha |\text{out}(X) \cap \text{out}(C)|$ , so the overall potential decreases by at least  $2\alpha |\text{out}(X) \cap \text{out}(C)| \geq 1/n$ .

The second case is when  $h_1 \geq k_1$ . Assume that  $h_1 \in L_{i'}$ . Then  $n_{i'} \leq 3h_1/2$ , and, since  $i' \leq i - 1$  must hold,  $\varphi(h) \geq \frac{4\alpha k_1}{n_{i'}} + \varphi(n_{i'-1}) = \frac{4\alpha k_1}{n_{i'}} + \varphi(h_1) \geq \frac{8\alpha k_1}{3h_1} + \varphi(h_1)$ . So the potential of each edge in  $\text{out}(X) \cap \text{out}(C)$  decreases by at least  $\frac{8\alpha k_1}{3h_1}$ , and the total decrease in potential due to these edges is at least  $\frac{8\alpha k_1}{3h_1} \cdot |\text{out}(X) \cap \text{out}(C)| \geq \frac{4\alpha k_1}{3}$ , since  $|\text{out}(X) \cap \text{out}(C)| \geq h_1/2$ . The total increase in the potential due to the edges in  $E(X, Y)$  is bounded by  $1.1|E(X, Y)| < 0.55\alpha k_1$ , since  $|E(X, Y)| \leq \alpha k_1/2$ . Overall, the total potential decreases by at least  $\frac{2\alpha k_1}{3} \geq 1/n$ .  $\square$

**Action 2: Separating a large cluster.** Let  $\mathcal{C}$  be the current acceptable clustering, and let  $C \in \mathcal{C}$  be a large cluster in  $\mathcal{C}$ . Assume further that we are given a partition  $(A, B)$  of  $V(G)$ , with  $C \subseteq A$ ,  $\mathcal{T} \subseteq B$ , and  $|E_G(A, B)| < k_1/2$ . We perform the following operation, that we denote by  $\text{SEPARATE}(C, A)$ .

Consider a cluster  $S \in \mathcal{C}$ . If  $S \setminus A \neq \emptyset$ , and  $|\text{out}(S \setminus A)| > |\text{out}(S)|$ , then we modify  $A$  by removing all vertices of  $S$  from it. Notice that in this case, the number of edges in  $E(S)$  that originally contributed to the cut  $(A, B)$ ,  $|E(S \cap A, S \cap B)| > |\text{out}(S) \cap E(A)|$  must hold, so  $|\text{out}(A)|$  only goes down as a result of this modification. We assume from now on that if  $S \setminus A \neq \emptyset$ , then  $|\text{out}(S \setminus A)| \leq |\text{out}(S)|$ . In particular, if  $S$  is a small cluster, and  $S \setminus A \neq \emptyset$ , then  $S \setminus A$  is also a small cluster.

We build a new partition  $\mathcal{C}'$  of  $V(G)$  as follows. First, we add every connected component of  $G[A]$  to  $\mathcal{C}$ . Notice that all these clusters are small, as  $|\text{out}(A)| < k_1/2$ . Next, for every cluster  $S \in \mathcal{C}$ , such that  $S \setminus A \neq \emptyset$ , we add every connected component of  $G[S \setminus A]$  to  $\mathcal{C}'$ . Notice that every terminal  $t \in \mathcal{T}$  is added as a separate cluster to  $\mathcal{C}'$ . So far we have defined a new partition  $\mathcal{C}'$  of  $V(G)$ . This partition may not be acceptable, since we are not guaranteed that every small cluster of  $\mathcal{C}'$  has the  $\alpha_{\text{BW}}$ -bandwidth property. In our final step, we perform a bandwidth decomposition of every small cluster of  $\mathcal{C}'$ , using Theorem 4.1, and obtain the final acceptable clustering  $\mathcal{C}''$  of vertices of  $G$ . Notice that if  $S \in \mathcal{C}''$  is a large cluster, then there must be some **large** cluster  $S'$  in the original partition  $\mathcal{C}$  with  $S \subseteq S'$ .

**Claim 4.3** *Let  $\mathcal{C}''$  be the outcome of operation  $\text{SEPARATE}(C, A)$ . Then  $\varphi(\mathcal{C}'') \leq \varphi(\mathcal{C}) - 1$ .*

**Proof:** In order to prove the claim, it is enough to prove that  $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - 1$ , since, from Theorem 4.1, bandwidth decompositions of small clusters do not increase the potential.

We now show that  $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - 2$ . We can bound the changes in the potential as follows:

- Every edge in  $\text{out}(A)$  contributes at most 1.1 to the potential of  $\mathcal{C}''$ , and there are at most  $\frac{k_1-1}{2}$  such edges.
- Every edge in  $\text{out}(C)$  contributed at least 1 to the potential of  $\mathcal{C}'$ , and there are at least  $k_1$  such edges, since  $C$  is a large cluster.

For every other edge  $e$ , the potential of  $e$  does not increase. Indeed, let  $e = (u, v)$ , where  $u \in S_1$ ,  $v \in S_2$ , with  $S_1, S_2 \in \mathcal{C}'$ , and  $S_1, S_2 \not\subseteq A$ . Then there are clusters  $S'_1, S'_2 \in \mathcal{C}$ , with  $S_1 \subseteq S'_1$  and  $S_2 \subseteq S'_2$ . Notice that  $S'_1 \neq S'_2$ , since  $S_1$  and  $S_2$  correspond to connected components of  $S'_1$  and  $S'_2$ ,

respectively, and so no edge can connect them. From our construction of  $\mathcal{C}'$ ,  $|\text{out}(S_1)| \leq |\text{out}(S'_1)|$  and  $|\text{out}(S_2)| \leq |\text{out}(S'_2)|$ , so the potential of  $e$  cannot increase. Therefore, the decrease in the potential is at least  $k_1 - \frac{1.1(k_1-1)}{2} \geq 1$ .  $\square$

To summarize, given any acceptable clustering  $\mathcal{C}$  of the vertices of  $G$ , let  $E'$  be the set of edges whose endpoints belong to distinct clusters of  $\mathcal{C}$ . Then  $|E'| \leq \varphi(\mathcal{C}) \leq 1.1|E'|$ . So the potential is a good estimate on the number of edges connecting the different clusters. We have also defined two actions on large clusters of  $\mathcal{C}$ : **PARTITION**( $C, X, Y$ ), that replaces a large cluster  $C$  with a pair  $X, Y$  of clusters, where  $(X, Y)$  is a  $(k_1, \alpha)$ -violating partition of  $C$ , and **SEPARATE**( $C, A$ ), where  $A$  is a cut of size less than  $k/2$ , separating a large cluster  $C$  from the terminals. Each such action returns a new acceptable clustering, whose potential goes down by at least  $1/n$ .

## 5 The Algorithm

In this section we prove Theorem 3.1, by providing an efficient randomized algorithm, that with high probability either computes a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of  $k/\text{poly log } k$  demand pairs and their routing with congestion at most 2 in  $G$ , or finds a good congestion-2 crossbar in  $G$ .

We maintain, throughout the algorithm, a good clustering  $\mathcal{C}$  of  $G$ . Initially,  $\mathcal{C}$  is a partition of  $V(G)$ , where every vertex of  $G$  belongs to a distinct cluster, that is,  $\mathcal{C} = \{\{v\} \mid v \in V(G)\}$ . Clearly, this is a good clustering. We then perform a number of phases. In every phase, we start with some good clustering  $\mathcal{C}$  and the corresponding legal contracted graph  $H_{\mathcal{C}}$ . The phase output is one of the following: either (1) a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of  $k/\text{poly log } k$  pairs, with the routing of the pairs in  $\mathcal{M}'$  with congestion at most 2 in  $G$ , or (2) a good congestion-2 crossbar  $(\mathcal{S}^*, \mathcal{M}^*, \tau^*)$  in  $G$ ; or (3) another good clustering  $\mathcal{C}'$ , such that  $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - 1$ . In the first two cases, we terminate the algorithm, and output either the routing of the pairs in  $\mathcal{M}'$ , or the good crossbar. In the latter case, we continue to the next phase. After  $O(|E|)$  phases, our algorithm will successfully terminate with the required output. It is therefore enough to prove the following theorem.

**Theorem 5.1** *Then there is an efficient randomized algorithm that, given a good clustering  $\mathcal{C}$  of the vertices of  $G$  and be the corresponding legal contracted graph  $H_{\mathcal{C}}$ , with high probability computes one of the following:*

- (routing-of-many-pairs): *a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of  $k/\text{poly log } k$  pairs and a routing of the pairs in  $\mathcal{M}'$  in graph  $G$  with congestion at most 2; or*
- (good-crossbar): *a good congestion-2 crossbar  $(\mathcal{S}^*, \mathcal{M}^*, \tau^*)$  in  $G$ ; or*
- (better-clustering): *a new good clustering  $\mathcal{C}'$ , such that  $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - 1$ .*

The rest of this section is dedicated to proving Theorem 5.1.

We assume that we are given a good clustering  $\mathcal{C}$  of the vertices of  $G$ , and the corresponding legal contracted graph  $G' = H_{\mathcal{C}}$ .

Let  $m$  be the number of edges in  $G' \setminus \mathcal{T}$ . From Claim 4.1,  $m \geq k/3$ . As a first step, we randomly partition the vertices in  $G' \setminus \mathcal{T}$  into  $\gamma$  subsets  $X_1, \dots, X_\gamma$ , where each vertex  $v \in V(G') \setminus \mathcal{T}$  selects an index  $1 \leq j \leq \gamma$  independently uniformly at random, and is then added to  $X_j$ . We need the following claim, that appeared in [Chu12]. The proof appears in Section C of the Appendix for completeness.

**Claim 5.1** *With probability at least  $\frac{1}{2}$ , for each  $1 \leq j \leq \gamma$ ,  $|\text{out}_{G'}(X_j)| < \frac{10m}{\gamma}$ , while  $|E_{G'}(X_j)| \geq \frac{m}{2\gamma^2}$ .*

Given a partition  $X_1, \dots, X_\gamma$ , we can efficiently check whether the conditions of Claim 5.1 hold. If they do not hold, we repeat the randomized partitioning procedure. From Claim 5.1, we are guaranteed that

with high probability, after  $\text{poly}(n)$  iterations, we will obtain a partition with the desired properties. Assume now that we are given the partition  $X_1, \dots, X_\gamma$  of  $V(G') \setminus \mathcal{T}$ , for which the conditions of Claim 5.1 hold. Then for each  $1 \leq j \leq \gamma$ ,  $|E_{G'}(X_j)| > \frac{|\text{out}_{G'}(X_j)|}{20\gamma}$ . Let  $X'_j \subseteq V(G) \setminus \mathcal{T}$  be the set of vertices obtained from  $X_j$ , after we un-contract each cluster, that is, for each super-node  $v_C \in X_j$ , we replace  $v_C$  with the vertices of  $C$ . Notice that  $\{X'_j\}_{j=1}^\gamma$  is a partition of  $V(G) \setminus \mathcal{T}$ .

The plan for the rest of the proof is as follows. For each  $1 \leq j \leq \gamma$ , we will maintain an acceptable clustering  $\mathcal{C}_j$  of the vertices of  $G$ . That is, for each  $1 \leq j \leq \gamma$ ,  $\mathcal{C}_j$  is a partition of  $V(G)$ . In addition to being an acceptable clustering, it will have the following property:

P1. If  $C \in \mathcal{C}_j$  is a large cluster, then  $C \subseteq X'_j$ .

The initial partition  $\mathcal{C}_j$ , for  $1 \leq j \leq \gamma$  is obtained as follows. Recall that  $\mathcal{C}$  is the current good clustering of the vertices of  $G$ , and every cluster  $C \in \mathcal{C}$  is either contained in  $X'_j$ , or it is disjoint from it. First, we add to  $\mathcal{C}_j$  all clusters  $C \in \mathcal{C}$  with  $C \cap X'_j = \emptyset$ . Next, we add to  $\mathcal{C}_j$  all connected components of  $G[X'_j]$ . If any of these components is a small cluster, then we perform a bandwidth decomposition of this cluster, using Theorem 4.1, and update  $\mathcal{C}_j$  accordingly. Let  $\mathcal{C}_j$  be the resulting final partition. Clearly, it is an acceptable clustering, with property (P1). Moreover, we show that  $\varphi(\mathcal{C}_j) \leq \varphi(\mathcal{C}) - 1$ :

**Claim 5.2** For each  $1 \leq j \leq \gamma$ ,  $\varphi(\mathcal{C}_j) \leq \varphi(\mathcal{C}) - 1$ .

**Proof:** Let  $\mathcal{C}'_j$  be the partition of  $V(G)$ , obtained as follows: we add to  $\mathcal{C}'_j$  all clusters  $C \in \mathcal{C}$  with  $C \cap X_j = \emptyset$ , and we add all connected components of  $G[X_j]$  to  $\mathcal{C}'_j$  (that is,  $\mathcal{C}'_j$  is obtained like  $\mathcal{C}_j$ , except that we do not perform bandwidth decompositions of the small clusters). From Theorem 4.1, it is enough to prove that  $\varphi(\mathcal{C}'_j) \leq \varphi(\mathcal{C}) - 1$ . The changes of the potential from  $\mathcal{C}$  to  $\mathcal{C}'_j$  can be bounded as follows:

- The edges in  $E_{G'}(X_j)$  contribute at least 1 to  $\varphi(\mathcal{C})$  and contribute 0 to  $\varphi(\mathcal{C}'_j)$ .
- The potential of edges in  $\text{out}_G(X'_j)$  may increase. The increase is at most  $\varphi(n) \leq \frac{1}{2^8\gamma}$  per edge. So the total increase is at most  $\frac{|\text{out}_{G'}(X_j)|}{2^8\gamma} \leq \frac{|E_{G'}(X_j)|}{4}$ . These are the only edges whose potential may increase.

Overall, the decrease in the potential is at least  $\frac{|E_{G'}(X_j)|}{2} \geq \frac{m}{4\gamma^2} \geq \frac{k}{12\gamma^2} \geq 1$ .  $\square$

If any of the partitions  $\mathcal{C}_1, \dots, \mathcal{C}_\gamma$  is a good clustering, then we have found a good clustering  $\mathcal{C}'$  with  $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - 1$ . We terminate the algorithm and return  $\mathcal{C}'$ . Otherwise, we select an arbitrary large cluster  $S_j \in \mathcal{C}_j$ . We then consider the resulting collection  $S_1, \dots, S_\gamma$  of large clusters, and try to exploit them to construct a good crossbar. Since for each  $1 \leq j \leq \gamma$ ,  $S_j \subseteq X'_j$ , the sets  $S_1, \dots, S_\gamma$  are mutually disjoint and they do not contain terminals. Our algorithm performs a number of iterations, using the following theorem.

**Theorem 5.2** There is an efficient randomized algorithm, that, given a collection  $\{S_1, \dots, S_\gamma\}$  of disjoint large clusters containing non-terminal vertices, with high probability computes one of the following:

- (routing-of-many-pairs): a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of  $k/\text{poly log } k$  pairs and a routing of the pairs in  $\mathcal{M}'$  in graph  $G$  with congestion at most 2; or



- (good-crossbar): a good congestion-2 crossbar  $(\mathcal{S}^*, \mathcal{M}^*, \tau^*)$  in  $G$ ; or
- (sparse-partition): a  $(k_1, \alpha)$ -violating partition  $(X, Y)$  of  $S_j$ , for some  $1 \leq j \leq \gamma$ ; or
- (small-separation): a cut  $(A, B)$  in  $G$  with  $S_j \subseteq A$ ,  $\mathcal{T} \subseteq B$  and  $|E_G(A, B)| < k_1/2$ , for some  $1 \leq j \leq \gamma$ .

We provide the proof of Theorem 5.2 in the following section, and complete the proof of Theorem 5.1 here. Suppose we are given a good clustering  $\mathcal{C}$  of the vertices of  $G$ . For each  $1 \leq j \leq \gamma$ , we compute an acceptable clustering  $\mathcal{C}_j$  of  $V(G)$  as described above. If any of the partitions  $\mathcal{C}_j$  is a good clustering, then we terminate the algorithm and return  $\mathcal{C}_j$ . From the above discussion,  $\varphi(\mathcal{C}_j) \leq \varphi(\mathcal{C}) - 1$ . Otherwise, for each  $1 \leq j \leq \gamma$ , we select any large cluster  $S_j \in \mathcal{C}_j$ , and apply Theorem 5.2 to the current family  $\{S_1, \dots, S_\gamma\}$  of large clusters. If the outcome of Theorem 5.2 is either (routing-of-many-pairs) or (good-crossbar), then we terminate the algorithm and return the resulting routing or crossbar. We say that the iteration is successful if one of these two cases happen. Otherwise, we apply the appropriate action: PARTITION( $S_j, X, Y$ ), or SEPARATE( $S_j, A$ ) to the clustering  $\mathcal{C}_j$ . As a result, we obtain an acceptable clustering  $\mathcal{C}'_j$ , with  $\varphi(\mathcal{C}'_j) \leq \varphi(\mathcal{C}_j) - 1/n$ . Moreover, it is easy to see that this clustering also has Property (P1): if the PARTITION operation is performed, then we only partition a large cluster; if the SEPARATE operation is performed, then the only large clusters in the new partition  $\mathcal{C}'_j$  are subsets of large clusters in  $\mathcal{C}_j$ .

If all clusters in  $\mathcal{C}'_j$  are small, then we can again terminate the algorithm with (better-clustering): a good clustering  $\mathcal{C}'_j$ , with  $\varphi(\mathcal{C}'_j) \leq \varphi(\mathcal{C}) - 1$ . Otherwise, we select any large cluster  $S'_j \in \mathcal{C}'_j$ , and continue to the next iteration. Overall, as long as we do not complete a successful iteration, and we do not find a good clustering  $\mathcal{C}'$  of  $V(G)$  with  $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - 1$ , we make progress in each iteration by decreasing the potential of one of the partitions  $\mathcal{C}_j$  by at least  $1/n$ , by performing either a SEPARATE or a PARTITION operation on one of the large clusters of  $\mathcal{C}_j$ . After polynomially many iterations we are then guaranteed to complete a successful iteration, or find a good clustering  $\mathcal{C}'$  with  $\varphi(\mathcal{C}') \leq \varphi(\mathcal{C}) - 1$ , and finish the algorithm. Therefore, in order to complete the proof of Theorem 5.1 it is now enough to prove Theorem 5.2.

## 6 Proof of Theorem 5.2

Let  $\mathcal{R} = \{S_1, \dots, S_\gamma\}$ .

Throughout the algorithm, we will sometimes be interested in routing flow inside the clusters  $S_j \in \mathcal{R}$ . Specifically, given two subsets  $\Gamma, \Gamma' \subseteq \text{out}(S_j)$  of edges, with  $|\Gamma| = |\Gamma'| \leq k_1/2$ , we will be interested in routing the edges of  $\Gamma$  to the edges of  $\Gamma'$  inside  $S_j$ , with congestion at most  $1/\alpha$ . In other words, we will be looking for a flow  $F : \Gamma \xrightarrow{1:1}_{1/\alpha} \Gamma'$ . Notice that if such a flow does not exist, then we can find a  $(k_1, \alpha)$ -violating partition  $(X, Y)$  of  $S_j$ , by using the min-cut max-flow theorem. We can then return this partition and terminate the algorithm with (sparse-partition). Therefore, in order to simplify the exposition of the algorithm, we will assume that whenever the algorithm attempts to find such a flow, it always succeeds. From the integrality of flow, we can also find a set  $\mathcal{P} : \Gamma \xrightarrow{1:1}_{[1/\alpha]} \Gamma'$  of paths contained in  $S_j$ .

We start by verifying that for each  $1 \leq j \leq \gamma$ , the vertices of  $S_j$  can send  $k_1/2$  flow units with no edge-congestion to the terminals in graph  $G$ . If this is not the case for some set  $S_j$ , then there is a cut  $(A, B)$  with  $S_j \subseteq A$ ,  $\mathcal{T} \subseteq B$  and  $|E_G(A, B)| < k_1/2$ . We then return the partition  $(A, B)$  of  $G$  and finish the algorithm with (small-separation). From now on we assume that each set  $S_j$  can send  $k_1/2$  flow units with no edge-congestion to the terminals.

Let  $G'$  be the graph obtained from  $G$  by replacing every edge of  $G$  by two bi-directed edges. The rest of the proof consists of three steps. In the first step, we construct a degree-3 tree  $\hat{T}$ , whose

vertex set  $V(\tilde{T}) = \{v_S \mid S \in \mathcal{R}'\}$ , for a large enough family  $\mathcal{R}' \subseteq \mathcal{R}$  of vertex subsets, and each edge  $e = (v_S, v_{S'})$  in tree  $\tilde{T}$  corresponds to a collection  $\mathcal{P}_e$  of paths in graph  $G'$ , connecting the vertices of  $S$  to the vertices of  $S'$  or vice versa. Moreover, we will ensure that the paths in  $\bigcup_{e \in E(\tilde{T})} \mathcal{P}_e$  only cause congestion 1 in  $G'$ . In the second step, we find a subset  $\mathcal{M}^* \subseteq \mathcal{M}$  of the demand pairs, and route the terminals in  $\mathcal{T}(\mathcal{M}^*)$  to the vertices of  $S \cup S'$ , where  $(S, S')$  is some pair of vertex subsets in  $\mathcal{R}'$ . In the final third step, we construct a good congestion-2 crossbar.

## 6.1 Step 1: Constructing the Tree

This step is summarized in the following theorem.

**Theorem 6.1** *There is an efficient algorithm, that either computes a  $(k_1, \alpha)$ -violating partition of some set  $S \in \mathcal{R}$ , or finds a subset  $\mathcal{R}' \subseteq \mathcal{R}$  of size  $r = 8\gamma_{\text{CMG}}$ , a tree  $\tilde{T}$  of maximum degree at most 3 and vertex set  $V(\tilde{T}) = \{v_S \mid S \in \mathcal{R}'\}$ , and a collection  $\mathcal{P}_e$  of  $k_2 = \Omega\left(\frac{k_1 \alpha \cdot \alpha_{\text{BW}}}{\gamma^{3.5}}\right)$  paths in graph  $G'$  for every edge  $e \in E(\tilde{T})$ , such that:*

- For each edge  $e = (v_S, v_{S'}) \in E(\tilde{T})$ , every path  $P \in \mathcal{P}_e$  connects a vertex of  $S$  to a vertex of  $S'$ , or a vertex of  $S'$  to a vertex of  $S$ , and it does not contain the vertices of  $\bigcup_{S'' \in \mathcal{R}'} S''$  as inner vertices.
- The set  $\bigcup_{e \in E(\tilde{T})} \mathcal{P}_e$  of paths causes congestion 1 in  $G'$ .

**Proof:** Since the set  $\mathcal{T}$  of terminals is 1-well-linked, every pair  $(S_j, S_{j'})$  of vertex subsets can send  $k_1/2$  flow units to each other with congestion at most 3 (concatenate the flows from  $S_j$  to a subset  $\mathcal{T}_1$  of the terminals, from  $S_{j'}$  to a subset  $\mathcal{T}_2$  of the terminals, and the flow between the two subsets of the terminals). Equivalently, for each pair  $(S_j, S_{j'})$ , there are at least  $\lfloor k_1/6 \rfloor$  edge-disjoint paths connecting  $S_j$  to  $S_{j'}$  in  $G$ . We can assume that these paths do not contain any terminals, as the degree of every terminal in  $G$  is 1. We say that a path  $P$  is *direct* iff it does not contain the vertices of  $\mathcal{T} \cup \left(\bigcup_{j=1}^{\gamma} S_j\right)$  as its inner vertices.

We build the following graph  $Z$ . Let  $k' = \lfloor \frac{k_1}{6\gamma^2} \rfloor = \frac{k}{\text{poly log } k}$ . The vertices of  $Z$  are  $\{v_1, \dots, v_\gamma\}$ , where vertex  $v_j$  represents the set  $S_j$ . There is an edge  $(v_j, v_{j'})$  for  $j \neq j'$  in  $Z$  iff there are at least  $k'$  direct edge-disjoint paths connecting  $S_j$  to  $S_{j'}$  in  $G$ . (Notice that this can be checked efficiently).

It is easy to verify that graph  $Z$  is connected: indeed, assume otherwise. Let  $A$  be any connected component of  $Z$ , and let  $B$  contain the rest of the vertices. Let  $v_j \in A$ ,  $v_{j'} \in B$ . Since there are at least  $\lfloor k_1/6 \rfloor$  edge-disjoint paths connecting  $S_j$  and  $S_{j'}$  in  $G$ , and these paths do not contain the terminals, there must be at least  $\lfloor k_1/6 \rfloor$  direct edge-disjoint paths connecting the vertices of  $\bigcup_{v_i \in A} S_i$  to the vertices of  $\bigcup_{v_i \in B} S_i$ . Since  $|A| + |B| = \gamma$ , at least one pair  $(S_i, S_{i'})$  with  $v_i \in A$ ,  $v_{i'} \in B$ , has at least  $k'$  direct edge-disjoint paths connecting them.

Let  $T$  be any spanning tree of  $Z$ , and let  $\gamma' = \gamma^{1/4}$ . For every pair  $v_i, v_j$  of vertices in graph  $Z$ , let  $n(v_i, v_j)$  be the maximum number of direct edge-disjoint paths connecting  $S_i$  to  $S_j$  in  $G$ . For every edge  $e = (v_i, v_j)$  of the tree  $T$ , let  $w(e) = n(v_i, v_j)$ , and let  $w(T) = \sum_{e \in E(T)} w(e)$ . Root the tree  $T$  at any vertex. We now perform a number of iterations, that, informally, aim to maximize the number of leaves and the value  $w(T)$  of the tree  $T$ . In every iteration, one of the following two improvement steps is performed, whenever possible.

**Improvement Step 1:** Let  $v_j$  be any vertex of the tree  $T$ , whose parent  $v_{j'}$  is a non-root degree-2 vertex. If there is any non-leaf vertex  $v_{j''}$ , such that  $v_{j''}$  is not the descendant of  $v_j$  in the tree  $T$ , and

the edge  $(v_j, v_{j''})$  belongs to the graph  $Z$ , then we delete the edge  $(v_j, v_{j'})$  from the tree  $T$ , and add the edge  $(v_j, v_{j''})$  to it. Notice that this operation increases the number of leaves in  $T$  by 1.

**Improvement Step 2:** Let  $v_j$  be any non-root degree-2 vertex of the tree, whose parent  $v_{j'}$  is also a degree-2 non-root vertex. Let  $e = (v_j, v_{j'})$  be the edge of tree  $T$  connecting  $v_j$  to its parent. Consider the two connected components of  $T \setminus e$ , and let  $e' \in E(Z)$  be the edge connecting these two components, with the maximum value  $w(e')$ . If  $w(e') > w(e)$ , then we delete  $e$  from  $T$  and add  $e'$  to it. Notice that this operation does not decrease the number of leaves in  $T$ , and it increases  $w(T)$  by at least 1.

We perform one of the above two improvement steps, while possible. Since each such improvement step either increases the number of leaves in  $T$ , or leaves the number of leaves unchanged but increases  $w(T)$  by at least 1, after polynomially many improvement steps, we obtain a final tree  $T$ , on which none of these improvement steps can be executed. Let  $L$  denote the set of leaves of  $T$ .

We say that Case 1 happens, if  $T$  contains at least  $\gamma'$  leaves. Otherwise, we say that Case 2 happens. We consider each of the two cases separately below. Since Case 2 is easier to analyze, we consider it first.

**Case 2** A path  $P$  in the tree  $T$  is called a 2-path iff it does not contain the root of  $T$ , and all its vertices have degree 2 in  $T$ . A 2-path  $P$  is maximal iff it is not a strict sub-path of any other 2-path. Since the number of leaves in  $T$  is less than  $\gamma' = \gamma^{1/4}$ , the number of maximal 2-paths in  $T$  is at most  $2\gamma'$ . Therefore,  $T$  contains at least one 2-path of length at least  $\gamma'$ . We let  $P'$  be such a 2-path, that minimizes the size of the sub-tree rooted at the last (lowest) vertex of  $P'$ . Recall that  $r = 8\gamma_{\text{CMG}} \leq \lceil \gamma'/8 \rceil$ . Let  $P$  be the sub-path of  $P'$ , consisting of the last  $r$  vertices of  $P'$ . Assume without loss of generality that  $P = (v_1, \dots, v_r)$ , where  $v_r$  is the vertex that lies deepest in the tree  $T$ . We set  $\mathcal{R}' = \{S_1, \dots, S_r\}$ , and we will build a tree  $\tilde{T}$  over the set  $v_1, \dots, v_r$  of vertices as required. In fact in this case, the tree  $\tilde{T}$  will be a path,  $(v_{S_1}, \dots, v_{S_r})$ . Therefore, we only need to define, for each edge  $e = (v_{S_i}, v_{S_{i+1}})$  on this path, the corresponding set  $\mathcal{P}_e$  of paths in graph  $G'$ .

Since the number of leaves in the original tree  $T$  is at most  $\gamma'$ , and every 2-path in the sub-tree of  $v_r$  has length at most  $\gamma'$  (by the choice of  $P'$ ), the size of the sub-tree rooted at  $v_r$  is bounded by  $3\gamma'^2$ .

We need the following claim:

**Claim 6.1** *For each  $1 \leq i < r$ ,  $n(v_i, v_{i+1}) \geq \frac{k_1}{24\gamma^{3/2}}$ , and for each  $1 \leq i, i' \leq r$  with  $|i - i'| > 1$ ,  $n(v_i, v_{i'}) < \frac{k_1}{6\gamma^2}$ .*

**Proof:** Fix some  $1 \leq i < r$ , and consider the edge  $e = (v_i, v_{i+1})$ . Notice that the size of the sub-tree rooted at  $v_{i+1}$  is bounded by  $r + 3\gamma'^2 \leq 4\gamma'^2$ , since the size of the sub-tree rooted at  $v_r$  is bounded by  $3\gamma'^2$ . Let  $T_1, T_2$  be the two connected components of  $T \setminus \{e\}$ . Recall that there are at least  $\lfloor k_1/6 \rfloor$  direct paths connecting the sets corresponding to vertices of  $T_1$  to the sets corresponding to vertices of  $T_2$ . So there is at least one pair  $v \in T_1, v' \in T_2$  of vertices, with  $n(v, v') \geq \frac{k_1}{24\gamma \cdot \gamma'^2} = \frac{k_1}{24\gamma^{3/2}}$ . Since we could not perform the second improvement step, the weight of the edge  $(v_i, v_{i+1})$  must be at least  $\frac{k_1}{24\gamma^{3/2}}$ .

In order to prove the second assertion, consider some pair  $v_i, v_{i'}$  of vertices on path  $P$ , where  $|i - i'| > 1$ , and assume without loss of generality that  $v_i$  is a descendant of  $v_{i'}$ . Since we could not execute the first improvement step for vertex  $v_i$ , there is no edge  $(v_i, v_{i'})$  in graph  $Z$ . Therefore,  $n(v_i, v_{i'}) < \frac{k_1}{6\gamma^2}$ .  $\square$

For each  $1 \leq i < r$ , let  $\tilde{\mathcal{P}}_i$  be the set of  $\lceil \frac{k_1}{24\gamma^{3/2}} \rceil$  direct edge-disjoint paths, connecting  $S_i$  to  $S_{i+1}$ . While the paths in each set  $\tilde{\mathcal{P}}_i$  are edge-disjoint, the total congestion caused by the set  $\bigcup_i \tilde{\mathcal{P}}_i$  of paths may

be as high as  $(r - 1)$ . In order to overcome this difficulty, we perform edge-splitting on an auxiliary graph  $H$ , constructed as follows. Start with the set of all the vertices and edges lying on the paths in  $\bigcup_{i=1}^{r-1} \tilde{\mathcal{P}}_i$ . Next, for each  $1 \leq i \leq r$ , we contract the vertices of  $S_i$ , that belong to the current graph, into a single vertex  $v_i$ . Finally, we replace every edge by a pair of bi-directed edges. Let  $H$  be the final directed graph. Observe that  $H$  is a directed Eulerian graph, that has the following properties:

- Graph  $H$  does not contain any terminals, or vertices from the sets  $S \in \mathcal{R} \setminus \mathcal{R}'$ .
- For each  $1 \leq i < r$ , there are at least  $\lceil \frac{k_1}{24\gamma^{3/2}} \rceil$  edge-disjoint paths connecting  $v_i$  to  $v_{i+1}$ , that do not contain any vertices in  $\{v_1, v_2, \dots, v_r\}$  as inner vertices.
- For all  $1 \leq i \neq i' \leq r$ , with  $|i - i'| > 1$ , there are at most  $\frac{k_1}{6\gamma^2}$  paths connecting  $v_i$  to  $v_{i'}$ , that do not contain the vertices of  $\{v_1, v_2, \dots, v_r\}$  as intermediate vertices.

We use the following theorem to perform edge-splitting in graph  $H$ . Let  $D = (V, A)$  be any directed multigraph with no self-loops. For any pair  $(v, v') \in V$  of vertices, their connectivity  $\lambda(v, v'; D)$  is the maximum number of edge-disjoint paths connecting  $v$  to  $v'$  in  $D$ . Given a pair  $a = (u, v)$ ,  $b = (v, w)$  of edges, a splitting-off procedure replaces the two edges  $a, b$  by a single edge  $(u, w)$ . We denote by  $D^{a,b}$  the resulting graph. We use the extension of Mader's theorem [Mad78] to directed graphs, due to Frank [Fra89] and Jackson [Jac98]. Following is a simplified version of Theorem 3 from [Jac98]:

**Theorem 6.2** *Let  $D = (V, A)$  be an Eulerian digraph,  $v \in V$  and  $a = (v, u) \in A$ . Then there is an edge  $b = (w, v) \in A$ , such that for all  $y, y' \in V \setminus \{v\}$ :  $\lambda(y, y'; D) = \lambda(y, y'; D^{ab})$ .*

We iteratively perform edge-splitting in graph  $H$  using Theorem 6.2, by repeatedly choosing vertices  $v \notin \{v_1, \dots, v_r\}$ , until all such vertices become isolated. Let  $H'$  denote the resulting graph, after we discard all isolated vertices and make all edges undirected. Then  $V(H') = \{v_1, \dots, v_r\}$ , and for each  $1 \leq i \neq i' \leq r$ , each edge  $e = (v_i, v_{i'})$  corresponds to a **direct** path  $P_e$  in graph  $G'$ , connecting some vertex of  $S_i$  to some vertex of  $S_{i'}$  (that is,  $P_e$  does not contain any vertices of  $(\bigcup_{S \in \mathcal{R}} S) \cup \mathcal{T}$  as inner vertices). Moreover, the set  $\{P_e \mid e \in E(H')\}$  of paths causes congestion at most 1 in  $G'$ , and the edges  $e' \in E(G)$  whose both endpoints belong to the same set  $S_i$  do not lie on any such path  $P_e$ .

We need the following claim.

**Claim 6.2** *For each  $1 \leq i, i' \leq r$  with  $|i - i'| > 1$ , there are at most  $\lceil \frac{k_1}{3\gamma^2} \rceil$  parallel edges  $(v_i, v_{i'})$  in  $H'$ , and for each  $1 \leq i < r$ , there are at least  $\frac{k_1}{48\gamma^{3/2}}$  parallel edges  $(v_i, v_{i+1})$  in  $H'$ .*

**Proof:** Assume for contradiction that the first assertion is false, and let  $v_i, v_{i'}$  be a pair of vertices in  $H'$  with  $1 \leq i, i' \leq r$  and  $|i - i'| > 1$ , such that there are more than  $\lceil \frac{k_1}{3\gamma^2} \rceil$  parallel edges  $(v_i, v_{i'})$  in  $H'$ . Let  $\mathcal{P}'$  be the set of paths in graph  $G$ , containing, for each such parallel edge  $e$ , the corresponding path  $P_e$ . Then  $\mathcal{P}'$  contains more than  $\lceil \frac{k_1}{3\gamma^2} \rceil$  direct paths, connecting the vertices of  $S_i$  to the vertices of  $S_{i'}$  in  $G$ . Moreover, the paths in  $\mathcal{P}'$  cause congestion at most 2 in  $G$ . Therefore, there is a flow  $F$  between  $S_i$  and  $S_{i'}$  of value at least  $\frac{1}{2} \lceil \frac{k_1}{3\gamma^2} \rceil + \frac{1}{2}$ , with no congestion in graph  $G$ , where all flow-paths are direct paths. From the integrality of flow, there is a set of at least  $\frac{k_1}{6\gamma^2}$  direct edge-disjoint paths connecting  $S_i$  to  $S_{i'}$  in  $G$ , contradicting Claim 6.1.

We now turn to prove the second assertion. Assume otherwise, and let  $1 \leq i < r$  be some index, such that there are fewer than  $\frac{k_1}{48\gamma^{3/2}}$  parallel edges  $(v_i, v_{i+1})$  in graph  $H'$ . Consider the following partition  $(A, B)$  of  $V(H')$ :  $A = \{v_1, \dots, v_i\}$ ,  $B = \{v_{i+1}, \dots, v_r\}$ . Recall that graph  $H$  contained at least  $\frac{k_1}{24\gamma^{3/2}}$

edge-disjoint paths connecting  $v_i$  to  $v_{i+1}$ . Since the edge splitting operations preserve the connectivity between  $v_i$  and  $v_{i'}$ , graph  $H'$  must contain at least  $\frac{k_1}{24\gamma^{3/2}}$  edge-disjoint paths connecting  $v_i$  to  $v_{i+1}$ . In particular,  $|E_{H'}(A, B)| \geq \frac{k_1}{24\gamma^{3/2}}$ . Let  $E' \subseteq E_{H'}(A, B)$  contain all edges between vertex sets  $A$  and  $B$ , except for the set of parallel edges  $(v_i, v_{i+1})$ . Since there are fewer than  $\frac{k_1}{48\gamma^{3/2}}$  such parallel edges,  $|E'| > \frac{k_1}{48\gamma^{3/2}}$ .

Notice however that there are at most  $\frac{r^2}{4} - 1$  pairs of vertices  $(v_j, v'_j) \in A \times B$ , where  $(v_j, v'_j) \neq (v_i, v_{i+1})$ . Each such pair contributes at most  $\lceil \frac{k_1}{3\gamma^2} \rceil$  edges to  $E'$ , from the first assertion. Therefore,

$$|E'| \leq \left( \frac{r^2}{4} - 1 \right) \cdot \lceil \frac{k_1}{3\gamma^2} \rceil \leq \frac{r^2 \cdot k_1}{12\gamma^2} < \frac{k_1}{48\gamma^{3/2}},$$

as  $r \leq \lceil \frac{\gamma'}{8} \rceil \leq \frac{\gamma'}{4} = \frac{\gamma^{1/4}}{4}$ , a contradiction.  $\square$

Our final tree  $\tilde{T}$  is simply the path connecting the vertices  $(v_{S_1}, v_{S_2}, \dots, v_{S_r})$  in this order. For each edge  $e = (v_{S_i}, v_{S_{i+1}})$  on this path, we let  $\mathcal{P}_e$  be the set of paths in graph  $G'$  corresponding to the set of parallel edges connecting  $v_i$  to  $v_{i+1}$  in graph  $H'$ . We discard paths from each set  $\mathcal{P}_e$  until  $|\mathcal{P}_e| = k_2$  holds. It is immediate to verify that  $\bigcup_{e \in E(\tilde{T})} \mathcal{P}_e$  cause congestion at most 1 in  $G'$ , and moreover, each such path does not contain the vertices of  $\bigcup_{S \in \mathcal{R}'} S$  as inner vertices.

**Case 1** Recall that in this case,  $T$  contains at least  $\gamma'$  leaves. Let  $\mathcal{R}' \subseteq \mathcal{R}$  be any set of  $r < \gamma'$  vertex subsets, corresponding to the leaves in tree  $T$ . For simplicity of notation, we will assume that  $\mathcal{R}' = \{S_1, \dots, S_r\}$ . Our first step is to select, for each  $S \in \mathcal{R}'$ , a subset  $\Gamma(S) \subseteq \text{out}(S)$  of  $k_3 = \Omega(k' \alpha \cdot \alpha_{\text{BW}}/r^3)$  edges, such that for each  $S, S' \in \mathcal{R}'$ , there is a set  $\mathcal{P}(S, S') : \Gamma(S) \xrightarrow{1:1} \Gamma(S')$  of paths in graph  $G$ . We do so in the following claim.

**Claim 6.3** *There is an efficient algorithm, that either computes a  $(k_1, \alpha)$ -violating partition for some set  $S \in \mathcal{R}$ , or computes, for each set  $S \in \mathcal{R}'$ , a subset  $\Gamma(S) \subseteq \text{out}(S)$  of  $k_3 = \Omega(k' \alpha \cdot \alpha_{\text{BW}}/r^3)$  edges, such that for each  $S, S' \in \mathcal{R}'$ , there is a set  $\mathcal{P}(S, S') : \Gamma(S) \xrightarrow{1:1} \Gamma(S')$  of paths in graph  $G$ , that do not contain the vertices of  $\bigcup_{S'' \in \mathcal{R}'} S''$  as inner vertices.*

**Proof:** Let  $v_{j^*}$  be the root of the tree  $T$ , and let  $S_{j^*} \in \mathcal{R}$  be the corresponding vertex subset. We claim that for each set  $S_j \in \mathcal{R}'$ , there is a flow  $F_j$  of value  $k'$ , connecting the vertices of  $S_{j^*}$  to the vertices of  $S_j$  in  $G$ , with congestion at most  $\frac{1}{\alpha} + \gamma \leq \frac{2}{\alpha}$ , such that the paths in  $\mathcal{P}_j$  do not contain any vertices of  $\bigcup_{S \in \mathcal{R}'} S$  as inner vertices. Indeed, consider the path  $(v_j^* = v_{i_1}, v_{i_2}, \dots, v_{i_h} = v_j)$  in the tree  $T$ , connecting  $v_{j^*}$  to  $v_j$ . For each edge  $e_z = (v_{i_z}, v_{i_{z+1}})$  on this path, there is a set  $\mathcal{Q}_z$  of  $k'$  direct edge-disjoint paths in  $G$ , connecting the vertices of  $S_{i_z}$  to the vertices of  $S_{i_{z+1}}$ , from the definition of the graph  $Z$ . Denote

$$\Gamma_z^2 = \{e' \in \text{out}(S_{i_z}) \mid e' \text{ is the first edge on some path in } \mathcal{Q}_z\},$$

and similarly

$$\Gamma_{z+1}^1 = \{e' \in \text{out}(S_{i_{z+1}}) \mid e' \text{ is the last edge on some path in } \mathcal{Q}_z\}.$$

For each  $1 < z < h$ , we now have two subsets  $\Gamma_z^1, \Gamma_z^2 \subseteq \text{out}(S_{i_z})$  of edges of cardinality  $k' < k_1/2$  each. From our assumption, there is a flow  $F_z' : \Gamma_z^1 \xrightarrow{1:1} \Gamma_z^2$  inside set  $S_z$ . Concatenating the flows  $(\mathcal{Q}_1, F_1', \mathcal{Q}_2, F_2', \dots, F_{h-1}', \mathcal{Q}_{h-1})$ , we obtain the desired flow  $F_j$  of value  $k'$ . The total congestion caused

by paths in  $\bigcup_{z=1}^{h-1} \mathcal{Q}_z$  is at most  $\gamma$  (since  $h \leq \gamma$ , and the paths in each set  $\mathcal{Q}_z$  are edge-disjoint), while each flow  $F'_z$  causes congestion at most  $1/\alpha$  inside the graph  $G[S_{i_z}]$ . Therefore, the total congestion due to flow  $F_j$  is bounded by  $\frac{1}{\alpha} + \gamma \leq \frac{2}{\alpha}$ .

Scaling all flows  $F_j$ , for  $S_j \in \mathcal{R}'$  down by factor  $2r/\alpha$ , we obtain a new flow  $F$ , where every set  $S_j \in \mathcal{R}'$  sends  $\frac{k'\alpha}{2r}$  flow units to  $S_{j^*}$ , and the total congestion due to  $F$  is at most 1. From the integrality of flow, there is a collection  $\{\mathcal{P}_j\}_{j=1}^r$  of path sets, where for each  $1 \leq j \leq r$ , set  $\mathcal{P}_j$  contains  $\lfloor \frac{k'\alpha}{2r} \rfloor$  paths connecting  $S_j$  to  $S_{j^*}$ , the paths in  $\bigcup_{j=1}^r \mathcal{P}_j$  are edge-disjoint, and they do not contain the vertices of  $\bigcup_{j=1}^r S_j$  as inner vertices. We also assume without loss of generality that the paths in  $\bigcup_{j=1}^r \mathcal{P}_j$  do not contain the vertices of  $S_{j^*}$  as inner vertices.

For each  $1 \leq j \leq r$ , let  $E_j \subseteq \text{out}(S_{j^*})$  be the set of edges  $e$ , such that  $e$  is the last edge on some path in  $\mathcal{P}_j$ , and let  $E' = \bigcup_{j=1}^r E_j$ . Notice that  $|E'| = r \cdot \lfloor \frac{k'\alpha}{2r} \rfloor < k_1$ . We set up an instance of the sparsest cut problem as follows. First, partition every edge  $e \in E'$  by a vertex  $t_e$ , and let  $\tilde{\mathcal{T}}$  be the set of the resulting new vertices. Consider the sub-graph  $G^*$  of the resulting graph induced by  $S_{j^*} \cup \tilde{\mathcal{T}}$ , with the set  $\tilde{\mathcal{T}}$  of terminals. We apply the algorithm  $\mathcal{A}_{\text{ARV}}$  to the resulting instance of the sparsest cut problem. If the algorithm returns a cut whose sparsity is less than  $\alpha$ , then we have found a  $(k_1, \alpha)$ -violating partition of  $S_{j^*}$ . We return this partition, and terminate the algorithm. Otherwise, we are guaranteed that  $S_{j^*}$  has the  $\alpha_{\text{BW}}$ -bandwidth property with respect to  $E'$ .

We then apply Theorem 2.2 to the same graph  $G^*$ , with the set  $\tilde{\mathcal{T}}$  of terminals, where the terminal subsets are defined to be  $\mathcal{T}_j = \{t_e \mid e \in E_j\}$  for  $1 \leq j \leq r$ . Notice that for all  $j$ ,  $|\mathcal{T}_j| = k/\text{poly log } k$ , while  $r = \text{poly log } k$  and  $\alpha_{\text{BW}} = 1/\text{poly log } k$ , so the conditions of Theorem 2.2 hold. From Theorem 2.2, we obtain, for each  $1 \leq j \leq r$ , a subset  $\mathcal{T}'_j \subseteq \mathcal{T}_j$  of  $\Omega\left(\frac{k' \cdot \alpha \cdot \alpha_{\text{BW}}}{r^3}\right) = k_3$  terminals, such that  $\bigcup_j \mathcal{T}'_j$  is 1-well-linked in  $G^*$ . Let  $E'_j \subseteq E_j$  be the set of edges corresponding to the vertices in  $\mathcal{T}'_j$ , let  $\mathcal{P}'_j \subseteq \mathcal{P}_j$  be the set of paths terminating at the edges of  $E'_j$ , and let  $\Gamma(S_j) \subseteq \text{out}(S_j)$  be the set of edges where the paths in  $\mathcal{P}'_j$  start, for each  $1 \leq j \leq r$ . Then for each  $1 \leq j \leq r$ ,  $|\Gamma(S_j)| = |\mathcal{T}'_j| = k_3$ . Moreover, we claim that for all  $1 \leq j, j' \leq r$ , there is a set  $\mathcal{P}(S_j, S_{j'}) : \Gamma(S_j) \overset{1:1}{\rightsquigarrow} \Gamma(S_{j'})$  of paths in graph  $G$ , that do not contain the vertices of  $\bigcup_{S \in \mathcal{R}'} S$  as inner vertices. Indeed, since  $\bigcup_j \mathcal{T}'_j$  is 1-well-linked in  $G^*$ , set  $S_{j^*}$  has the 1-bandwidth property with respect to  $\bigcup_i E'_i$ , and so there is a set  $\mathcal{P}_{j,j'} : E_j \overset{1:1}{\rightsquigarrow} E_{j'}$  of paths contained in  $S_{j^*}$ . The set  $\mathcal{P}(S_j, S_{j'})$  of paths is obtained by concatenating  $\mathcal{P}'_j, \mathcal{P}_{j,j'}$  and  $\mathcal{P}'_{j'}$ .  $\square$

Our next step is to perform edge splitting, similarly to Case 2. We build an auxiliary graph  $H$  from graph  $G$ , as follows. We delete from  $G$  all terminals, and for each  $1 \leq j \leq r$ , we delete all edges in  $\text{out}(S_j) \setminus \Gamma(S_j)$ . Next, we contract each set  $S_j$  into a super-node  $v_j$ , for all  $1 \leq j \leq r$ . Finally, we replace every edge in the resulting graph by a pair of bi-directed edges. Let  $H$  be this final graph. Notice that in-degree and the out-degree of every super-node  $v_j$ , for  $1 \leq j \leq r$ , in graph  $H$  is exactly  $k_3$ , and for every pair  $v_j, v_{j'}$  of such super-nodes, there are  $k_3$  edge-disjoint paths connecting  $v_j$  to  $v_{j'}$ , and  $k_3$  edge-disjoint paths connecting  $v_{j'}$  to  $v_j$  in  $H$ . We iteratively perform edge-splitting in graph  $H$  using Theorem 6.2, by repeatedly choosing vertices  $v \notin \{v_1, \dots, v_r\}$ , until all such vertices become isolated. Let  $H'$  denote the resulting graph, after we discard all isolated vertices. Then  $V(H') = \{v_1, \dots, v_r\}$ , and for each  $1 \leq i \neq i' \leq r$ , each edge  $e = (v_i, v_{i'})$  corresponds to a path  $P_e$  in graph  $G'$ , connecting some vertex of  $S_i$  to some vertex of  $S_{i'}$ , and  $P_e$  does not contain any vertices of  $\bigcup_{S \in \mathcal{R}'} S$  as inner vertices (recall that  $G'$  is the directed graph obtained from  $G$  by replacing each edge with a pair of bi-directed edges). Moreover, the set  $\{P_e \mid e \in E(H')\}$  of paths causes congestion at most 1 in  $G'$ . For each vertex  $v_j$  in graph  $H'$ , the in-degree and the out-degree of  $v_j$  are equal to  $k_3$ , and for each pair  $v_j, v_{j'}$  of vertices, there are exactly  $k_3$  edge-disjoint paths connecting  $v_j$  to  $v_{j'}$  in  $H'$ , and  $k_3$  edge-disjoint paths connecting  $v_{j'}$  to  $v_j$ .

We now construct a graph  $Z'$ , whose vertex set is  $\{v_1, \dots, v_r\}$ , and there is an edge  $(v_i, v_j)$  in graph  $Z'$  iff the total number of parallel edges  $(v_i, v_j)$  and  $(v_j, v_i)$  in graph  $H'$  is at least  $2k_3/r^3$ . The proof

of the next claim is identical to the proof of Claim 3.14 in [Chu12] and is omitted here.

**Claim 6.4** *There is an efficient algorithm to find a spanning tree  $\tilde{T}$  of maximum vertex degree at most 3 in graph  $Z'$ .*

We output the tree  $\tilde{T}$  as our final tree. Notice that each edge  $e = (v_i, v_j) \in E(\tilde{T})$  corresponds to a collection  $\mathcal{P}_e$  of paths, where each path connects either a vertex of  $S_i$  to a vertex of  $S_j$ , or vice versa in graph  $G'$  (these are the paths that correspond to the edges of  $H'$  in the original graph  $G'$ , obtained by the edge splitting procedure). Moreover, the paths in  $\bigcup_{e \in E(\tilde{T})} \mathcal{P}_e$  cause congestion at most 1 in  $G'$ , and these paths do not contain the vertices of  $\bigcup_{j=1}^r S_j$  as inner vertices. The number of paths in each set is:

$$|\mathcal{P}_e| \geq \frac{2k_3}{r^3} = \Omega\left(\frac{k'\alpha \cdot \alpha_{\text{BW}}}{r^6}\right) = \Omega\left(\frac{k_1\alpha \cdot \alpha_{\text{BW}}}{\gamma^2 r^6}\right) = \Omega\left(\frac{k_1\alpha \cdot \alpha_{\text{BW}}}{\gamma^{3.5}}\right) = k_2.$$

□

## 6.2 Step 2: connecting the terminals

We assume without loss of generality that  $\mathcal{R}' = \{S_1, \dots, S_r\}$ . In this step we connect a subset of terminals to two subsets  $S, S' \in \{S_1, \dots, S_r\}$ , using the following theorem. Intuitively, we find a large enough subset  $\mathcal{M}_1 \subseteq \mathcal{M}$  of demand pairs, and two sets of edge-disjoint paths: set  $\mathcal{P}'_1$  connecting the sources of the pairs in  $\mathcal{M}_1$  to  $S$ , and set  $\mathcal{P}''_1$  connecting their destinations to  $S'$  (up to switching the sources and the destinations). We would also like to ensure that the paths in  $\mathcal{P}'_1 \cup \mathcal{P}''_1$  are disjoint from the paths  $\mathcal{P}_e$  representing the edges  $e \in E(\tilde{T})$  of the tree. We will achieve this property, only with respect to somewhat smaller subsets  $\mathcal{P}'_e \subseteq \mathcal{P}_e$  of paths.

**Theorem 6.3** *There is an efficient algorithm, that either finds a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of  $k/\text{poly log } k$  demand pairs, and a routing of the pairs in  $\mathcal{M}'$  via edge-disjoint paths in  $G$ , or finds a subset  $\mathcal{M}_1 \subseteq \mathcal{M}$  of  $k_4 = \Omega(k_2/r^2)$  demand pairs, two subsets  $\mathcal{P}'_1, \mathcal{P}''_1$  of paths in  $G$ , and for each edge  $e \in E(\tilde{T})$ , a subset  $\mathcal{P}'_e \subseteq \mathcal{P}_e$  of  $\lfloor k_2/2 \rfloor$  paths, such that:*

- We are given a partition of  $\mathcal{T}(\mathcal{M}_1)$  into two subsets  $\mathcal{T}'_1, \mathcal{T}''_1$ , where for each pair  $(s, t) \in \mathcal{M}_1$ ,  $s \in \mathcal{T}'_1, t \in \mathcal{T}''_1$ , or vice versa.
- There are two sets  $S, S' \in \mathcal{R}'$  (with possibly  $S = S'$ ), such that  $\mathcal{P}'_1 : \mathcal{T}'_1 \rightsquigarrow \text{out}(S)$ ,  $\mathcal{P}''_1 : \mathcal{T}''_1 \rightsquigarrow \text{out}(S')$ .
- The paths in  $\mathcal{P}'_1 \cup \mathcal{P}''_1 \cup \left(\bigcup_{e \in E(\tilde{T})} \mathcal{P}'_e\right)$  do not contain any vertices of  $\bigcup_{j=1}^r S_j$  as inner vertices, and they cause congestion at most 2 in  $G$ .

**Proof:** We start with the following lemma, whose proof uses standard techniques.

**Lemma 6.1** *Let  $S \in \mathcal{R}'$  be any set. Then there is an efficient algorithm, that either finds a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of  $k/\text{poly log } k$  demand pairs and a routing of  $\mathcal{M}'$  via edge-disjoint paths in  $G$ , or finds a subset  $\mathcal{M}_0 \subseteq \mathcal{M}$  of  $\Omega(k_1)$  pairs, and a set  $\mathcal{P} : \mathcal{T}(\mathcal{M}_0) \rightsquigarrow_1 \text{out}(S)$  of paths in graph  $G$ , connecting every terminal in  $\mathcal{T}(\mathcal{M}_0)$  to some edge of  $\text{out}(S)$  via edge-disjoint paths.*

**Proof:** Recall that we have assumed that there is a flow  $F$  of value  $k_1/2$  and no congestion in graph  $G$  between  $\text{out}(S)$  and  $\mathcal{T}$ . From the integrality of flow, there is a set  $\mathcal{P}^*$  of  $\lfloor k_1/2 \rfloor$  edge-disjoint paths

connecting the terminals in  $\mathcal{T}$  to the edges in  $\text{out}(S)$ . Recall that the degree of every terminal in  $G$  is 1, so every terminal is an endpoint of at most one path in  $\mathcal{P}^*$ . Let  $\mathcal{T}_0 \subseteq \mathcal{T}$  be the subset of  $\lfloor k_1/2 \rfloor$  terminals that serve as endpoints of paths in  $\mathcal{P}^*$ . We partition the remaining terminals into  $h = \lceil \frac{k}{\lfloor k_1/2 \rfloor} \rceil = O(\gamma^3 \log \gamma)$  subsets of at most  $\lfloor k_1/2 \rfloor$  terminals each. Since the set  $\mathcal{T}$  of terminals is 1-well-linked, for each  $1 \leq j \leq h$ , there is a flow  $\mathcal{F}_j : \mathcal{T}_j \rightsquigarrow_1 \mathcal{T}_0$ . Concatenating flow  $\mathcal{F}_j$  with the paths in  $\mathcal{P}^*$ , and taking the union of all resulting flows over all  $1 \leq j \leq h$ , we obtain a flow  $F'$ , where every terminal in  $\mathcal{T}$  sends one flow unit to some edge in  $\text{out}(S)$ , and the congestion caused by  $F'$  is at most  $2h = O(\gamma^3 \log \gamma)$ .

Our next step is to perform a grouping of the terminals, using Theorem 2.1, with the parameter  $q = 20h$ . Let  $\mathcal{G}$  be the resulting partition of the terminals. Recall that each set  $U \in \mathcal{G}$  contains at least  $q$  and at most  $3q$  terminals, and it is associated with a tree  $T_U$ , containing the terminals of  $U$ , such that the trees  $\{T_U\}_{U \in \mathcal{G}}$  are edge-disjoint.

We partition  $\mathcal{M}$  into two subsets:  $\mathcal{M}'_1$  containing all pairs  $(s, t)$  where both  $s$  and  $t$  belong to the same group  $U \in \mathcal{G}$ , and  $\mathcal{M}'_2$  containing all remaining pairs.

Assume first that  $|\mathcal{M}'_1| \geq |\mathcal{M}|/2$ . We then select a subset  $\mathcal{M}' \subseteq \mathcal{M}'_1$  of demand pairs as follows: for each group  $U \in \mathcal{G}$ , we select one arbitrary pair  $(s, t) \in \mathcal{M}'_1$  with  $s, t \in U$ , if such a pair exists, and add it to  $\mathcal{M}'$ . Notice that since every group contains at most  $60h$  terminals, the number of pairs in  $\mathcal{M}'$  is at least  $\frac{k}{120h} = \Omega\left(\frac{k}{\gamma^3 \log \gamma}\right) = \Omega\left(\frac{k}{\text{poly log } k}\right)$ . Every pair  $(s, t) \in \mathcal{M}'_1$  can be connected by some path contained in the tree  $T_U$ , where  $s, t \in U$ . Since the trees  $\{T_U\}_{U \in \mathcal{G}}$  are edge-disjoint, the resulting routing of the pairs in  $\mathcal{M}'$  is via edge-disjoint paths. We assume from now on that  $|\mathcal{M}'_2| \geq |\mathcal{M}|/2$ . We need the following simple claim.

**Claim 6.5** *We can efficiently find a subset  $\mathcal{M}'' \subseteq \mathcal{M}'_2$  of at least  $k_1/960$  demand pairs, such that there is a flow  $F^* : \mathcal{T}(\mathcal{M}'') \rightsquigarrow_{1.1} \text{out}(S)$  in  $G$ .*

**Proof:** We start with  $\mathcal{M}'' = \emptyset$ . While  $\mathcal{M}'_2 \neq \emptyset$ , select any pair  $(s, t) \in \mathcal{M}'_2$  and add it to  $\mathcal{M}''$ . Assume that  $s \in U, t \in U'$ . Remove from  $\mathcal{M}'_2$  all pairs  $(s', t')$ , where either  $s'$  or  $t'$  belong to  $U \cup U'$ . Notice that for each pair added to  $\mathcal{M}''$ , at most  $6q = 120h$  pairs are deleted from  $\mathcal{M}'_2$ . Therefore, at the end of this procedure,  $|\mathcal{M}''| \geq \frac{k}{240h} \geq \frac{k}{240} \cdot \frac{k_1}{4k} = \frac{k_1}{960}$ .

We now show that there is a flow  $F^* : \mathcal{T}(\mathcal{M}'') \rightsquigarrow_{1.1} \text{out}(S)$  in  $G$ . Let  $F''$  be the flow  $F'$ , scaled down by factor  $20h$ . Then every terminal in  $\mathcal{T}$  sends  $1/(20h)$  flow units to the edges of  $\text{out}(S)$ , and the total congestion caused by  $F''$  is at most  $1/10$ . In order to define the flow  $F^*$ , consider some terminal  $t \in \mathcal{T}(\mathcal{M}'')$ , and let  $U_t \in \mathcal{G}$  be the group to which it belongs. Let  $U'_t \subseteq U_t$  be any subset of  $20h$  terminals in  $U_t$ . Terminal  $t$  then sends  $1/(20h)$  flow units to each terminal in  $U'_t$ , inside the tree  $T_{U_t}$ . This flow is then concatenated with the flow that leaves each terminal in  $U'_t$  in  $F''$ . Taking the union of this flow for each  $t \in \mathcal{T}(\mathcal{M}'')$ , we obtain a flow  $F^*$ , where every terminal in  $\mathcal{T}(\mathcal{M}'')$  sends one flow unit to the edges of  $\text{out}(S)$ . In order to bound the edge congestion, we observe that flow  $F''$  contributes at most  $1/10$  congestion to every edge. Since the connected components  $\{T_{U_t}\}_{t \in \mathcal{T}(\mathcal{M}'')}$  are edge-disjoint, the flow inside these components contributes at most 1 to the total edge congestion. Overall, the congestion caused by the flow  $F^*$  is at most 1.1.  $\square$

From the integrality of flow, there is a collection  $\mathcal{P}$  of  $\lfloor \frac{2|\mathcal{M}''|}{1.1} \rfloor \geq 1.8|\mathcal{M}''|$  of edge-disjoint paths connecting the terminals in  $\mathcal{T}(\mathcal{M}'')$  to the edges of  $\text{out}(S)$ . Let  $\mathcal{T}' \subseteq \mathcal{T}(\mathcal{M}'')$  be the set of terminals where these paths originate, and let  $\mathcal{M}_0 \subseteq \mathcal{M}''$  be the subset of the demand pairs contained in  $\mathcal{T}'$ . Since  $|\mathcal{T}'| \geq 1.8|\mathcal{M}''|$ , we get that  $|\mathcal{M}_0| \geq |\mathcal{M}''|/2$ . Thus, we have obtained a set  $\mathcal{M}_0$  of  $\Omega(k_1)$  demand pairs, such that all terminals in  $\mathcal{T}(\mathcal{M}_0)$  can be connected to the edges of  $\text{out}(S)$  via edge-disjoint paths.  $\square$



In order to complete the proof of Theorem 6.3, we start by applying Lemma 6.1 to the set  $S = S_1$ . If the outcome is a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of  $k/\text{poly log } k$  demand pairs and a routing of pairs in  $\mathcal{M}'$  via edge-disjoint paths in  $G$ , then we return  $\mathcal{M}'$  together with this routing, and terminate the algorithm.

We therefore assume from now on, that Lemma 6.1 returns a set  $\mathcal{M}_0 \subseteq \mathcal{M}$  of  $\Omega(k_1)$  pairs, together with the set  $\mathcal{P} : \mathcal{T}(\mathcal{M}_0) \rightsquigarrow_1 \text{out}(S_1)$  of edge-disjoint paths in graph  $G$ . We discard terminal pairs from  $\mathcal{M}_0$ , and the corresponding paths from  $\mathcal{P}$  until  $|\mathcal{M}_0| = \lfloor k_2/4 \rfloor$ , so  $|\mathcal{P}| = |\mathcal{T}(\mathcal{M}_0)| \leq k_2/2$  holds.

For each path  $P \in \mathcal{P}$ , we direct the path from the terminal towards  $S_1$ , and we truncate it at the first vertex that belongs to  $\bigcup_{i=1}^r S_i$ . The new collection  $\mathcal{P}$  of paths then connects every terminal in  $\mathcal{T}(\mathcal{M}_0)$  to some vertex of  $\bigcup_{i=1}^r S_i$ . Moreover, the paths in  $\mathcal{P}$  do not contain the vertices of  $\bigcup_{i=1}^r S_i$  as inner vertices.

We now combine the two sets of paths,  $\mathcal{P}$  and  $\bigcup_{e \in E(\tilde{T})} \mathcal{P}_e$ , to ensure that they cause congestion 1 in  $G'$  altogether, using the following simple claim.

**Claim 6.6** *We can efficiently find, for each edge  $e \in E(\tilde{T})$ , a subset  $\mathcal{P}'_e \subseteq \mathcal{P}_e$  of  $\lfloor k_2/2 \rfloor$  paths, and a collection  $\mathcal{P}' : \mathcal{T}(\mathcal{M}_0) \rightsquigarrow_1 \bigcup_{i=1}^r S_i$  of paths in  $G'$ , such that the paths in set  $\left(\bigcup_{e \in E(\tilde{T})} \mathcal{P}'_e\right) \cup \mathcal{P}'$  cause congestion 1 in  $G'$ .*

**Proof:** The proof is very similar to the arguments used by Conforti et al. [CHR03]. We re-route the paths in  $\mathcal{P}$  via some paths in  $\bigcup_{e \in E(\tilde{T})} \mathcal{P}_e$ , by setting up an instance of the stable matching problem in a multi-graph. In this problem, we are given a complete bipartite multigraph  $\tilde{G} = (A, B, E)$ , where  $|A| = |B|$ . Each vertex  $v \in A \cup B$  specifies an ordering  $R_v$  of the edges adjacent to  $v$  in  $\tilde{G}$ . A complete matching  $M$  between the vertices of  $A$  and  $B$  is called stable iff, for every edge  $e = (a, b) \in E \setminus M$ , the following holds. Let  $e_a, e_b$  be the edges adjacent to  $a$  and  $b$  respectively that belong to  $M$ . Then either  $a$  prefers  $e_a$  over  $e$ , or  $b$  prefers  $e_b$  over  $e$ . Conforti et al. [CHR03], generalizing the famous theorem of Gale and Shapley [GS62], show an efficient algorithm to find a complete stable matching  $M$  in any such multigraph.

Let  $\mathcal{Q} = \bigcup_{e \in E(\tilde{T})} \mathcal{P}_e$ . We set up an instance of the stable matching problem as follows. Set  $A$  contains a vertex  $v(P)$  for each path  $P \in \mathcal{P}$ , and set  $B$  contains a vertex  $v(Q)$  for each path  $Q \in \mathcal{Q}$ . In order to ensure that  $|A| = |B|$ , we add dummy vertices to  $A$  as needed.

For each pair  $P \in \mathcal{P}$ ,  $Q \in \mathcal{Q}$  of paths, for each edge  $e$  that these paths share, we add an edge  $(v(P), v(Q))$ , and we think of this edge as representing the edge  $e$ . We add additional dummy edges as needed to turn the graph into a complete bipartite graph.

Finally, we define preference lists for vertices in  $A$  and  $B$ . For each vertex  $v(P) \in A$ , the edges incident to  $v(P)$  are ordered according to the order in which they appear on the path  $P$ . The dummy edges incident to  $v(P)$  are ordered arbitrarily at the end of the list. For each vertex  $v(Q) \in B$ , its adjacent edges are appear in the reverse order of their appearance on the path  $Q$ , with the dummy edges appearing at the end of the preference list. The preference lists of the dummy vertices are arbitrary.

Let  $M$  be a stable matching in the resulting graph  $\tilde{G}$ . For each edge  $e \in E(\tilde{T})$ , we let  $\mathcal{P}'_e \subseteq \mathcal{P}_e$  be the set of paths  $Q$ , whose vertex  $v(Q)$  is matched to a dummy vertex in  $M$ . Since  $|\mathcal{P}| \leq k_2/2 \leq |\mathcal{P}_e|/2$ , this ensures that  $|\mathcal{P}'_e| \geq k_2/2$ .

In order to construct the set  $\mathcal{P}'$  of paths, consider some path  $P \in \mathcal{P}$ . If  $P$  participates in the matching  $M$  via a dummy edge, then we add  $P$  to  $\mathcal{P}'$ . Otherwise, assume that  $(v(P), v(Q)) \in M$ , and the edge  $(v(P), v(Q))$  corresponds to some edge  $e$  shared by both  $P$  and  $Q$ . We then replace  $P$  with a path  $P'$  that consists of two subpaths: a segment of  $P$  from the beginning of  $P$  until the edge  $e$ , and a segment of  $Q$  from the edge  $e$  to the end of  $Q$ . Notice that  $P'$  starts at the same terminal as  $P$ , and terminates

at some vertex of  $\bigcup_{i=1}^r S_i$ . Path  $P'$  is then added to set  $\mathcal{P}'$ .

This finishes the definition of sets  $\mathcal{P}'$ ,  $\{\mathcal{P}'_e\}_{e \in E(\tilde{T})}$  of paths. Observe that the paths in  $\mathcal{P}'$  connect every terminal in  $\mathcal{T}(\mathcal{M}_0)$  to some vertex of  $\bigcup_{j=1}^r S_j$ , and they do not contain the vertices of  $\bigcup_{j=1}^r S_j$  as inner vertices. It now only remains to show that the set  $\left(\bigcup_{e \in E(\tilde{T})} \mathcal{P}'_e\right) \cup \mathcal{P}'$  of paths causes congestion 1 in  $G'$ .

Clearly, the paths in  $\bigcup_{e \in E(\tilde{T})} \mathcal{P}'_e$  are edge-disjoint from each other in graph  $G'$ , since the paths in  $\bigcup_{e \in E(\tilde{T})} \mathcal{P}_e$  were edge-disjoint from each other.

Assume now for contradiction that some pair of paths  $P', P'' \in \mathcal{P}'$  share an edge. Assume without loss of generality that  $P'$  was obtained from concatenating a segment of some path  $P_1 \in \mathcal{P}$  and some path  $Q_1 \in \mathcal{Q}$  via some edge  $e_1$  (where possibly  $Q_1$  is empty and  $P' = P_1$ ), and similarly  $P''$  was obtained from concatenating a segment of some path  $P_2 \in \mathcal{P}$  and some path  $Q_2 \in \mathcal{Q}$  via some edge  $e_2$ . The paths  $P', P''$  may share an edge only if the first segment of  $P'$  shares an edge with the second segment of  $P''$ , or vice versa. Assume without loss of generality that it is the former, that is, there is an edge  $e'$ , that appears on  $P_1$  before  $e_1$ , and  $e'$  appears on  $Q_2$  after  $e_2$  (notice that it is impossible that  $e' = e_1$  or  $e' = e_2$ , since  $Q_1, Q_2$  cannot share an edge). But then both  $P_1$  and  $Q_2$  prefer the edge  $e'$  to their current matching in  $M$ , contradicting the fact that  $M$  is a stable matching.

Finally, if some pair of paths  $P' \in \mathcal{P}'$  and  $P'' \in \bigcup_{e \in E(\tilde{T})} \mathcal{P}'_e$  share an edge, we reach a contradiction exactly as before.  $\square$

Let  $\mathcal{T}_0 = \mathcal{T}(\mathcal{M}_0)$ . Let  $S \in \mathcal{R}'$  be the set such that at least  $|\mathcal{T}_0|/r$  paths in  $\mathcal{P}'$  terminate at  $S$ . Let  $\mathcal{P}^1 \subseteq \mathcal{P}'$  be the subset of paths that terminate at  $S$ , and let  $\mathcal{T}^1 \subseteq \mathcal{T}_0$  be the subset of terminals where these paths originate. Let  $\mathcal{M}^1 \subseteq \mathcal{M}_0$  be the subset of pairs  $(s, t)$  where both  $s, t \in \mathcal{T}^1$ , and let  $\mathcal{M}^2 \subseteq \mathcal{M}_0$  be the subset of pairs  $(s, t)$  where  $s \in \mathcal{T}^1, t \notin \mathcal{T}^1$ .

Assume first that  $|\mathcal{M}^1| \geq |\mathcal{T}^1|/4$ . Then  $\mathcal{M}^1$  contains  $\Omega\left(\frac{k_2}{r}\right) > k_4$  demand pairs, and we have a set  $\mathcal{P}^1 : \mathcal{T}(\mathcal{M}^1) \rightsquigarrow S$  of paths connecting the terminals participating in pairs in  $\mathcal{M}^1$  to the vertices of  $S$ . We discard arbitrary pairs from  $\mathcal{M}^1$ , until  $|\mathcal{M}^1| = k_4$  holds, and we discard the corresponding paths from  $\mathcal{P}^1$ . We then partition the resulting set  $\mathcal{T}(\mathcal{M}^1)$  of terminals into two subsets,  $\mathcal{T}'_1, \mathcal{T}''_1$ , such that for each demand pair  $(s, t) \in \mathcal{M}^1$ , exactly one terminal belongs to each of the subsets. This defines a corresponding partition  $(\mathcal{P}'_1, \mathcal{P}''_1)$  of the set  $\mathcal{P}^1$  of paths. We set  $S = S'$  in this case.

From now on we assume that  $|\mathcal{M}^1| < |\mathcal{T}^1|/4$ , and so  $|\mathcal{M}^2| = \Omega\left(\frac{k_2}{r}\right)$ . Let  $\mathcal{T}^2 = \{t \mid (s, t) \in \mathcal{M}^2, s \in \mathcal{T}^1\}$ , and let  $\mathcal{P}^2 \subseteq \mathcal{P}'$  be the subset of paths originating at the terminals of  $\mathcal{T}^2$ . Let  $S' \in \mathcal{R}'$  be the set where at least  $|\mathcal{T}^2|/r$  of the paths in  $\mathcal{P}^2$  terminate. We let  $\mathcal{P}''_1 \subseteq \mathcal{P}^2$  be the subset of paths terminating at the vertices of  $S'$ , and we let  $\mathcal{T}''_1 \subseteq \mathcal{T}^2$  be the subset of terminals where these paths originate. We let  $\mathcal{T}'_1 = \{s \mid (s, t) \in \mathcal{M}^2, t \in \mathcal{T}''_1\}$ ,  $\mathcal{P}'_1 \subseteq \mathcal{P}^1$  the set of paths originating at the terminals of  $\mathcal{T}'_1$ , and  $\mathcal{M}_1 \subseteq \mathcal{M}^2$  the set of pairs  $(s, t)$  with  $s \in \mathcal{T}'_1, t \in \mathcal{T}''_1$ . Notice that  $|\mathcal{M}_1| = \Omega\left(\frac{k_2}{r^2}\right) = k_4$  as required.  $\square$

Let  $v$  and  $v'$  be the vertices corresponding to the sets  $S$  and  $S'$  respectively in the tree  $\tilde{T}$ . Assume first that  $v \neq v'$ . Since the maximum vertex degree in  $\tilde{T}$  is at most 3, if we remove  $v$  and  $v'$  from  $\tilde{T}$ , we obtain at most 5 sub-trees. Let  $T_0$  denote the sub-tree incident on both  $v$  and  $v'$ ,  $T_1, T_2$  the sub-trees incident on  $v$ , and  $T_3, T_4$  the sub-trees incident on  $v'$  (possibly some of the sub-trees are empty). Let  $T \in \{T_1, \dots, T_4\}$  be the sub-tree containing the most vertices, and assume without loss of generality that  $T = T_1$ , so it is incident on  $v$ . Let  $T^*$  be the tree obtained from  $\tilde{T}$ , by deleting the sub-trees  $T_2, T_3, T_4$  from it. It is easy to see that  $T^*$  must contain at least  $r/4 \geq 2\gamma_{\text{CMG}}$  vertices, and it contains both vertices  $v$  and  $v'$ . Moreover, the degree of  $v$  is at most 2, and the degree of  $v'$  is 1 in  $T^*$ . We will view  $v$  as the root of  $T^*$ .

If  $v = v'$ , then we root the tree  $\tilde{T}$  at vertex  $v$ , and consider the sub-trees rooted at the children of  $v$ . We delete all but one of these sub-trees, leaving the sub-tree containing the most vertices. We denote by  $T^*$  the resulting tree. In this case, the degree of  $v$  is 1, and  $T^*$  again contains at least  $2\gamma_{\text{CMG}}$  vertices.

Let  $\mathcal{R}'' \subseteq \mathcal{R}'$  be the collection of sets  $S_j$  whose corresponding vertex  $v_j \in V(T^*)$ , so  $|\mathcal{R}''| \geq 2\gamma_{\text{CMG}}$ .

For each terminal  $t \in \mathcal{T}(\mathcal{M}_1)$ , let  $P_t \in \mathcal{P}'_1 \cup \mathcal{P}''_1$  be the unique path originating at  $t$ . Recall that every edge  $e \in \text{out}(S) \cup \text{out}(S')$  may participate in up to two paths in  $\mathcal{P}'_1 \cup \mathcal{P}''_1$  (and in this case  $S = S'$  must hold, and  $e$  must be the last edge on these paths). It will be convenient to ensure that each such edge participates in at most one such path. In order to achieve this, we construct two new subsets  $\mathcal{M}'_1, \mathcal{M}''_1 \subseteq \mathcal{M}_1$  of demand pairs, as follows. Start with  $\mathcal{M}'_1 = \mathcal{M}''_1 = \emptyset$ . While  $\mathcal{M}_1 \neq \emptyset$ , select any pair  $(s, t) \in \mathcal{M}_1$ , and delete it from  $\mathcal{M}_1$ . Let  $e$  be the last edge on  $P_s$ , and let  $e'$  be the last edge on  $P_t$ . If  $e = e'$ , then we add pair  $(s, t)$  to  $\mathcal{M}'_1$ , and continue to the next iteration. Otherwise, we add  $(s, t)$  to  $\mathcal{M}''_1$ , and we delete from  $\mathcal{M}_1$  every pair  $(s', t') \in \mathcal{M}_1$ , where either  $P_{s'}$  or  $P_{t'}$  contain either of the edges  $e$  or  $e'$ . If, at the end of this procedure,  $|\mathcal{M}'_1| \geq |\mathcal{M}_1|/2$ , then we can route every pair  $(s, t) \in \mathcal{M}'_1$  by concatenating the paths  $P_s$  and  $P_t$ . This gives a routing of the paths in  $\mathcal{M}'_1$  with congestion at most 2. From now on we assume that  $|\mathcal{M}'_1| < |\mathcal{M}_1|/2$ . Observe that for each pair that we add to  $\mathcal{M}''_1$ , we delete at most three pairs from  $\mathcal{M}_1$ . Therefore,  $|\mathcal{M}''_1| \geq |\mathcal{M}_1|/6$ . We delete demand pairs from  $\mathcal{M}''_1$ , until  $|\mathcal{M}''_1| = \lfloor k_4/6 \rfloor$  holds, and we delete from  $\mathcal{T}'_1, \mathcal{T}''_1$  all terminals that do not participate in the pairs in  $\mathcal{M}''_1$ , and we delete their corresponding paths from  $\mathcal{P}'_1$  and  $\mathcal{P}''_1$ . For convenience, we will refer to  $\mathcal{M}''_1$  as  $\mathcal{M}_1$  from now on. Observe that now every edge in  $\text{out}(S) \cup \text{out}(S')$  participates in at most one path in  $\mathcal{P}'_1 \cup \mathcal{P}''_1$ .

To summarize this step, we now have a tree  $T^*$  with maximum degree at most 3, and a subset  $\mathcal{R}'' \subseteq \mathcal{R}$  of at least  $2\gamma_{\text{CMG}}$  vertex subsets, such that  $V(T^*) = \{v_j \mid S_j \in \mathcal{R}''\}$ . Tree  $T^*$  contains two special vertices, vertex  $v$  corresponding to the set  $S \in \mathcal{R}''$  - the root of the tree, whose degree is at most 2, and vertex  $v'$  corresponding to the set  $S' \in \mathcal{R}''$ , whose degree is 1. It is possible that  $S = S'$  and so  $v = v'$ . For each edge  $e = (v_i, v_j) \in E(T^*)$ , we are given a set  $\mathcal{P}'_e$  of  $\lfloor k_2/2 \rfloor$  paths in graph  $G$ , connecting the vertices of  $S_i$  to the vertices of  $S_j$ . We are also given a subset  $\mathcal{M}_1$  of  $\lfloor k_4/6 \rfloor < \lfloor k_2/2 \rfloor$  demand pairs, and two subset  $\mathcal{T}'_1, \mathcal{T}''_1$  of terminals, such that for each pair  $(s, t) \in \mathcal{M}_1$ ,  $s \in \mathcal{T}'_1$ ,  $t \in \mathcal{T}''_1$  or vice versa. Additionally, we have two sets  $\mathcal{P}'_1 : \mathcal{T}'_1 \rightsquigarrow \text{out}(S)$ ,  $\mathcal{P}''_1 : \mathcal{T}''_1 \rightsquigarrow \text{out}(S')$  of paths. The paths in  $\mathcal{P}'_1 \cup \mathcal{P}''_1 \cup \left( \bigcup_{e \in E(T^*)} \mathcal{P}'_e \right)$  do not contain the vertices of  $\bigcup_{S'' \in \mathcal{R}''} S''$  as inner vertices, and the total congestion they cause in graph  $G$  is at most 2. Each edge in  $\text{out}(S) \cup \text{out}(S')$  participates in at most one path in  $\mathcal{P}'_1 \cup \mathcal{P}''_1$ .

### 6.3 Step 3: Building the Good Crossbar

We assume without loss of generality that  $\mathcal{R}'' = \{S_1, \dots, S_{r''}\}$ , where  $r'' \geq 2\gamma_{\text{CMG}}$ . We say that a path  $P$  in graph  $G$  is good iff it does not contain the vertices of  $\mathcal{T} \cup \left( \bigcup_{j=1}^{r''} S_j \right)$  as inner vertices. Consider some set  $S_j \in \mathcal{R}''$ , and let  $\mathcal{P}$  be any collection of good paths in  $G$ . We denote by  $\Gamma_j(\mathcal{P}) \subseteq \text{out}(S_j)$  the multi-set of edges, that appear as the first or the last edge on any path in  $\mathcal{P}$ . That is,

$$\Gamma_j(\mathcal{P}) = \{e \in \text{out}(S_j) \mid e \text{ is the first or the last edge on some path } P \in \mathcal{P}\}$$

Our first step is to select, for each edge  $e \in E(T^*)$ , a large subset  $\mathcal{P}''_e \subseteq \mathcal{P}'_e$  of paths, and two large subsets  $\mathcal{P}'_2 \subseteq \mathcal{P}'_1, \mathcal{P}''_2 \subseteq \mathcal{P}''_1$  of paths, such that, on the one hand, for each  $1 \leq j \leq r''$ , set  $S_j$  has the 1-bandwidth property with respect to the subset  $\Gamma_j(\mathcal{P}')$  of edges, where  $\mathcal{P}' = \mathcal{P}'_2 \cup \mathcal{P}''_2 \cup \left( \bigcup_{e \in E(T^*)} \mathcal{P}''_e \right)$ , while on the other hand the endpoints of paths in  $\mathcal{P}'_2$  and  $\mathcal{P}''_2$  form source-sink pairs in set  $\mathcal{M}_1$ . This is done in the following theorem.

**Theorem 6.4** *There is an efficient algorithm, that either computes a subset  $\mathcal{M}' \subseteq \mathcal{M}_1$  of  $k/\text{poly log } k$  demand pairs and their routing with congestion at most 2 in  $G$ , or a  $(k_1, \alpha)$ -violating partition of some set  $S_j \in \mathcal{R}''$ , or it computes, for each edge  $e \in E(T^*)$ , a subset  $\mathcal{P}_e'' \subseteq \mathcal{P}'_e$  of  $k_5 = \Omega(\alpha_{\text{BW}}^2 k_2)$  paths, and two subsets  $\mathcal{P}'_2 \subseteq \mathcal{P}'_1, \mathcal{P}''_2 \subseteq \mathcal{P}''_1$ , such that:*

- *Each set  $S_j \in \mathcal{R}''$  has the 1-bandwidth property with respect to the set  $\Gamma_j(\mathcal{P}')$  of edges, where  $\mathcal{P}' = \mathcal{P}'_2 \cup \mathcal{P}''_2 \cup \left( \bigcup_{e \in E(T^*)} \mathcal{P}''_e \right)$ .*
- *There is a subset  $\mathcal{M}_2 \subseteq \mathcal{M}_1$  of  $\Omega(\alpha_{\text{BW}}^2) \cdot |\mathcal{M}_1| = \Omega(\alpha_{\text{BW}}^2 \cdot k_4)$  demand pairs, and two subsets  $\mathcal{T}'_2 \subseteq \mathcal{T}'_1, \mathcal{T}''_2 \subseteq \mathcal{T}''_1$  of terminals, such that for each pair  $(s, t) \in \mathcal{M}_2$ ,  $s \in \mathcal{T}'_2, t \in \mathcal{T}''_2$  or vice versa. Moreover,  $\mathcal{P}'_2$  connects every terminal of  $\mathcal{T}'_2$  to some edge in  $\text{out}(S)$ , and  $\mathcal{P}''_2$  connects every terminal of  $\mathcal{T}''_2$  to some edge in  $\text{out}(S')$ .*

**Proof:** We start with  $\mathcal{P} = \mathcal{P}'_1 \cup \mathcal{P}''_1 \cup \left( \bigcup_{e \in E(T^*)} \mathcal{P}'_e \right)$ , and we consider the sets  $S_j$  one-by-one. For each such set  $S_j$ , we compute the instance  $\text{SC}(G, S_j, \Gamma_j)$  of the sparsest cut problem, where  $\Gamma_j = \Gamma_j(\mathcal{P})$ , and run the algorithm  $\mathcal{A}_{\text{ARV}}$  on it. Let  $(A, B)$  be the output of the algorithm. If the sparsity of the cut  $(A, B)$  is less than  $\alpha$ , then  $(A, B)$  defines a  $(k_1, \alpha)$ -violating partition of  $S_j$ . We then stop the algorithm and return this partition. Otherwise, we are guaranteed that set  $S_j$  has the  $\alpha/\alpha_{\text{ARV}} = \alpha_{\text{BW}}$ -bandwidth property with respect to  $\Gamma_j(\mathcal{P})$ .

From now on, we assume that for each  $1 \leq j \leq r''$ , set  $S_j$  has the  $\alpha_{\text{BW}}$ -bandwidth property with respect to  $\Gamma_j(\mathcal{P})$ . Our next step is to boost this bandwidth property for every subset  $S_j$  to 1, by selecting subsets of paths in each set  $\mathcal{P}'_e$  (and in sets  $\mathcal{P}'_1, \mathcal{P}''_1$ ). This is done by iteratively applying Theorems 2.2 and 2.3 to the sets  $S_1, \dots, S_{r''}$ .

We start with the case where  $v' = v$ . We process the vertices of the tree  $T^*$  in the bottom-up fashion. Let  $v_j$  be some vertex of  $T^*$ . We maintain the invariant that once vertex  $v_j$  is processed, for every edge  $e$  in the sub-tree of  $v_j$ , we have computed a subset  $\mathcal{P}''_e \subseteq \mathcal{P}'_e$  of at least  $k_5$  paths, such that the following holds. Let  $\mathcal{Q}$  denote the union of all paths in the sets  $\mathcal{P}''_e$ , where  $e$  belongs to the sub-tree of  $v_j$ . Then for every vertex  $v_{j'}$  in the sub-tree of  $v_j$ , set  $S_{j'}$  has the 1-bandwidth property with respect to  $\Gamma_{j'}(\mathcal{Q})$ . Additionally, if  $e$  is the edge connecting  $v_j$  to its parent, then we have also computed a subset  $\tilde{\mathcal{P}}_e \subseteq \mathcal{P}'_e$  of at least  $\tilde{k} = \Omega(\alpha_{\text{BW}} k_2)$  paths, such that set  $S_j$  has the 1-bandwidth property with respect to the set  $\Gamma_j(\mathcal{Q} \cup \tilde{\mathcal{P}}_e)$  of edges.

We now show how to process each vertex, so that the above invariant is preserved. Assume first that  $v_j$  is a leaf vertex, and let  $e$  be the unique edge of  $T^*$  incident to  $v_j$ . Let  $\Gamma = \Gamma_j(\mathcal{P}_e)$ . Consider the instance  $\text{SC}(G, S_j, \Gamma)$  of the sparsest cut problem. Recall that it is defined on a graph  $G_j(\Gamma)$ , with a set  $\mathcal{T}(\Gamma)$  of terminals, where each terminal  $t_e \in \mathcal{T}(\Gamma)$  corresponds to an edge  $e \in \Gamma$ . We apply Theorem 2.2 to graph  $G_j(\Gamma)$  with a single set  $\mathcal{T}(\Gamma)$  of terminals. Since  $S_j$  has the  $\alpha_{\text{BW}}$ -bandwidth property with respect to  $\Gamma$ , set  $\mathcal{T}(\Gamma)$  is  $\alpha_{\text{BW}}$ -well-linked in graph  $G_j(\Gamma)$ , and so we will obtain a subset  $\mathcal{T}'(\Gamma) \subseteq \mathcal{T}(\Gamma)$  of  $\Omega(\alpha_{\text{BW}} k_2) \geq \tilde{k}$  vertices. Let  $\Gamma'_j \subseteq \Gamma$  be the subset of edges corresponding to the vertices in  $\mathcal{T}'(\Gamma)$ . We define  $\tilde{\mathcal{P}}_e \subseteq \mathcal{P}'_e$  to be the subset of paths whose last edge belongs to  $\Gamma'_j$ . Notice that each edge  $e' \in \Gamma'_j$  may serve as the last edge for at most two paths in  $\mathcal{P}'_e$ . If two such paths exist, then only one of them is added to  $\tilde{\mathcal{P}}_e$ . This defines the required set  $\tilde{\mathcal{P}}_e$ , with  $|\tilde{\mathcal{P}}_e| \geq \tilde{k}$ . Observe that  $S_j$  has the 1-bandwidth property with respect to  $\Gamma_j(\tilde{\mathcal{P}}_e)$ .

Consider now some non-leaf vertex  $v_j \neq v$ . We assume that  $v_j$  has two children. The case where  $v_j$  has one child is handled similarly. Assume that  $e$  is the edge connecting  $v_j$  to its parent, and  $e_1, e_2$  are the edges connecting  $v_j$  to its two children in  $T^*$ . We define three subsets of edges in  $\text{out}(S_j)$ :  $\Gamma_0$ , containing all edges of  $\text{out}(S_j)$  that belong to paths in  $\mathcal{P}'_e$ , and  $\Gamma_1, \Gamma_2$ , containing the edges of  $\text{out}(S_j)$

that belong to the paths in  $\tilde{\mathcal{P}}_{e_1}, \tilde{\mathcal{P}}_{e_2}$ , respectively. Let  $\Gamma = \Gamma_0 \cup \Gamma_1 \cup \Gamma_2$ . Recall that  $S_j$  has the  $\alpha_{\text{BW}}$ -bandwidth property with respect to  $\Gamma$ . As before, we build an instance  $\text{SC}(G, S_j, \Gamma)$  of the sparsest cut problem, with the graph  $G_j(\Gamma)$  and the set  $\mathcal{T}(\Gamma)$  of terminals. Let  $\mathcal{T}_0(\Gamma), \mathcal{T}_1(\Gamma), \mathcal{T}_2(\Gamma) \subseteq \mathcal{T}(\Gamma)$  be the subsets of these new terminals corresponding to the edges in the sets  $\Gamma_0, \Gamma_1$ , and  $\Gamma_2$ , respectively. Then  $\mathcal{T}(\Gamma)$  is  $\alpha_{\text{BW}}$ -well-linked in graph  $G_j(\Gamma)$ , and moreover  $|\mathcal{T}_0(\Gamma)| \geq \lfloor k_2/2 \rfloor$  and  $|\mathcal{T}_1(\Gamma)|, |\mathcal{T}_2(\Gamma)| \geq \tilde{k}/2$  (since each edge in  $\text{out}(S_j)$  may participate in up to two paths in  $\mathcal{P}$ ). We apply Theorem 2.2 to graph  $G_j(\Gamma)$  and three subsets  $\mathcal{T}_0(\Gamma), \mathcal{T}_1(\Gamma), \mathcal{T}_2(\Gamma)$  of terminals. Let  $\mathcal{T}'_0(\Gamma) \subseteq \mathcal{T}_0(\Gamma)$ ,  $\mathcal{T}'_1(\Gamma) \subseteq \mathcal{T}_1(\Gamma)$ , and  $\mathcal{T}'_2(\Gamma) \subseteq \mathcal{T}_2(\Gamma)$  be the output of the theorem. Recall that  $|\mathcal{T}'_0(\Gamma)| \geq \Omega(\alpha_{\text{BW}} k_2) \geq \tilde{k}$ , while  $|\mathcal{T}'_1(\Gamma)|, |\mathcal{T}'_2(\Gamma)| \geq \Omega(\alpha_{\text{BW}} \tilde{k}) = \Omega(\alpha_{\text{BW}}^2 k_2) \geq k_5$ . Moreover,  $\mathcal{T}'_0(\Gamma) \cup \mathcal{T}'_1(\Gamma) \cup \mathcal{T}'_2(\Gamma)$  is 1-well-linked in graph  $G_j(\Gamma)$ , and these vertex sets are mutually disjoint. As before, we let  $\Gamma'_0 \subseteq \Gamma_0$  be the subset of edges corresponding to the vertices of  $\mathcal{T}'_0(\Gamma)$ , and we let  $\Gamma'_1, \Gamma'_2$  be the subsets of edges corresponding to the vertices of  $\mathcal{T}'_1(\Gamma)$  and  $\mathcal{T}'_2(\Gamma)$ , respectively. Finally, we set  $\tilde{\mathcal{P}}_e \subseteq \mathcal{P}'_e$  be the subset of paths that contain an edge in  $\Gamma'_0$  (if some edge in  $\Gamma'_0$  is contained in two such paths, only one such path is added to  $\tilde{\mathcal{P}}_e$ ). We define  $\mathcal{P}''_{e_1} \subseteq \tilde{\mathcal{P}}_{e_1}$  and  $\mathcal{P}''_{e_2} \subseteq \tilde{\mathcal{P}}_{e_2}$  similarly, using the sets  $\Gamma_1, \Gamma_2$  of edges. Observe that  $|\tilde{\mathcal{P}}_e| \geq \tilde{k}$ , and  $|\mathcal{P}_{e_1}|, |\mathcal{P}_{e_2}| \geq k_5$  as required. It is easy to verify that the invariant is preserved.

In our final step we process the root vertex  $v = v_j$ . Recall that since we have assumed that  $v = v'$ , vertex  $v$  has degree 1. Let  $e$  be the unique edge incident on  $v$ , and let  $S_j \in \mathcal{R}''$  be the set corresponding to  $v$ . We again define three subsets  $\Gamma_0, \Gamma_1, \Gamma_2 \subseteq \text{out}(S_j)$  of edges, where  $\Gamma_0$  contains all edges participating in the paths in  $\tilde{\mathcal{P}}_e$ , and  $\Gamma_1$  and  $\Gamma_2$  contain all edges participating in the paths in  $\mathcal{P}'_1$  and  $\mathcal{P}''_1$ , respectively. Notice that  $\mathcal{M}_1$  defines a matching  $\tilde{\mathcal{M}}$  over the set  $\Gamma_1 \cup \Gamma_2$  of edges, where for each pair  $(s, t) \in \mathcal{M}_1$ , if  $P, P' \in \mathcal{P}'_1 \cup \mathcal{P}''_1$  are the two paths originating at  $s$  and  $t$  respectively, and  $e', e''$  are the last edges on these two paths, then  $(e, e'') \in \tilde{\mathcal{M}}$ . As before, we set up an instance  $\text{SC}(G, S_j, \Gamma)$  of the sparsest cut problem, with graph  $G_j(\Gamma)$  and the set  $\mathcal{T}(\Gamma)$  of terminals, where  $\Gamma = \Gamma_0 \cup \Gamma_1 \cup \Gamma_2$ . Let  $\mathcal{T}_0(\Gamma), \mathcal{T}_1(\Gamma)$  and  $\mathcal{T}_2(\Gamma)$  be the subsets of terminals in  $\mathcal{T}(\Gamma)$  corresponding to the edges in sets  $\Gamma_0, \Gamma_1$ , and  $\Gamma_2$ , respectively. We apply Theorem 2.3 to graph  $G_j(\Gamma)$ , where the first set of terminals is  $\mathcal{T}_1(\Gamma) \cup \mathcal{T}_2(\Gamma)$ , together with the matching  $\tilde{\mathcal{M}}$ , while the second set of terminals is  $\mathcal{T}_0(\Gamma)$ . Let  $\mathcal{T}'_0(\Gamma) \subseteq \mathcal{T}_0(\Gamma)$ ,  $\mathcal{T}'_1(\Gamma) \subseteq \mathcal{T}_1(\Gamma)$ , and  $\mathcal{T}'_2(\Gamma) \subseteq \mathcal{T}_2(\Gamma)$  be the output of the theorem, and let  $\tilde{\mathcal{M}}' \subseteq \tilde{\mathcal{M}}$  be the corresponding matching over the vertices of  $\mathcal{T}'_1(\Gamma) \cup \mathcal{T}'_2(\Gamma)$ . As before, we define subsets  $\Gamma'_0, \Gamma'_1, \Gamma'_2$  of edges corresponding to the terminal sets  $\mathcal{T}'_0(\Gamma), \mathcal{T}'_1(\Gamma)$  and  $\mathcal{T}'_2(\Gamma)$ , respectively. The subset  $\mathcal{P}''_e \subseteq \tilde{\mathcal{P}}_e$  of paths is defined exactly as before, from the subset  $\Gamma'_0$  of edges.

Finally, let  $\mathcal{P}'_2 \subseteq \mathcal{P}'_1, \mathcal{P}''_2 \subseteq \mathcal{P}''_1$  be the subsets of paths containing edges in  $\Gamma'_1$  and  $\Gamma'_2$ , respectively (recall that each edge in  $\Gamma'_1 \cup \Gamma'_2$  may belong to at most one such path). Let  $\mathcal{T}'_2 \subseteq \mathcal{T}'_1$  be the subset of terminals where the paths of  $\mathcal{P}'_2$  originate, and similarly let  $\mathcal{T}''_2 \subseteq \mathcal{T}''_1$  be the subset of terminals where the paths of  $\mathcal{P}''_2$  originate. Let  $\mathcal{M}_2 = \{(s, t) \in \mathcal{M}_1 \mid s \in \mathcal{T}'_2, t \in \mathcal{T}''_2\}$ . From the above discussion,  $|\mathcal{M}_2| = |\mathcal{M}'| \geq \Omega(\alpha_{\text{BW}})|\mathcal{M}_1|$ . This finishes the proof of the theorem for the case where  $v = v'$ .

Assume now that  $v \neq v'$ . The algorithm is exactly as before, except that the vertices  $v$  and  $v'$  are processed differently. In order to process vertex  $v'$ , assume that  $v' = v_j$ , and let  $e$  be the unique edge incident on  $v'$  in tree  $T^*$ . We consider the two sets  $\mathcal{P}''_1, \mathcal{P}'(e)$  of paths, and the two corresponding subsets  $\Gamma_1, \Gamma_2 \subseteq \text{out}(S_j)$  of edges, that lie on the paths in  $\mathcal{P}''_1$  and  $\mathcal{P}'(e)$ , respectively. We then process vertex  $v'$  as any other degree-2 inner vertex in tree  $T^*$ , using the sets  $\mathcal{P}''_1, \mathcal{P}'(e)$  of paths, and the corresponding subsets  $\Gamma_1, \Gamma_2$  of edges. The output of this iteration is a subset  $\tilde{\mathcal{P}}(e) \subseteq \mathcal{P}'(e)$  of at least  $\tilde{k}$  paths, and a subset  $\tilde{\mathcal{P}}''_1 \subseteq \mathcal{P}''_1$  of at least  $\Omega(\alpha_{\text{BW}})|\mathcal{M}_1|$  paths, such that set  $S_j$  has the 1-bandwidth property with respect to  $\Gamma_j(\tilde{\mathcal{P}}''_1 \cup \tilde{\mathcal{P}}(e))$ .

Let  $\tilde{\mathcal{T}}''_1 \subseteq \mathcal{T}''_1$  be the subset of terminals  $t$ , where  $P_t \in \tilde{\mathcal{P}}''_1$ . Next, we discard from  $\mathcal{M}_1$  all pairs  $(s, t)$  that do not contain a terminal in  $\tilde{\mathcal{P}}''_1$ , obtaining a new set  $\tilde{\mathcal{M}} \subseteq \mathcal{M}_1$  of demand pairs. We also construct a set  $\tilde{\mathcal{T}}'_1 \subseteq \mathcal{T}'_1$  of terminals, that participate in the pairs in  $\tilde{\mathcal{M}}$ , and a set  $\tilde{\mathcal{P}}'_1 \subseteq \mathcal{P}'_1$  of paths originating from these terminals.

Vertex  $v$  is processed as follows. Let  $e, e'$  be two edges adjacent to  $v$  in  $T^*$  (the case where the degree of  $v$  is 1 is handled similarly). We process vertex  $v$  exactly like all other inner vertices of tree  $T^*$ , with the corresponding three subsets of paths:  $\tilde{\mathcal{P}}(e), \tilde{\mathcal{P}}(e')$ , and  $\tilde{\mathcal{P}}_1$ . As a result, we obtain three new subsets of paths:  $\mathcal{P}''(e) \subseteq \tilde{\mathcal{P}}(e)$  and  $\mathcal{P}''(e') \subseteq \tilde{\mathcal{P}}(e')$  containing at least  $k_5$  paths each, and a subset  $\mathcal{P}'_2 \subseteq \tilde{\mathcal{P}}_1$  of  $\Omega(\alpha_{\text{BW}}^2)|\mathcal{M}_1|$  paths. Let  $\mathcal{T}'_2 \subseteq \tilde{\mathcal{T}}'_1$  be the subset of terminals  $t$  where  $P_t \in \mathcal{P}'_2$ , and let  $\mathcal{M}_2 \subseteq \mathcal{M}_1$  be the set of all pairs  $(s, t)$  that contain a terminal in  $\mathcal{T}'_2$ . We set  $\mathcal{T}''_2 \subseteq \tilde{\mathcal{T}}''_1$  be the set of all terminals participating in pairs in  $\mathcal{M}_2$ , and we let  $\mathcal{P}''_2 \subseteq \tilde{\mathcal{P}}''_1$  be the set of paths originating at the terminals of  $\mathcal{P}''_2$ .  $\square$

Let  $k_6 = |\mathcal{M}_2| = \Omega(\alpha_{\text{BW}}^2)k_4$ . Notice that  $k_6 < k_5/2$ . If  $v = v'$ , then for each edge  $e \in E(T^*)$ , we discard arbitrary paths from  $\mathcal{P}''(e)$ , until  $|\mathcal{P}''(e)| = 2k_6$  holds. Otherwise, we do the same to ensure that  $|\mathcal{P}''(e)| = k_6$  for each edge  $e \in E(T^*)$ . Recall that  $|\mathcal{P}'_2|, |\mathcal{P}''_2| = k_6$ .

We are now ready to define the good crossbar in graph  $G$ . Let  $\mathcal{S}^*$  contain all sets  $S_j$ , where the degree of vertex  $v_j$  in tree  $T^*$  is either 1 or 2 (excluding the set  $S$  corresponding to the root vertex  $v$ ). Notice that at least half the vertices of  $T^*$  must have this property, and therefore,  $|\mathcal{S}^*| \geq \gamma_{\text{CMG}}$ . If  $|\mathcal{S}^*| > \gamma_{\text{CMG}}$ , then we discard vertex subsets from  $\mathcal{S}^*$  arbitrarily, until  $|\mathcal{S}^*| = \gamma_{\text{CMG}}$  holds.

Consider some set  $S_j \in \mathcal{S}^*$ , and let  $e$  be any one of the edges adjacent to the vertex  $v_j$  in tree  $T^*$ . We then define  $\Gamma_j^* \subseteq \text{out}(S_j)$  to be the subset of edges that participate in the paths in  $\mathcal{P}''(e)$ . Recall that from Theorem 6.4, set  $S_j$  has the 1-bandwidth property with respect to  $\Gamma_j^*$ .

Consider any set  $S_j \in \mathcal{R}''$ . If the degree of  $v_j$  in tree  $T^*$  is 2, then let  $e, e'$  be the two edges incident to  $v_j$  in  $T^*$ , and let  $\Gamma, \Gamma' \subseteq \text{out}(S_j)$  be the sets of edges lying on the paths  $\mathcal{P}''(e)$  and  $\mathcal{P}''(e')$ , respectively. Since  $S_j$  has the 1-bandwidth property with respect to  $\Gamma \cup \Gamma'$ , we can find a set  $\mathcal{Q}_j : \Gamma \xrightarrow{1:1} \Gamma'$  of paths contained in  $S_j$ .

Assume now that the degree of  $v_j$  in tree  $T^*$  is 3 (in this case,  $S_j \notin \mathcal{S}^*$ ), and let  $e, e'$  and  $e''$  be the three edges incident to  $v_j$  in  $T^*$ , where  $e$  connects  $v_j$  to its parent. Let  $\Gamma, \Gamma', \Gamma'' \subseteq \text{out}(S_j)$  be the sets of edges lying on the paths in  $\mathcal{P}''(e), \mathcal{P}''(e')$ , and  $\mathcal{P}''(e'')$ , respectively. Since  $S_j$  has the 1-bandwidth property with respect to  $\Gamma \cup \Gamma' \cup \Gamma''$ , we can find two sets  $\mathcal{Q}_j^1 : \Gamma \xrightarrow{1:1} \Gamma'$  and  $\mathcal{Q}_j^2 : \Gamma \xrightarrow{1:1} \Gamma''$  of paths inside  $S_j$ . Let  $\mathcal{Q}_j = \mathcal{Q}_j^1 \cup \mathcal{Q}_j^2$ .

Assume now that  $v = v'$ , and let  $S$  be the set corresponding to the vertex  $v$ . Let  $e$  be the unique edge incident on  $v$  in tree  $T^*$ , and let  $\Gamma \subseteq \text{out}(S)$  be the subset of edges lying on the paths in  $\mathcal{P}''(e)$ . Similarly, let  $\Gamma' \subseteq \text{out}(S)$  be the subset of edges lying on the paths in  $\mathcal{P}'_2 \cup \mathcal{P}''_2$ . We can then find a set  $\mathcal{Q}(v) : \Gamma \xrightarrow{1:1} \Gamma'$  of paths contained in set  $S$ .

We are now ready to complete the construction of the good crossbar for the case where  $v = v'$ . We let  $\mathcal{M}^* = \mathcal{M}_2$ . For each terminal  $t \in \mathcal{T}(\mathcal{M}^*)$ , we construct a tree  $T_t$ , which is then added to  $\tau^*$ , as follows. We start with the path  $P_t \in \mathcal{P}'_2 \cup \mathcal{P}''_2$ , and let  $\tilde{e}$  be the last edge on this path. We then add to  $T_t$  the unique path in  $\mathcal{Q}(S)$  that originates from  $\tilde{e}$ . Let  $\tilde{e}'$  be the last edge on this path, let  $e'$  be the edge of  $T^*$  connecting  $v$  to its child, and let  $P$  be the unique path in  $\mathcal{P}''_{e'}$  originating at  $\tilde{e}'$ . We then add  $P$  to the tree  $T_t$ .

In general, we process the vertices of the tree  $T^*$  in the top-bottom manner. We maintain the invariant that when a vertex  $v_j$  is processed, for each edge  $e$  adjacent to  $v_j$  in tree  $T^*$ , one path of  $\mathcal{P}''_e$  has been added to tree  $T_t$ . Vertex  $v_j$  is processed as follows. Let  $e$  be the edge connecting  $v_j$  to its parent, and let  $P \in \mathcal{P}''_e$  be the path that belongs to tree  $T_t$ . Let  $\tilde{e} \subseteq \text{out}(S_j)$  be the last edge on path  $P$ . If the degree of  $v_j$  is 2 in  $T^*$ , then we add to  $T_t$  the unique path in  $\mathcal{Q}_j$  that originates at  $\tilde{e}$ . Let  $\tilde{e}'$  be the last edge on this path, and let  $e'$  be the edge in  $T^*$  connecting  $v_j$  to its child. We add the unique path in  $\mathcal{P}''_{e'}$  originating at  $\tilde{e}'$  to the tree  $T_t$ .

Otherwise, if the degree of  $v_j$  is 3, then we add to  $T_t$  two paths:  $P_1 \in \mathcal{Q}_j^1$ , and  $P_2 \in \mathcal{Q}_j^2$ , that originate

at  $\tilde{e}$ . Let  $e', e''$  denote the two edges of  $T^*$  connecting the vertex  $v_j$  to its two children. As before, we add a unique path  $P_3 \in \mathcal{P}_{e'}''$ , originating at the last edge of  $P_1$ , and  $P_4 \in \mathcal{P}_{e''}''$ , originating at the last edge of  $P_2$  to tree  $T_t$ .

If the degree of  $v_j$  is 1, then we do nothing. Once all vertices of  $T^*$  are processed, we obtain a tree  $T_t$ , containing the terminal  $t$ , and for each set  $S_j \in \mathcal{S}^*$ , a unique edge in  $\Gamma^*$ . We let  $\tau^*$  denote the collection of all trees  $T_t$ , for all  $t \in \mathcal{T}(\mathcal{M}^*)$ . It is immediate to verify that the trees in  $\tau^*$  cause congestion at most 2 in graph  $G$ . Moreover, since the sets  $S_j \in \mathcal{S}^*$  correspond to vertices of degree 1 or 2 in  $T^*$ , each edge of  $G[S_j]$  for such sets  $S_j$  belongs to at most one tree.

Assume now that  $v \neq v'$ . Let  $e, e'$  be the two edges incident to the vertex  $v$  in tree  $T^*$  (if the degree of  $v$  is 1, the argument is similar). As before, we denote by  $\Gamma', \Gamma'' \subseteq \text{out}(S)$  the subsets of edges that lie on the paths in  $\mathcal{P}_e'', \mathcal{P}_{e'}''$ , respectively, and we let  $\Gamma \subseteq \text{out}(S)$  denote the subset of edges lying on the paths in  $\mathcal{P}_2'$ . We then compute two sets  $\mathcal{Q}_1(S) : \Gamma \overset{1:1}{\rightsquigarrow} \Gamma'$ , and  $\mathcal{Q}_2(S) : \Gamma \overset{1:1}{\rightsquigarrow} \Gamma''$  of paths contained in set  $S$ .

Let  $e$  be the unique edge incident on the vertex  $v'$  in tree  $T^*$ , and let  $\Gamma \subseteq \text{out}(S')$  be the subset of edges participating in the paths in  $\mathcal{P}_e''$ . Let  $\Gamma' \subseteq \text{out}(S')$  be the subset of edges participating in the paths in  $\mathcal{P}_2''$ . We compute a set  $\mathcal{Q}(S') : \Gamma \overset{1:1}{\rightsquigarrow} \Gamma'$  of paths contained in  $S'$ .

For each terminal  $t \in \mathcal{T}(\mathcal{M}_2)$ , we now build a tree  $T_t$ . This tree is constructed exactly as before, and it contains the path  $P_t \in \mathcal{P}_2'$ , and some additional path  $P \in \mathcal{P}_2''$ . As before, for each edge  $e \in T^*$ , the tree contains exactly one path in  $\mathcal{P}_e''$ . Notice that now each tree  $T \in \tau^*$  contains two terminals: one in  $\mathcal{T}_2'$ , and one in  $\mathcal{T}_2''$ . We partition  $\mathcal{M}_2$  into two subsets,  $\mathcal{M}_2'$  and  $\mathcal{M}_2''$ , where  $\mathcal{M}_2'$  contains all pairs  $(s, t)$  where both  $s$  and  $t$  belong to the same tree in  $\tau^*$ , and  $\mathcal{M}_2''$  contains all remaining pairs. If  $|\mathcal{M}_2'| \geq |\mathcal{M}_2|/2$ , then we can route every pair  $(s, t) \in \mathcal{M}_2'$  along the tree  $T \in \tau^*$  containing  $s$  and  $t$ . This gives a routing of the pairs in  $\mathcal{M}_2'$  with congestion at most 2. We now assume that  $|\mathcal{M}_2''| \geq |\mathcal{M}_2|/2$ .

We construct the set  $\mathcal{M}^* \subseteq \mathcal{M}_2''$  of demand pairs as follows. Start with  $\mathcal{M}^* = \emptyset$ . While  $\mathcal{M}_2'' \neq \emptyset$ , select any pair  $(s, t) \in \mathcal{M}_2''$  and move it to  $\mathcal{M}^*$ . Let  $T_s, T_t \in \tau^*$  be the trees in which  $s$  and  $t$  participate, respectively. Remove from  $\mathcal{M}^*$  all pairs  $(s', t')$ , where either  $s'$  or  $t'$  participate in either  $T_s$  or  $T_t$ . Observe that for each pair that we add to  $\mathcal{M}^*$ , at most three pairs are removed from  $\mathcal{M}_2''$ . Therefore,  $|\mathcal{M}^*| \geq |\mathcal{M}_2''|/3 \geq |\mathcal{M}_2|/6 \geq k_6/6$ . Finally, we remove from  $\tau^*$  all trees that do not contain terminals in  $\mathcal{T}(\mathcal{M}^*)$ . Observe that now each tree in  $\tau^*$  contains exactly one terminal in  $\mathcal{T}(\mathcal{M}^*)$ . For each set  $S_j \in \mathcal{S}^*$ , we also discard from  $\Gamma_j^*$  all edges that do not participate in trees in  $\tau^*$ . Observe that for each  $S_j \in \mathcal{S}^*$ , we now have  $|\Gamma_j^*| = |\tau^*| \geq k_6/6$ . Setting  $k^* = |\mathcal{M}^*| = |\tau^*|$  finishes the construction.

## 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.
- [AGLR94] Baruch Awerbuch, Rainer Gawlick, Tom Leighton, and Yuval Rabani. On-line admission control and circuit routing for high performance computing and communication. In *Proc. 35th IEEE Symp. on Foundations of Computer Science*, pages 412–423, 1994.
- [And10] Matthew Andrews. Approximation algorithms for the edge-disjoint paths problem via Raecke decompositions. In *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science, FOCS '10*, pages 277–286, Washington, DC, USA, 2010. IEEE Computer Society.

- [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.
- [AR06] Yossi Azar and Oded Regev. Combinatorial algorithms for the unsplittable flow problem. *Algorithmica*, 44(1):49–66, 2006.
- [ARV09] Sanjeev Arora, Satish Rao, and Umesh V. Vazirani. Expander flows, geometric embeddings and graph partitioning. *J. ACM*, 56(2), 2009.
- [AZ06] Matthew Andrews and Lisa Zhang. Logarithmic hardness of the undirected edge-disjoint paths problem. *Journal of the ACM (JACM)*, 53(5):745–761, 2006.
- [AZ07] Matthew Andrews and Lisa Zhang. Hardness of the undirected congestion minimization problem. *SIAM J. Comput.*, 37(1):112–131, 2007.
- [BFSU94] Andrei Z. Broder, Alan M. Frieze, Stephen Suen, and Eli Upfal. Optimal construction of edge-disjoint paths in random graphs. In *Proc. 5th ACM-SIAM SODA*, pages 603–612, 1994.
- [BFU94] Andrei Z. Broder, Alan M. Frieze, and Eli Upfal. Existence and construction of edge-disjoint paths on expander graphs. *SIAM J. Comput.*, pages 976–989, 1994.
- [BS00] Alok Baveja and Aravind Srinivasan. Approximation algorithms for disjoint paths and related routing and packing problems. *Mathematics of Operations Research*, 25:2000, 2000.
- [CC14] Chandra Chekuri and Julia Chuzhoy. Polynomial bounds for the grid-minor theorem. In David B. Shmoys, editor, *STOC*, pages 60–69. ACM, 2014.
- [CE13] Chandra Chekuri and Alina Ene. Poly-logarithmic approximation for maximum node disjoint paths with constant congestion. In *Proc. of ACM-SIAM SODA*, 2013.
- [CHR03] M. Conforti, R. Hassin, and R. Ravi. Reconstructing flow paths. *Operations Research Letters*, 31:273–276, 2003.
- [Chu12] Julia Chuzhoy. Routing in undirected graphs with constant congestion. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 855–874. ACM, 2012. full version at <http://ttic.uchicago.edu/~cjulia/publications.html>.
- [Chu16] Julia Chuzhoy. Improved bounds for the excluded grid theorem. *arXiv preprint arXiv:1602.02629*, 2016.
- [CKS04a] C. Chekuri, S. Khanna, and F. B. Shepherd. Edge-disjoint paths in planar graphs. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '04, pages 71–80, Washington, DC, USA, 2004. IEEE Computer Society.
- [CKS04b] Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd. The all-or-nothing multicommodity flow problem. In László Babai, editor, *STOC*, pages 156–165. ACM, 2004. A full version at <http://www.math.mcgill.ca/~bshepherd/PS/all.pdf>.
- [CKS05] Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd. Multicommodity flow, well-linked terminals, and routing problems. In *STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 183–192, New York, NY, USA, 2005. ACM.



- [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.
- [CKS09] Chandra Chekuri, Sanjeev Khanna, and F Bruce Shepherd. Edge-disjoint paths in planar graphs with constant congestion. *SIAM Journal on Computing*, 39(1):281–301, 2009.
- [CMS07] Chandra Chekuri, Marcelo Mydlarz, and F. Bruce Shepherd. Multicommodity demand flow in a tree and packing integer programs. *ACM Trans. Algorithms*, 3, August 2007.
- [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.
- [Fra89] A. Frank. On connectivity properties of Eulerian digraphs. *Ann. Discrete Math.*, 41, 1989.
- [Fri00] Alan M. Frieze. Edge-disjoint paths in expander graphs. *SIAM Journal On Computing*, 30:2001, 2000.
- [GS62] David Gale and Lloyd Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 1:9–14, 1962.
- [GVY97] Naveen Garg, Vijay V. Vazirani, and Mihalis Yannakakis. Primal-dual approximation algorithms for integral flow and multicut in trees. *Algorithmica*, 18(1):3–20, 1997.
- [HS70] A. Hajnal and E. Szemerédi. Proof of a conjecture of P. Erdos. *Combinatorial Theory and its Applications*, pages 601–623, 1970.
- [Jac98] Bill Jackson. Some remarks on arc-connectivity, vertex splitting, and orientation in graphs and digraphs. *J. Graph Theory*, 12:429–436, 1998.
- [Kar72] R. Karp. Reducibility among combinatorial problems. In R. Miller and J. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.
- [KK11] Ken-ichi Kawarabayashi and Yusuke Kobayashi. Breaking  $O(n^{1/2})$ -approximation algorithms for the edge-disjoint paths problem with congestion two. In Lance Fortnow and Salil P. Vadhan, editors, *STOC*, pages 81–88. ACM, 2011.
- [KK13] Ken-ichi Kawarabayashi and Yusuke Kobayashi. An  $O(\log n)$ -approximation algorithm for the edge-disjoint paths problem in Eulerian planar graphs. *ACM Transactions on Algorithms*, 9(2):16, 2013.
- [Kle05] Jon M. Kleinberg. An approximation algorithm for the disjoint paths problem in even-degree planar graphs. In *FOCS'05*, pages 627–636, 2005.
- [KR96] Jon Kleinberg and Ronitt Rubinfeld. Short paths in expander graphs. In *In Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, pages 86–95, 1996.
- [KRV09] Rohit Khandekar, Satish Rao, and Umesh V. Vazirani. Graph partitioning using single commodity flows. *J. ACM*, 56(4), 2009.
- [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 *FOCS*, pages 52–61. IEEE Computer Society, 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.
- [LR99] F. T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM*, 46:787–832, 1999.
- [Mad78] W. Mader. A reduction method for edge connectivity in graphs. *Ann. Discrete Math.*, 3:145–164, 1978.
- [OSVV08] Lorenzo Orecchia, Leonard J. Schulman, Umesh V. Vazirani, and Nisheeth K. Vishnoi. On partitioning graphs via single commodity flows. In *Proceedings of the 40th annual ACM symposium on Theory of computing*, STOC '08, pages 461–470, New York, NY, USA, 2008. ACM.
- [Räc02] Harald Räcke. Minimizing congestion in general networks. In *In Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 43–52, 2002.
- [RS86] Neil Robertson and P D Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, August 1986.
- [RS90] N. Robertson and P. D. Seymour. Outline of a disjoint paths algorithm. In *Paths, Flows and VLSI-Layout*. Springer-Verlag, 1990.
- [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 Rafail Ostrovsky, editor, *FOCS*, pages 200–209. IEEE, 2011.

## A Proof of Theorems 2.2 and 2.3

We start with several definitions and observations that are common to the proofs of both theorems. Both proofs closely follow the arguments of [CKS04b]. Assume that we are given a connected graph  $G$  with maximum vertex degree at most 4, and a set  $\mathcal{T} \subseteq V(G)$  of terminals, such that  $\mathcal{T}$  is  $\alpha$ -well-linked in  $G$ , for some  $0 < \alpha < 1$ . Let  $q$  be the smallest even integer with  $q \geq 4/\alpha$ , so  $4/\alpha \leq q \leq 8/\alpha$ . We assume that  $|\mathcal{T}| \geq q$ .

We use the following simple theorem, whose proof is very similar to the proof of Theorem 2.1, to find an initial clustering of the vertices of  $G$ .

**Theorem A.1** *There is an efficient algorithm that finds a partition  $\mathcal{C}$  of the vertices of  $G$ , where each set  $C \in \mathcal{C}$  contains at least  $q$  and at most  $4q$  terminals. Moreover, each cluster  $C \in \mathcal{C}$  induces a connected sub-graph in  $G$ .*

**Proof:** Let  $T$  be any spanning tree of  $G$ . Our algorithm is iterative. We start with  $\mathcal{C} = \emptyset$ . During the algorithm, we remove some vertices from  $T$ , and we maintain the invariant that  $T$  contains at least  $q$  terminals.

If  $T$  contains at most  $4q$  terminals, then we add all the vertices of  $T$  into  $\mathcal{C}$  as a single cluster, and terminate the algorithm. Otherwise, we find the lowest vertex  $v$  in  $T$ , such that the sub-tree rooted at  $v$  contains at least  $q$  terminals. Notice that since we have assumed that the maximum vertex degree in  $G$  is 4,  $v$  has at most three children, and the subtree of each such child contains fewer than  $q$  terminals. Therefore, the total number of terminals contained in the subtree of  $v$  is at most  $3q$ . We add all the vertices in the sub-tree of  $v$  as a new cluster to  $\mathcal{C}$ , and remove them from  $T$ . Notice that since we assumed that  $T$  contained at least  $4q$  terminals at the beginning of the current iteration, it contains at least  $q$  terminals at the end of the current iteration. We then continue to the next iteration.  $\square$

Let  $\mathcal{H}$  be the set of all sub-graphs induced by the vertices in the subsets in  $\mathcal{C}$ , that is,  $\mathcal{H} = \{G[C] \mid C \in \mathcal{C}\}$ . We will refer to the sub-graphs  $H \in \mathcal{H}$  as *clusters*. Notice that all clusters in  $\mathcal{H}$  are mutually vertex-disjoint, and every vertex belongs to exactly one cluster. For each cluster  $H \in \mathcal{H}$ , we let  $\mathcal{T}(H) \subseteq \mathcal{T}$  be the subset of terminals contained in  $H$ , and for each terminal  $t \in \mathcal{T}$ , we let  $H(t) \in \mathcal{H}$  denote the unique cluster to which  $t$  belongs.

Let  $E' \subseteq E(G)$  be the set of all cut edges of graph  $G$  (recall that  $e \in E(G)$  is a cut edge iff its removal disconnects the graph  $G$ ). We need the following definition.

**Definition 7.** Let  $H$  be any sub-graph of  $G$ ,  $\mathcal{T}(H)$  any subset of terminals contained in  $H$ , and let  $v \in \mathcal{T}(H)$ . We say that  $v$  is a *center* of  $H$  with respect to  $\mathcal{T}(H)$ , iff  $v$  can send one flow unit to the vertices of  $\mathcal{T}(H)$  in graph  $H$ , such that every vertex of  $\mathcal{T}(H)$  receives at most  $1/q$  flow units, and the flow on every edge  $e \in E(H)$  is at most  $1/2$ .

We say that  $v$  is a *pseudo-center* of  $H$  with respect to  $\mathcal{T}(H)$ , iff  $v$  can send one flow unit to the vertices of  $\mathcal{T}(H)$ , with the same restrictions as above, except that the edges  $e \in E'$  are allowed to carry up to one flow unit.

We need the following claim.

**Claim A.1** *Assume that we are given a collection  $\mathcal{G}$  of **edge-disjoint** sub-graphs of  $G$ , a subset  $\mathcal{T}(H) \subseteq \mathcal{T} \cap V(H)$  of terminals for each sub-graph  $H \in \mathcal{G}$ , and a subset  $\mathcal{T}' \subseteq \mathcal{T}$  of terminals, such that:*

- For  $H \neq H'$ ,  $\mathcal{T}(H) \cap \mathcal{T}(H') = \emptyset$ .

- For each  $t \in \mathcal{T}'$ , there is a cluster  $H(t) \in \mathcal{G}$ , such that  $t \in \mathcal{T}(H)$ . Moreover,  $t$  is a pseudo-center of  $H$  with respect to  $\mathcal{T}(H)$ .
- For each  $H \in \mathcal{G}$ ,  $|\mathcal{T}(H) \cap \mathcal{T}'| \leq 1$ .

Then  $\mathcal{T}'$  is 1-well-linked in  $G$ .

**Proof:** Let  $(A, B)$  be any partition of  $V(G)$ . Denote  $\mathcal{T}_A = A \cap \mathcal{T}'$ ,  $\mathcal{T}_B = B \cap \mathcal{T}'$ , and assume without loss of generality that  $|\mathcal{T}_A| \leq |\mathcal{T}_B|$ . In order to show that  $|E(A, B)| \geq |\mathcal{T}_A|$ , it is enough to prove that there is a flow  $F : \mathcal{T}_A \rightsquigarrow_1 \mathcal{T}_B$  in  $G$ . Let  $\mathcal{T}'_B \subseteq \mathcal{T}_B$  be any subset of terminals of size  $|\mathcal{T}_A|$ .

We compute the flow  $F$  in two steps. In the first step, we find a flow  $F' : \mathcal{T}_A \overset{1:1}{\rightsquigarrow} \mathcal{T}'_B$ , where the flow on every edge  $e \in E(G) \setminus E'$  is at most 1, and the flow on every edge  $e \in E'$  is at most 1.5. In the second step, we transform  $F'$  into the desired flow  $F : \mathcal{T}_A \rightsquigarrow_1 \mathcal{T}_B$ .

We start by showing the existence of the flow  $F'$ . Recall that each terminal  $t \in \mathcal{T}'$  is associated with a cluster  $H(t)$ , and there is a flow  $F_t$  inside  $H(t)$ , where  $t$  sends one flow unit to the vertices of  $\mathcal{T}(H(t))$ , and each such vertex receives at most  $1/q$  flow units. The congestion due to flow  $F_t$  on every edge  $e \in E(G) \setminus E'$  is at most  $\frac{1}{2}$ , and the congestion on every edge  $e \in E'$  is at most 1. Since we have selected  $q$  to be an even integer, from the integrality of flow, we can assume that the flow  $F_t$  is  $1/q$ -integral (to do so, set the capacity of every edge  $e \notin E'$  to be  $q/2$ , the capacity of every edge  $e \in E'$  to  $q$ , add a source  $s$ , and connect it to every vertex  $v \in \mathcal{T}(H(t))$  with an edge of capacity 1. The integral  $s$ - $t$  flow of value  $q$  in this network, scaled down by factor  $q$ , gives the desired flow  $F_t$ ). Every vertex  $v \in \mathcal{T}(H(t))$  receives either 0 or  $1/q$  flow units in  $F_t$ . Let  $S'_t \subseteq \mathcal{T}(H(t))$  be the set of exactly  $q$  vertices that receive  $1/q$  flow units in  $F_t$ .

Let  $A' = \bigcup_{t \in \mathcal{T}_A} S'_t$ , and let  $B' = \bigcup_{t \in \mathcal{T}'_B} S'_t$ . Notice that  $|A'| = q|\mathcal{T}_A| = q|\mathcal{T}'_B| = |B'|$ , since the terminal sets  $\mathcal{T}(H), \mathcal{T}(H')$  are disjoint for  $H \neq H'$ . Since both sets  $A'$  and  $B'$  only contain terminals, and the terminals are  $\alpha$ -well-linked in  $G$ , there is a flow  $F'' : A' \overset{1:1}{\rightsquigarrow}_{2/\alpha} B'$  in graph  $G$ . Scaling this flow down by factor  $q \geq 4/\alpha$ , we obtain a flow  $F^*$ , where every vertex  $v \in A'$  sends  $1/q$  flow units, every vertex  $v \in B'$  receives  $1/q$  flow units, and the total edge congestion is at most  $1/2$ . Let flow  $F'$  be the concatenation of  $\bigcup_{t \in \mathcal{T}_A} F_t$ , the flow  $F^*$ , and  $\bigcup_{t \in \mathcal{T}'_B} F_t$ . Then every vertex  $v \in \mathcal{T}_A$  sends one flow unit in  $F'$ , every vertex  $v \in \mathcal{T}'_B$  receives one flow unit, the flow on every edge  $e \notin E'$  is at most 1, and the flow on every edge  $e \in E'$  is at most 1.5. We assume without loss of generality that flow  $F'$  is half-integral. That is, we can decompose  $F'$  into flow-paths, each of which carries exactly  $1/2$  flow units. Notice that every vertex is an endpoint of either 0 or 2 such flow-paths.

We now transform the flow  $F'$  into a flow  $F : \mathcal{T}_A \rightsquigarrow_1 \mathcal{T}_B$ . This part is similar to the proof of Lemma 3.10 in [CKS04b]. The transformation replaces every flow-path in  $F'$  by a corresponding simple flow-path, by removing all cycles from the path, if such exist. Let  $F$  be the resulting flow. It is now enough to show that every edge  $e \in E'$  carries at most one flow unit. Let  $e \in E'$  be any cut edge, and assume for contradiction that more than one flow unit is sent via  $e$ . Let  $X, Y$  be the two connected components of  $G \setminus \{e\}$ . Since the amount of flow on  $e$  is bounded by 1.5 and we assume that all flow-paths are simple and carry exactly  $1/2$  flow units each, there are three flow-paths that use edge  $e$ , and each such path connects a vertex of  $X$  to a vertex of  $Y$ .

For every vertex  $v \in X$ , let  $n_1(v)$  be the total number of flow-paths originating or terminating at  $v$ , and let  $n_1 = \sum_{v \in X} n_1(v)$ . Since  $n_1(v) \in \{0, 2\}$  for each  $v$ ,  $n_1$  is an even integer. Let  $n_2(v)$  be the total number of flow-paths originating or terminating at  $v$ , whose other endpoint also belongs to  $X$ , and let  $n_2 = \sum_{v \in X} n_2(v)$ . Since every flow-path contained in  $X$  contributes 2 to  $n_2$ ,  $n_2$  is also an even integer. However, since there are exactly three flow-paths that start at  $X$  and terminate at  $Y$ ,  $n_1 - n_2 = 3$ , which is impossible.  $\square$

## Proof of Theorem 2.2

The proof roughly follows the arguments of [CKS04b], and consists of three steps. In the first step, we define an initial clustering of the vertices of  $G$  and select one representative terminal from each cluster. If we could claim at this point that each selected representative is a pseudo-center of its cluster, then we would be done by Claim A.1. However, this is not necessarily the case. To overcome this difficulty, we use a procedure suggested by [CKS04b] called *tagging*, where each cluster  $C$  “tags” some other cluster  $C'$ , such that the representative of  $C$  is a pseudo-center of the merged cluster  $C \cup C'$ . The tagging procedure is performed in Step 2. Finally, in Step 3, we merge some pairs of clusters and select one pseudo-center for each such merged cluster.

**Step 1: Initial Clustering and Representative Selection** We use Theorem A.1 to find an initial clustering  $\mathcal{H}$  of the vertices of  $G$ . Next, we select a representative for each cluster  $H \in \mathcal{H}$ , using the following lemma.

**Lemma A.1** *There is an efficient randomized algorithm that with high probability finds, for each  $1 \leq j \leq r$ , a subset  $\tilde{\mathcal{T}}_j \subseteq \mathcal{T}_j$  of terminals, such that  $|\tilde{\mathcal{T}}_j| \geq \Omega(\alpha|\mathcal{T}_j|/r)$ , the sets  $\{\tilde{\mathcal{T}}_j\}_{j=1}^r$  are mutually disjoint, and for each cluster  $H \in \mathcal{H}$ ,  $\mathcal{T}(H) \cap \left(\bigcup_{j=1}^r \tilde{\mathcal{T}}_j\right) = 1$ .*

**Proof:** We start by selecting, for each  $1 \leq j \leq r$ , at most one representative  $t \in \mathcal{T}(H) \cap \mathcal{T}_j$  for every cluster  $H$ , and we will later ensure that at most one representative is chosen for each cluster overall.

Fix some  $1 \leq j \leq r$ . We construct an initial set  $\tilde{\mathcal{T}}'_j \subseteq \mathcal{T}_j$  of representatives, as follows. Start with  $\tilde{\mathcal{T}}'_j = \emptyset$ . While  $\mathcal{T}_j \neq \emptyset$ , select any terminal  $t \in \mathcal{T}_j$  and add it to  $\tilde{\mathcal{T}}'_j$ . Remove from  $\mathcal{T}_j$  all terminals that belong to  $\mathcal{T}(H(t))$ , and continue to the next iteration. Let  $\tilde{\mathcal{T}}'_j$  be the final subset of terminals. Then  $|\tilde{\mathcal{T}}'_j| \geq \frac{|\mathcal{T}_j|}{4q}$ , since for each terminal added to  $\tilde{\mathcal{T}}'_j$ , at most  $4q$  terminals are removed from  $\mathcal{T}_j$ . Moreover, for each cluster  $H \in \mathcal{H}$ , at most one terminal from  $\mathcal{T}(H)$  belongs to  $\tilde{\mathcal{T}}'_j$ .

Consider now some cluster  $H \in \mathcal{H}$ , and let  $\mathcal{R}(H)$  be a multi-set, containing, for each  $1 \leq j \leq r$ , the unique terminal in  $\mathcal{T}(H) \cap \tilde{\mathcal{T}}'_j$ , if it exists (if some terminal  $t \in \mathcal{T}(H)$  belongs to several sets  $\tilde{\mathcal{T}}'_j$ , then it appears several times in  $\mathcal{R}(H)$ ). Since  $\mathcal{R}(H)$  contains at most one representative from each set  $\tilde{\mathcal{T}}'_j$ ,  $|\mathcal{R}(H)| \leq r$ . Cluster  $H$  then selects one terminal  $t \in \mathcal{R}(H)$  uniformly at random, and terminal  $t$  becomes the *representative of the cluster  $H$* .

For each  $1 \leq j \leq r$ , let  $\tilde{\mathcal{T}}_j \subseteq \tilde{\mathcal{T}}'_j$  be the subset of terminals that serve as representatives of clusters in  $\mathcal{H}$ . Then the sets  $\{\tilde{\mathcal{T}}_j\}_{j=1}^r$  are mutually disjoint, and each cluster  $H \in \mathcal{H}$  has exactly one terminal in  $\mathcal{T}(H) \cap \left(\bigcup_{j=1}^r \tilde{\mathcal{T}}_j\right)$ . It now only remains to show that for each  $1 \leq j \leq r$ ,  $|\tilde{\mathcal{T}}_j| \geq \Omega(\alpha|\mathcal{T}_j|/r)$  with high probability.

Fix some  $1 \leq j \leq r$ , and let  $\mu_j$  denote the expectation of  $|\tilde{\mathcal{T}}'_j|$ . Then  $\mu_j \geq |\tilde{\mathcal{T}}'_j|/r \geq |\mathcal{T}_j|/(4qr) = \Omega(\alpha|\mathcal{T}_j|/r) = \Omega(r \log r)$ . From the Chernoff bound,  $\Pr \left[|\tilde{\mathcal{T}}'_j| < \mu_j/2\right] \leq e^{-\mu_j/12} \leq 1/r^r$ . Using the union bound, the probability that  $|\tilde{\mathcal{T}}'_j| \geq \Omega(\alpha|\mathcal{T}_j|/r)$  for all  $1 \leq j \leq r$  is at least  $(1 - 1/r^{r-1})$ . By repeating the procedure a polynomial number of times, we can ensure that this algorithm succeeds with high probability.  $\square$

For each cluster  $H \in \mathcal{H}$ , the unique terminal in  $\mathcal{T}(H) \cap \left(\bigcup_{j=1}^r \tilde{\mathcal{T}}_j\right)$  is called the *representative of  $H$* . So far we have constructed the sets  $\tilde{\mathcal{T}}_j$ , for  $1 \leq j \leq r$  of terminals that have all desired properties, except that we are not guaranteed that  $\bigcup_{j=1}^r \tilde{\mathcal{T}}_j$  is 1-well linked in  $G$ . If we could guarantee that each terminal

$t \in \bigcup_{j=1}^r \tilde{\mathcal{T}}_j$  is a pseudo-center of its cluster  $H(t)$ , then this property would follow from Claim A.1. However, we cannot guarantee this property at this stage. We overcome this difficulty in the next step, using the tagging technique of [CKS04b]. For each  $1 \leq j \leq r$ , we denote  $\tilde{k}_j = |\tilde{\mathcal{T}}_j| = \Omega(\alpha|\mathcal{T}_j|/r)$ .

**Step 2: Tagging** The idea of the tagging technique is that we merge some pairs of clusters in  $\mathcal{H}$ , and select one representative for each such merged cluster (from among the two representatives of the merged clusters), in a way that ensures that these representatives are pseudo-centers of their new clusters. We use the following result of [CKS04b].

**Theorem A.2** (Lemma 3.13 in [CKS04b]) *Let  $H \in \mathcal{H}$  be any cluster, and let  $s \in V(H)$  be any vertex that is not a pseudo-center of  $H$ . Then there is an edge  $e = (x, y)$  with  $x \in V(H)$ ,  $y \notin V(H)$ , such that for any path  $P$ , connecting  $H$  to another cluster  $H' \in \mathcal{H}$ , where the first edge of  $P$  is  $e$ , and  $P$  is edge-disjoint from  $H$ , vertex  $s$  is a pseudo-center of the cluster  $H \cup P \cup H'$  with respect to  $\mathcal{T}(H) \cup \mathcal{T}(H')$ . Moreover, the edge  $e = (x, y)$  can be found efficiently.*

For every terminal  $t \in \bigcup_{j=1}^r \tilde{\mathcal{T}}_j$ , if  $t$  is not a pseudo-center of  $H(t)$ , let  $e_t = (x_t, y_t)$  be the corresponding edge guaranteed by Theorem A.2, and let  $H'(t)$  be the cluster of  $\mathcal{H}$  containing  $y_t$  (such a cluster exists since every vertex of  $G$  belongs to one of the clusters in  $\mathcal{H}$ , from Theorem A.1). We then say that the terminal  $t$  and the cluster  $H(t)$  tag the cluster  $H'(t)$ .

Intuitively, we can now create a new cluster, by merging  $H(t)$  and  $H'(t)$ . From Theorem A.2, we are guaranteed that  $t$  is a pseudo-center of the new cluster. However, since we require that the clusters are disjoint and every cluster only has one representative, the representative of the cluster  $H'(t)$  will need to be discarded. While in general we only expect to discard a constant fraction of the terminals, it is possible that many of these terminals belong to one of the sets  $\tilde{\mathcal{T}}_j$ , and we will end up discarding too many terminals from one such set. Another difficulty is when many terminals tag the same cluster  $H'$ . In this case, following the approach of [CKS04b], we can partition all such terminals into pairs, and connect their corresponding clusters by edge-disjoint paths that are contained in  $H'$ . For each such new merged cluster, we can select one of the two representative terminals that is guaranteed to be a pseudo-center by Theorem A.2. Again, we will need to discard some clusters (such as the cluster  $H'$  in this case), and we need to ensure that we do not discard too many representatives from any set  $\tilde{\mathcal{T}}_j$ .

We start by building a collection  $\mathcal{H}^* \subseteq \mathcal{H}$  of clusters, such that for any pair  $(H, H')$  of clusters, where  $H$  tags  $H'$ , at most one of these clusters may belong to  $\mathcal{H}^*$ . We also ensure that among the representatives of the clusters in  $\mathcal{H}^*$ , enough terminals belong to each set  $\tilde{\mathcal{T}}_j$ .

For convenience, for each  $1 \leq j \leq r$ , we say that all terminals in  $\tilde{\mathcal{T}}_j$  are of color  $j$ . For each terminal  $t \in \tilde{\mathcal{T}}_j$ , the color of the corresponding cluster  $H(t)$  is also  $j$ . Therefore, we have at least  $\tilde{k}_j$  clusters of each color  $j$ . Following the arguments of [CKS04b], we build a directed graph  $D$ , whose vertex set corresponds to the clusters in  $\mathcal{H}$ , that is,  $V(D) = \{v_H \mid H \in \mathcal{H}\}$ , and there is an edge  $(v_H, v_{H'})$  iff  $H$  tags  $H'$ . We say that a vertex  $v_H$  has color  $j$  iff the color of  $H$  is  $j$ . As in [CKS04b], we need to find a large independent set  $\mathcal{I}$  in graph  $D$ . However, we also need to ensure that  $\mathcal{I}$  contains  $\Omega(\tilde{k}_j/r)$  vertices of each color  $j$ .

**Claim A.2** *There is an efficient algorithm that finds an independent set  $\mathcal{I}$  in graph  $D$ , containing at least  $\Omega(\tilde{k}_j/r)$  vertices of each color  $1 \leq j \leq r$ .*

**Proof:** For each  $1 \leq j \leq r$ , let  $k'_j$  denote the number of vertices of color  $j$  in graph  $D$ . We assume without loss of generality that  $k'_2 = k'_3 = \dots = k'_r = \Omega(\alpha k_2/r)$ , and  $k'_1 = \Omega(\alpha k_1/r)$ , by discarding, if necessary, some vertices from the graph  $D$ . We consider two cases.

Assume first that  $k'_1 \leq rk'_2$ . Let  $\beta = rk'_2/k'_1$ . Notice that the out-degree of every vertex in  $D$  is at most 1. We say that a vertex  $v_H$  of color 1 is good, iff the total number of edges adjacent to  $v_H$ , whose other endpoint's color is different from 1, is at most  $8\beta$ . Since the total number of vertices in the graph  $D$  is at most  $rk'_2 + k'_1 \leq 2rk'_2 = 2\beta k'_1$ , the total number of edges is also at most  $2\beta k'_1$ . Therefore, at least half the vertices of color 1 are good. Let  $D'$  be the sub-graph of  $D$  induced by the good color-1 vertices. Since the out-degree of every vertex in  $D'$  is at most 1, and  $D'$  contains at least  $k'_1/2$  vertices, we can efficiently find an independent set  $\mathcal{I}_1$  of size exactly  $\lfloor \frac{k'_1}{16r} \rfloor$  in graph  $D'$ , by iteratively greedily selecting a vertex with minimum total degree. The selected vertex  $v$  is then added to  $\mathcal{I}_1$ , and it is deleted from  $D'$ , together with all its neighbors. Let  $\mathcal{I}_1$  be the resulting independent set of size  $\lfloor \frac{k'_1}{16r} \rfloor$ , and let  $V'$  be the subset of the vertices of colors  $2, \dots, r$ , whose neighbors belong to  $\mathcal{I}_1$ . Then  $|V'| \leq \lfloor \frac{k'_1}{16r} \rfloor \cdot 8\beta \leq \frac{k'_2}{2}$ .

Let  $D''$  be the sub-graph of  $D$  induced by the vertices of colors  $2, \dots, r$ , that do not belong to  $V'$ . Then for each  $2 \leq j \leq r$ ,  $D''$  contains at least  $k'_2/2$  vertices of color  $j$ . Next we construct, for each  $2 \leq j \leq r$ , a set  $\mathcal{I}_j$  of  $\lfloor \frac{k'_2}{6r} \rfloor$  vertices of color  $j$ , such that  $\bigcup_{j=1}^r \mathcal{I}_j$  is an independent set in graph  $D$ . We start with  $\mathcal{I}_j = \emptyset$  for all  $2 \leq j \leq r$ , and then iterate. In every iteration, we select a minimum-degree vertex  $v$  from  $D''$ , and add it to the set  $\mathcal{I}_j$ , where  $j$  is the color of  $v$ . We then delete  $v$  and all its neighbors from  $D''$ . If, at any point of the algorithm execution, the size of the set  $\mathcal{I}_j$ , for any  $2 \leq j \leq r$ , reaches  $\lfloor \frac{k'_2}{6r} \rfloor$ , then we delete all the remaining vertices of color  $j$  from  $D''$ , and continue. Notice that since the out-degree of every vertex is at most 1 in any induced sub-graph of  $D''$ , we can always find a vertex whose degree is at most 2 in the current graph. Therefore, in each iteration, at most 3 vertices are deleted from  $D''$ .

It is easy to see that at the end of the algorithm, for each  $2 \leq j \leq r$ ,  $|\mathcal{I}_j| = \lfloor \frac{k'_2}{6r} \rfloor$ . Indeed, assume otherwise, and assume that for some  $j^*$ , set  $\mathcal{I}_{j^*}$  contains fewer than  $\lfloor \frac{k'_2}{6r} \rfloor$  vertices upon the termination of the algorithm. But the original graph  $D''$  contained at least  $k'_2/2$  vertices of color  $j^*$ , the number of iterations is bounded by  $(r-1) \cdot \lfloor \frac{k'_2}{6r} \rfloor < \frac{k'_2}{6}$ , and in each iteration at most 3 vertices of color  $j^*$  are deleted from  $D''$ . Therefore, upon the termination of the algorithm, at least one vertex of color  $j^*$  must be present in the graph, a contradiction. The output of the algorithm is  $\bigcup_{j=1}^r \mathcal{I}_j$ .

Consider now the second case, where  $k'_1 > rk'_2$ . Let  $\beta = \frac{k'_1}{rk'_2}$ . We say that a vertex of colors  $2, \dots, r$  is good if it has at most  $8\beta r$  neighbors of color 1. Since the total number of edges in the graph is bounded by  $k'_1 + rk'_2 \leq 2k'_1 = 2\beta rk'_2$ , for every color  $2 \leq j \leq r$ , at least half the vertices of color  $j$  are good. We construct a sub-graph  $D'$  of graph  $D$ , induced by all good vertices of colors  $2, \dots, r$ . Since graph  $D'$  contains at least  $k'_2/2$  vertices of each color, using the same algorithm as before, we can find, for each  $2 \leq j \leq r$ , a collection  $\mathcal{I}_j$  of  $\lfloor \frac{k'_2}{16r} \rfloor$  good color- $j$  vertices, such that the set  $\mathcal{I}' = \bigcup_{j=1}^r \mathcal{I}_j$  is an independent set in  $D'$ . Let  $V'$  be the subset of color-1 vertices that have neighbors in  $\mathcal{I}'$ . Then, since all vertices in  $\mathcal{I}'$  are good,  $|V'| \leq r \cdot \frac{k'_2}{16r} \cdot 8\beta r = \frac{\beta k'_2 r}{2} = \frac{k'_1}{2}$ . Let  $V''$  be the subset of color-1 vertices that do not belong to  $V'$ , and let  $D''$  be the subgraph of  $D$  induced by  $V''$ . Then  $D''$  contains at least  $k'_1/2$  vertices, and using the same algorithm as before, we can find an independent set  $\mathcal{I}_1$  of size at least  $k'_1/8$  in graph  $D''$ . We then output  $\bigcup_{j=1}^r \mathcal{I}_j$ , which is guaranteed to be an independent set.  $\square$

Let  $\mathcal{H}^* \subseteq \mathcal{H}$  contain the set of clusters  $H$  whose corresponding vertex  $v_H \in \mathcal{I}$ . For each  $1 \leq j \leq r$ , let  $\tilde{\mathcal{T}}_j^* \subseteq \tilde{\mathcal{T}}_j$  denote the sets of terminals  $t$  where  $H(t) \in \mathcal{H}^*$ . Then for each  $1 \leq j \leq r$ ,  $|\tilde{\mathcal{T}}_j^*| = \Omega(\alpha |\mathcal{T}_j|/r^2)$  with high probability.

**Step 3: Merging the Clusters** In this step we merge some pairs of clusters, and select a representative for each new merged cluster, such that, on the one hand, this representative is a pseudo-center of the merged cluster, and on the other hand, we obtain enough representatives from each set  $\tilde{\mathcal{T}}_j^*$  of

terminals.

We construct a new collection  $\mathcal{G}$  of merged clusters, and the final collection  $\mathcal{T}' = \bigcup_{j=1}^r \mathcal{T}'_j$  of representative terminals, where for each  $1 \leq j \leq r$ ,  $\mathcal{T}'_j \subseteq \tilde{\mathcal{T}}_j^*$ . We start with  $\mathcal{G} = \emptyset$  and  $\mathcal{T}' = \emptyset$ .

We add to  $\mathcal{G}$  all clusters  $H \in \mathcal{H}^*$  that do not tag any cluster. Recall that in this case, the unique terminal  $t \in \mathcal{T}(H) \cap \left(\bigcup_{j=1}^r \tilde{\mathcal{T}}_j^*\right)$  is a pseudo-center of  $H$ . We add  $t$  to  $\mathcal{T}'$ .

Consider now some cluster  $H \in \mathcal{H}$ , and let  $\mathcal{G}(H) \subseteq \mathcal{H}^*$  be the set of all clusters in  $\mathcal{H}^*$  that tag  $H$ . If  $|\mathcal{G}(H)| = 1$ , then let  $H' \in \mathcal{H}^*$  be the unique cluster that tags  $H$ . Let  $t' \in \bigcup_{j=1}^r \tilde{\mathcal{T}}_j^*$  be the representative of  $H'$ , that is,  $t' \in \mathcal{T}(H')$ . We then create a new merged cluster,  $H^* = H \cup H' \cup e_t(H')$ , which is added to  $\mathcal{G}$ . We set  $\mathcal{T}(H^*) = \mathcal{T}(H) \cup \mathcal{T}(H')$ . By Theorem A.2, we are guaranteed that  $t'$  is a pseudo-center of  $H^*$ . We add  $t'$  to  $\mathcal{T}'$ .

Assume now that  $|\mathcal{G}(H)| > 1$ . We use the following lemma from [CKS04b].

**Lemma A.2** (Lemma 3.14 in [CKS04b]) *Let  $T$  be a tree and  $A$  some even multi-set of vertices in  $V(T)$ . Then we can efficiently find  $|A|/2$  edge-disjoint paths in  $T$ , such that each vertex  $v \in A$  is the endpoint of exactly  $n_v$  of these paths, where  $v$  occurs  $n_v$  times in  $A$ .*

Let  $\mathcal{G}(H) = \{H_1, \dots, H_p\}$ , and for each  $1 \leq i \leq p$ , let  $t_i \in \mathcal{T}(H_i)$  be the unique terminal that belongs to  $\bigcup_{j=1}^r \tilde{\mathcal{T}}_j^*$ . If  $p$  is odd, then we choose one of the clusters  $H_1, \dots, H_p$  uniformly at random, and discard it from  $\mathcal{H}^*$ . Notice that  $p \geq 3$  must hold in this case, so the probability that any cluster is discarded is at most  $1/3$ . We assume from now on that  $p$  is even. Let  $A = \{y_{t_1}, y_{t_2}, \dots, y_{t_p}\}$ . We use Lemma A.2 to find a collection  $\mathcal{P}$  of edge-disjoint paths contained in  $H$  (by first building a spanning tree  $T$  of  $H$ ), where every path connects a distinct pair  $y_{t_i}, y_{t_j}$  of vertices. For each such path  $P \in \mathcal{P}$ , if  $P$  connects  $y_{t_i}$  to  $y_{t_j}$ , then we construct a new cluster  $H^* = H_i \cup P \cup H_j$ , and add it to  $\mathcal{G}$ . We set  $\mathcal{T}(H^*) = \mathcal{T}(H_i) \cup \mathcal{T}(H_j)$ . Notice that from Theorem A.2, both  $t_i$  and  $t_j$  are pseudo-centers of  $H^*$ . We select one of these two terminals uniformly at random, and add it to  $\mathcal{T}'$ .

Consider the final set  $\mathcal{G}$  of clusters. From the construction, we are guaranteed that all these clusters are edge-disjoint, since the original clusters  $\mathcal{H}$  were vertex-disjoint. Every terminal  $t \in \mathcal{T}'$  is guaranteed to be a pseudo-center of its cluster in  $\mathcal{G}$ . It is now enough to show that for each  $1 \leq j \leq r$ , at least  $|\tilde{\mathcal{T}}_j^*|/12 = \Omega(\alpha|\mathcal{T}_j|/r^2)$  terminals from  $\tilde{\mathcal{T}}_j^*$  belong to  $\mathcal{T}'$  with high probability.

Fix some  $1 \leq j \leq r$ , let  $t \in \tilde{\mathcal{T}}_j^*$  be any terminal, and let  $H(t) \in \mathcal{H}^*$  be the cluster that  $t$  represents. Recall that  $H(t)$  may be discarded from  $\mathcal{H}^*$  with probability at most  $1/3$  (if  $H(t)$  tags some other cluster  $H'$ , for which the parameter  $p$  is an odd number greater than 1), and additionally, if  $H(t)$  is merged with another cluster  $H'$ , then  $t$  is added to  $\mathcal{T}'$  with probability at least  $\frac{1}{2}$ . Using the standard Chernoff bound, the probability that more than a  $2/3$ -fraction of clusters of color  $j$  are discarded from  $\mathcal{H}^*$  is at most  $1/\text{poly}(r)$ , and similarly, the probability that fewer than  $\frac{|\tilde{\mathcal{T}}_j^*|}{12}$  terminals from  $\tilde{\mathcal{T}}_j^*$  are added to  $\mathcal{T}'$  is bounded by  $1/\text{poly}(r)$ , since  $|\tilde{\mathcal{T}}_j^*| \geq \Omega(\alpha|\mathcal{T}_j|/r^2) = \Omega(\log r)$ . Using the union bound over all  $1 \leq j \leq r$ , with probability at least  $\frac{1}{2}$ , set  $\mathcal{T}'$  contains at least  $|\tilde{\mathcal{T}}_j^*|/12$  terminals from  $\tilde{\mathcal{T}}_j$ , for all  $1 \leq j \leq r$ . By repeating this procedure a polynomial number of times, we can ensure that it succeeds with high probability. Applying Claim A.1 to the set  $\mathcal{G}$  of clusters and the set  $\mathcal{T}'$  of terminals, we conclude that  $\mathcal{T}'$  is 1-well-linked in  $G$ .

### Proof of Theorem 2.3

The proof closely follows the proof of Theorem 2.2, except that now we are only given two subsets  $\mathcal{T}_1, \mathcal{T}_2$  of terminals, and we need to ensure that the terminals in the set  $\mathcal{T}'_1$  are chosen so that for every pair  $(s, t) \in \mathcal{M}$ , either both  $s, t$  belong to  $\mathcal{T}'_1$ , or none of them.



**Step 1: Initial Clustering and Representative Selection** As before, we use Theorem A.1 to find an initial clustering  $\mathcal{H}$ . We then select representatives for the clusters  $H \in \mathcal{H}$ , using the following analogue of Lemma A.1.

**Lemma A.3** *There is an efficient algorithm, that either routes a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of  $\Omega(\alpha k_1)$  pairs on edge-disjoint paths in  $G$ , or finds two subsets  $\tilde{\mathcal{T}}_1 \subseteq \mathcal{T}_1$  and  $\tilde{\mathcal{T}}_2 \subseteq \mathcal{T}_2$ , with  $|\tilde{\mathcal{T}}_1| = \Omega(\alpha k_1)$ ,  $|\tilde{\mathcal{T}}_2| = \Omega(\alpha k_2)$ , such that:*

- $\tilde{\mathcal{T}}_1 \cap \tilde{\mathcal{T}}_2 = \emptyset$ ,
- for each cluster  $H \in \mathcal{H}$ ,  $|\mathcal{T}(H) \cap (\tilde{\mathcal{T}}_1 \cup \tilde{\mathcal{T}}_2)| \leq 1$ , and
- for each pair  $(s, t) \in \mathcal{M}$ , either both  $s, t$  belong to  $\tilde{\mathcal{T}}_1$ , or neither of them does.

**Proof:** We start by selecting a subset  $\tilde{\mathcal{T}}_2' \subseteq \mathcal{T}_2$  of representatives from the set  $\mathcal{T}_2$ , exactly as in the proof of Lemma A.1. We construct a collection  $\tilde{\mathcal{T}}_2' \subseteq \mathcal{T}_2$  of at least  $\frac{k_2}{4q}$  terminals, such that for each cluster  $H \in \mathcal{H}$ , at most one terminal of  $\mathcal{T}(H)$  belongs to  $\tilde{\mathcal{T}}_2'$ .

The selection of representatives for  $\mathcal{T}_1$  is performed slightly differently. We construct a new matching  $\mathcal{M}_1 \subseteq \mathcal{M}$ , that consists of the selected representatives. We start with  $\mathcal{M}_1 = \emptyset$ . While  $\mathcal{M} \neq \emptyset$ , let  $(s, t) \in \mathcal{M}$  be any pair of terminals in  $\mathcal{M}$ . Add  $(s, t)$  to  $\mathcal{M}_1$ , and remove from  $\mathcal{T}_1$  all terminals contained in  $\mathcal{T}(H(s)) \cup \mathcal{T}(H(t))$ . Remove from  $\mathcal{M}$  all pairs that are not contained in the current set  $\mathcal{T}_1$  of terminals. Let  $\mathcal{M}_1$  be the final matching that the algorithm has constructed. It is easy to verify that  $|\mathcal{M}_1| \geq \frac{k_1}{16q}$ , since for every pair that we add to  $\mathcal{M}_1$ , we remove at most  $8q$  pairs from  $\mathcal{M}$ . Let  $\tilde{\mathcal{T}}_1' \subseteq \mathcal{T}_1$  be the subset of terminals that participate in pairs in  $\mathcal{M}_1$ . Then for every cluster  $H \in \mathcal{H}$ , either  $|\mathcal{T}(H) \cap \tilde{\mathcal{T}}_1'| \in \{0, 1\}$ , or  $|\mathcal{T}(H) \cap \tilde{\mathcal{T}}_1'| = 2$ . In the latter case, the two terminals  $s, t \in \mathcal{T}(H) \cap \tilde{\mathcal{T}}_1'$  belong to  $\mathcal{M}$  as a pair.

Partition  $\mathcal{M}_1$  into two subsets  $\mathcal{M}'_1$  and  $\mathcal{M}''_1$ , where  $\mathcal{M}'_1$  contains all pairs  $(s, t)$  where both  $s, t \in \mathcal{T}(H)$  for some cluster  $\mathcal{T}(H)$ , and  $\mathcal{M}''_1$  contains all remaining pairs. If  $|\mathcal{M}'_1| \geq \frac{1}{2}|\mathcal{M}_1|$ , then we have found a collection  $\mathcal{M}'_1 \subseteq \mathcal{M}$  of  $\Omega(\alpha k_1)$  pairs, where each pair  $(s, t) \in \mathcal{M}'_1$  is associated with a distinct cluster  $H \in \mathcal{H}$ , where  $s, t \in \mathcal{T}(H)$ . Since each cluster  $H \in \mathcal{H}$  is connected, and the clusters are mutually edge-disjoint, we can route the pairs in  $\mathcal{M}'_1$  on edge-disjoint paths. Therefore, we assume from now on, that  $|\mathcal{M}''_1| \geq |\mathcal{M}_1|/2 \geq \frac{k_1}{32q}$ . To simplify the notation, we denote  $\mathcal{M}''_1$  by  $\mathcal{M}_1$  from now on. We denote by  $\tilde{\mathcal{T}}_1'$  the set of terminals participating in the pairs in  $\mathcal{M}_1$ , so  $|\tilde{\mathcal{T}}_1'| = \Omega(\alpha k_1)$ .

Notice that now for every cluster  $H \in \mathcal{H}$ , at most one terminal  $t \in \mathcal{T}(H)$  belongs to  $\tilde{\mathcal{T}}_1'$ , and at most one terminal  $t' \in \mathcal{T}(H)$  belongs to  $\tilde{\mathcal{T}}_2'$ , while it is possible that  $t = t'$ . If  $\mathcal{T}(H)$  contains both a terminal of  $\tilde{\mathcal{T}}_1'$  and a terminal of  $\tilde{\mathcal{T}}_2'$ , then we say that  $H$  is a mixed cluster. Let  $\mathcal{H}'$  be the set of all mixed clusters.

We partition  $\mathcal{H}'$  into two subsets  $\mathcal{H}'_1, \mathcal{H}'_2$ , such that  $|\mathcal{H}'_1|, |\mathcal{H}'_2| \geq |\mathcal{H}'|/2 - 1$ . Our partition ensures that if, for some pair  $(s, t) \in \mathcal{M}_1$ , there are two clusters  $H, H' \in \mathcal{H}'$  with  $s \in \mathcal{T}(H)$  and  $t \in \mathcal{T}(H')$ , then either both  $H, H' \in \mathcal{H}'_1$ , or both  $H, H' \in \mathcal{H}'_2$ . The partition is performed by using the following greedy procedure. Start with  $\mathcal{H}'_1, \mathcal{H}'_2 = \emptyset$ . While  $\mathcal{H}' \neq \emptyset$ , let  $H \in \mathcal{H}'$  be any cluster, and let  $s$  be the unique terminal in  $\mathcal{T}(H) \cap \tilde{\mathcal{T}}_1'$ . Let  $t$  be the terminal such that  $(s, t) \in \mathcal{M}_1$ , and let  $H' \in \mathcal{H}'$  be the cluster for which  $t \in \mathcal{T}(H')$ . If  $H' \notin \mathcal{H}'$ , then we add  $H$  to either  $\mathcal{H}'_1$  or  $\mathcal{H}'_2$  - whichever currently contains fewer clusters. Otherwise, we add both  $H$  and  $H'$  to one of the two sets  $\mathcal{H}'_1, \mathcal{H}'_2$ , that contains fewer clusters. The ties are broken arbitrarily. Let  $(\mathcal{H}'_1, \mathcal{H}'_2)$  be the resulting partition of  $\mathcal{H}'$ .

For each cluster  $H \in \mathcal{H}'_1$ , we remove the unique terminal  $t \in \mathcal{T}(H) \cap \tilde{\mathcal{T}}_2'$  from  $\tilde{\mathcal{T}}_2'$ . For each cluster  $H \in \mathcal{H}'_2$ , let  $t$  be the unique terminal in  $\mathcal{T}(H) \cap \tilde{\mathcal{T}}_1'$ , and let  $s \in \tilde{\mathcal{T}}_1'$  such that  $(s, t) \in \mathcal{M}_1$ . We remove

$(s, t)$  from  $\mathcal{M}_1$ , and we remove both  $s$  and  $t$  from  $\tilde{\mathcal{T}}_1'$ . Notice that since  $|\mathcal{H}'_1| \geq |\mathcal{H}'_2| - 2$ , we only remove a constant fraction of the terminals from  $\tilde{\mathcal{T}}_1'$ . Therefore, at the end of this procedure, we obtain a subset  $\tilde{\mathcal{M}} \subseteq \mathcal{M}$  of pairs of terminals, with  $|\tilde{\mathcal{M}}| = \Omega(\alpha k_1)$ , a set  $\tilde{\mathcal{T}}_1$  of terminals participating in the pairs in  $\tilde{\mathcal{M}}$ , and another set  $\tilde{\mathcal{T}}_2 \subseteq \mathcal{T}_2$  of terminals, with  $|\tilde{\mathcal{T}}_2| \geq \Omega(\alpha k_2)$ , such that  $\tilde{\mathcal{T}}_1 \cap \tilde{\mathcal{T}}_2 = \emptyset$ , and for every cluster  $H \in \mathcal{H}$ ,  $\mathcal{T}(H)$  contains at most one terminal  $t \in \tilde{\mathcal{T}}_1 \cup \tilde{\mathcal{T}}_2$ .  $\square$

From now on, we denote by  $\tilde{\mathcal{M}} \subseteq \mathcal{M}$  the collection of terminal pairs contained in  $\tilde{\mathcal{T}}_1$ . For each cluster  $H \in \mathcal{H}$ , we say that the unique terminal  $t \in \mathcal{T}(H) \cap (\tilde{\mathcal{T}}_1 \cup \tilde{\mathcal{T}}_2)$  is the *representative* of  $H$ . Notice that some clusters  $H \in \mathcal{H}$  may have no representative.

**Step 2: Tagging** The tagging is performed exactly as before. Consider some cluster  $H \in \mathcal{H}$ , that has a representative  $t \in \mathcal{T}(H) \cap (\tilde{\mathcal{T}}_1 \cup \tilde{\mathcal{T}}_2)$ . If  $t$  is not a pseudo-center of  $H$ , then we can find an edge  $e_t = (x_t, y_t)$  with  $x_t \in V(H), y_t \notin V(H)$ , guaranteed by Theorem A.2. Let  $H' \in \mathcal{H}$  be a cluster containing the vertex  $y_t$  (again, such a cluster must exist since every vertex belongs to some cluster by Theorem A.2). We say that cluster  $H$  and terminal  $t$  *tag* the cluster  $H'$ .

We say that a pair  $(s, t) \in \tilde{\mathcal{M}}$  is good if  $H(s)$  tags  $H(t)$  or vice versa. Let  $\tilde{\mathcal{M}}' \subseteq \tilde{\mathcal{M}}$  be the set of all good pairs, and let  $\tilde{\mathcal{M}}'' = \tilde{\mathcal{M}} \setminus \tilde{\mathcal{M}}'$ . If  $|\tilde{\mathcal{M}}'| \geq |\tilde{\mathcal{M}}|/2$ , then we can route all pairs in  $\tilde{\mathcal{M}}'$  via edge-disjoint paths, where each pair  $(s, t) \in \tilde{\mathcal{M}}'$  is routed inside the connected component  $H(s) \cup H(t) \cup e_s$ , if  $H(s)$  tagged  $H(t)$ , or inside the component  $H(s) \cup H(t) \cup e_t$  if  $t$  tagged  $s$ . It is easy to see that all resulting paths are edge-disjoint. Therefore, we assume from now on that  $|\tilde{\mathcal{M}}''| \geq |\tilde{\mathcal{M}}|/2$ . To simplify notation, we will refer to  $\tilde{\mathcal{M}}''$  as  $\tilde{\mathcal{M}}$ , and we discard from  $\tilde{\mathcal{T}}_1$  terminals that do not participate in the new set  $\tilde{\mathcal{M}}$ .

We say that the terminals of  $\tilde{\mathcal{T}}_1$  are red and the terminals of  $\tilde{\mathcal{T}}_2$  are blue. If  $H \in \mathcal{H}$  has a representative terminal in  $\tilde{\mathcal{T}}_1$ , then we say that  $H$  is a red cluster; if it has a blue representative, then it is a blue cluster. Notice that it is possible that  $H$  has no representative, in which case it has no color.

As before, we would like to find a subset  $\mathcal{H}^* \subseteq \mathcal{H}$  of clusters, such that whenever  $H$  tags  $H'$ , only one of these two clusters belongs to  $\mathcal{H}^*$ . However, we have an additional restriction: for every pair  $(s, t) \in \tilde{\mathcal{M}}$ , we would like to ensure that either both  $H(s)$  and  $H(t)$  belong to  $\mathcal{H}^*$ , or none of them.

In order to achieve this, we construct a directed graph  $D$ . The set of vertices of  $D$  consists of two subsets: the set of vertices representing the terminals in  $\tilde{\mathcal{T}}_2$ ,  $V_2 = \{v_t \mid t \in \tilde{\mathcal{T}}_2\}$ , and the set of vertices representing the pairs of terminals in  $\tilde{\mathcal{M}}$ ,  $V_1 = \{v_{s,t} \mid (s, t) \in \tilde{\mathcal{M}}\}$ .

We say that a vertex  $v_t \in V_2$  represents the cluster  $H(t)$ , and a vertex  $v_{s,t} \in V_1$  represents the clusters  $H(s)$  and  $H(t)$ . We let  $V(D) = V_1 \cup V_2$ , and we add an edge  $(u, u')$  to  $D$  iff at least one of the clusters represented by  $u$  tags a cluster represented by  $u'$ . We denote  $|V_1| = \tilde{k}_1$  and  $|V_2| = \tilde{k}_2$ . Notice that  $\tilde{k}_1 = \Omega(\alpha k_1)$  and  $\tilde{k}_2 = \Omega(\alpha k_2)$ . We use the following variation of Claim A.2.

**Claim A.3** *There is an independent set  $\mathcal{I}$  in  $D$ , containing at least  $\lfloor \tilde{k}_1/24 \rfloor$  vertices of  $V_1$  and at least  $\lfloor \tilde{k}_2/24 \rfloor$  vertices of  $V_2$ .*

**Proof:** For simplicity, we treat  $D$  as a general directed graph, where the out-degree of every vertex is at most 2. Therefore, we can assume without loss of generality that  $\tilde{k}_1 \leq \tilde{k}_2$ , and we denote  $\beta = \tilde{k}_2/\tilde{k}_1$ . We say that a vertex  $v \in V_1$  is good iff it is adjacent to at most  $8\beta$  vertices of  $V_2$ . Since the total number of vertices in the graph is bounded by  $2\beta k_1$ , the number of edges is bounded by  $4\beta k_1$ , and so at least half the vertices of  $V_1$  are good.

Let  $D'$  be the sub-graph of  $D$  induced by the good vertices in  $V_1$ . Since the out-degree of every vertex in  $D'$  is at most 2, and  $|V(D)| \geq \tilde{k}_1/2$ , we can find an independent set  $\mathcal{I}_1$  in  $D'$  of size exactly  $\lfloor \tilde{k}_1/24 \rfloor$ , using the standard greedy algorithm as before.

Let  $V' \subseteq V_2$  be the subset of vertices adjacent to the vertices of  $\mathcal{I}_1$ . Since all vertices in  $\mathcal{I}_1$  are good,  $|V'| \leq 8\beta \lfloor \frac{\tilde{k}_1}{24} \rfloor \leq \frac{\tilde{k}_2}{3}$ . Let  $D''$  be the sub-graph of  $D$  induced by the vertices of  $V_2 \setminus V'$ . Since the out-degree of every vertex in  $D''$  is at most 2, and  $|V(D'')| \geq |V_2|/2$ , we can find an independent set  $\mathcal{I}_2$  of size  $\lfloor \tilde{k}_2/24 \rfloor$  in  $D''$ . We output  $\mathcal{I}_1 \cup \mathcal{I}_2$  as our final solution.  $\square$

Let  $\mathcal{M}^* \subseteq \tilde{\mathcal{M}}$  denote the set of pairs  $(s, t)$  with  $v_{s,t} \in \mathcal{I}$ , let  $\mathcal{T}_1^* \subseteq \tilde{\mathcal{T}}_1$  be the set of all terminals participating in pairs in  $\mathcal{M}^*$ , and let  $\mathcal{T}_2^*$  be the set of terminals  $t$  with  $v_t \in \mathcal{I}$ . Finally, let  $\mathcal{H}^* = \{H(t) \mid t \in \mathcal{T}_1^* \cup \mathcal{T}_2^*\}$ . Then for every cluster  $H \in \mathcal{H}^*$ ,  $\mathcal{T}(H)$  has exactly one representative in  $\mathcal{T}_1^* \cup \mathcal{T}_2^*$ , and for any pair  $H, H'$  of clusters, where  $H$  tags  $H'$ , at most one of the two clusters may belong to  $\mathcal{H}^*$ . Notice that we are also guaranteed that  $|\mathcal{T}_1^*| = \Omega(\tilde{k}_1) = \Omega(\alpha k_1)$ , and  $|\mathcal{T}_2^*| = \Omega(\tilde{k}_2) = \Omega(\alpha k_2)$ .

**Step 3: Merging the Clusters** In this step, we create the final set  $\mathcal{G}$  of clusters, by merging some pairs of clusters in  $\mathcal{H}^*$ . For each cluster  $H \in \mathcal{G}$ , we first select up to two terminals  $t, t' \in \mathcal{T}(H) \cap (\mathcal{T}_1^* \cup \mathcal{T}_2^*)$  that we call *potential pseudo-centers* for  $H$ . In the end, at most one such terminal will be selected for every cluster.

If  $H \in \mathcal{H}^*$  is a cluster that does not tag any other cluster, then we add  $H$  to  $\mathcal{G}$ . Let  $t \in \mathcal{T}(H)$  be the unique representative of  $H$  in  $\mathcal{T}_1^* \cup \mathcal{T}_2^*$ . Then we say that  $t$  is the potential pseudo-center of  $H$ . Notice that from Theorem A.2,  $t$  is indeed a pseudo-center of  $H$ .

Consider now some cluster  $H \in \mathcal{H}$ , and let  $\mathcal{G}(H) \subseteq \mathcal{H}^*$  be the set of clusters  $H'$  that tag  $H$ . Assume that  $\mathcal{G}(H) = \{H_1, \dots, H_p\}$ . For each  $1 \leq i \leq p$ , let  $t_i \in \mathcal{T}(H_i)$  be the representative of  $H_i$  in  $\mathcal{T}_1^* \cup \mathcal{T}_2^*$ . If  $p = 1$ , then we create one new cluster  $H^* = H \cup H_1 \cup \{e_{t_1}\}$  and add it to  $\mathcal{G}$ . We set  $\mathcal{T}(H^*) = \mathcal{T}(H_1) \cup \mathcal{T}(H)$ , and we let  $t_1$  be its potential pseudo-center. Notice that  $t_1$  is indeed a pseudo-center of the new cluster, from Theorem A.2.

Assume now that  $p > 1$ , and it is even. As before, we use Lemma A.2 to find a collection  $\mathcal{P}$  of edge-disjoint paths contained in  $H$ , where every path connects a distinct pair  $y_{t_i}, y_{t_j}$  of vertices. For each such path  $P \in \mathcal{P}$ , if  $P$  connects  $y_{t_i}$  to  $y_{t_j}$ , then we construct a new cluster  $H^* = H_i \cup P \cup H_j$ , and add it to  $\mathcal{G}$ . We set  $\mathcal{T}(H^*) = \mathcal{T}(H_i) \cup \mathcal{T}(H_j)$ . Notice that from Theorem A.2, both  $t_i$  and  $t_j$  are pseudo-centers of the new cluster. We say that  $t_i$  and  $t_j$  are *potential pseudo-centers* for the new cluster, and we will later select at most one of them to represent this cluster.

Finally, assume that  $p$  is odd. If the number of blue clusters in  $\mathcal{G}(H)$  is at least two, then we discard one arbitrary blue cluster from  $\mathcal{G}(H)$ , and continue as in the case where  $p$  is even. Otherwise,  $\mathcal{G}(H)$  must contain at least two red clusters. In this case we say that  $H$  is an *odd cluster*.

For each such odd cluster  $H$ , we would like to select one red cluster  $H' \in \mathcal{G}(H)$  to discard. However, once such a cluster is discarded, if  $s \in \mathcal{T}(H')$  is the representative of  $H'$  in  $\mathcal{T}_1^* \cup \mathcal{T}_2^*$ , and  $t$  is the terminal with  $(s, t) \in \mathcal{M}^*$ , we will need to discard the terminal  $t$  as well. That is,  $t$  will not be able to serve as a pseudo-center of its cluster  $H(t)$ . In particular, if  $H(t) \in \mathcal{G}(H')$  of some other odd cluster  $H'$ , we need to ensure that  $H(t)$  is the cluster that we discard from  $\mathcal{G}(H')$  in this case.

We proceed in two stages. In the first stage, while there is a pair  $(s, t) \in \mathcal{M}^*$ , such that  $H(s) \in \mathcal{G}(H)$ ,  $H(t) \in \mathcal{G}(H')$ , with  $H \neq H'$ , and both  $H$  and  $H'$  are odd clusters, we discard  $H(s)$  from  $\mathcal{G}(H)$ ,  $H(t)$  from  $\mathcal{G}(H')$ , and  $(s, t)$  from  $\mathcal{M}^*$ . The clusters  $H$  and  $H'$  now stop being odd, as  $|\mathcal{G}(H)|, |\mathcal{G}(H')|$  are now even. We process these two clusters as in the case where  $p$  is even. When no such pair  $(s, t) \in \mathcal{M}^*$  remains, the first stage finishes and the second stage begins. We process the remaining odd clusters one-by-one. For each such cluster  $H$ , there must be at least one red cluster  $H' \in \mathcal{G}(H)$ , such that, if  $t \in \mathcal{T}(H')$  is the representative of  $H'$  in  $\mathcal{T}_1^*$ , and  $s$  is the terminal with  $(s, t) \in \mathcal{M}^*$ , then  $H(s)$  does not belong to the set  $\mathcal{G}(H'')$  of any odd cluster  $H''$ . We then discard  $(s, t)$  from  $\mathcal{M}^*$ , and discard  $H'$  from  $\mathcal{G}(H)$ . If  $s$  is a potential pseudo-center for any cluster in  $\mathcal{G}$ , then it stops being a potential pseudo-center for that cluster.

Let  $\mathcal{M}^{**} \subseteq \mathcal{M}^*$  be the set of all surviving pairs. We need the following simple claim.

**Claim A.4**  $|\mathcal{M}^{**}| \geq |\mathcal{M}^*|/4$ .

**Proof:** Consider the first stage of the algorithm, where we have removed a collection of pairs  $(s, t)$ , with  $H(s) \in \mathcal{G}(H)$ ,  $H(t) \in \mathcal{G}(H')$ , for odd clusters  $H$  and  $H'$ . For each such odd cluster  $H^*$ , we have removed at most one such pair. However, since  $\mathcal{G}(H^*)$  contained at least two red clusters, there is another pair  $(s', t') \in \mathcal{M}^*$ , such that  $H(t') \in \mathcal{G}(H^*)$ , and  $(s', t')$  was not removed from  $\mathcal{M}^*$  during the first step. It is easy to verify that at most half the pairs in  $\mathcal{M}^*$  were removed during the first stage. Let  $\mathcal{M}'$  be the collection of pairs that survive after the first stage.

In the second stage, for each remaining odd cluster  $H$ , we select one red cluster  $H' \in \mathcal{G}(H)$ , and delete the pair  $(s, t)$  with  $H(s) = H'$  from  $\mathcal{M}'$ . However, since  $H$  still remains an odd cluster, we are guaranteed that there is another pair  $(s', t') \in \mathcal{M}'$ , such that  $H(s') \in \mathcal{G}(H)$ , and  $H(t')$  does not belong to any current odd cluster. (Otherwise we could have continued the first phase for another step). We can therefore charge the pair  $(s, t)$  to the pair  $(s', t')$ , that will not be removed during the second step. So overall, at most half the pairs are removed from  $\mathcal{M}'$  during the second stage.  $\square$

We conclude that  $|\mathcal{M}^{**}| = \Omega(\alpha k_1)$ . Let  $\mathcal{T}_1^{**}$  be the set of all red terminals participating in the pairs in  $\mathcal{M}^{**}$ , and let  $\mathcal{T}_2^{**}$  be the set of all blue terminals that serve as potential pseudo-centers for clusters in  $\mathcal{G}$ . It is easy to see that  $|\mathcal{T}_2^{**}| \geq |\mathcal{T}_2^*|/2 = \Omega(\alpha k_2)$ , since every blue cluster discarded from  $\mathcal{H}^*$  can be charged to another blue cluster that remains in  $\mathcal{H}^*$ .

To summarize, we have defined a collection  $\mathcal{H}^*$  of edge-disjoint clusters, a set  $\mathcal{M}^{**} \subseteq \mathcal{M}$  of  $\Omega(\alpha k_1)$  pairs of terminals, a subset  $\mathcal{T}_1^{**} \subseteq \mathcal{T}_1$  of terminals that participate in the pairs in  $\mathcal{M}^{**}$ , and a subset  $\mathcal{T}_2^{**} \subseteq \mathcal{T}_2$  of  $\Omega(\alpha k_2)$  terminals, with  $\mathcal{T}_1^{**} \cap \mathcal{T}_2^{**} = \emptyset$ . Moreover, every terminal  $t \in \mathcal{T}_1^{**} \cup \mathcal{T}_2^{**}$  is a potential pseudo-center for one of the clusters in  $\mathcal{G}$ , and every cluster  $H \in \mathcal{G}$  now has either 0, 1, or 2 potential pseudo-centers. The sets  $\{\mathcal{T}(H)\}_{H \in \mathcal{G}}$  of terminals are mutually disjoint.

Notice that it is possible that for some clusters  $H \in \mathcal{G}$ , we have two potential pseudo-centers  $s, t$  for  $H$ , which form a pair in  $\mathcal{M}^{**}$ . Let  $\mathcal{M}_1^{**} \subseteq \mathcal{M}^{**}$  be the subset of all such pairs of terminals. In other words,  $(s, t) \in \mathcal{M}_1^{**}$  iff both  $s$  and  $t$  are potential pseudo-centers of the same cluster  $H \in \mathcal{G}$ . If  $|\mathcal{M}_1^{**}| \geq |\mathcal{M}^{**}|/2$ , then every pair in  $\mathcal{M}_1^{**}$  can be routed inside its own cluster, via edge-disjoint paths. In this case we terminate the algorithm and return this collection of paths. From now on we assume that this is not the case. To simplify notation, we denote by  $\mathcal{M}^{**}$  the set of pairs in  $\mathcal{M}^{**} \setminus \mathcal{M}_1^{**}$ , and by  $\mathcal{T}_1^{**}$  the set of terminals participating in pairs in  $\mathcal{M}^{**}$ .

**Step 4: Selecting the Pseudo-Centers** In this step we select a subset  $\mathcal{M}' \subseteq \mathcal{M}^{**}$  of terminal pairs, a set  $\mathcal{T}'_1 \subseteq \mathcal{T}_1^{**}$  of terminals participating in pairs in  $\mathcal{M}'$ , and a subset  $\mathcal{T}'_2 \subseteq \mathcal{T}_2^{**}$ , such that for each cluster  $H \in \mathcal{G}$ , at most one potential pseudo-center belongs to  $\mathcal{T}'_1 \cup \mathcal{T}'_2$ .

Our first step is to ensure that every cluster in  $\mathcal{G}$  contains at most one red potential pseudo-center. In order to do so, we start with  $\mathcal{M}' = \emptyset$ . While  $\mathcal{M}^{**} \neq \emptyset$ , we select an arbitrary pair  $(s, t) \in \mathcal{M}^{**}$ , remove it from  $\mathcal{M}^{**}$  and add it to  $\mathcal{M}'$ . Assume that  $s$  is a potential pseudo-center of  $H$ , and  $t$  is a potential pseudo-center of  $H'$ . We also remove from  $\mathcal{M}^{**}$  any other pair  $(s', t')$ , where either  $s'$  or  $t'$  are potential pseudo-centers of  $H$  or  $H'$ . Notice that for each pair added to  $\mathcal{M}'$ , at most three pairs are deleted from  $\mathcal{M}^{**}$ . Therefore,  $|\mathcal{M}'| \geq |\mathcal{M}^{**}|/3 = \Omega(\alpha k_1)$ . Let  $\mathcal{T}'_1$  be the set of all terminals participating in pairs in  $\mathcal{M}^{**}$ .

Next, we ensure that every cluster in  $\mathcal{G}$  contains at most one blue potential pseudo-center. In order to do so, we start with  $\mathcal{T}'_2 = \emptyset$ . For every cluster  $H \in \mathcal{G}$  that contains two blue potential pseudo-center, we select one of these two blue terminals arbitrarily and add it to  $\mathcal{T}'_2$ . Notice that  $|\mathcal{T}'_2| \geq |\mathcal{T}_2^{**}|/2 = \Omega(\alpha k_2)$ .

Finally, we need to take care of clusters  $H \in \mathcal{G}$  that contain one red and one blue potential pseudo-

centers. We call such clusters *mixed clusters*, and we take care of them like in Step 1. Let  $\mathcal{G}' \subseteq \mathcal{G}$  be the set of all mixed clusters. We partition  $\mathcal{G}'$  into two subsets  $\mathcal{G}'_1, \mathcal{G}'_2$ , such that  $|\mathcal{G}'_1|, |\mathcal{G}'_2| \geq |\mathcal{G}'|/2 - 1$ . As before, we ensure that if, for some pair  $(s, t) \in \mathcal{M}'$ , there are two clusters  $H, H' \in \mathcal{G}'$ , where  $s$  is a potential pseudo-center of  $H$ , and  $t$  is a potential pseudo-center of  $H'$ , then either both  $H, H' \in \mathcal{G}'_1$ , or both  $H, H' \in \mathcal{G}'_2$ . The partition is performed exactly as in Step 1.

For each cluster  $H \in \mathcal{G}'_1$ , we remove its blue potential pseudo-center from  $\mathcal{T}'_2$ . For each cluster  $H \in \mathcal{G}'_2$ , let  $t$  be its red potential pseudo-center, and let  $s \in \mathcal{T}'_1$  such that  $(s, t) \in \mathcal{M}'$ . We remove  $(s, t)$  from  $\mathcal{M}'$ , and we remove both  $s$  and  $t$  from  $\mathcal{T}'_1$ . Since  $|\mathcal{G}'_1| \geq |\mathcal{G}'_2| - 2$ , we only remove a constant fraction of the terminals from  $\mathcal{T}'_1$ . Therefore, at most a constant fraction of the pairs is removed from  $\mathcal{M}'$ , and at most a constant fraction of terminals is removed from  $\mathcal{T}'_2$ .

Let  $\mathcal{M}', \mathcal{T}'_1, \mathcal{T}'_2$  be the resulting subsets of terminal pairs and terminals. Then  $|\mathcal{T}'_1| = \Omega(\alpha k_1)$ ,  $|\mathcal{T}'_2| = \Omega(\alpha k_2)$ . For every cluster  $H \in \mathcal{G}$ , at most one potential pseudo-center for  $H$  belongs to  $\mathcal{T}'_1 \cup \mathcal{T}'_2$ . Using Claim A.1, set  $\mathcal{T}' = \mathcal{T}'_1 \cup \mathcal{T}'_2$  of terminals is 1-well-linked in graph  $G$ .

## B Proof of Theorem 1.1

In this section we complete the proof of Theorem 1.1, using Theorem 3.1.

As in previous work on the EDPwC problem [RZ10, And10, Chu12], we follow the outline of [CKS05], by first partitioning the graph  $G$  into a collection of sub-instances, where in each sub-instance, the corresponding set of terminals is well-linked. We then solve each such sub-instance separately, by embedding an expander into it. The specific type of embedding that we use is similar to the one proposed in [Chu12], and is summarized in the following definition.

**Definition 8.** Let  $\mathcal{M}' \subseteq \mathcal{M}$  be a subset of the demand pairs, and let  $\mathcal{T}' = \mathcal{T}(\mathcal{M}')$  be the set of terminals participating in the pairs in  $\mathcal{M}'$ . Let  $H$  be a graph whose vertex set is  $\mathcal{T}'$ . An *embedding* of  $H$  into  $G$  maps every vertex  $t \in V(H)$  into a connected sub-graph  $C_t \subseteq G$  with  $t \in C_t$ , and every edge  $e = (t, t') \in E(H)$  into a path  $P_e$  in graph  $G$ , connecting some vertex  $v \in C_t$  to some vertex  $v' \in C_{t'}$ . Given an edge  $e' \in E(G)$ , the *load* on edge  $e'$  is the total number of sub-graphs  $\{C_t\}_{t \in \mathcal{T}'}$  and paths  $\{P_e\}_{e \in E(H)}$  to which edge  $e'$  belongs. The *congestion* of the embedding is the maximum load on any edge  $e' \in E(G)$ .

Our algorithm uses the cut-matching game of [KRV09], summarized in Theorem 2.4, to embed an expander into the graph  $G$ . To simplify notation, we say that an  $N$ -vertex (multi)-graph  $X$  is a *good expander*, iff  $N$  is even, the expansion of  $X$  is at least  $\alpha_{\text{CMG}}(N)$ , and the degree of every vertex of  $X$  is exactly  $\gamma_{\text{CMG}}(N)$ . Once we embed a good expander into the graph  $G$ , we will need to route demand pairs across  $X$  via vertex-disjoint paths.

There are many algorithms for routing on expanders, e.g. [LR99, BFU94, BFSU94, KR96, Fri00], that give different types of guarantees. We use the following theorem, due to Rao and Zhou [RZ10] (see also a proof in [Chu12]).

**Theorem B.1 (Theorem 7.1 in [RZ10])** *Let  $G = (V, E)$  be any  $n$ -vertex  $d$ -regular  $\alpha$ -expander. Assume further that  $n$  is even, and that the vertices of  $G$  are partitioned into  $n/2$  disjoint demand pairs  $\mathcal{M} = \{(s_1, t_1), \dots, (s_{n/2}, t_{n/2})\}$ . Then there is an efficient algorithm that routes  $\Omega\left(\frac{\alpha n}{\log n \cdot d^2}\right)$  of the demand pairs on vertex-disjoint paths in  $G$ .*

**Corollary B.1** *Let  $X$  be a good  $N$ -vertex expander, and assume that the vertices of  $X$  are partitioned into  $N/2$  disjoint demand pairs  $\mathcal{M} = \{(s_1, t_1), \dots, (s_{N/2}, t_{N/2})\}$ . Then there is an efficient algorithm that routes  $\Omega\left(\frac{N \cdot \alpha_{\text{CMG}}(N)}{\log N \cdot \gamma_{\text{CMG}}^2(N)}\right) = \Omega\left(\frac{N}{\log^4 N}\right)$  of the demand pairs via vertex-disjoint paths in  $X$ .*

The following theorem allows us to construct a good expander  $H$  and embed it into  $G$  with congestion at most 2, given a good congestion-2 crossbar  $(\mathcal{S}^*, \mathcal{M}^*, \tau^*)$  in graph  $G$ .

**Theorem B.2** *Assume that we are given an undirected graph  $G = (V, E)$  with vertex degrees at most 4, and a set  $\mathcal{M}$  of  $k$  demand pairs, defined over a set  $\mathcal{T}$  of terminals. Assume further that the degree of every terminal is 1, every terminal participates in exactly one pair in  $\mathcal{M}$ , and  $\mathcal{T}$  is 1-well-linked in  $G$ . Then there is an efficient randomized algorithm, that with high probability outputs one of the following:*

- *Either a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of  $k/\text{poly log } k$  demand pairs and the routing of the pairs in  $\mathcal{M}'$  with congestion at most 2 in  $G$ ;*
- *Or a good expander  $H$  together with the embedding of  $H$  into  $G$  with congestion at most 2.*

**Proof:** We apply Theorem 3.1 to graph  $G$ . If the outcome is a subset  $\mathcal{M}' \subseteq \mathcal{M}$  of  $k/\text{poly log } k$  demand pairs and their routing in  $\mathcal{M}'$  with congestion at most 2, then we return this routing, and terminate the algorithm. We assume from now on that the algorithm in Theorem 3.1 produces a good crossbar  $(\mathcal{S}^*, \mathcal{M}^*, \tau^*)$ .

The good expander  $H$  is constructed as follows. The set of vertices of  $H$ ,  $V(H) = \mathcal{T}(\mathcal{M}^*)$ . The embedding of each vertex  $t \in V(H)$  is the unique tree  $T \in \tau^*$ , to which terminal  $t$  belongs. In order to compute the set of edges of  $H$  and their embedding into  $G$ , we perform the cut-matching game, using Theorem 2.4. Recall that this game consists of  $\gamma_{\text{CMG}}$  iterations, where in iteration  $j$ , for  $1 \leq j \leq \gamma_{\text{CMG}}$ , the cut player computes a partition  $(A_j, B_j)$  of  $V(H)$ , and the matching player responds with a perfect matching between the vertices of  $A_j$  and the vertices of  $B_j$ . Given the partition  $(A_j, B_j)$  of  $V(H)$ , computed by the cut player, we find a corresponding partition  $(A'_j, B'_j)$  of  $\Gamma_j^*$ , as follows. For each  $t_i \in V(H)$ , if  $t_i \in A_j$ , then we add  $e_{i,j}$  to  $A'_j$ , and otherwise we add it to  $B'_j$ . Since the set  $S_j^*$  has the 1-bandwidth property with respect to  $\Gamma_j^*$ , we can find a collection  $\mathcal{Q}_j^* : A'_j \overset{1:1}{\rightsquigarrow} B'_j$  of edge-disjoint paths contained in  $S_j^*$ . These paths define a matching  $M_j$  between the edges of  $A'_j$  and the edges of  $B'_j$ , which in turn defines a matching between the terminals of  $A_j$  and the terminals of  $B_j$ . We then add the edges of the matching to the graph  $H$ . For each such edge  $e \in M_j$ , its embedding into  $G$  is the corresponding path  $Q \in \mathcal{Q}_j^*$ . From Theorem 2.4, after  $\gamma_{\text{CMG}}$  iterations, we obtain a good expander  $H$ , together with its embedding into  $G$ . It is easy to see that the embedding causes congestion at most 2 in graph  $G$ . Indeed, consider any edge  $e \in E$ . If edge  $e$  belongs to any sub-graph  $G[S_j]$ , for  $S_j \in \mathcal{S}^*$ , then it belongs to at most one tree in  $\tau^*$ , and to at most one path in  $\{P_{e'} \mid e' \in E(H)\}$ . If  $e$  does not belong to any such sub-graph, then it is contained in at most two trees in  $\tau^*$ , and it is not contained in any path in  $\{P_{e'} \mid e' \in E(H)\}$ . Therefore, the total congestion of this embedding is at most 2.  $\square$

We are now ready to complete the proof of Theorem 1.1. From the discussion in Section 2, it is enough to consider graphs  $G$ , where the maximum vertex degree is at most 4, the degree of every terminal is 1, and each terminal participates in exactly one demand pair. Our starting point is similar to that used in previous work on EDP [CKS04b, CKS05, RZ10, And10, Chu12]. We use the standard multicommodity flow LP-relaxation for the EDPwC problem: we have an indicator variable  $x_i$  for each  $1 \leq i \leq k$ , for whether the pair  $(s_i, t_i)$  is routed. Let  $\mathcal{P}_i$  be the set of all paths connecting  $s_i$  to  $t_i$  in  $G$ . The LP relaxation is defined as follows.

$$(LP1) \quad \max \quad \sum_{i=1}^k x_i$$

$$\text{s.t.} \quad \sum_{P \in \mathcal{P}_i} f(P) \geq x_i \quad \forall 1 \leq i \leq k \quad (1)$$

$$\sum_{P: e \in P} f(P) \leq 2 \quad \forall e \in E \quad (2)$$

$$0 \leq x_i \leq 1 \quad \forall 1 \leq i \leq k \quad (3)$$

$$f(P) \geq 0 \quad \forall 1 \leq i \leq k, \forall P \in \mathcal{P}_i \quad (4)$$

While this LP has exponentially many variables, it can be efficiently solved using standard techniques, e.g. by using an equivalent polynomial-size edge-based flow formulation. We denote by  $\text{OPT}$  the value of the optimal solution to the LP. Clearly, the value of the optimal solution to the EDPwC problem instance with congestion 2 is at most  $\text{OPT}$ .

Notice that if we replace Constraint (2) with the following constraint:

$$\sum_{P: e \in P} f(P) \leq 1 \quad \forall e \in E,$$

then we obtain the standard multicommodity LP relaxation for the EDP problem itself, where no congestion is allowed. Let  $\text{OPT}'$  denote the optimal solution value of this new LP. Then clearly  $\text{OPT}' \leq \text{OPT} \leq 2\text{OPT}'$ .

The next theorem follows from the work of Chekuri, Khanna and Shepherd [CKS04b, CKS05], and we provide a short proof sketch for completeness.

**Theorem B.3** *Suppose we are given a graph  $G = (V, E)$  and a set  $\mathcal{M}$  of  $k$  source-sink pairs in  $G$ . Then we can efficiently partition  $G$  into a collection  $G_1, \dots, G_\ell$  of vertex-disjoint induced sub-graphs, and compute, for each  $1 \leq i \leq \ell$ , a collection  $\mathcal{M}_i \subseteq \mathcal{M}$  of source-sink pairs contained in  $G_i$ , such that  $\sum_{i=1}^{\ell} |\mathcal{M}_i| = \Omega(\text{OPT}/\log^2 k)$ , and moreover, if for each  $1 \leq i \leq \ell$ ,  $\mathcal{T}_i$  denotes the set of terminals participating in pairs in  $\mathcal{M}_i$ , then  $\mathcal{T}_i \subseteq V(G_i)$ , and either (i) the algorithm computes a set  $\mathcal{P}_i$  of edge-disjoint paths connecting the pairs in  $\mathcal{M}_i$ , or (ii)  $\mathcal{T}_i$  is 1-well-linked in  $G_i$ .*

**Proof:** We need the following definition.

**Definition 9.** Given a graph  $G = (V, E)$ , and a subset  $\mathcal{T} \subseteq V$  of vertices called terminals, we say that  $\mathcal{T}$  is *flow-well-linked* in  $G$ , iff any matching  $\mathcal{M}$  on  $\mathcal{T}$  can be fractionally routed with congestion at most 2 in  $G$ .

The next theorem follows from the work of Chekuri, Khanna and Shepherd [CKS04b, CKS05], and its proof also appears in [Chu12].

**Theorem B.4** *Suppose we are given a graph  $G = (V, E)$  and a set  $\mathcal{M}$  of  $k$  source-sink pairs in  $G$ . Then we can efficiently partition  $G$  into a collection  $G_1, \dots, G_\ell$  of vertex-disjoint induced sub-graphs, and compute, for each  $1 \leq i \leq \ell$ , a collection  $\mathcal{M}_i \subseteq \mathcal{M}$  of source-sink pairs contained in  $G_i$ , such that  $\sum_{i=1}^{\ell} |\mathcal{M}_i| = \Omega(\text{OPT}/\log^2 k)$ , and moreover, if for each  $1 \leq i \leq \ell$ ,  $\mathcal{T}_i$  denotes the set of terminals participating in pairs in  $\mathcal{M}_i$ , then  $\mathcal{T}_i \subseteq V(G_i)$ , and  $\mathcal{T}_i$  is flow-well-linked in  $G_i$ .*

Fix some  $1 \leq i \leq \ell$ . From the min-cut max-flow theorem, since  $\mathcal{T}_i$  is flow-well-linked in  $G_i$ , then it must be 1/2-well-linked in  $G_i$ . We apply the grouping technique from Theorem 2.3 to graph  $G_i$ , where the first set of terminals is  $\mathcal{T}_i$ , together with the matching  $\mathcal{M}_i$  defined over  $\mathcal{T}_i$ , and the second set of

terminals is  $\emptyset$ . Notice that  $\alpha = 1/2$  in this case. If the outcome is a subset  $\mathcal{M}'_i \subseteq \mathcal{M}_i$  of  $\Omega(|\mathcal{M}_i|)$  pairs, together with their routing on edge-disjoint paths, then we let  $\mathcal{P}_i$  denote this routing. Otherwise, we obtain a subset  $\mathcal{M}'_i \subseteq \mathcal{M}_i$  of  $\Omega(|\mathcal{M}_i|)$  pairs, where  $\mathcal{T}(\mathcal{M}'_i)$  is 1-well-linked in  $G_i$ .  $\square$

We now proceed to solve the problem on each one of the graphs  $G_i$  separately. For each  $i$ , let  $k_i = |\mathcal{M}'_i|$ . Notice that in order to complete the proof of Theorem 1.1, it is now enough to show that there is an efficient randomized algorithm that for each  $1 \leq i \leq \ell$ , routes  $\Omega(k_i / \text{poly log } k_i)$  pairs in  $\mathcal{M}'_i$  with congestion at most 2 in graph  $G_i$  with high probability.

Fix some  $1 \leq i \leq \ell$ . If we already have a set  $\mathcal{P}_i$  of edge-disjoint paths routing the pairs in  $\mathcal{M}'_i$ , then we return this routing. Otherwise, set  $\mathcal{T}(\mathcal{M}'_i)$  is 1-well-linked in  $G_i$ , and we apply Theorem B.2 to graph  $G_i$  and the set  $\mathcal{M}'_i$  of the demand pairs. If the output of Theorem B.2 is a collection  $\mathcal{M}''_i \subseteq \mathcal{M}'_i$  of  $k_i / \text{poly log } k_i$  demand pairs, together with their routing with congestion at most 2 in  $G_i$ , then we terminate the algorithm and return this routing for graph  $G_i$ . Otherwise, we obtain a collection  $\mathcal{M}''_i \subseteq \mathcal{M}'_i$  of  $k_i / \text{poly log } k_i$  demand pairs, a good expander  $H$  whose vertex set is  $V(H) = \mathcal{T}(\mathcal{M}''_i)$ , and an embedding of  $H$  into the graph  $G_i$  with congestion 2. Using Theorem B.1, we can compute a collection  $\tilde{\mathcal{M}}_i \subseteq \mathcal{M}''_i$  of  $\Omega\left(\frac{|\mathcal{M}''_i|}{(\log |\mathcal{M}''_i|)^4}\right) = \Omega(k_i / \text{poly log } k_i)$  demand pairs, and a routing  $\mathcal{P}$  of the pairs in  $\tilde{\mathcal{M}}_i$  in the expander  $H$  via vertex-disjoint paths. We now construct a routing of the pairs in  $\tilde{\mathcal{M}}_i$  in the graph  $G_i$ , with congestion at most 2.

In order to construct the routing, consider any path  $P \in \mathcal{P}$ , and assume that it connects a pair  $(s, t) \in \tilde{\mathcal{M}}_i$ . We transform the path  $P$  into a path  $Q$ , connecting  $s$  to  $t$  in graph  $G_i$ , as follows. Assume that  $P = (t_0 = s, t_1, \dots, t_r = t)$ , and let  $e_j$  be the edge connecting  $t_j$  to  $t_{j+1}$ , for  $0 \leq j < r$ . For each such edge  $e_j$ , let  $P_{e_j}$  be the path via which  $e_j$  is embedded into  $G$ . Let  $v_j$  be the first vertex on  $P_{e_j}$ , and let  $u_{j+1}$  be the last vertex on  $P_{e_j}$ . Then  $v_j$  belongs to the connected sub-graph  $C_{t_j}$ , and  $u_{j+1}$  belongs to the connected sub-graph  $C_{t_{j+1}}$  of  $G$ .

For each  $1 \leq j < r$ , let  $Q_j$  be any path connecting  $u_j$  to  $v_j$  inside the connected graph  $C(t_j)$ . Let  $Q_0$  be any path connecting  $t_0$  to  $v_0$  inside  $C_{t_0}$ , and let  $Q_r$  be any path connecting  $u_r$  to  $t_r$  inside  $C_{t_r}$ . Path  $Q$  is then obtained by concatenating  $Q_0, P_{e_0}, Q_1, \dots, P_{e_{r-1}}, Q_r$ . Notice that path  $Q$  connects the original pair  $(s, t)$  of terminals to each other.

Let  $\mathcal{Q}$  be the final set of paths, obtained after processing all paths  $P \in \mathcal{P}$ . From the above discussion, every pair  $(s, t) \in \tilde{\mathcal{M}}_i$  has a path connecting  $s$  to  $t$  in  $\mathcal{Q}$ . In order to bound the congestion caused by  $\mathcal{Q}$ , recall that the paths in  $\mathcal{P}$  are vertex-disjoint. Therefore, the paths in  $\mathcal{Q}$  traverse every graph  $C_t$  for  $t \in V(H)$  at most once, and use every path  $P_e$  for  $e \in E(H)$  at most once. Since the congestion of the embedding of  $H$  is bounded by 2, the paths in  $\mathcal{Q}$  cause congestion at most 2 overall.

To conclude, for each  $1 \leq i \leq \ell$ , we have computed a collection  $\mathcal{M}^*_i \subseteq \mathcal{M}_i$  of  $k_i / \text{poly log } k_i$  demand pairs, and a routing  $\mathcal{Q}_i$  of the pairs in  $\mathcal{M}^*_i$  with congestion at most 2 in  $G_i$ . Since  $\sum_{i=1}^{\ell} k_i = \Omega(\text{OPT} / \log^2 k)$ , overall  $\bigcup_{i=1}^{\ell} \mathcal{Q}_i$  is a  $(\text{poly log } k)$ -approximate solution for the EDPwC instance  $(G, \mathcal{M})$  with congestion 2.

## C Proof of Claim 5.1

Let  $H = G' \setminus \mathcal{T}$ . Fix some  $1 \leq j \leq \gamma$ . Let  $\mathcal{E}_1(j)$  be the bad event that  $\sum_{v \in X_j} d_H(v) > \frac{2m}{\gamma} \cdot \left(1 + \frac{1}{\gamma}\right)$ . In order to bound the probability of  $\mathcal{E}_1(j)$ , we define, for each vertex  $v \in V(H)$ , a random variable  $x_v$ , whose value is  $\frac{d_H(v)}{k_1}$  if  $v \in X_j$  and 0 otherwise. Notice that  $x_v \in [0, 1]$ , and the random variables  $\{x_v\}_{v \in V(H)}$  are pairwise independent. Let  $B = \sum_{v \in V(H)} x_v$ . Then the expectation of  $B$ ,  $\mu_1 = \sum_{v \in V(H)} \frac{d_H(v)}{\gamma k_1} = \frac{2m}{\gamma k_1}$ . Using the standard Chernoff bound (see e.g. Theorem 1.1 in [DP09]),



$$\Pr [\mathcal{E}_1(j)] = \Pr [B > (1 + 1/\gamma) \mu_1] \leq e^{-\mu_1/(3\gamma^2)} = e^{-\frac{2m}{3\gamma^3 k_1}} < \frac{1}{6\gamma}$$

since  $m \geq k/3$  and  $k_1 = \frac{k}{192\gamma^3 \log \gamma}$ .

For each terminal  $t \in \mathcal{T}$ , let  $e_t$  be the unique edge adjacent to  $t$  in graph  $G'$ , and let  $u_t$  be its other endpoint. Let  $U = \{u_t \mid t \in \mathcal{T}\}$ . For each vertex  $u \in U$ , let  $w(u)$  be the number of terminals  $t$ , such that  $u = u_t$ . Notice that  $w(u) \leq k_1$  must hold. We say that a bad event  $\mathcal{E}_2(j)$  happens iff  $\sum_{u \in U \cap X_j} w(u) \geq \frac{k}{\gamma} \cdot \left(1 + \frac{1}{\gamma}\right)$ . In order to bound the probability of the event  $\mathcal{E}_2(j)$ , we define, for each  $u \in U$ , a random variable  $y_u$ , whose value is  $w(u)/k_1$  if  $u \in X_j$ , and it is 0 otherwise. Notice that  $y_u \in [0, 1]$ , and the variables  $y_u$  are independent for all  $u \in U$ . Let  $Y = \sum_{u \in U} y_u$ . The expectation of  $Y$  is  $\mu_2 = \frac{k}{k_1 \gamma}$ , and event  $\mathcal{E}_2(j)$  holds iff  $Y \geq \frac{k}{k_1 \gamma} \cdot \left(1 + \frac{1}{\gamma}\right) \geq \mu_2 \cdot \left(1 + \frac{1}{\gamma}\right)$ . Using the standard Chernoff bound again, we get that:

$$\Pr [\mathcal{E}_2(j)] \leq e^{-\mu_2/(3\gamma^2)} \leq e^{-k/(3k_1\gamma^3)} \leq \frac{1}{6\gamma}$$

since  $k_1 = \frac{k}{192\gamma^3 \log \gamma}$ . Notice that if events  $\mathcal{E}_1(j), \mathcal{E}_2(j)$  do not hold, then:

$$|\text{out}_{G'}(X_j)| \leq \sum_{v \in X_j} d_H(v) + \sum_{u \in U \cap X_j} w(u) \leq \left(1 + \frac{1}{\gamma}\right) \left(\frac{2m}{\gamma} + \frac{k}{\gamma}\right) < \frac{10m}{\gamma}$$

since  $m \geq k/3$ .

Let  $\mathcal{E}_3(j)$  be the bad event that  $|E_{G'}(X_j)| < \frac{m}{2\gamma^2}$ . We next prove that  $\Pr [\mathcal{E}_3(j)] \leq \frac{1}{6\gamma}$ . We say that two edges  $e, e' \in E(G' \setminus \mathcal{T})$  are *independent* iff they do not share any endpoints. Our first step is to compute a partition  $U_1, \dots, U_r$  of the set  $E(G' \setminus \mathcal{T})$  of edges, where  $r \leq 2k_1$ , such that for each  $1 \leq i \leq r$ ,  $|U_i| \geq \frac{m}{4k_1}$ , and all edges in set  $U_i$  are mutually independent. In order to compute such a partition, we construct an auxiliary graph  $Z$ , whose vertex set is  $\{v_e \mid e \in E(G')\}$ , and there is an edge  $(v_e, v_{e'})$  iff  $e$  and  $e'$  are not independent. Since the maximum vertex degree in  $G'$  is at most  $k_1$ , the maximum vertex degree in  $Z$  is bounded by  $2k_1 - 2$ . Using the Hajnal-Szemerédi Theorem [HS70], we can find a partition  $V_1, \dots, V_r$  of the vertices of  $Z$  into  $r \leq 2k_1$  subsets, where each subset  $V_i$  is an independent set, and  $|V_i| \geq \frac{|V(Z)|}{r} - 1 \geq \frac{m}{4k_1}$ . The partition  $V_1, \dots, V_r$  of the vertices of  $Z$  gives the desired partition  $U_1, \dots, U_r$  of the edges of  $G' \setminus \mathcal{T}$ . For each  $1 \leq i \leq r$ , we say that the bad event  $\mathcal{E}_3^i(j)$  happens iff  $|U_i \cap E(X_j)| < \frac{|U_i|}{2\gamma^2}$ . Notice that if  $\mathcal{E}_3(j)$  happens, then event  $\mathcal{E}_3^i(j)$  must happen for some  $1 \leq i \leq r$ . Fix some  $1 \leq i \leq r$ . The expectation of  $|U_i \cap E(X_j)|$  is  $\mu_3 = \frac{|U_i|}{\gamma^2}$ . Since all edges in  $U_i$  are independent, we can use the standard Chernoff bound to bound the probability of  $\mathcal{E}_3^i(j)$ , as follows:

$$\Pr [\mathcal{E}_3^i(j)] = \Pr [ |U_i \cap E(X_j)| < \mu_3/2 ] \leq e^{-\mu_3/8} = e^{-\frac{|U_i|}{8\gamma^2}}$$

Since  $|U_i| \geq \frac{m}{4k_1}$ ,  $m \geq k/3$ ,  $k_1 = \frac{k}{192\gamma^3 \log \gamma}$ , and  $\gamma = \Theta(\log^4 k)$ , this is bounded by  $\frac{1}{12k_1\gamma}$ . We conclude that  $\Pr [\mathcal{E}_3^i(j)] \leq \frac{1}{12k_1\gamma}$ , and by using the union bound over all  $1 \leq i \leq r$ ,  $\Pr [\mathcal{E}_3(j)] \leq \frac{1}{6\gamma}$ .

Using the union bound over all  $1 \leq j \leq \gamma$ , with probability at least  $\frac{1}{2}$ , none of the events  $\mathcal{E}_1(j), \mathcal{E}_2(j), \mathcal{E}_3(j)$  for  $1 \leq j \leq \gamma$  happen, and so for each  $1 \leq j \leq \gamma$ ,  $|\text{out}_{G'}(X_j)| < \frac{10m}{\gamma}$ , and  $|E_{G'}(X_j)| \geq \frac{m}{2\gamma^2}$  must hold.

## D Integrality Gap of the Multi-Commodity Flow Relaxation for EDP

Recall that in the multi-commodity flow relaxation of the EDP problem, the goal is to maximize the total amount of flow sent between the demand pairs, with no congestion, with at most one flow unit being sent between any pair. The gap construction for this relaxation, due to Garg, Vazirani and Yannakakis [GVY97] starts from a sub-graph of the  $(k + 1) \times (k + 1)$  grid depicted in Figure 2. Every degree-4 vertex in this graph is then replaced by the gadget depicted in the figure. This gadget ensures that whenever two paths traverse a degree-4 vertex of the grid, one vertically and one horizontally, then they must share an edge.




Figure 2: The gap example for the multicommodity flow relaxation

Notice that every pair  $(s_i, t_i)$  can send  $\frac{1}{2}$  flow unit from  $s_i$  to  $t_i$  along the path that contains the vertical segment of the grid incident to  $s_i$  and the horizontal segment incident to  $t_i$ . Therefore, the value of the multi-commodity flow relaxation is  $k/2 = \Omega(\sqrt{n})$ . However, it is easy to see that the value of the optimal integral solution is 1. Indeed, for any path  $P$  connecting any pair  $(s_i, t_i)$ , if we remove all edges of  $P$  from the graph, then all other demand pairs become disconnected. Therefore, the integrality gap of this example is  $\Omega(\sqrt{n})$ .

Notice also that in this example we can route  $k = \Omega(\sqrt{n})$  demand pairs integrally with congestion 2, while at most one demand pair can be routed with congestion 1. This shows a polynomial gap between routing with congestion 1 and routing with congestion 2 and higher.

Figure 3 depicts a wall graph. Obtaining sub-polynomial approximation algorithms for EDP on wall graphs remains an interesting open problem.




Figure 3: A wall graph

## E Table of Parameters

$\gamma_{\text{CMG}}$	$\Theta(\log^2 k)$	Parameter for the cut-matching game, from Theorem 2.4.
$\gamma$	$2^{24}\gamma_{\text{CMG}}^4 = O(\log^8 k)$	Number of sets in the initial partition, $\mathcal{R} = \{S_1, \dots, S_\gamma\}$
$\gamma'$	$\gamma^{1/4} = 2^6\gamma_{\text{CMG}}$	Threshold for the number of leaves for Cases 1 and 2
$\alpha_{\text{ARV}}(k)$	$O(\sqrt{\log k})$	Approximation factor of the algorithm of [ARV09] for Sparsest Cut.
$\alpha$	$\frac{1}{2^{11}\gamma \log k} = \Omega\left(\frac{1}{\log^9 k}\right)$	Parameter for bandwidth decomposition.
$\alpha_{\text{BW}}$	$\alpha/\alpha_{\text{ARV}}(k) = \Omega\left(\frac{1}{\log^{9.5} k}\right)$	Bandwidth property parameter.
$r$	$8\gamma_{\text{CMG}}$	Number of sets in $\mathcal{R}'$
$k_1$	$\frac{k}{192\gamma^3 \log \gamma} = \Omega\left(\frac{k}{\log^{24} k \log \log k}\right)$	Threshold for large clusters.
$k'$	$\left\lfloor \frac{k_1}{6\gamma^2} \right\rfloor = \Omega\left(\frac{k}{\text{poly log } k}\right)$	Parameter in the construction of graph $Z$ in Step 1.
$k_2$	$\Omega\left(\frac{k_1 \alpha \cdot \alpha_{\text{BW}}}{\gamma^{3.5}}\right) = \Omega\left(\frac{k}{\text{poly log } k}\right)$	Number of paths for every edge in $\tilde{T}$ in Step 1.
$k_4$	$\Omega(k_2/r^2) = \Omega\left(\frac{k}{\text{poly log } k}\right)$	Number of demand pairs routed to $S, S'$ in Step 2.
$k_6$	$\Omega(\alpha_{\text{BW}}^2 k_4) = \Omega\left(\frac{k}{\text{poly log } k}\right)$	Number of trees in $\tau^*$ is at least $k_6/2$ .