

New Algorithms for Maximum Disjoint Paths Based on Tree-Likeness*

Krzysztof Fleszar¹, Matthias Mnich², and Joachim Spoerhase³

- 1 Universität Würzburg, Würzburg, Germany
krzysztof.fleszar@uni-wuerzburg.de
- 2 Universität Bonn, Bonn, Germany
mmnich@uni-bonn.de
- 3 Universität Würzburg, Würzburg, Germany
joachim.spoerhase@uni-wuerzburg.de

Abstract

We study the classical NP-hard problems of finding maximum-size subsets from given sets of k terminal pairs that can be routed via edge-disjoint paths (MAXEDP) or node-disjoint paths (MAXNDP) in a given graph. The approximability of MAXEDP/NDP is currently not well understood; the best known lower bound is $\Omega(\log^{1/2-\varepsilon} n)$, assuming $\text{NP} \not\subseteq \text{ZPTIME}(n^{\text{poly} \log n})$. This constitutes a significant gap to the best known approximation upper bound of $\mathcal{O}(\sqrt{n})$ due to Chekuri et al. (2006) and closing this gap is currently one of the big open problems in approximation algorithms. In their seminal paper, Raghavan and Thompson (Combinatorica, 1987) introduce the technique of randomized rounding for LPs; their technique gives an $\mathcal{O}(1)$ -approximation when edges (or nodes) may be used by $\mathcal{O}\left(\frac{\log n}{\log \log n}\right)$ paths.

In this paper, we strengthen the above fundamental results. We provide new bounds formulated in terms of the *feedback vertex set number* r of a graph, which measures its vertex deletion distance to a forest. In particular, we obtain the following.

- For MAXEDP, we give an $\mathcal{O}(\sqrt{r} \cdot \log^{1.5} kr)$ -approximation algorithm. As $r \leq n$, up to logarithmic factors, our result strengthens the best known ratio $\mathcal{O}(\sqrt{n})$ due to Chekuri et al.
- Further, we show how to route $\Omega(\text{OPT})$ pairs with congestion $\mathcal{O}\left(\frac{\log kr}{\log \log kr}\right)$, strengthening the bound obtained by the classic approach of Raghavan and Thompson.
- For MAXNDP, we give an algorithm that gives the optimal answer in time $(k+r)^{\mathcal{O}(r)} \cdot n$. This is a substantial improvement on the run time of $2^{kr^{\mathcal{O}(r)}} \cdot n$, which can be obtained via an algorithm by Scheffler.

We complement these positive results by proving that MAXEDP is NP-hard even for $r = 1$, and MAXNDP is W[1]-hard for parameter r . This shows that neither problem is fixed-parameter tractable in r unless $\text{FPT} = \text{W}[1]$ and that our approximability results are relevant even for very small constant values of r .

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases disjoint paths, approximation algorithms, feedback vertex set

Digital Object Identifier 10.4230/LIPIcs.ESA.2016.42

1 Introduction

In this paper, we study disjoint paths routing problems. In this setting, we are given an undirected graph G and a collection of source-destination pairs $\mathcal{M} = \{(s_1, t_1), \dots, (s_k, t_k)\}$.

* This research was supported by ERC Starting Grant 306465 (BeyondWorstCase).



The goal is to select a maximum-sized subset $\mathcal{M}' \subseteq \mathcal{M}$ of the pairs that can be *routed*, where a routing of \mathcal{M}' is a collection \mathcal{P} of paths such that, for each pair $(s_i, t_i) \in \mathcal{M}'$, there is a path in \mathcal{P} connecting s_i to t_i . In the MAXIMUM EDGE DISJOINT PATHS (MAXEDP) problem, a routing \mathcal{P} is *feasible* if its paths are pairwise edge-disjoint, and in the MAXIMUM NODE DISJOINT PATHS (MAXNDP) problem the paths in \mathcal{P} must be pairwise vertex-disjoint.

Disjoint paths problems are fundamental problems with a long history and significant connections to optimization and structural graph theory. The decision versions EDP of MAXEDP and NDP of MAXNDP ask whether all of the pairs can be routed. Karp [27] showed that, when the number of pairs is part of the input, the decision problem is NP-complete. In undirected graphs, MAXEDP and MAXNDP are solvable in polynomial time when the number of pairs is a fixed constant; this is a very deep result of Robertson and Seymour [40] that builds on several fundamental results in structural graph theory from their graph minors project.

In this paper, we consider the optimization problems MAXEDP and MAXNDP when the number of pairs are part of the input. In this setting, the best approximation ratio for MAXEDP is achieved by an $\mathcal{O}(\sqrt{n})$ -approximation algorithm [12, 33], where n is the number of nodes, whereas the best hardness for undirected graphs is only $\Omega(\log^{1/2-\varepsilon} n)$ [3]. Bridging this gap is a fundamental open problem that seems quite challenging at the moment.

Most of the results for routing on disjoint paths use a natural multi-commodity flow relaxation as a starting point. A well-known integrality gap instance due to Garg et al. [24] shows that this relaxation has an integrality gap of $\Omega(\sqrt{n})$, and this is the main obstacle for improving the $\mathcal{O}(\sqrt{n})$ -approximation ratio in general graphs. The integrality instance on an $n \times n$ grid (of treewidth $\Theta(\sqrt{n})$) exploits a topological obstruction in the plane that prevents a large integral routing; see Fig. 1. This led Chekuri et al. [15] to studying the approximability of MAXEDP with respect to the *tree-width* of the underlying graph. In particular, they pose the following conjecture:

► **Conjecture 1** ([13]). *The integrality gap of the standard multi-commodity flow relaxation for MAXEDP is $\Theta(w)$, where w is the treewidth of the graph.*

Recently, Ene et al. [21] showed that MAXEDP admits an $\mathcal{O}(w^3)$ -approximation algorithm on graphs of treewidth at most w . This is the best known approximation ratio in terms of w , improving on an earlier $\mathcal{O}(w \cdot 3^w)$ -approximation algorithm due to Chekuri et al. This shows that the problem seems more amenable on “tree-like” graphs.

However, for $w = \omega(n^{1/6})$, the bound is weaker than the bound of $\mathcal{O}(\sqrt{n})$. In fact, EDP remains NP-hard even for graphs of *constant* treewidth, namely treewidth $w = 2$ [37]. This further rules out the existence of a fixed-parameter algorithm for MAXEDP parameterized by w , assuming $P \neq NP$. Therefore, to obtain fixed-parameter tractability results as well as better approximation guarantees, one needs to resort to parameters stronger than treewidth.

Another route to bridge the large gap between approximation lower and upper bounds for MAXEDP is to allow the paths to have *low congestion* c : that is, instead of requiring the routed paths to be pairwise disjoint, at most c paths can use an edge. In their groundbreaking work, Raghavan and Thompson [38] introduced the technique of randomized rounding of LPs to obtain polynomial-time approximation algorithms for combinatorial problems. Their approach allows to route $\Omega(\text{OPT})$ pairs of paths with congestion $\mathcal{O}\left(\frac{\log n}{\log \log n}\right)$. This extensive line of research [2, 18, 28] has culminated in a $\log^{\mathcal{O}(1)} k$ -approximation algorithm with congestion 2 for MAXEDP [20]. A slightly weaker result also holds for MAXNDP [11].

1.1 Motivation and Contribution

The goal of this work is to study disjoint paths problems under another natural measure for how “far” a graph is from being a tree. In particular, we propose to examine MAXEDP and MAXNDP under the *feedback vertex set number*, which for a graph G denotes the smallest size r of a set R of G for which $G - R$ is a forest. Note that the treewidth of G is at most $r + 1$. Therefore, given the NP-hardness of EDP for $w = 2$ and the current gap between the best known upper bound $\mathcal{O}(w^3)$ and the linear upper bound suggested by Conjecture 1, it is interesting to study the stronger restriction of bounding the feedback vertex set number r of the input graph. Our approach is further motivated by the fact that MAXEDP is efficiently solvable on trees by means of the algorithm of Garg, Vazirani and Yannakakis [24]. Similarly, MAXNDP is easy on trees (see Theorem 4).

Our main insight is that one can in fact obtain bounds in terms of r that either strengthen the best known bounds or are almost tight (see Table 1). It therefore seems that parameter r correlates quite well with the “difficulty” of disjoint paths problems.

Our first result allows the paths to have small congestion: in this setting, we strengthen the result, obtained by the classic randomized LP-rounding approach of Raghavan and Thompson [38], that one can always route $\Omega(\text{OPT})$ pairs with congestion $\mathcal{O}\left(\frac{\log n}{\log \log n}\right)$ with constant probability.

► **Theorem 2.** *For any instance (G, \mathcal{M}) of MAXEDP, one can efficiently find a routing of $\Omega(\text{OPT})$ pairs with congestion $\mathcal{O}\left(\frac{\log kr}{\log \log kr}\right)$ with constant probability; in other words, there is an efficient $\mathcal{O}(1)$ -approximation algorithm for MAXEDP with congestion $\mathcal{O}\left(\frac{\log kr}{\log \log kr}\right)$.*

Our second main result builds upon Theorem 2 and uses it as a subroutine. We show how to use a routing for MAXEDP with low congestion to obtain a polynomial-time approximation algorithm for MAXEDP *without congestion* that performs well in terms of r .

► **Theorem 3.** *The integrality gap of the multi-commodity flow relaxation for MAXEDP with k terminal pairs is $\mathcal{O}(\sqrt{r} \cdot \log^{1.5} rk)$ for graphs with feedback vertex set number r . Moreover, there is a polynomial time algorithm that, given a fractional solution to the relaxation of value opt , it constructs an integral routing of size $\text{opt}/\mathcal{O}(\sqrt{r} \cdot \log^{1.5} rk)$.*

In particular, our algorithm strengthens the best known approximation algorithm for MAXEDP on general graphs [12] as always $r \leq n$, and indeed it matches that algorithm’s performance up to polylogarithmic factors. Substantially improving upon our bounds would also improve the current state of the art of MAXEDP. Conversely, the result implies that it suffices to study graphs with close to linear feedback vertex set number in order to improve the currently best upper bound of $\mathcal{O}(\sqrt{n})$ on the approximation ratio [12].

Our algorithmic approaches harness the forest structure of $G - R$ for any feedback vertex set R . However, the technical challenge comes from the fact that the edge set running between $G - R$ and R is unrestricted. Therefore, the “interaction” between R and $G - R$ is non-trivial, and flow paths may run between the two parts in an arbitrary manner and multiple times. In fact, we show that MAXEDP is already NP-hard if R consists of a *single node* (Theorem 6); this contrasts the efficient solvability on forests [24].

In order to overcome the technical hurdles we propose several new concepts, which we believe could be of interest in future studies of disjoint paths or routing problems.

In the randomized rounding approach of Raghavan and Thompson [38], it is shown that the probability that the congestion on any fixed edge is larger than $c \frac{\log n}{\log \log n}$ for some constant c is at most $1/n^{\mathcal{O}(1)}$. Combining this with the fact that there are at most n^2 edges,

yields that every edge has bounded congestion w.h.p. The number of edges in the graph may, however, be unbounded in terms of r and k . Hence, in order to prove Theorem 2, we propose a non-trivial *pre-processing step* of the optimum LP solution that is applied prior to the randomized rounding. In this step, we aggregate the flow paths by a careful rerouting so that the flow “concentrates” in $O(kr^2)$ nodes (so-called *hot spots*) in the sense that if all edges incident on hot spots have low congestion then so have all edges in the graph. Unfortunately, for any such hot spot the number of incident edges carrying flow may still be unbounded in terms of k and r . We are, however, able to give a refined probabilistic analysis that suitably relates the probability that the congestion bound is exceeded to the amount of flow on that edge. Since the total amount of flow on each hot spot is bounded in terms of k , the probability that *all* edges incident on the same hot spot have bounded congestion is inverse polynomial in r and k .

The known $\mathcal{O}(\sqrt{n})$ -approximation algorithm for MAXEDP by Chekuri et al. [12] employs a clever LP-rounding approach. If there are many long paths then there must be a single node carrying a significant fraction of the total flow and a good fraction of this flow can be realized by integral paths by solving a single-source flow problem. If the LP solution contains many short flow paths then greedily routing these short paths yields the bound since each such path blocks a bounded amount of flow. In order to prove Theorem 3, it is natural to consider the case where there are many paths visiting a large number of nodes in R . In this case, we reduce to a single-source flow problem, similarly to the approach of Chekuri et al. The case where a majority of the flow paths visit only a few nodes in R turns out more challenging, since any such path may still visit an unbounded number of edges in terms of k and r . We use two main ingredients to overcome these difficulties. First, we apply our Theorem 2 as a building block to obtain a solution with logarithmic congestion while losing only a constant factor in the approximation ratio. Second, we introduce the concept of *irreducible routings with low congestion* which allows us exploit the structural properties of the graph and the congestion property to identify a sufficiently large number of flow paths blocking only a small amount of flow.

Note that the natural greedy approach of always routing the shortest conflict-free path gives only $\mathcal{O}(\sqrt{m})$ for MAXEDP. We believe that it is non-trivial to obtain our bounds via a more direct or purely combinatorial approach.

Our third result is a fixed-parameter algorithm for MAXNDP in $k + r$.

► **Theorem 4.** MAXNDP can be solved in time $(8k + 8r)^{2r+2} \cdot \mathcal{O}(n)$ on graphs with feedback vertex set number r and k terminal pairs.

This run time is polynomial for constant r . We also note that for small r , our algorithm is asymptotically significantly faster than the fastest known algorithm for NDP, by Kawarabayashi and Wollan [29], which requires time at least *quadruple-exponential* in k [1]. Namely, if r is at most triple-exponential in k , our algorithm is asymptotically faster than theirs. We achieve this result by the idea of so-called *essential pairs* and *realizations*, which characterizes the “interaction” between the feedback vertex set R and the paths in an optimum solution. Note that in our algorithm of Theorem 4, parameter k does not appear in the exponent of the run time at all. Hence, for small values of r , our algorithm is also faster than reducing MAXNDP to NDP by guessing the subset of pairs to be routed (at an expense of 2^k in the run time) and using Scheffler’s [41] algorithm for NDP with run time $2^{\mathcal{O}(r \log r)} \cdot \mathcal{O}(n)$.

Once a fixed-parameter algorithm for a problem has been obtained, the existence of a polynomial-size kernel comes up. Here we note that MAXNDP does not admit a polynomial kernel for parameter $k + r$, unless $\text{NP} \subseteq \text{coNP}/\text{poly}$ [8].

■ **Table 1** Summary of results obtained in this paper.

const.	param.	EDP	MAXEDP	NDP	MAXNDP
$r = 0$		poly [24]	poly [24]	poly [41]	poly (Thm. 4)
$r = 1$		<i>open</i>	NP-hard (Thm. 6)	poly [41]	poly (Thm. 4)
$r \geq 2$		NP-hard (Thm. 6)	NP-hard (Thm. 6)	poly [41]	poly (Thm. 4)
	r	para-NP-hard (Thm. 6)		FPT [41]	W[1]-hard (Thm. 5)
		$\mathcal{O}(\sqrt{r} \cdot \log^{1.5} kr)$ -approx (Thm. 3)		exact $(k+r)^{\mathcal{O}(r)}$ (Thm. 4)	
		$\mathcal{O}(1)$ -approx. w.cg. $\mathcal{O}\left(\frac{\log kr}{\log \log kr}\right)$ (Thm. 2)			

Another natural question is whether the run time $f(k, r) \cdot n$ in Theorem 4 can be improved to $f(r) \cdot n^{\mathcal{O}(1)}$. We answer this question in the negative, ruling out the existence of a fixed-parameter algorithm for MAXNDP parameterized by r (assuming $\text{FPT} \neq \text{W}[1]$):

► **Theorem 5.** *MAXNDP in unit-capacity graphs is W[1]-hard parameterized by r .*

This contrasts the known result that NDP is fixed-parameter tractable in r [41]—which further stresses the relevance of understanding this parameter. We prove Theorem 5 in the full version of the paper [22].

For MAXEDP, we prove that the situation is, in a sense, even worse:

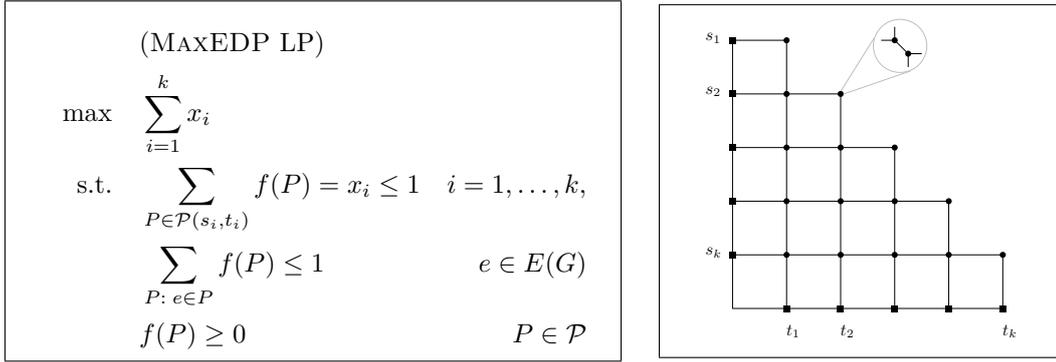
► **Theorem 6.** *MAXEDP is NP-hard for unit-capacity graphs with $r = 1$ and EDP is NP-hard for unit-capacity graphs with $r = 2$.*

This theorem also shows that our algorithms are relevant for small values of r , and they nicely complement the NP-hardness for MAXEDP in capacitated trees [24].

Our results are summarized in Table 1.

Related Work. Our study of the feedback vertex set number is in line with the general attempt to obtain bounds for MaxEDP (or related problems) that are independent of the input size. Besides the above-mentioned works that provide bounds in terms of the *tree-width* of the input graph, Günlük [25] and Chekuri et al. [17] give bounds on the *flow-cut gap* for the closely related integer multicommodity flow problem that are logarithmic with respect to the *vertex cover number* of a graph. This improved upon earlier bounds of $\mathcal{O}(\log n)$ [34] and $\mathcal{O}(\log k)$ [5, 35]. As every feedback vertex set is in particular a vertex cover of a graph, our results generalize earlier work for disjoint path problems on graphs with bounded vertex cover number. Bodlaender et al. [8] showed that NDP does not admit a polynomial kernel parameterized by vertex cover number *and* the number k of terminal pairs, unless $\text{NP} \subseteq \text{coNP}/\text{poly}$; therefore, NDP is unlikely to admit a polynomial kernel in $r+k$ either. Ene et al. [21] showed that MAXNDP is W[1]-hard parameterized by treedepth, which is another restriction of treewidth that is incomparable to the feedback vertex set number.

The basic gap in understanding the approximability of MaxEDP has led to several improved results for special graph classes, and also our results can be seen in this light. For example, polylogarithmic approximation algorithms are known for graphs whose global minimum cut value is $\Omega(\log^5 n)$ [39], for bounded-degree expanders [10, 9, 30, 34, 23], and for Eulerian planar or 4-connected planar graphs [28]. Constant factor approximation algorithms are known for capacitated trees [24, 14], grids and grid-like graphs [4, 6, 31, 32]. For planar graphs, there is a constant-factor approximation algorithm with congestion 2 [42]. Very



■ **Figure 1** Multi-commodity flow relaxation for MAXEDP. Right: $\Omega(\sqrt{n})$ integrality gap for MAXEDP [24]: any integral routing routes at most one pair, whereas a multi-commodity flow can send $1/2$ unit of flow for each pair (s_i, t_i) along the canonical path from s_i to t_i in the grid.

recently, Chuzhoy et al. [19] gave a $\tilde{\mathcal{O}}(n^{9/19})$ -approximation algorithm for MaxNDP on *planar* graphs. However, improving the $\mathcal{O}(\sqrt{n})$ -approximation algorithm for MAXEDP remains elusive even for *planar* graphs.

2 Preliminaries

We use standard graph theoretic notation. For a graph G , let $V(G)$ denote its vertex set and $E(G)$ its edge set. Let G be a graph. A *feedback vertex set* of G is a set $R \subseteq V(G)$ such that $G - R$ is a forest. A *minor* of G is a graph H that is obtained by successively contracting edges from a subgraph of G (and deleting any occurring loops). A class \mathcal{G} of graphs is *minor-closed* if for any graph in \mathcal{G} also all its minors belong to \mathcal{G} .

For an instance (G, \mathcal{M}) of MAXEDP/MAXNDP, we refer to the vertices participating in the pairs \mathcal{M} as *terminals*. It is convenient to assume that \mathcal{M} forms a matching on the terminals; this can be ensured by making several copies of a terminal and attaching them as leaves.

Multi-commodity flow relaxation. We use the following standard multi-commodity flow relaxation for MAXEDP (there is an analogous relaxation for MAXNDP). We use $\mathcal{P}(u, v)$ to denote the set of all paths in G from u to v , for each pair (u, v) of nodes. Since the pairs \mathcal{M} form a matching, the sets $\mathcal{P}(s_i, t_i)$ are pairwise disjoint. Let $\mathcal{P} = \bigcup_{i=1}^k \mathcal{P}(s_i, t_i)$. The LP has a variable $f(P)$ for each path $P \in \mathcal{P}$ representing the amount of flow on P . For each pair $(s_i, t_i) \in \mathcal{M}$, the LP has a variable x_i denoting the total amount of flow routed for the pair (in the corresponding IP, x_i denotes whether the pair is routed or not). The LP imposes the constraint that there is a flow from s_i to t_i of value x_i . Additionally, the LP has constraints that ensure that the total amount of flow on paths using a given edge (resp. node for MAXNDP) is at most 1.

It is well-known that the relaxation MAXEDP LP can be solved in polynomial time, since there is an efficient separation oracle for the dual LP (alternatively, one can write a compact relaxation). We use (f, \mathbf{x}) to denote a feasible solution to MAXEDP LP for an instance (G, \mathcal{M}) of MAXEDP. For each terminal v , let $x(v)$ denote the total amount of flow routed for v and we refer to $x(v)$ as the *marginal value* of v in the multi-commodity flow f .

We will use the following result by Chekuri et al. [12, Sect. 3.1]; see also Proposition 3.3 of Chekuri et al. [16].

► **Proposition 7.** *Let (f, \mathbf{x}) be a fractional solution to the LP relaxation of a MAXEDP instance (G, \mathcal{M}) . If some node v is contained in all flow paths of f , then we can find an integral routing of size at least $\frac{1}{12} \sum_i x_i$ in polynomial time.*

3 Bi-Criteria Approximation for MaxEDP with Low Congestion

We present a randomized rounding algorithm that will lead to the proof of Theorem 2.

3.1 Algorithm

Consider an instance (G, \mathcal{M}) of MAXEDP. Let R be a 2-approximate minimum feedback vertex set of G and let $r = |R|$; note that such a set R can be obtained in polynomial time [7].

For the sake of easier presentation, we will assume in this section that the feedback vertex set R contains all terminal nodes from \mathcal{M} . This can be achieved by temporarily adding the set of terminals to the feedback vertex set R . Also note that this assumption increases the bound of Theorem 2 by at most a constant factor.

First, solve the corresponding MAXEDP LP. We obtain an optimal extreme point solution (f, \mathbf{x}) . For each $(s_i, t_i) \in \mathcal{M}$, this gives us a set $\mathcal{P}'(s_i, t_i) = \{P \in \mathcal{P}(s_i, t_i) \mid f(P) > 0\}$ of positive weighted paths that satisfy the LP constraints.

Since we have an extreme point solution, the number of tight constraints is not smaller than the number of variables. As the number of constraints that are not of type $f(P) \geq 0$ is polynomially bounded in the input size, the same holds for the cardinality of the set $\mathcal{P}' = \bigcup_{i=1}^k \mathcal{P}'(s_i, t_i)$. In what follows, we will modify \mathcal{P}' and then select an (unweighted) subset \mathcal{S} of \mathcal{P}' that will form our integral solution.

Each $P \in \mathcal{P}'$ has the form $(r_1, \dots, r_2, \dots, r_\ell)$ where r_1, \dots, r_ℓ are the nodes in R that are traversed by P in this order. The paths (r_j, \dots, r_{j+1}) with $j = 1, \dots, \ell - 1$ are called *subpaths* of P . For every subpath P' of P , we set $f(P') = f(P)$. Let \mathcal{J} be the multi-set of all subpaths of all paths in \mathcal{P}' . Let $F = G - R$ be the forest obtained by removing R .

We now modify some paths in \mathcal{P}' , one by one, and at the same time construct a subset H_0 of nodes that we will call “hot spots”. At the end, every subpath in \mathcal{J} will contain at least one hot spot.

Initially, let $H_0 = \emptyset$. Consider any tree T in F and fix any of its nodes as a root. Then let \mathcal{J}_T be the multi-set of all subpaths in \mathcal{J} that, excluding the endpoints, are contained in T . For each subpath $P \in \mathcal{J}_T$, define its *highest node* $h(P)$ as the node on P closest to the root. Note that $P \cap T = P \cap F$ is a path. Now, pick a subpath $P \in \mathcal{J}_T$ that does not contain any node in H_0 and whose highest node $h(P)$ is *farthest away* from the root. Consider the multi-set $\mathcal{J}[P]$ of all subpaths in \mathcal{J}_T that are identical to P (but may be subpaths of different flow paths in \mathcal{P}'). Note that the weight $f(\mathcal{J}[P]) := \sum_{P' \in \mathcal{J}[P]} f(P')$ of $\mathcal{J}[P]$ is at most 1 by the constraints of the LP. Let $u, v \in R$ be the endpoints of P . We define \mathcal{J}_{uv} as the set of all subpaths in $\mathcal{J} \setminus \mathcal{J}[P]$ that have u and v as their endpoints and that do not contain any node in H_0 .

Intuitively speaking, we now aggregate flow on P by rerouting as much flow as possible from \mathcal{J}_{uv} to P . To this end, we repeatedly perform the following operation as long as $f(\mathcal{J}[P]) < 1$ and $\mathcal{J}_{uv} \neq \emptyset$. We pick a path P' in \mathcal{J} that contains a subpath in \mathcal{J}_{uv} . We reroute flow from P' by creating a new path P'' that arises from P' by replacing its subpath between u and v with P , and assign it the weight $f(P'') = \min\{f(P'), 1 - f(\mathcal{J}[P])\}$. Then we set the weight of (the original path) P' to $\max\{0, f(P') + f(\mathcal{J}[P]) - 1\}$. We update the sets \mathcal{P}' , $\mathcal{P}'(s_i, t_i)$, \mathcal{J} , \mathcal{J}_T , $\mathcal{J}[P]$ and \mathcal{J}_{uv} accordingly.

As soon as $f(\mathcal{J}[P]) = 1$ or $\mathcal{J}_{uv} = \emptyset$, we add $h(P)$ to H_0 . Then, we proceed with the next $P \in \mathcal{J}_T$ not containing a hot spot and whose highest node $h(P)$ is farthest away from the root. If no such P is left, we consider the next tree T in F .

At the end, we create our solution \mathcal{S} by randomized rounding: We route every terminal pair (s_i, t_i) with probability x_i . In case (s_i, t_i) is routed, we randomly select a path from $\mathcal{P}'(s_i, t_i)$ and add it to \mathcal{S} where the probability that path P is taken is $f(P)/x_i$.

3.2 Analysis

First, observe that \mathbf{x} did not change during our modifications of the paths, as the total flow between any terminal pair did not change. Thus, the expected number of pairs routed in our solution is $\sum_{i=1}^k x_i \geq \text{OPT}$. Using the Chernoff bound, the probability that we route less than $\text{OPT}/2$ pairs is at most $e^{-1/8 \text{OPT}} < 1/2$, assuming that $\text{OPT} > 8$. Secondly, we bound the congestion of our solution—our second criterion.

► **Lemma 8.** *The congestion of flow f is at most 2.*

Proof. In our algorithm, we increase the flow only along flow subpaths that are pairwise edge-disjoint. To see this, consider two distinct flow subpaths P and P' on which we increase the flow. Assume, without loss of generality, that P was considered before P' by the algorithm. If there was an edge e lying on P and P' , then both subpaths traverse the same tree in forest F . Hence, the path from e to $h(P')$ would visit $h(P)$, and $h(P)$ would be an internal node of P' . This yields a contradiction, as $h(P)$ was already marked as a hot spot when P' was considered. This shows that we increased the flow along any edge by at most one unit, and, hence, f has congestion at most 2. ◀

We now bound the congestion of the integral solution obtained by randomized rounding. In the algorithm, we constructed a set H_0 of hot spots. As a part of the analysis, we will now extend this set to a set H as follows. Initially, $H = H_0$. We build a sub-forest F' of F consisting of all edges of F that lie on a path connecting two hot spots. Then we add to H all nodes that have degree at least 3 in F' . Since the number of nodes of degree 3 in any forest is at most its number of leaves and since every leaf of F' is a hot spot, it follows that this can at most double the size of H to $2|H_0|$. Finally, we add the set R of all feedback vertex nodes to H . In the following, all nodes in H are called hot spots.

► **Lemma 9.** *The number $|H|$ of hot spots is $\mathcal{O}(kr^2)$.*

Proof. It suffices to show that $|H_0| \in \mathcal{O}(kr^2)$. To this end, fix two nodes $u, v \in R$ and consider the set of flow subpaths P with end nodes u and v for which we added $h(P)$ to H_0 . Due to the aggregation of flows in our algorithm, all except possibly one of the subpaths are saturated, that is, they carry precisely one unit of flow. Since no two of these subpaths are contained in a same flow path of f and since the flow value of f is bounded from above by k , we added only $\mathcal{O}(k)$ hot spots for the pair u, v . Since there are at most r^2 pairs in R , the claim follows. ◀

► **Definition 10.** A hot spot $u \in H$ is *good* if the congestion on any edge incident on u is bounded by $c \cdot \frac{\log kr}{\log \log kr}$, where c is a sufficiently large constant; otherwise, u is *bad*.

► **Lemma 11.** *Let $u \in H$ be a hot spot. Then the probability that u is bad is at most $1/(k^2 r^3)$.*

Proof. Let $e_1 = uv_1, \dots, e_\ell = uv_\ell$ be the edges incident on u and let f_i be the total flow on edge uv_i for $i = 1, \dots, \ell$. By Lemma 8, we have that $f_i \leq 2$. Since any flow path visits at

most two of the edges incident on u , the total flow $\sum_{i=1}^{\ell} f_i$ on the edges incident on u is at most $2k$.

For any $i = 1, \dots, \ell$, we have that $f_i = \sum_{P: P \ni e_i} f(P)$, where P runs over the set of all paths connecting some terminal pair and containing e_i . Let $f_{ij} = \sum_{P \in \mathcal{P}(s_j, t_j): P \ni e_i} f(P)$ be the total amount of flow sent across e_i by terminal pair (s_j, t_j) . Recall that x_j is the total flow sent for terminal pair (s_j, t_j) . The probability that the randomized rounding procedure picks path P with $P \in \mathcal{P}(s_j, t_j)$ is precisely $x_j \cdot \frac{f(P)}{x_j} = f(P)$. Given the disjointness of the respective events, the probability that pair (s_j, t_j) routes a path across e_i is precisely f_{ij} . Let X_{ij} be the binary random variable indicating whether pair (s_j, t_j) routes a path across e_i . Then $\mathbb{P}[X_{ij} = 1] = f_{ij}$. Let $X_i = \sum_j X_{ij}$ be the number of paths routed across e_i by the algorithm. By linearity of expectation, we have that $\mathbb{E}[X_i] = \sum_j \mathbb{E}[X_{ij}] = \sum_j f_{ij} = f_i$.

Fix any edge e_i . Set $\delta = c \cdot \frac{\log kr}{\log \log kr}$ and $\delta' = 2\frac{\delta}{f_i} - 1$. Note that for fixed i , the variables X_{ij} are independent. Hence, by the Chernoff bound, we have that

$$\begin{aligned} \mathbb{P}\left[X_i \geq c \cdot \frac{\log kr}{\log \log kr}\right] &\leq \mathbb{P}[X_i \geq (1 + \delta')f_i] < \left(\frac{e^{\delta'}}{(1 + \delta')^{1 + \delta'}}\right)^{f_i} \\ &\leq \left(\frac{f_i}{2}\right)^{2\delta} \cdot \left(\frac{\delta}{e}\right)^{-2\delta} \leq f_i e^{-c' \log \log kr \cdot \frac{\log kr}{\log \log kr}} \leq \frac{f_i}{2k^3 r^3}. \end{aligned}$$

Here, we use that $f_i \leq 2$ for the second last inequality and for the last inequality we pick c' sufficiently large by making c and k sufficiently large. (Note that MAXEDP can be solved efficiently for constant k .)

Now, using the union bound, we can infer that the probability that any of the edges incident on u carries more than δ paths is at most $\sum_i f_i / (2k^3 r^3) \leq (2k) / (2k^3 r^3) = 1 / (k^2 r^3)$. ◀

► **Lemma 12.** *Assume that every hot spot is good. Then the congestion on any edge is bounded by $2c \frac{\log kr}{\log \log kr}$.*

Proof. Consider an arbitrary edge $e = uv$ that is not incident on any hot spot. In particular, this means that e lies in the forest $F = G - R$. A hot spot z in F is called *direct to e* if the path in F from z to e excluding e does not contain any hot spot other than z .

We claim that there can be at most two distinct hot spots z, z' direct to e . If there were a third hot spot z'' direct to e , then consider the unique node $z_0 \in V(F)$ such that no two of the hot spots z, z', z'' are connected in $F - z_0$. Such a node z_0 exists since z, z', z'' cannot lie on a common path in F since they are all direct to e . The node z_0 , however, would be added as a hot spot at the latest when H was built. Now, this is a contradiction, because then one of the paths connecting z, z' or z'' to e would contain z_0 and thus one of these hot spots would not be direct to e .

Now let P be an arbitrary path that is routed by our algorithm and that traverses e , and let $P' \in \mathcal{J}$ be the subpath of P in F visiting e . Moreover, let $P_z, P_{z'}$ be the paths in F connecting z, z' to e excluding e , and let $e_z, e_{z'}$ be the edges on these paths incident on z, z' , respectively. By our construction, P' must visit a hot spot in F . If P' visited neither z nor z' , then P' would contain a hot spot direct to u or to v that is distinct from z and z' —a contradiction. Hence P' and thus also P visit e_z or $e_{z'}$. The claim now follows from the fact that this holds for any path traversing e , that z and z' are good, and that therefore altogether at most $2c \frac{\log kr}{\log \log kr}$ paths visit e_z or $e_{z'}$. ◀

► **Theorem 13.** *The algorithm from Sect. 3.1 produces—with constant probability—a routing with $\Omega(\text{OPT})$ paths, such that the congestion is $\mathcal{O}\left(\frac{\log kr}{\log \log kr}\right)$.*

Proof. As argued above, we route less than $\text{OPT}/2$ paths with probability at most $1/2$. By Lemma 9, there are $\mathcal{O}(kr^2)$ hotspots. The probability that at least one of these hot spots is bad is $\mathcal{O}(kr^2/(k^2r^3)) = \mathcal{O}(1/(kr))$, by Lemma 11. Hence, with constant probability, we route at least $\text{OPT}/2$ pairs with congestion at most $2c_{\frac{\log kr}{\log \log kr}}$, by Lemma 12. ◀

4 Refined Approximation Bound for MaxEDP

In this section, we provide an improved approximation guarantee for MAXEDP *without* congestion, thereby proving Theorem 3. (In contrast to the previous section, we do not assume here that all terminals are contained in the feedback vertex set.)

4.1 Irreducible Routings with Low Congestion

We first develop the concept of *irreducible routings with low congestion*, which is (besides Theorem 2) a key ingredient of our strengthened bound on the approximability of MAXEDP based on the feedback vertex number.

Consider any multigraph G and any set \mathcal{P} of (not necessarily simple) paths in G with congestion c . We say that an edge e is *redundant in \mathcal{P}* if there is an edge $e' \neq e$ such that the set of paths in \mathcal{P} covering (containing) e is a subset of the set of paths in \mathcal{P} covering e' . Thus, any edge that is not covered by any path in \mathcal{P} is redundant in \mathcal{P} if G contains at least two edges.

► **Definition 14.** Set \mathcal{P} is called an *irreducible routing with congestion c* if each edge belongs to at most c paths of \mathcal{P} and there is no edge redundant in \mathcal{P} .

In contrast to a feasible routing of an MAXEDP instance, we do not require an irreducible routing to connect a set of terminal pairs. If there is an edge e redundant in \mathcal{P} , we can apply the following *reduction rule*: We contract e in G and we contract e in every path of \mathcal{P} that covers e . By this, we obtain a minor G' of G and a set \mathcal{P}' of paths that consists of all the contracted paths and of all paths in \mathcal{P} that were not contracted. Thus, there is a one-to-one correspondence between the paths in \mathcal{P} and \mathcal{P}' .

We make the following observation about \mathcal{P} and \mathcal{P}' .

► **Observation 15.** *A subset of paths in \mathcal{P}' is edge-disjoint in G' if and only if the corresponding subset of paths in \mathcal{P} is edge-disjoint in G .*

As applying the reduction rule strictly decreases the number of redundant edges, an iterative application of this rule yields an irreducible routing on a minor of the original graph.

► **Theorem 16.** *Let \mathcal{G} be a minor-closed class of multigraphs and let $p_{\mathcal{G}} > 0$. If for each graph $G \in \mathcal{G}$ and every non-empty irreducible routing \mathcal{S} of G with congestion c there exists a path in \mathcal{S} of length at most $p_{\mathcal{G}}$, then the average length of the paths in \mathcal{S} is at most $c \cdot p_{\mathcal{G}}$.*

Proof. Take a path P_0 of length at most $p_{\mathcal{G}}$. Contract all edges of P_0 in G and obtain a minor $G' \in \mathcal{G}$ of G . For each path in \mathcal{S} contract all edges shared with P_0 to obtain a set \mathcal{S}' of paths. Remove P_0 along with all degenerated paths from \mathcal{S}' , thus $|\mathcal{S}'| < |\mathcal{S}|$. Note that \mathcal{S}' is an irreducible routing of G' with congestion c . We repeat this reduction procedure recursively on G' and \mathcal{S}' until \mathcal{S}' is empty which happens after at most $|\mathcal{S}|$ steps. At each step we decrease the total path length by at most $c \cdot p_{\mathcal{G}}$. Hence, the total length of paths in \mathcal{S} is at most $|\mathcal{S}| \cdot c \cdot p_{\mathcal{G}}$. ◀

As a consequence of Theorem 16, we get the following result for forests.

► **Lemma 17.** *Let F be a forest and let \mathcal{S} be a non-empty irreducible routing of F with congestion c . Then the average path length in \mathcal{S} is at most $2c$.*

Proof. We show that \mathcal{S} contains a path of length at most 2 . The lemma follows immediately by applying Theorem 16.

Take any tree in F , root it with any node and consider a leaf v of maximum depth. If v is adjacent to the root, then the tree is a star and every path in the tree has length at most 2 . Otherwise, let e_1 and e_2 be the first two edges on the path from v to the root. By definition of irreducible routing, the set of all paths covering e_1 is not a subset of the paths covering e_2 ; hence, e_1 is covered by a path which does not cover e_2 . Since all other edges incident to e_1 end in a leaf, this path has length at most 2 . ◀

Note that the bound provided in Lemma 17 is actually tight up to a constant. Let $c \geq 1$ be an arbitrary integer. Consider a graph that is a path of length $c - 1$ with a star of $c - 1$ leaves attached to one of its end points. The $c - 1$ paths of length c together with the $2c - 2$ paths of length 1 form an irreducible routing with congestion c . The average path length is $((c - 1)c + (2c - 2))/(3c - 3) = (c + 2)/3$.

4.2 Approximation Algorithm

Consider an instance (G, \mathcal{M}) of MAXEDP, and let r be the size of a feedback vertex set R in G . Using our result of Sect. 3, we can efficiently compute a routing \mathcal{P} with congestion $c := \mathcal{O}\left(\frac{\log kr}{\log \log kr}\right)$ containing $\Omega(\text{OPT})$ paths.

Below we argue how to use the routing \mathcal{P} to obtain a feasible routing of cardinality $\Omega(|\mathcal{P}|/(c^{1.5}\sqrt{r}))$, which yields an overall approximation ratio of $\mathcal{O}(\sqrt{r} \cdot \log^{1.5} rk)$; that will prove Theorem 3.

Let $r' = \sqrt{r/c}$. We distinguish the following cases.

Case 1: At least half of the paths in \mathcal{P} visit at most r' nodes of the feedback vertex set R . Let $\overline{\mathcal{P}}$ be the subset of these paths. Initialize \mathcal{P}' with $\overline{\mathcal{P}}$. As long as there is an edge e not adjacent to R that is redundant in \mathcal{P}' , we iteratively apply the reduction rule from Sect. 4.1 on e . Let G' be the obtained minor of G with forest $F' = G' - R$. The obtained set \mathcal{P}' is a set of (not necessarily simple) paths in G' corresponding to $\overline{\mathcal{P}}$. By (iterated application of) Observation 15 to path sets $\overline{\mathcal{P}}$ and \mathcal{P}' , it suffices to show that there is a subset $\mathcal{P}'_0 \subseteq \mathcal{P}'$ of pairwise edge-disjoint paths of size $|\mathcal{P}'_0| = \Omega(|\mathcal{P}'|/(cr'))$ in order to obtain a feasible routing for (G, \mathcal{M}) of size $\Omega(|\mathcal{P}'|/(cr'))$.

To obtain \mathcal{P}'_0 , we first bound the total path length in \mathcal{P}' . Removing R from G' “decomposes” the set \mathcal{P}' into a set $\mathcal{S} := \{S \text{ is a connected component of } P \cap F' \mid P \in \mathcal{P}'\}$ of subpaths lying in F' . Observe that \mathcal{S} is an irreducible set of F' with congestion c , as the reduction rule is not applicable anymore. (Note that a single path in \mathcal{P}' may lead to many paths in the cover \mathcal{S} which are considered distinct.) Thus, by Lemma 17, the average path length in \mathcal{S} is at most $2c$.

Let P be an arbitrary path in \mathcal{P}' . Each edge on P that is *not* in a subpath in \mathcal{S} is incident on a node in R , and each node in R is incident on at most two edges in P . Together with the fact that P visits at most r' nodes in R and that the average length of the subpaths in \mathcal{S} is at most $2c$, we can upper bound the total path length $\sum_{P \in \mathcal{P}'} |P|$ by $|\mathcal{P}'|r'(2c + 2)$. Let \mathcal{P}'' be the set of the $|\mathcal{P}'|/2$ shortest paths in \mathcal{P}' . Hence, each path in \mathcal{P}'' has length at most $4r'(c + 1)$.

We greedily construct a feasible solution \mathcal{P}'_0 by iteratively picking an arbitrary path P from \mathcal{P}'' adding it to \mathcal{P}'_0 and removing all paths from \mathcal{P}'' that share some edge with P (including P itself). We stop when \mathcal{P}'' is empty. As \mathcal{P}'' has congestion c , we remove at most $4r'c(c+1)$ paths from \mathcal{P}'' per iteration. Thus, $|\mathcal{P}'_0| \geq |\mathcal{P}''|/(4r'c(c+1)) = \Omega(|\mathcal{P}|/(c^{1.5}\sqrt{r}))$.

Case 2: At least half of the paths in \mathcal{P} visit at least r' nodes of the feedback vertex set R . Let \mathcal{P}' be the subset of these paths. Consider each path in \mathcal{P}' as a flow of value $1/c$ and let f be the sum of all these flows. Note that f provides a feasible solution to the MAXEDP LP relaxation for (G, M) of value at least $|\mathcal{P}|/(2c)$. Note that each such flow path contributes $1/c$ unit of flow to each of the r' nodes in R it visits. Since every flow path in f has length at least r' , the total inflow of the nodes in R is at least $|f|r'$. By averaging, there must be a node $v \in R$ of inflow at least $r'|f|/r = |f|/r'$. Let f' be the subflow of f consisting of all flow paths visiting v . This subflow corresponds to a feasible solution (f', \mathbf{x}') of the LP relaxation of value at least $|f|/r' \geq |\mathcal{P}|/(2cr')$. Using Proposition 7, we can recover an integral feasible routing of size at least $\frac{1}{12} \sum_i x'_i \geq |\mathcal{P}|/(24cr') = \Omega(|\mathcal{P}|/(c^{1.5}\sqrt{r}))$.

This completes the proof of Theorem 3. ◀

5 Fixed-Parameter Algorithm for MaxNDP

We give a fixed-parameter algorithm for MAXNDP with run time $(k+r)^{\mathcal{O}(r)} \cdot n$, where r is the size of a minimum feedback vertex set in the given instance (G, \mathcal{M}) . A feedback vertex set R of size r can be computed in time $2^{\mathcal{O}(r)} \cdot \mathcal{O}(n)$ [36]. By the matching assumption, each terminal in \mathcal{M} is a leaf. We can thus assume that none of the terminals is contained in R .

Consider an optimal routing \mathcal{P} of the given MAXNDP instance. Let $\mathcal{M}_R \subseteq \mathcal{M}$ be the set of terminal pairs that are connected via \mathcal{P} by a path that visits at least one node in R . Let $P \in \mathcal{P}$ be a path connecting a terminal pair $(s_i, t_i) \in \mathcal{M}_R$. This path has the form $(s_i, \dots, r_1, \dots, r_2, \dots, r_\ell, \dots, t_i)$, where r_1, \dots, r_ℓ are the nodes in R that are traversed by P in this order. The pairs (s_i, r_1) , (r_ℓ, t_i) and (r_j, r_{j+1}) with $j = 1, \dots, \ell - 1$ are called *essential* pairs for P . A node pair is called *essential* if it is essential for some path in \mathcal{P} . Let \mathcal{M}_e be the set of essential pairs.

Let F be the forest that arises when deleting R from the input graph G . Let (u, v) be an essential pair. A u - v path P in G is said to *realize* (u, v) if all internal nodes of P lie in F . A set \mathcal{P}' of paths is said to *realize* \mathcal{M}_e if every pair in \mathcal{M}_e is realized by some path in \mathcal{P}' and if two paths in \mathcal{P}' can only intersect at their end nodes. Note that the optimal routing \mathcal{P} induces a natural realization of \mathcal{M}_e , by considering all maximal subpaths of paths in \mathcal{P} whose internal nodes all lie in F . Conversely, for any realization \mathcal{P}' of \mathcal{M}_e , we can concatenate paths in \mathcal{P}' to obtain a feasible routing that connects all terminal pairs in \mathcal{M}_R . Therefore, we consider \mathcal{P}' (slightly abusing notation) also as a feasible routing for \mathcal{M}_R .

In our algorithm, we first guess the set \mathcal{M}_e (and thus \mathcal{M}_R). Then, by a dynamic program, we construct two sets of paths, \mathcal{P}_e and \mathcal{P}_F where \mathcal{P}_e realizes \mathcal{M}_e and \mathcal{P}_F connects in F a subset of $\overline{\mathcal{M}}_R := \mathcal{M} \setminus \mathcal{M}_R$. In our algorithm, the set $\mathcal{P}_e \cup \mathcal{P}_F$ forms a feasible routing that maximizes $|\mathcal{P}_F|$ and routes all pairs in \mathcal{M}_R . (Recall that we consider the realization \mathcal{P}_e of \mathcal{M}_e as a feasible routing for \mathcal{M}_R .)

Now assume that we know set \mathcal{M}_e . We will describe below a dynamic program that computes an optimum routing in time $2^{\mathcal{O}(r)}(k+r)^{\mathcal{O}(1)}n$. For the sake of easier presentation, we only describe how to compute the cardinality of such a routing.

We make several technical assumptions that help to simplify the presentation. First, we modify the input instance as follows. We subdivide every edge incident on a node in R by

introducing a single new node on this edge. Note that this yields an instance equivalent to the input instance. As a result, every neighbor of a node in R that lies in F , that is, every node in $N_G(R)$, is a leaf in F . Moreover, the set R is an independent set in G . Also recall that we assumed that every terminal is a leaf. Therefore, we may assume that R does not contain any terminal. We also assume that forest F is a rooted tree, by introducing a dummy node (which plays the role of the root) and arbitrarily connecting this node to every connected component of F by an edge. In our dynamic program, we will take care that no path visits this root node. We also assume that F is an ordered tree by introducing an arbitrary order among the children of each node.

For any node v , let F_v be the subtree of F rooted at v . Let $c_v := \deg_F(v) - 1$ be the number of children of v and let v_1, \dots, v_{c_v} be the (ordered) children of v . Then, for $i = 1, \dots, c_v$, let F_v^i denote the subtree of F_v induced by the union of v with the subtrees F_{v_1}, \dots, F_{v_i} . For leaves v , we define F_v^0 as $F_v = v$.

We introduce a dynamic programming table T . It contains an entry for every F_v^i and every subset \mathcal{M}'_e of \mathcal{M}_e . Roughly speaking, the value of such an entry is the solution to the subproblem, where we restrict the forest to F_v^i , and the set of essential pairs to \mathcal{M}'_e . More precisely, table T contains five parameters. Parameters v and i describing F_v^i , parameter \mathcal{M}'_e , and two more parameters u and b . Parameter u is either a terminal, or a node in R , and b is in one of the three states: *free*, *to-be-used*, or *blocked*. The value $T[v, i, \mathcal{M}'_e, u, b]$ is the maximum cardinality of a set \mathcal{P}_F of paths with the following properties:

1. \mathcal{P}_F is a feasible routing of some subset of $\overline{\mathcal{M}}_R$.
2. \mathcal{P}_F is completely contained in F_v^i .
3. There is an additional set \mathcal{P}_e of paths with the following properties:
 - a. \mathcal{P}_e is completely contained in $F_v^i \cup R$ and node-disjoint from the paths in \mathcal{P}_F .
 - b. \mathcal{P}_e is a realization of $\mathcal{M}'_e \cup \{(u, v)\}$ if $b = \textit{to-be-used}$. Else, it is a realization of \mathcal{M}'_e .
 - c. There is no path in $\mathcal{P}_e \cup \mathcal{P}_F$ visiting v if $b = \textit{free}$.

If no such set \mathcal{P}_F exists then $T[v, i, \mathcal{M}'_e, u, b]$ is $-\infty$.

Note that the parameter u is only relevant when $b = \textit{to-be-used}$ (otherwise, it can just be ignored). Observe that $T[v, i, \mathcal{M}'_e, u, \textit{blocked}] \geq T[v, i, \mathcal{M}'_e, u, \textit{free}] \geq T[v, i, \mathcal{M}'_e, u, \textit{to-be-used}]$. Below, we describe how to compute the entries of T in a bottom-up manner.

In the base case v is a leaf. We set $T[v, 0, \emptyset, u, \textit{free}] = 0$. Then we set $T[v, 0, \mathcal{M}'_e, u, \textit{blocked}] = 0$ if \mathcal{M}'_e is either empty, consists of a single pair of nodes in $R \cap N_G(v)$, or consists of a single pair where one node is v and the other one is in $R \cap N_G(v)$. Finally, we set $T[v, 0, \emptyset, u, \textit{to-be-used}] = 0$ if $u = v$ or u is in $R \cap N_G(v)$. For all other cases where v is a leaf, we set $T[v, i, \mathcal{M}'_e, u, b] = -\infty$.

For the inductive step, we consider the two cases $i = 1$ and $i > 1$. Let $i = 1$. It holds that $T[v, 1, \mathcal{M}'_e, u, \textit{to-be-used}] = T[v_1, c_v, \mathcal{M}'_e, u, \textit{to-be-used}]$ since the path in \mathcal{P}_e realizing (u, v) has to start at a leaf node of F_{v_1} . It also holds that $T[v, 1, \mathcal{M}'_e, u, \textit{blocked}]$ and $T[v, 1, \mathcal{M}'_e, u, \textit{free}]$ are equal to $T[v_1, c_v, \mathcal{M}'_e, u, \textit{blocked}]$.

Now, let $i > 1$. In a high level view, we guess which part of \mathcal{M}'_e is realized in $F_v^{i-1} \cup R$ and which part is realized in $F_{v_i} \cup R$. For this, we consider every tuple $(\mathcal{M}'_{e1}, \mathcal{M}'_{e2})$ such that $\mathcal{M}'_{e1} \uplus \mathcal{M}'_{e2}$ is a partition of \mathcal{M}'_e . By our dynamic programming table, we find a tuple that maximizes our objective. In the following, we assume that we guessed $(\mathcal{M}'_{e1}, \mathcal{M}'_{e2})$ correctly. Let us consider the different cases of b in more detail.

For $b = \textit{free}$, node v is not allowed to be visited by any path, especially by any path in $F_v^{i-1} \cup R$. Hence, $T[v, i, \mathcal{M}'_e, u, \textit{free}]$ is equal to

$$T[v, i-1, \mathcal{M}'_{e1}, u, \textit{free}] + T[v_i, c_{v_i}, \mathcal{M}'_{e2}, u, \textit{blocked}] .$$

In the case of $b = \text{to-be-used}$, we have to realize (u, v) in $F_v^i \cup R$. For this, there are two possibilities: either (u, v) is realized by a path in $F_v^{i-1} \cup R$, or there is a realizing path that first goes through $F_{v_i} \cup R$ and then reaches v via the edge (v_i, v) . Hence, for the first case, we consider

$$T[v, i-1, \mathcal{M}'_{e_1}, u, \text{to-be-used}] + T[v_i, c_{v_i}, \mathcal{M}'_{e_2}, u, \text{blocked}],$$

for the second case, we consider

$$T[v, i-1, \mathcal{M}'_{e_1}, u, \text{free}] + T[v_i, c_{v_i}, \mathcal{M}'_{e_2}, u, \text{to-be-used}].$$

Maximizing over both, we obtain $T[v, i, \mathcal{M}'_e, u, \text{to-be-used}]$.

For the case of $b = \text{blocked}$, we will consider two subcases. In the first subcase, there is no path in $\mathcal{P}_e \cup \mathcal{P}_F$ going through edge (v_i, v) , hence, we get

$$T[v, i-1, \mathcal{M}'_{e_1}, u, \text{blocked}] + T[v_i, c_{v_i}, \mathcal{M}'_{e_2}, u, \text{blocked}].$$

In the second subcase, there is a path P in $\mathcal{P}_e \cup \mathcal{P}_F$ going through edge (v_i, v) . Since P is connecting two leaves in F_v^i , a part of P is in $F_v^{i-1} \cup R$ and the other part is in $F_{v_i} \cup R$. If $P \in \mathcal{P}_e$, then it is realizing a pair of \mathcal{M}'_e . Hence, for every pair $(u_1, u_2) \in \mathcal{M}'_e$, we have to consider the term

$$T[v, i-1, \mathcal{M}'_{e_1} - (u_1, u_2), u_1, \text{to-be-used}] + T[v_i, c_{v_i}, \mathcal{M}'_{e_2} - (u_1, u_2), u_2, \text{to-be-used}]$$

and the symmetric term where we swap u_1 and u_2 . If $P \in \mathcal{P}_F$, then it is realizing a terminal pair of $\overline{\mathcal{M}}_R$. Hence, for every pair $(u_1, u_2) \in \overline{\mathcal{M}}_R$ we get the term

$$1 + T[v, i-1, \mathcal{M}'_{e_1}, u_1, \text{to-be-used}] + T[v_i, c_{v_i}, \mathcal{M}'_{e_2}, u_2, \text{to-be-used}]$$

and the symmetric term where we swap u_1 and u_2 . Note that we count the path realizing (u_1, u_2) in our objective. Maximizing over all the terms of the two subcases, we obtain $T[v, i, \mathcal{M}'_e, u, \text{to-be-used}]$.

Let us analyze the run time of algorithm described in Sect. 5. In order to guess \mathcal{M}_e , we enumerate all potential sets of essential pairs. There are at most $(2k + r + 1)^{2r}$ candidate sets to consider, since each pair contains a node in R , and each node in R is paired with at most two other nodes each of which is either a terminal or another node in R . For each particular guess \mathcal{M}_e , we run the above dynamic program. The number of entries in T —as specified by the five parameters v, i, \mathcal{M}'_e, u and b —for each fixed \mathcal{M}_e is at most $(\sum_{v \in V(F)} \deg_F(v)) \times 2^{2r} \times (2k + r) \times 3$. For the computation of each such entry, we consider all combinations of at most 2^{2r} partitions of \mathcal{M}'_e with either at most r essential pairs in \mathcal{M}'_e , or with at most k terminal pairs in $\overline{\mathcal{M}}_R$. Altogether, this gives a run time of $(8k + 8r)^{2r+2} \cdot \mathcal{O}(n)$. This finishes the proof of Theorem 4.

6 Hardness of Edge-Disjoint Paths in Almost-Forests

In this section we show that EDP (and hence MAXEDP) is NP-hard already in graphs that are forests after deleting two nodes. That is, we prove Theorem 6.

Proof of Theorem 6. We first show NP-hardness of EDP for $r = 2$. We reduce from the problem EDGE 3-COLORING in cubic graphs, which is NP-hard [26]. Given a cubic graph H , we construct a complete bipartite graph G , where one of the two partite classes of $V(G)$ consists of three nodes $\{v_1, v_2, v_3\}$, and the other partite class consists of $V(H)$. As terminal

pairs, we create the set $\mathcal{M} = \{(s, t) \mid \{s, t\} \in E(H)\}$; in words, we want to connect a pair of nodes by a path in G if and only if they are connected by an edge in H . This completes the construction of the instance (G, \mathcal{M}) of MAXEDP. Notice that G has a feedback vertex set of size $r = 2$, since removing any size-2 subset of $\{v_1, v_2, v_3\}$ from G yields a forest.

Regarding correctness of the reduction, we show that H is 3-edge-colorable if and only if *all* pairs in \mathcal{M} can be routed in G .

In the forward direction, suppose that H is 3-edge-colorable. Let $\varphi : E(H) \rightarrow \{1, 2, 3\}$ be a proper 3-edge-coloring of H . For $c = 1, 2, 3$, let $E_c \subseteq E(H)$ be the set of edges that receive color c under φ . Then there is a routing in G that routes all terminal pairs $\{(s, t) \in \mathcal{M} \mid \{s, t\} \in E_c\}$ exclusively via the node v_c (and thus via paths of length 2). Notice that this routing indeed yields edge-disjoint paths, for if there are distinct vertices $s, t_1, t_2 \in V(H)$ and edges $e_1 = \{s, t_1\}, e_2 = \{s, t_2\} \in E(H)$, then e_1, e_2 receive distinct colors under φ (as φ is proper), and so the two terminal pairs $\{s, t_1\}, \{s, t_2\}$ are routed via distinct nodes $c_1, c_2 \in \{v_1, v_2, v_3\}$, and thus also via edge-disjoint paths.

In the backward direction, suppose that all terminal pairs in \mathcal{M} can be routed in G . Since H is cubic, any node $s \in V(H)$ is contained in three terminal pairs. Therefore, no path of the routing can have a node in $V(H)$ as an internal node and thus all paths in the routing have length 2. Then this routing naturally corresponds to a proper 3-edge-coloring φ of H , where any terminal pair $\{s, t\}$ routed via c means that we color the edge $\{s, t\} \in E(H)$ with color c under φ .

In order to show NP-hardness of MAXEDP for $r = 1$, we also reduce from EDGE 3-COLORING in cubic graphs and perform a similar construction as described above: This time, we construct a bipartite graph G with one subset of the partition being $\{v_1, v_2\}$, the other being $V(H)$, and the set \mathcal{M} of terminal pairs being again specified by the edges of H . This completes the reduction. The resulting graph G has a feedback vertex set of size $r = 1$.

We claim that H is 3-colorable if and only if we can route $n = |V(H)|$ pairs in G .

In the forward direction, suppose that H is 3-edge-colorable. Let $\varphi : E(H) \rightarrow \{1, 2, 3\}$ be a proper 3-edge-coloring of H . For $c = 1, 2, 3$, let $E_c \subseteq E(H)$ be the set of edges that receive color c under φ . Then there is a routing in G that routes all $\{(s, t) \in \mathcal{M} \mid \{s, t\} \in E_c\}$ exclusively via the node v_c (and thus via paths of length 2) for the colors $c = 1, 2$. (The terminals corresponding to edges receiving color 3 remain unrouted.)

The reasoning that the resulting routing is feasible is analogous to the case of $r = 2$. Since for each of the n terminals exactly two of the three terminal pairs are routed, this means that precisely n terminal pairs are routed overall.

In the backward direction, suppose that n terminal pairs in \mathcal{M} can be routed in G . Since any terminal v in G is a node in $V(H)$ has therefore has degree two in G , this means that at most two paths can be routed for v . As n terminal pairs are realized, this also means that *exactly* two paths are routed for each terminal. Hence, none of the paths in the routing has length more than two. Otherwise, it would contain an internal node in $V(H)$, which then could not be part of two other paths in the routing. Then this routing naturally corresponds to a partial edge-coloring of H , where any terminal pair $\{s, t\}$ routed via c means that we color the edge $\{s, t\} \in E(H)$ with color c . Since each terminal v in $V(H)$ is involved in exactly two paths in the routing, exactly one terminal pair for v remains unrouted. Hence, exactly one edge incident on v in H remains uncolored in the partial coloring. We color all uncolored edges in H by color 3 to obtain a proper 3-coloring. ◀

Thus, we almost close the complexity gap for EDP with respect to the size of a minimum feedback vertex set, only leaving the complexity of the case $r = 1$ open.

References

- 1 Isolde Adler, Stavros G. Kolliopoulos, Philipp Klaus Krause, Daniel Lokshantov, Saket Saurabh, and Dimitrios Thilikos. Tight bounds for linkages in planar graphs. In *Proc. ICALP 2011*, volume 6755 of *Lecture Notes Comput. Sci.*, pages 110–121, 2011.
- 2 Matthew Andrews. Approximation algorithms for the edge-disjoint paths problem via Racke decompositions. In *Proc. FOCS 2010*, pages 277–286, 2010.
- 3 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.
- 4 Yonatan Aumann and Yuval Rabani. Improved bounds for all optical routing. In *Proc. SODA 1995*, pages 567–576, 1995.
- 5 Yonatan Aumann and Yuval Rabani. An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM J. Comput.*, 27(1):291–301, 1998.
- 6 Baruch Awerbuch, Rainer Gawlick, Tom Leighton, and Yuval Rabani. On-line admission control and circuit routing for high performance computing and communication. In *Proc. FOCS 1994*, pages 412–423, 1994.
- 7 Vineet Bafna, Piotr Berman, and Toshihiro Fujito. A 2-approximation algorithm for the undirected feedback vertex set problem. *SIAM J. Discrete Math.*, 12(3):289–297 (electronic), 1999.
- 8 Hans L. Bodlaender, Stephan Thomasse, and Anders Yeo. Kernel bounds for disjoint cycles and disjoint paths. *Theoret. Comput. Sci.*, 412(35):4570–4578, 2011. doi:10.1016/j.tcs.2011.04.039.
- 9 Andrei Z. Broder, Alan M. Frieze, Stephen Suen, and Eli Upfal. Optimal construction of edge-disjoint paths in random graphs. *SIAM J. Comput.*, 28(2):541–573 (electronic), 1999.
- 10 Andrei Z. Broder, Alan M. Frieze, and Eli Upfal. Existence and construction of edge-disjoint paths on expander graphs. *SIAM J. Comput.*, 23(5):976–989, 1994.
- 11 Chandra Chekuri and Alina Ene. Poly-logarithmic approximation for maximum node disjoint paths with constant congestion. In *Proc. SODA 2013*, pages 326–341, 2013.
- 12 Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd. An $O(\sqrt{n})$ approximation and integrality gap for disjoint paths and unsplittable flow. *Theory Comput.*, 2:137–146, 2006. doi:10.4086/toc.2006.v002a007.
- 13 Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd. A note on multiflows and treewidth. *Algorithmica*, 54(3):400–412, 2009.
- 14 Chandra Chekuri, Marcelo Mydlarz, and F. Bruce Shepherd. Multicommodity demand flow in a tree and packing integer programs. *ACM Trans. Algorithms*, 3(3):Art. 27, 23, 2007.
- 15 Chandra Chekuri, Guylsain Naves, and F. Bruce Shepherd. Maximum edge-disjoint paths in k -sums of graphs. In *Proc. ICALP 2013*, volume 7965 of *Lecture Notes Comput. Sci.*, pages 328–339, 2013.
- 16 Chandra Chekuri, Guylsain Naves, and F. Bruce Shepherd. Maximum edge-disjoint paths in k -sums of graphs, 2013. URL: <http://arxiv.org/abs/1303.4897>.
- 17 Chandra Chekuri, F. Bruce Shepherd, and Christophe Weibel. Flow-cut gaps for integer and fractional multiflows. *J. Comb. Theory, Ser. B*, 103(2):248–273, 2013.
- 18 Julia Chuzhoy. Routing in undirected graphs with constant congestion. In *Proc. STOC 2012*, pages 855–874, 2012.
- 19 Julia Chuzhoy, David H. K. Kim, and Shi Li. Improved approximation for node-disjoint paths in planar graphs. In *Proc. STOC 2016*, 2016. to appear.
- 20 Julia Chuzhoy and Shi Li. A polylogarithmic approximation algorithm for edge-disjoint paths with congestion 2. In *Proc. FOCS 2012*, pages 233–242, 2012.

- 21 Alina Ene, Matthias Mních, Marcin Pilipczuk, and Andrej Risteski. On routing disjoint paths in bounded treewidth graphs. In *Proc. SWAT 2016*, LIPIcs, 2016. to appear.
- 22 Krzysztof Fleszar, Matthias Mních, and Joachim Spoerhase. New algorithms for maximum disjoint paths based on tree-likeness, 2016. URL: <http://arxiv.org/abs/1603.01740>.
- 23 Alan M. Frieze. Edge-disjoint paths in expander graphs. *SIAM J. Comput.*, 30(6):1790–1801 (electronic), 2001.
- 24 Naveen Garg, Vijay V. Vazirani, and Mihalis Yannakakis. Primal-dual approximation algorithms for integral flow and multicut in trees. *Algorithmica*, 18(1):3–20, 1997. doi: 10.1007/BF02523685.
- 25 Oktay Günlük. A new min-cut max-flow ratio for multicommodity flows. *SIAM J. Discrete Math.*, 21(1):1–15, 2007.
- 26 Ian Holyer. The NP-completeness of edge-coloring. *SIAM J. Comput.*, 10(4):718–720, 1981.
- 27 Richard M. Karp. On the computational complexity of combinatorial problems. *Networks*, 5:45–68, 1975.
- 28 Ken-ichi Kawarabayashi and Yusuke Kobayashi. Breaking $\mathcal{O}(n^{1/2})$ -approximation algorithms for the edge-disjoint paths problem with congestion two. In *Proc. STOC 2011*, pages 81–88, 2011.
- 29 Ken-ichi Kawarabayashi and Paul Wollan. A shorter proof of the graph minor algorithm: the unique linkage theorem. In *Proc. STOC 2010*, pages 687–694, 2010.
- 30 Jon Kleinberg and Ronitt Rubinfeld. Short paths in expander graphs. In *Proc. FOCS 1996*, pages 86–95, 1996.
- 31 Jon Kleinberg and Éva Tardos. Disjoint paths in densely embedded graphs. In *Proc. FOCS 1995*, pages 52–61, 1995.
- 32 Jon Kleinberg and Éva Tardos. Approximations for the disjoint paths problem in high-diameter planar networks. *J. Comput. System Sci.*, 57(1):61–73, 1998.
- 33 Stavros G. Kolliopoulos and Clifford Stein. Approximating disjoint-path problems using packing integer programs. *Math. Program.*, 99(1):63–87, 2004.
- 34 Tom Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999.
- 35 Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.
- 36 Daniel Lokshtanov, M. S. Ramanujan, and Saket Saurabh. Linear time parameterized algorithms for subset feedback vertex set. In *Proc. ICALP 2015*, pages 935–946, 2015.
- 37 Takao Nishizeki, Jens Vygen, and Xiao Zhou. The edge-disjoint paths problem is NP-complete for series-parallel graphs. *Discrete Appl. Math.*, 115(1-3):177–186, 2001. doi: 10.1016/S0166-218X(01)00223-2.
- 38 Prabhakar Raghavan and Clark D. Tompson. Randomized rounding: A technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7(4):365–374, 1987.
- 39 Satish Rao and Shuheng Zhou. Edge disjoint paths in moderately connected graphs. *SIAM J. Comput.*, 39(5):1856–1887, 2010.
- 40 Neil Robertson and P. D. Seymour. Graph minors. XIII. The disjoint paths problem. *J. Combin. Theory Ser. B*, 63(1):65–110, 1995. doi:10.1006/jctb.1995.1006.
- 41 Petra Scheffler. A practical linear time algorithm for disjoint paths in graphs with bounded tree-width. Technical Report TR 396/1994, FU Berlin, Fachbereich 3 Mathematik, 1994.
- 42 Loïc Séguin-Charbonneau and F. Bruce Shepherd. Maximum edge-disjoint paths in planar graphs with congestion 2. In *Proc. FOCS 2011*, pages 200–209, 2011.