Disjoint Paths and Unsplittable Flow
(Version 2.7)

Stavros G. Kolliopoulos
Department of Informatics and Telecommunications
National and Kapodistrian University of Athens
Athens 157 84, Greece
E-mail: sgk@di.uoa.gr

October 4, 2016

1.1 Introduction

Finding disjoint paths in graphs is a problem that has attracted considerable attention from at least three perspectives: graph theory, VLSI design and network routing/flow. The corresponding literature is extensive. In this chapter we focus mostly on results on offline approximation algorithms for problems on general graphs as influenced from the network flow perspective. Surveys examining the underlying graph theory, combinatorial problems in VLSI, and disjoint paths on special graph classes can be found in [70, 71, 139, 141, 120, 132, 119, 88]. We sporadically mention some results on fixed-parameter tractability, but this is not an aspect this survey covers at any depth.

An instance of *edge-disjoint (vertex-disjoint) paths* consists of a graph $G = (V, E)$ and a multiset $\mathcal{T} = \{(s_i, t_i) : s_i \in V, t_i \in V,\ i = 1, \ldots, k\}$ of k source-sink pairs. Any source or sink is called a *terminal.* An element of $\mathcal{T}$ is also called a *commodity.* In the decision problem, one asks whether there is a a set of edge-disjoint (vertex-disjoint) paths $P_1, P_2, \ldots, P_k$, where P_i is an s_i-t_i path, $i = 1, \ldots, k$. The graph G can be either directed or undirected. Typically, a terminal may appear in more than one pair in $\mathcal{T}$. For vertex-disjoint paths one requires that the terminal pairs are mutually disjoint. We abbreviate the edge-disjoint paths problem by Edp and vertex-disjoint paths by Vdp. The notation introduced will be used throughout the chapter to refer to an input instance. We will also denote $|V|$ by n and $|E|$ by m for the corresponding graph.

Based on whether G is directed or undirected and the edge- or vertex-disjointness condition one obtains four basic problem versions. The following polynomial-time reductions exist among them. Any undirected problem can be reduced to its directed counterpart by replacing an undirected edge with an appropriate gadget; both reductions maintain planarity. See [124] and [141, Chapter 70] for details. An edge-disjoint problem can be reduced to its vertex-disjoint counterpart by replacing G with its line graph (or digraph as the case may be). Directed vertex-disjoint paths reduce to directed edge-disjoint paths by replacing every vertex with a pair of new vertices connected by an edge. There is no known reduction from a directed to an undirected problem. These transformations can serve for translating approximation guarantees or hardness results from the edge-disjoint to the vertex-disjoint setting and vice versa.

The *unsplittable flow* problem (UFP) is the generalization of EDP where every edge $e \in E$ has a positive capacity u_e, and every commodity i has a demand $d_i > 0$. The demand from s_i to t_i has to be routed in an unsplittable manner, i.e., along a single path from s_i to t_i. For every edge e the total demand routed through that edge should be at most u_e. We will sometimes refer to a capacitated graph as a *network*. In a similar manner a vertex-capacitated generalization of vertex-disjoint paths can be defined. UFP was introduced in the PhD thesis of Kleinberg [88]. Versions of the problem had been studied before though not under the UFP moniker (see, e.g., [53, 12]).

If one relaxes the requirement that every commodity should use exactly one path, one obtains the *multicommodity flow* problem which is solvable in polynomial time, for example through linear programming. When all the sources of a multicommodity flow instance coincide at a vertex s and all the sinks at a vertex t, we obtain the classical *s-t flow* problem, whose maximization version is the well-studied *maximum flow* problem. The relation between UFP and multicommodity flow is an important one to which we shall return often in this survey. We will denote a solution to either problem as a flow vector f, defined on the edges or the paths of G as appropriate.

Complexity of disjoint-path problems. For general k all four basic problems are NP-complete. The undirected VDP was shown to be NP-complete by Knuth in 1974 (see [83]), via a reduction from SAT, and by Lynch [111]. This implies the NP-completeness of directed VDP and directed EDP. Even, Itai and Shamir [65] showed that both problems remain NP-complete on directed acyclic graphs (DAGs). In the same paper the undirected EDP was shown NP-complete even when the multiset $\mathcal{T}$ contains only two distinct pairs of terminals. In the case when $s_1 = s_2 = \ldots = s_k$ all four versions are in P as special cases of maximum flow. For planar graphs Lynch's reduction [111] shows NP-completeness for undirected VDP; Kramer and van Leeuwen [104] show that undirected EDP is NP-complete. The NP-completeness of the directed planar versions follows.

For fixed k, the directed versions are NP-complete even for the case of two pairs with opposing source-sinks, i.e., (s, t) and (t, s) [69][1]. Undirected VDP, and by implication EDP as well, can be solved in polynomial time [133]. This is an outcome of the celebrated project of Robertson and Seymour on graph minors. See [23] for an informal description of the highly impractical Robertson-Seymour algorithm, which runs in $f(k) \cdot n^3$ for an immense, but computable, function $f(k)$. By an easy reduction, the Robertson-Seymour algorithm also works when terminal pairs are not mutually disjoint and one wants the output to consist of internally-disjoint paths. A decision problem is *fixed-parameter tractable* (FPT) parameterized by t if it can be solved in time $f(t)n^{O(1)}$ for a computable function f that depends only on t. Therefore, the Robertson-Seymour algorithm implies that the decision version of disjoint paths in undirected graphs is FPT when parameterized by k. It is notable that for fixed k, VDP, and by consequence EDP, can be solved on DAGs by a fairly simple polynomial-time algorithm [69]. Earlier polynomial-time algorithms for $k = 2$ include the one by Perl and Shiloach on DAGs [124] and the ones derived independently by Seymour [144], Shiloach [148] and Thomassen [161] for VDP on general undirected graphs.

[1] The NP-completeness proof holds for a sparse graph with $m = \Theta(n)$; this observation has consequences for hardness of approximation proofs in [78, 13, 147]

For planar graphs and fixed k the directed VDP is in P [140] while the complexity of the edge-disjoint case is open, even for $k = 2$. Schrijver's algorithm runs in $n^{O(k)}$ but Cygan et al. [56] showed that the problem is in fact FPT with an algorithm that runs in $2^{2^{O(k)}} \cdot n^{O(1)}$. These algorithms solve the decision problems. A few polynomial-time algorithms are known for optimization versions. When the input graph is a tree, Garg, Vazirani and Yannakakis gave a polynomial-time algorithm to maximize the number of pairs that can be connected by edge-disjoint paths [77]. The algorithm extends for vertex-disjoint paths [N. Garg, personal communication, July 2005]. By total unimodularity, maximizing the number of pairs that can be connected by edge-disjoint paths is polynomial-time solvable on *di-trees* as well, i.e., directed graphs in which there is a unique directed path from s_i to t_i, for all i; a reduction to a minimum-cost circulation problem is also possible in this case (cf. [54]). Reducing directed vertex-disjoint paths to EDP maintains the di-tree property, hence the maximization version of the former problem is polynomial-time solvable as well. Observe that directed out- and in-trees are special cases of di-trees. Additional complexity results for special graph classes can be found in [118, 164, 123, 115]. For a comprehensive complexity classification up to 2009, see [121].

Optimization versions. Two basic NP-hard optimization problems are associated with unsplittable flow and hence with EDP. Given a UFP instance an *unsplittable flow solution* or simply a *routing* is a selection of $k' \leq k$ paths, one each for a subset $\mathcal{T}' \subseteq \mathcal{T}$ of k' commodities. For every commodity $(s_i, t_i) \in \mathcal{T}'$ demand d_i is routed along the corresponding path. Any routing can be expressed as a flow vector f; the flow f_e through edge e equals the sum of the demands using e. A *feasible* routing is one that respects the capacity constraints. In the *maximum-demand* optimization problem one seeks a feasible routing of a subset $\mathcal{T}'$ of commodities such that $\sum_{i \in \mathcal{T}'} d_i$ is maximized. The *congestion* of a routing f is defined as $\max_{e \in E}\{\max\{f_e/u_e, 1\}\}$. Note that the events $f_e < u_e$ and $f_e = u_e$ are equivalent for this definition. In the *minimum-congestion* optimization problem one seeks a routing of all k commodities that minimizes the congestion, i.e., one seeks the minimum $\lambda \geq 1$ such that all k commodities can be feasibly routed if all the capacities are multiplied by λ. From now on, we use MEDP to denote the maximum-demand optimization version of EDP. Similarly, MVDP for the vertex-disjoint maximization problem. Some other objective functions of interest will be defined in Section 1.3.

Fractional solutions and flow-cut gaps. We present now some background on multicommodity flow.

LP-rounding algorithms. As mentioned, multicommodity flow is an efficiently-solvable relaxation of EDP. Hence it is no accident that multicommodity flow theory has played such an important part in developing algorithms for disjoint-path problems. This brings us to the standard linear programming formulation for the optimization version of multicommodity flow. Let $\mathcal{P}_i$ denote the set of paths from s_i to t_i. Set $\mathcal{P} := \bigcup_{i=1}^{k} \mathcal{P}_i$. Consider the following linear program (LP) formulation for *maximum multicommodity flow:*

$$\text{maximize} \sum_{P \in \mathcal{P}} f_P \qquad \text{(LP-MCF)}$$

$$\sum_{P \in \mathcal{P}_i} f_P \leq d_i \qquad \text{for } i = 1, \ldots, k$$

$$\sum_{P \in \mathcal{P} \,:\, P \ni e} f_P \leq u_e \qquad \text{for } e \in E$$

$$f_P \geq 0 \qquad \text{for } P \in \mathcal{P}$$

The number of variables in the LP is exponential in the size of the graph. By using flow variables defined on the edges one can write an equivalent LP of polynomial size. We choose to deal with the more elegant flow-path formulation. Observe that adding the constraint $f_P \in \{0, d_i\}, \forall P \in \mathcal{P}_i$, to (LP-MCF) turns it into an exact formulation for maximum-demand Ufp. We call an LP solution for the optimization problem of interest *fractional.* A similar LP, corresponding to the *concurrent flow* problem, can be written for minimizing congestion. See [163] for details. Our presentation in this section focuses on edge-capacitated graphs where the dual problems concern the computation of minimum (fractional) edge cuts. Similarly one can write the multicommodity flow LP for vertex-capacitated graphs where the dual objects of interest are vertex cuts.

Several early approximation algorithms for Ufp, and more generally integer multicommodity flow, work in two stages. First a fractional solution f is computed. Then f is rounded to an unsplittable solution $\hat{f}$ through procedures of varying intricacy, most commonly by randomized rounding as shown by Raghavan and Thompson [128]. The randomized rounding stage can usually be derandomized using the method of conditional probabilities [62, 153, 127]. The derandomization component has gradually become very important in the literature for two reasons. First, through the key work of Srinivasan [156, 154] on pessimistic estimators, good deterministic approximation algorithms were designed even in cases where the success probability of the randomized experiment was small. See [155, 22] for applications to disjoint paths. Second, in some cases the above two-stage scheme can be implemented rather surprisingly without solving first the linear program. Instead one designs directly a suitable Langrangean relaxation algorithm implementing the derandomization part. See the work of Young [165] and Chapter 4 in this volume.

We note that some of the approximation ratios obtained through the LP-rounding method can nowadays be matched (or surpassed) by simple combinatorial algorithms. By combinatorial one usually means algorithms restricted to ordered ring operations as opposed to ordered field ones. Two distinct greedy algorithms for MEdp were given by Kleinberg [88] (see also [92]), and Kolliopoulos and Stein [99] (see also [95]). A lot of the subsequent work on combinatorial algorithms for general graphs uses these two approaches as a basis. Still the influence of rounding methods on the development of algorithms for disjoint-path problems can hardly be overstated. See Chapters 6 and 7 in this volume for further background on LP-based approximation algorithms.

Approximate max-flow/min-cut theorems. One of the first results on disjoint paths and in fact one of the cornerstones of graph theory is Menger's Theorem [117]: an undirected graph is k vertex-connected if and only if there are k vertex-disjoint paths between any two vertices. The edge analogue holds as well and the min-max relation behind the theorem has resurfaced in a number of guises, most notably as the max-flow/min-cut theorem for s-t flows. Let $G = (V, E)$ be undirected. For $S \subseteq V$, define $\delta_G(S) := \{\{u, v\} \in E : u \in S \text{ and } v \in V \setminus S\}$. Similarly $\text{dem}(S)$ is the sum of all demands over commodities which are separated by the cut $\delta_G(S)$. A necessary condition for the existence of a feasible fractional solution to (LP-MCF) that satisfies all demands is the *cut condition:*

$$\sum_{e \in \delta_G(S)} u_e \geq \text{dem}(S), \text{ for each } S \subset V.$$

Define the quantity $\frac{\sum_{e \in \delta_G(S)} u_e}{\text{dem}(S)}$ as the *sparsity* of the cut $\delta_G(S)$, and let λ be the sparsity of the minimum sparsity cut. The cut condition is equivalent to $\lambda \geq 1$. For s-t flows, the max-flow/min-cut theorem [68, 74, 60] yields that the cut condition is also sufficient. For undirected multicommodity flow, Hu showed that the cut condition is sufficient for $k = 2$ [80]. It fails in general for $k \geq 3$. For directed multicommodity flows there are simple examples with $k = 2$, for which the directed analogue of the cut condition holds but the demands cannot be satisfied fractionally (see, e.g., [141]). For undirected EDP, already for $k = 2$ the cut condition is not sufficient for a solution to exist [70]. Let $\phi > 0$ be the largest quantity so that a (fractional) multicommodity flow exists that respects capacities and routes the scaled demand values ϕd_i for all i. The *concurrent-flow/cut gap* of an instance is the quantity $\frac{\lambda}{\phi}$. Starting with the seminal work of Leighton and Rao [107] a lot of effort has been spent on determining the concurrent-flow/cut gap in a variety of settings. Among several other results an optimal upper bound of $O(\log k)$ has been established for general undirected graphs both for edge as well as vertex cuts [11, 110, 67].

A *multicut* in an undirected capacitated graph $G = (V, E)$ is a subset of edges $F \subseteq E$, such that if all edges in F are deleted no two vertices from a pair (s_i, t_i) $i = 1, \ldots, k$ are in the same connected component of the remaining graph. The *maximum-multiflow/cut gap* of an instance is the ratio between the values of the minimum-capacity multicut and the maximum (fractional) multicommodity flow. The latter quantity is simply the optimum of the (LP-MCF) formulation. Garg, Vazirani and Yannakakis [76] showed that the maximum-multiflow/cut gap in an undirected graph is $O(\log k)$ and this is existentially tight. See [149, 163, 55] for surveys of the results in this area and their applications to approximation algorithms. In contrast to the undirected case, for both types of flow-cut gaps there are polynomial lower bounds in directed graphs [136, 49]. Polylogarithmic bounds exist for the problem where one wants to separate sets as opposed to pairs of vertices [87]. This includes the case where demands are symmetric, i.e., there are commodities for both ordered pairs (s_i, t_i) and (t_i, s_i).

The outline of this chapter is as follows. In Section 1.2 we present hardness of approximation results and (mostly greedy) algorithms for MEDP and MVDP. In Section 1.3 we examine algorithms for the various optimization versions of UFP, properties of the fractional relaxation, packing integer programs and UFP-specific hardness results. Section 1.4 outlines developments specific to undirected graphs that have for the most part occurred at the intersection of approximation algorithms and structural graph theory. In Section 1.5 we present results on some variants of the

Problem	Approximation Ratio	Integrality Gap	Hardness
Directed MEDP	$O(\min\{\sqrt{m}, n^{2/3}\log^{1/3} n\})$ [89, 155, 99, 162]	$\Omega(\sqrt{n})$ [77]	$\Omega(n^{1/2-\varepsilon})$ [78]
Undirected MEDP	$O(\sqrt{n})$ [38]	$\Omega(\sqrt{n})$ [77]	$\log^{1/2-\varepsilon} n$ [8]
Directed MVDP	$O(\sqrt{n})$ [22, 99]	$\Omega(\sqrt{n})$ [77]	$\Omega(n^{1/2-\varepsilon})$ [78]
Undirected MVDP	$O(\sqrt{n})$ [22, 99]	$\Omega(\sqrt{n})$ [77]	$\log^{1/2-\varepsilon} n$ [8]

Table 1.1: Known upper and lower bounds for disjoint paths on general graphs in terms of m and n. The best upper bounds on the integrality gap for MEDP are $O(\min\{\sqrt{m}, n^{4/5}\})$ on directed graphs [155, 36] and $O(\sqrt{n})$ on undirected graphs [38]. For MVDP the corresponding upper bound is $O(\sqrt{n})$ [99] on both types of graphs.

basic problems. Unless mentioned otherwise, all of the algorithms we describe for directed graphs work also on undirected graphs.

1.2 Disjoint paths

In this section we examine the problem of finding a maximum-size set of edge-disjoint paths, mostly from the perspective of combinatorial algorithms. We defer the discussion of the LP-rounding algorithms and the integrality gaps of the linear relaxations until Section 1.3, where we examine them in the more general context of UFP. Similarly for some key results on expander graphs and hardness bounds particular to UFP. Table 1.1 summarizes the known positive and negative results for MEDP and MVDP on general graphs, including bounds on the integrality gap of the multicommodity flow relaxation.

Hardness results. Ma and Wang [112] showed via the PCP theorem that MVDP and MEDP on directed graphs cannot be approximated within $2^{O(\log^{1-\varepsilon} n)}$, unless $\mathsf{NP} = \mathsf{DTIME}(2^{\text{polylog}(n)})$. Guruswami et al. [78] showed that on directed graphs it is NP-hard to obtain an $O(n^{1/2-\varepsilon})$ approximation for any fixed $\varepsilon > 0$. They gave a gap-inducing reduction from the two-pair decision problem to EDP on a sparse graph with $\Theta(n)$ edges. Since this EDP problem reduces to a vertex-disjoint path instance on a graph with $N = \Theta(n)$ vertices, we obtain that is NP-hard to approximate MVDP on graphs with N vertices within $O(N^{1/2-\varepsilon})$, for any fixed $\varepsilon > 0$. See Chapter 17 in this volume for background on the PCP theorem and the theory of inapproximability. Chalermsook et al. [27] showed that it is NP-hard to obtain an $n^{1/2-\varepsilon}$ approximation for MEDP on DAGs. MEDP on undirected graphs is much less understood. It was first shown MAX SNP-hard in [77]. Improving upon an earlier result [6], Andrews et al. [8] showed that, for any constant $\varepsilon > 0$, there is no $\log^{1/2-\varepsilon} n$ approximation algorithm unless $\mathsf{NP} \subseteq \mathsf{ZPTIME}(n^{\text{polylog}(n)})$. $\mathsf{ZPTIME}(n^{\text{polylog}(n)})$ is the set of languages that have randomized algorithms that always give the correct answer in expected running time $n^{\text{polylog}(n)}$. The same hardness result holds for MVDP on undirected graphs. Even when congestion $1 \le c \le \alpha \log\log n / \log\log\log n$ is allowed for some constant $\alpha > 0$, [8] shows that MEDP and MVDP are $\log^{\Omega(1/c)} n$-hard to approximate. For directed graphs, there is a constant $0 < \lambda < 1/4$ such that approximating MEDP while allowing congestion $1 \le c \le \log^{\lambda} n$ is hard to approximate within a factor $\Omega(n^{1/c})$, unless $\mathsf{NP} \subseteq \mathsf{ZPTIME}(n^{\text{polylog}(n)})$ [48]. We return to disjoint paths with congestion in Section 1.4.

Greedy algorithms. The first approximation algorithm analyzed in the literature for MEDP on general graphs seems to be the online *Bounded Greedy Algorithm (BGA)* in the PhD thesis of Kleinberg [88]; see also [92]. The algorithm is parameterized by a quantity L. The terminal pairs are examined in one pass. When (s_i, t_i) is considered, check if s_i can be connected to t_i by a path of length at most L. If so, route (s_i, t_i) on such a path P_i. Delete P_i from G and iterate. To simplify the analysis we assume that the last terminal pair is always routed if all the previous pairs have been rejected.

The idea behind the analysis of BGA [88] is very simple but it has influenced later work such as [99], [101] and [36]. Informally it states that *in any graph there cannot be too many long paths that are edge-disjoint.* In [88] the algorithm was shown to achieve a $(2L+1)$-approximation if $L = \max\{diam(G), \sqrt{m}\}$. Several people quickly realized that the analysis can be slightly altered to obtain an $O(\sqrt{m})$-approximation. We provide such an analysis with $L = \sqrt{m}$. The first published $O(\sqrt{m})$ approximation for MEDP was given by Srinivasan using LP-rounding methods [155].

Let $\mathcal{O}$ be a maximum-cardinality set of edge-disjoint paths connecting pairs of $\mathcal{T}$. Let $\mathcal{B}$ be the set of paths output by BGA and $\mathcal{O}_u \subset \mathcal{O}$ be the set of paths corresponding to terminal pairs unrouted by the BGA. We have that

$$|\mathcal{O}| - |\mathcal{O}_u| = |\mathcal{O} \setminus \mathcal{O}_u| \leq |\mathcal{B}|. \tag{1.1}$$

One tries to relate $|\mathcal{O}_u|$ to $|\mathcal{B}|$. This is done by observing that a commodity l routed in $\mathcal{O}_u$ was not routed in $\mathcal{B}$ because one of two things happened: (i) no path of length shorter than L exists or (ii) the existing paths from s_l to t_l were blocked by (intersect on at least one edge with) paths selected earlier in $\mathcal{B}$. The paths in $\mathcal{O}_u$ can thus be partitioned into the two corresponding subsets $\mathcal{O}_1$ and $\mathcal{O}_2$. $\mathcal{O}_1$ contains paths blocked by a path in $\mathcal{B}$ and has size at most $L|\mathcal{B}|$, since the elements of $\mathcal{B}$ are edge-disjoint paths of length at most L. The second set $\mathcal{O}_2 := \mathcal{O}_u \setminus \mathcal{O}_1$, consists of disjoint paths longer than L, hence $|\mathcal{O}_2| < m/L$. Therefore

$$|\mathcal{O}_u| < \frac{m}{L} + L|\mathcal{B}| = \sqrt{m} + \sqrt{m}|\mathcal{B}| \leq 2\sqrt{m}|\mathcal{B}|. \tag{1.2}$$

Adding inequalities (1.1) and (1.2) yields that the BGA is an $O(\sqrt{m})$-approximation algorithm. In Section 1.3 we return to the performance of the BGA on expander graphs.

The astute reader has noticed that the idea used in the analysis above is an old one. It goes back to the blocking flow method of Dinitz [57] for the maximum s-t flow problem as applied to unit-capacity networks by Even and Tarjan [66]. A *blocking flow* is a flow that cannot be augmented without rerouting. The blocking flow method iterates over the residual graph. In every iteration a blocking flow is found over the subgraph of the residual graph that contains the edges on a shortest path from s to t. At the end of an iteration the distance from s to t in the new residual graph can be shown to have increased by at least one. When the distance becomes larger than L, the number of edge-disjoint paths from s to t is $O(\min\{m/L, n^2/L^2\})$ and this bounds also the remaining number of augmentations required by the algorithm [66].

Kolliopoulos and Stein [99] made the connection with the blocking flow idea explicit and proposed the offline Greedy_Path algorithm, from now on called simply *the Greedy algorithm.* The motivation behind the Greedy algorithm was the following: what amount of residual flow has survived if one is never allowed to reroute the flow sent along shortest paths at a given iteration? In every iteration, Greedy picks the unrouted (s_i, t_i) pair such that the length of the shortest path P_i from s_i to t_i is minimized. The pair is routed using P_i; the edges of P_i are deleted from the graph. In the original paper on the Greedy algorithm, it was shown to output a solution of size $\Omega(\max\{OPT/\sqrt{m_0}, OPT^2/m_0, OPT/d_0\})$, where OPT is the optimum, m_0 is the minimum number of edges used in an optimal solution, and d_0 is the minimum average length of the paths in an optimal solution [99]. In particular, when the terminals are disjoint and there exists an acyclic optimal solution one can show using a result in [82] that $m_0 = O(n^{3/2})$. Using the BGA notation and analysis from above we obtain the following. See also [36].

Lemma 1.1 *Consider the restriction of the Greedy algorithm that as soon as the minimum shortest path length among the unrouted pairs exceeds L selects one more path and terminates. The approximation guarantee is at most* $\max\{L, |\mathcal{O}_2|\}$.

The analysis in [99] used the fact that $|\mathcal{O}_2| \leq m/L$. This was extended by Chekuri and Khanna [36]:

Theorem 1.1 *[36] Using the notation defined above $|\mathcal{O}_2| = O(n^2/L^2)$ for undirected simple graphs and $|\mathcal{O}_2| = O(n^4/L^4)$ for the directed case.*

The theorem together with Lemma 1.1 and [99] yield immediately that the Greedy algorithm achieves an $O(\min\{\sqrt{m}, n^{2/3}\})$-approximation for undirected MEdp and an $O(\min\{\sqrt{m}, n^{4/5}\})$ for directed MEdp. Varadarajan and Venkataraman [162] improved the bound for directed graphs to $O(\min\{\sqrt{m}, n^{2/3}\log^{1/3} n\})$, again for the Greedy algorithm. Their argument shows the existence of a cut of size $O((n^2/L^2)\log(n/L))$ that separates all terminal pairs (s_i, t_i) lying at distance $L \geq \log n$ or more. This brings us almost full circle back to the Even-Tarjan bound [66] for s-t flows. The latter argument demonstrates the existence of a cut of size $O(n^2/L^2)$ when the source is at distance L or more from the sink. [36] demonstrates an infinite family of undirected and directed acyclic instances on which the approximation ratio achieved by the Greedy algorithm is $\Omega(n^{2/3})$. New ideas are thus required to bring the approximation for directed graphs down to $O(\sqrt{n})$ which in [36] is conjectured to be possible. This conjecture is still open.

Because of the d_0-approximation outlined earlier, one can assume without loss of generality that all shortest $s_i - t_i$ paths have length $\Omega(\sqrt{n})$. Then a counting argument shows that there is a vertex u such that at least $\Omega(OPT/\sqrt{n})$ paths in the optimal solution go through this "congested" vertex u. We guess u and concentrate on finding an $O(1)$-approximation to the maximum-size *u-solution,* to our original Edp instance: this consists only of paths going through u. Using this approach Chekuri, Khanna and Shepherd [38] and independently Nguyen [122] obtained $O(\sqrt{n})$-approximation algorithms for Edp on undirected graphs and DAGs. The results of [38] establish matching tight bounds on the integrality gap.

We now sketch the proof of Theorem 1.1 for the undirected case as given by Chekuri and Khanna [36]. The theorem holds for the fractional solution as well, i.e., the value ν of the maximum fractional multicommodity flow connecting terminals at distance more than L. Call a vertex of G *high-degree* if its degree is more than $6n/L$ and *low-degree* otherwise. The total capacity incident to low-degree vertices is $O(n^2/L)$. We claim that every $s_i - t_i$ path, $(s_i, t_i) \in \mathcal{T}$, must contain at least $L/6$ of the low-degree vertices. Therefore ν, the sum of flow values over the paths used in the fractional solution, is $O(n^2/L^2)$. To prove the claim consider a breadth-first search tree rooted at s_i and let *layer* L_j be the set of vertices at distance j from s_i. We will show something stronger: there are at least $L/6$ layers among the first L consisting only of low-degree vertices. Partition the layers into blocks of three contiguous layers and let B_j denote the block made up of layers $L_{3j+1}, L_{3j+2}, L_{3j+3}$. Discard the blocks which contain at least one layer consisting entirely of low-degree vertices. If $L/6$ or more blocks are discarded, we are done. So assume that we are left with at least $L/6$ blocks. The blocks are disjoint so at least one of the remaining blocks, call it B_*, must contain at most $6n/L$ vertices. Consider a high-degree vertex in the middle layer of B_*. By the breadth-first search property all its neighbors must be within B_* itself, a contradiction. This completes the proof of Theorem 1.1.

Vertex-disjoint paths. The Greedy algorithm, with the obvious modification, connects a set of terminal pairs of size $\Omega(\max\{OPT/\sqrt{n_0}, OPT^2/n_0, OPT/d_0\})$ [99]. Here n_0 denotes the minimum size of a set of vertices used in the optimal solution and d_0 the minimum average path length in an optimal solution. By the hardness result of [78] this result is asymptotically tight on directed graphs, unless $\mathsf{P} = \mathsf{NP}$.

1.3 Unsplittable flow

We start with some additional definitions. We assume that a UFP instance satisfies the *no-bottleneck assumption (*NBA*)*[2]*:* $d_{\max} := \max_{i=1,\ldots,k} d_i \leq u_{\min} := \min_{e \in E} u_e$, i.e., any commodity can be routed through any of the edges. This assumption is common in the literature and we will mention it explicitly when NBA is not met. In the *weighted* UFP, commodity i has an associated weight (profit) $w_i > 0$; one wants to route feasibly a subset of commodities with maximum total weight. Note that maximizing demand reduces to maximizing the weight: simply set $w_i := d_i$, $i = 1, \ldots, k$. Another objective function of interest in addition to maximizing demand and minimizing congestion is routing in the *minimum number of rounds.* A round corresponds to a set of commodities that can be routed feasibly, hence one seeks a minimum-size partition of the set of commodities into feasible unsplittable flow solutions. A *uniform capacity unsplittable flow problem* (UCUFP) is a UFP in which all edges of the input graph have the same capacity value.

Randomized rounding and UFP. Some of the approximation ratios achieved by LP-rounding that we are about to present are currently also obtainable with simple greedy algorithms. See the paragraph on combinatorial algorithms below. Nevertheless LP-rounding algorithms have the advantage that they are analyzed with respect to the existentially weak optima of the linear relaxations. In addition their analysis yields upper bounds on the respective integrality gaps. An implementation study comparing the actual performance of the LP-based vs. the more combinatorially-

[2] In early literature it was called the *balance condition.*

flavored algorithms would be of interest. For an in-depth survey of randomization for routing problems see [158].

Minimizing congestion. The best known algorithm for congestion is also perhaps the best known example of the randomized rounding method of Raghavan and Thompson [128]. A fractional solution f to the concurrent flow problem is computed and then one path is selected independently for every commodity from the following distribution: commodity i is assigned to path $P \in \mathcal{P}_i$ with probability f_P/d_i. An application of the Chernoff bound [44] shows that with high probability the resulting congestion is $O(\log n/\log\log n)$ times the fractional optimum. The process can be derandomized using the method of conditional probabilities [127]. Young [165] shows how to construct the derandomized algorithm without having first obtained the fractional solution.

The analysis of the performance guarantee cannot be improved. Leighton, Rao and Srinivasan [105] provide an instance on a directed graph on which a fractional solution routes at most $1/\log^c n$ flow per edge, for any constant $c > 0$, while any unsplittable solution incurs congestion $\Omega(\log n/\log\log n)$. If the unsplittable solution uses only paths with nonzero fractional flow the lower bound holds for both undirected and directed instances with optimal UFP congestion 1 [105, 114]. Trivially, it is NP-hard to approximate congestion within better than 2 in the case of EDP; this would solve the decision problem. For directed graphs, the hardness results were improved in [45, 9] and finally Chuzhoy et al. [48] showed a tight $\Omega(\log n/\log\log n)$ bound, assuming that $\mathsf{NP} \not\subseteq \mathsf{BPTIME}(n^{O(\log\log n)})$. For undirected graphs, Andrews and Zhang [5] show that congestion cannot be approximated within $(\log\log n)^{1-\varepsilon}$, for any constant $\varepsilon > 0$, unless $\mathsf{NP} \subseteq \mathsf{ZPTIME}(n^{\mathrm{polylog}(n)})$. Improving upon this bound, Andrews et al. [8] showed that minimizing congestion in undirected graphs is hard to approximate within $\Omega(\log\log n/\log\log\log n)$, assuming $\mathsf{NP} \not\subseteq \mathsf{ZPTIME}(n^{\mathrm{polylog}(n)})$.

Maximizing demand. Srinivasan published the first $O(\sqrt{m})$-approximation for MEDP and more generally maximum-demand UCUFP in [155]. The first non-trivial $O(\sqrt{m}\log m)$-approximation for (weighted) UFP was published in the IPCO version of [99]. Simultaneously and independently, Baveja and Srinivasan refined the results in [155] to obtain an $O(\sqrt{m})$-approximation for weighted UFP; this work was published in [22]. The Baveja-Srinivasan methods extend the earlier key work of Srinivasan on LP-rounding methods for approximating Packing Integer programs [156, 154]. We outline now some of the ideas in [156, 155, 22]. The algorithm computes first a fractional solution f to the weighted modification of the (LP-MCF) relaxation, that has the same constraints as (LP-MCF) and objective function $\sum_{i=1}^k w_i \sum_{P\in\mathcal{P}_i} f_P$. We call this relaxation from now on (LP-WMCF). The rounding method has two phases. First, a randomized rounding experiment is analyzed to show that it produces with positive probability a near-optimal feasible unsplittable solution. Second, the experiment is derandomized yielding a deterministic polynomial-time algorithm for computing a feasible near-optimal solution. Let y_* be the fractional optimum.

One starts by scaling down every variable f_P by an appropriate parameter $\alpha > 1$. This is done to boost the probability that after randomized rounding all edge capacities are met. Let B_i denote the event that in the unsplittable solution, the capacity of the edge $e_i \in E$ is violated. Let B_{m+1} denote the event that the routed demand will be less than $y_*/(\beta\alpha)$, for some $\beta > 1$. The quantity $\beta\alpha$ is the targeted approximation ratio. The randomized rounding method of Raghavan and Thompson in the context of UFP works by bounding the probability of the "bad" event

$\bigcup_{i=1}^{m+1} B_i$ by $\sum_{i=1}^{m+1} Pr(B_i)$. Srinivasan [155] and later Srinivasan and Baveja [22] exploit the fact that the events $\overline{B_i}$ are *positively correlated:* if it is given that a routing respects the capacities of the edges in some $S \subset E$, the conditional probability that for $e_i \in E \setminus S$, $\overline{B_i}$ occurs, is at least $Pr(\overline{B_i})$. Mathematically this is expressed via the FKG inequality due to Fortuin, Ginibre and Kasteleyn (see [3, Chapter 6]). Using the positive correlation property, Baveja and Srinivasan obtain a better upper bound on $Pr(\bigcup_{e_i \in E} B_i)$ than the naive union bound and therefore can prove the existence of an unsplittable solution while using a better, i.e., smaller, $\beta\alpha$ scaling factor than traditional randomized rounding. The second ingredient of Srinivasan's method in [156, 154] is to design an appropriate pessimistic estimator to constructively derandomize the method. Such an estimator is shown for UFP as well in [22]. The by-now standard derandomization approach of Raghavan [127] fails since it relies precisely on the probability $Pr(\bigcup_{i=1}^{m+1} B_i)$ being upper-bounded by $\sum_{i=1}^{m+1} Pr(B_i)$.

Let d denote the *dilation* of the optimal fractional solution f, i.e., the maximum number of edges on any flow-carrying path. The Baveja-Srinivasan algorithm computes a solution to weighted UFP of value

$$\Omega(\max\{(y_*)^2/m, y_*/\sqrt{m}, y_*/d\}), \tag{1.3}$$

The corresponding upper bounds on the integrality gap of (LP-WMCF) follow. The analysis of [156] was simplified by Srinivasan in [157] by using randomized rounding followed by alteration. Here the outcome of the random experiment is allowed to violate some constraints. It is then altered in a greedy manner to achieve feasibility. The problem-dependent alteration step should be analyzed to quantify the potential degradation of the performance guarantee. This method was applied to UFP in [26].

Integrality gaps for weighted MEDP and MVDP. In *weighted* MEDP or MVDP one wants to maximize the total weight of the paths that can be feasibly routed, For weighted MVDP the bounds of (1.3) hold with n in place of m [22, 99]. In LP-based algorithms where the selection of the fractional paths that respect capacities and the rounding are two distinct stages, it is possible in the vertex-disjoint case to accommodate the more general setting where different commodities may share terminals, i.e., when one requires that the internal vertices of a chosen path are not used in any other. We show how one can obtain an $O(\sqrt{n+\mu})$-approximation in this case, where μ denotes the sum of the multiplicities of the terminals that belong to more than one pairs. Clearly, $\mu \leq 2k$. We restrict the discussion to undirected graphs, but it is easy to extend it to the directed case. Let G be the input graph and $\mathcal{F}$ be the set of paths that support a feasible fractional solution for weighted MVDP. We will set up the rounding stage on a modified graph G' with a modified set of paths $\mathcal{F}'$. Consider a terminal s that has multiplicity $l := \mu(s)$, i.e., belongs to the pairs of $l > 1$ commodities. We create a new graph G' where s is replaced by l vertices $s^1, \ldots, s^l$ where each is connected to the neighborhood of s in G. Fractional paths in $\mathcal{F}$ that correspond to one of the l commodities originating at s are each mapped to a path in $\mathcal{F}'$ that originates at the corresponding s^i. Let P be a path in $\mathcal{F}$ that uses s as an internal vertex. In $\mathcal{F}'$ obtain P' from P by splicing out s and replacing it by a subpath through $s^1, \ldots, s^l$. For each i, set up the vertex-capacity constraint for s^i so that either one can route a commodity originating from s^i or one can use it as an internal vertex in at most one path in $\mathcal{F}'$. Producing $G', \mathcal{F}'$ and enforcing the new vertex capacity constraints requires simply writing the corresponding packing integer program where the columns correspond to the paths in $\mathcal{F}'$, see the rounding algorithms for VDP in [99] and the upcoming discussion of packing integer programs below.

In addition to the upper bounds on the integrality gap of (LP-WMCF) given by (1.3), the integrality gap for weighted MEDP is $O(\sqrt{n})$ on undirected graphs [38] and $O(n^{4/5}\log n)$ on directed graphs [36]. The gap is known to be at least $k/2$ for the unit-weight case by an example in a grid-like planar graph with $k = \Theta(\sqrt{n})$, for both MEDP and MVDP [77]. The best known lower bounds bounds for maximum-demand and weighted UFP are summarized in Table 2.2. We provide some additional negative results for UFP in the upcoming paragraph on combinatorial algorithms.

Minimizing the number of rounds. Aumann and Rabani [10] (see also [88]) show that a ρ-approximation for maximum demand translates to an $O(\rho \log n)$ guarantee for the number of rounds objective. [22] provides improvements when all edge capacities are unit. Let $\chi(\mathcal{T})$ be the minimum number of rounds. In deterministic polynomial-time one can feasibly "route in rounds", the number of rounds being the minimum of (i) $O(\chi(\mathcal{T})d^{\delta}\log n + d(y_* + \log n))$ for any fixed $\delta \in (0,1)$, (ii) $O(\eta^{-1}d(y_* + \log n))$, if for all i, $d_i \geq \eta$ and (iii) $O\left(\chi(\mathcal{T})\sqrt{m\,(1 + (\log n)/\chi(\mathcal{T}))}\right)$ [22]. Minimizing the number of rounds for UFP is related to wavelength assignment in optical networks. Connections routed in the same round can be viewed as being assigned the same wavelength. There is extensive literature dealing with *path coloring* as this problem is often called; usually the focus is on special graph classes. See [159, Chapter 2] for an introduction to this area.

UFP with small demands. Versions of UFP have been studied that impose stronger restrictions on $d_{\max}$ than just NBA. In UFP *with small demands* one assumes that $d_{\max} \leq u_{\min}/B$, for some $B > 1$. Various improved bounds that depend on B exist, some obtainable via combinatorial algorithms.

In the rather arbitrarily named *high-capacity* UFP, $B = \Omega(\log m)$. An optimal deterministic $O(\log n)$-competitive online algorithm was obtained by Awerbuch, Azar and Plotkin [12]. It maintains length functions for the edges that are exponential in the current load. This idea was introduced for multicommodity flow in [145] and heavily used thereafter (see, e.g., [106, 125, 165, 75]). Raghavan [127] showed that standard randomized rounding achieves with high probability an $O(1)$-approximation for maximum weight with respect to the fractional optimum. Similarly, one obtains that the high-capacity UFP admits an $O(1)$-approximation for congestion. For general $B > 1$ various bounds that depend on B exist, some obtainable via combinatorial algorithms. Baveja and Srinivasan obtained an $O(t^{1/\lfloor B \rfloor})$-approximation where t is the maximum total capacity used by a commodity along a path in a fractional solution. Under the standard assumption that demands have been scaled to lie in $[0, 1]$, $t \leq d$. Azar and Regev gave a combinatorial algorithm that achieves an $O(BD^{1/B})$-approximation where D is the maximum length of a path used in the optimal solution. Kolman and Scheideler [101] investigated the approximability in terms of a different network measure, cf. the upcoming discussion on the concept of the flow number. See [99, 22, 13, 101, 26] for further bounds and details. Interesting results for the *half-disjoint* case, i.e., when $B = 2$, include the following: a polynomial-time algorithm on undirected graphs for the decision version of UFP for fixed k [90] a polylogarithmic approximation for MEDP on undirected planar graphs [28] and an $O(\sqrt{n})$-approximation for MEDP on directed graphs [122]. See also the discussion on Disjoint Paths with congestion in Section 1.4.

Packing Integer Programs and UFP. Given $A \in [0,1]^{M\times N}$, $b \in [1,\infty)^M$ and $c \in [0,1]^N$ with $\max_j c_j = 1$, a *packing integer program (PIP)* $\mathcal{P} = (A, b, c)$ seeks to maximize $c^T \cdot x$ subject to $x \in \mathbb{Z}_+^N$ and $Ax \le b$. Let B and ζ denote respectively $\min_i b_i$, and the maximum number of non-zero entries in any column of A. The restrictions on the values in A, b, c are without loss of generality; arbitrary nonnegative values can be scaled appropriately [156]. When $A \in \{0,1\}^{M\times N}$, we say that we have a $(0,1)$-*PIP*. The best guarantees known for PIPs are due to Srinivasan; those for $(0,1)$-PIPs are better than those known for general PIPs by as much as an $\Omega(\sqrt{M})$ factor [156, 154].

As witnessed by the (LP-WMCF) relaxation, weighted UFP is a packing problem, albeit one with an exponential number of variables. Motivated by UFP, Kolliopoulos and Stein [99] defined the class of *column-restricted PIPs (CPIPs):* these are the PIPs in which all nonzero entries of column j of A have the same value ρ_j, for all j. Observe that a CPIP generalizes Knapsack. If one obtains the fractional solution f to (LP-WMCF), one can formulate, at a loss of a $\log m$ factor, the *rounding problem* as a polynomial-size CPIP where the columns of A correspond to the paths used in the fractional solution and the rows correspond to edges in the graph, hence to capacity constraints. The column value ρ_j equals the demand d_j of the commodity corresponding to the path represented by the column. In combination with improved bounds for CPIPs this approach yielded the $O(\sqrt{m}\log m)$-approximation for weighted UFP [99] mentioned above.

A result of independent interest in [99] shows that approximating a family of column-restricted PIPs can be reduced in an approximation-preserving fashion to approximating the corresponding family of $(0,1)$-PIPs. This result is obtained constructively via the *grouping-and-scaling* technique which first appeared in [98] in the context of single-source UFP. Let z_* be the fractional optimum. For a general CPIP the reduction of [99] using the bounds for $(0,1)$-PIPs in [156, 154] translates to the existence of an integral solution of value $\Omega\left(\max\left\{\frac{z_*}{M^{1/(\lfloor B\rfloor+1)}},\ \frac{z_*}{\zeta^{1/\lfloor B\rfloor}},\ z_*\left(\frac{z_*}{M}\right)^{1/\lfloor B\rfloor}\right\}\right)$.

Baveja and Srinivasan [21] improved the dilation bound for CPIPs to $\Omega(\frac{z_*}{t^{1/\lfloor B\rfloor}})$ where $t \le \zeta$ is the maximum column sum of A. Notably, Baveja and Srinivasan [21] treat CPIPs as a special case of an abstract generalization of UFP that they call *low-congestion routing problem (LCRP).* LCRP is a PIP where it is convenient to think of the columns of a given matrix A as corresponding to "paths", even though there is no underlying graph. Columns are partitioned into k groups: the set of variables is $\{z_{u,v} \colon u \in [k], v \in [l_u]\}$, where l_u are given integers. In addition to the column-restricted packing constraints $Az \le b$, at most one variable from each group can be set to 1, i.e., $\forall u, \sum_{v\in[l_u]} z_{u,v} \le 1$. The objective is to maximize $\sum_{u\in[k]} w_u \sum_{v\in[l_u]} z_{u,v}$, for a vector $w \in \mathbb{R}_+^k$. Let m denote the number of rows of A, z_* the fractional optimum of an LCRP, and ζ the maximum number of nonzero entries in any column of A, Baveja and Srinivasan [22] show how to obtain an integer solution to LCRP of value

$$\Omega(\max\{(z_*)^2/m, z_*/\sqrt{m}, z_*/\zeta\}) \tag{1.4}$$

i.e., LCRP can be approximated as well as its special case of weighted UFP. In [21] they showed the $O(t^{1/\lfloor B\rfloor})$-approximation mentioned above for LCRPs, and as a consequence for weighted UFP as well. Observe that the approach of [99] to UFP was the opposite to the one of [22, 21]. In [99] the authors reduced the rounding stage of a UFP algorithm to a CPIP, at a loss of a $\log m$

factor. Baveja and Srinivasan attacked directly LCRPs, saved the $\log m$ factor from the UFP approximation ratio, and then obtained results for CPIPS as corollaries.

Chekuri et al. [42] translated in a convenient way the grouping-and-scaling approach of [99] so that it can be used as a black box for UFP problems: the integrality gap for instances with NBA is within a constant factor of the integrality gap for unit-demand instances. Chekuri et al. [34] investigated CPIPs without NBA and among other results established an $O(L)$-approximation for CPIPs with at most L non-zero entries per column.

Combinatorial algorithms. For UFP with polynomially bounded demands and without NBA Guruswami et al. [78] gave a simple randomized algorithm that achieves an $O(\sqrt{m}\log^{3/2} m)$-approximation and generalized the Greedy algorithm for MEDP [99] (cf. Section 1.2) to UFP, to obtain an $O(\sqrt{m}\log^2 m)$-approximation. Azar and Regev [13] provided the first strongly-polynomial algorithm for weighted UFP with NBA that achieves an $O(\sqrt{m})$-approximation. For weighted UFP without NBA they obtained a strongly-polynomial $O(\sqrt{m}\log(2+\frac{d_{\max}}{u_{\min}}))$-approximation algorithm. By a reduction from the two-pair decision problem, it is NP-hard to obtain an $O(n^{1-\varepsilon})$-approximation for weighted UFP on directed graphs without NBA, for any fixed $\varepsilon > 0$ [13]. The lower bound applies with all the commodities sharing the same source and the weights being such that the objective function is the cardinality of the set of commodities that can be feasibly routed. To quantify the effect of $d_{\max}/u_{\min}$ Azar and Regev showed via a different reduction that it is NP-hard to obtain an $O(n^{1/2-\varepsilon}\sqrt{\log(2+\frac{d_{\max}}{u_{\min}})})$ ratio [13]. For weighted UFP without NBA the integrality gap of the multicommodity flow relaxation is $\Omega(n)$ even when the input graph is a path [26].

Guruswami et al. [78] considered the *integral splittable flow (*ISF*)* problem in which one allows the flow for a commodity to be split along more than one path but each of these paths must carry an integral amount of flow. The objective function is to maximize the total weight of the commodities for which the entire demand has been routed. This problem is NP-hard on both directed and undirected graphs, even with just two sources and sinks [65]. Hardness results for MEDP trivially carry over to ISF. Guruswami et al. [78] observe that there is an approximation-preserving reduction from Maximum Independent Set to ISF, therefore the latter problem cannot be approximated on undirected graphs within $m^{1/2-\varepsilon}$ unless $\mathsf{NP} = \mathsf{ZPP}$. Generalizing the Greedy algorithm of [99] they showed that ISF with polynomially-bounded demands is approximable within a factor of $O(\sqrt{md_{max}}\log^2 m)$ [78]. Another of the few known polynomial lower bounds for undirected graphs was also given in [78]: it is NP-hard to approximate the maximum-cardinality objective of vertex-capacitated UFP within a factor of $n^{1/2-\varepsilon}$.

Further progress in terms of greedy algorithms was achieved by Kolman and Scheideler [101] and Kolman [100]. Recall the BGA algorithm from Section 1.2. Kolman and Scheideler [101] proposed the *careful BGA,* parameterized by L. The commodities are ordered according to their demands, starting with the largest. Commodity i is accepted if there is a feasible path P for it such that, after routing i, the total flow is larger than half their capacity on at most L edges of P. Let $\mathcal{B}_1$ be the solution thus obtained and $\mathcal{B}_2$ the solution consisting simply of the largest demand routed on any path. The output is $\mathcal{B} := \max\{\mathcal{B}_1, \mathcal{B}_2\}$. In [101] the careful BGA is shown to achieve an $O(\sqrt{m})$-approximation for maximum-demand UFP without NBA. Generalizing

UFP Problem	NBA	Approximation Ratio	Integrality Gap	Hardness
Directed max-demand	w/o	$O(\min\{\sqrt{m}, n^{4/5}\})$ [100]	$\Omega(\sqrt{n})$ [77]*	$\Omega(n^{1/2-\varepsilon})$ [78]*
Undirected max-demand	w/o	$O(\min\{\sqrt{m}, n^{2/3}\})$ [100]	$\Omega(\sqrt{n})$ [77]*	$\Omega(n^{1/2-\varepsilon})$ [147]
Directed max-weight	w	$O(\min\{\sqrt{m}, n^{4/5}\log n\})$ [22, 13, 36]	$\Omega(\sqrt{n})$ [77]	$\Omega(n^{1/2-\varepsilon})$ [78]
Undirected max-weight	w	$O(\sqrt{n})$ [38]	$\Omega(\sqrt{n})$ [77]	$\log^{1/2-\varepsilon} n$ [8]
Directed max-weight	w/o	$O(\sqrt{m}\log(2+\frac{d_{\max}}{u_{\min}}))$ [13]	$\Omega(n)$ [26]	$\Omega(n^{1-\varepsilon})$ [13]
Undirected max-weight	w/o	$O(\sqrt{m}\log(2+\frac{d_{\max}}{u_{\min}}))$ [13]	$\Omega(n)$ [26]	$\Omega(n^{1-\varepsilon})$ [147]

Table 2.2: Known upper and lower bounds involving m and n for the maximization versions of UFP on general graphs, with (w) or without (w/o) NBA. The best known upper bound on the integrality gap for Directed Max-weight with NBA is $O(\min\{\sqrt{m}, n^{4/5}\log n\})$ [22, 36]. For Undirected Max-weight with NBA a matching $O(\sqrt{n})$ upper bound on the integrality gap was given in [38]. Lower bounds with a (*) apply also with NBA, in fact they carry over from MEDP.

Theorem 1.1 above to maximum-demand UFP, Kolman showed that the careful BGA achieves an $O(\min\{\sqrt{m}, n^{2/3}\})$-approximation on undirected networks and $O(\min\{\sqrt{m}, n^{4/5}\})$-approximation on directed networks, even without NBA [101]. Currently these are the best published bounds for maximum-demand UFP; previously they had been shown for maximum-demand UCUFP in [36]. Using the grouping-and-scaling translation from unit to arbitrary demands Chekuri et al. [38] obtained an LP-based $O(\sqrt{n})$-approximation for weighted UFP on undirected graphs and DAGs. Table 2.2 summarizes the known bounds that involve m and n for the maximization versions of UFP, including bounds on the integrality gap of the multicommodity-flow based relaxation.

Exploiting the network structure. Existing approximation guarantees for UFP are rather weak and on directed graphs one cannot hope for significant improvements, unless $\mathsf{P} = \mathsf{NP}$. A different line of work has aimed for approximation ratios depending on parameters other than n and m. This type of work was originally motivated in part by popular hypercube-derived interconnection networks (cf. [138]). Theoretical advances on these networks are typically facilitated by their rich expansion properties. A graph $G = (V, E)$ is an α-expander if for every set X of at most half the vertices, the number of edges leaving X is at least $\alpha|X|$. Concluding a long line of research, Frieze [72] showed that in any r-regular graph with sufficiently strong expansion properties and r a sufficiently large constant, *any* $\Omega(n/\log n)$ vertex pairs can be connected via edge-disjoint paths. See [72] for references on the long history of the topic and the precise underlying assumptions. In such an expander the median distance between pairs of vertices is $O(\log n)$, hence the result of Frieze is within a constant factor of optimal. This basic property, that expanders are rich in short edge-disjoint paths, has been exploited in various guises in the literature. Results for fractional multicommodity flows along short paths were first given by Leighton and Rao [107].

Kleinberg and Rubinfeld analyzed the BGA on expanders in [91]. In the light of Frieze's result above, the BGA achieves an $O(\log n)$-approximation. In [91] it was also shown that for UCUFP one can efficiently compute a fractional solution that routes at least half the maximum demand with dilation $d = O(\Delta^2\alpha^{-2}\log^3 n)$. Here Δ denotes the maximum degree of the (arbitrary) input graph.

The bound on d was improved in [101]. Kolman and Scheideler introduced a new network measure, the *flow number* $F_{G,u}$, and showed that (in undirected graphs) there is always a $(1+\varepsilon)$-approximate fractional flow of dilation $O(F_{G,u}/\varepsilon)$. The flow number is defined based on the solution to a product multicommodity flow problem on G and is computable in polynomial time. If $u_{\min} \geq 1$, $F_{G,u}$ is always $\Omega(\alpha^{-1})$ and $O(\Delta\alpha^{-1}\log n)$ [101]. The BGA examining the demands in nonincreasing order and with $L := 4F_{G,u}$ achieves an $O(F_{G,u})$ approximation for UFP on undirected graphs with NBA [101]. For UFP with small demands Kolman and Scheideler gave an $O(u_{\min}(F_{G,u}^{1/u_{\min}}-1))$ guarantee for integral $u_{\min} \geq 1$, which is $O(\log F_{G,u})$ if $u_{\min} \geq \log F_{G,u}$. It is NP-hard to approximate maximum-demand UFP on directed graphs with $F_{G,u} = n^{\gamma}$, $0 < \gamma \leq 1/2$, within $F_{G,u}^{1-\varepsilon}$ [101]. Chakrabarti et al. [26] provide an $O(F_G \log n)$-approximation for weighted undirected UFP where F_G is a definition of the flow number concept of [101] made independent of capacities. F_G and $F_{G,u}$ coincide on uniform-capacity networks but are otherwise incomparable. Notably [26] presents an $O(\sqrt{\Delta \log n})$-approximation for UCUFP on Δ-regular graphs with sufficiently strong, in the sense of [72], expansion properties.

A special case that has received considerable attention is UFP *on a path* where the input graph is a path. Commodities have weights and the objective is to maximize the total weight of the commodities that can be feasibly routed. This formulation can model bandwidth allocation on a single link where every user requests an amount of bandwidth for a given time window and the capacity of the link changes over time, under the assumption that the time breakpoints are integer quantities. UFP on a path remains NP-hard since Knapsack reduces to UFP on a single edge. The currently best polynomial-time approximation is a $(2+\varepsilon)$-approximation (without NBA), in time $n^{O(1/\varepsilon^4)}$ [4]. Bansal et al. [19] obtained a $(1+\varepsilon)$-approximation in quasi-polynomial time. An improved QPTAS and some better results for special cases are given in [20]. The first constant-factor approximations were given in [26] with NBA, and without NBA in [24]. UFP *on trees* has also been studied, but so far a constant-factor approximation remains elusive, see [34, 73]. The only hardness result known is that the problem is APX-hard on capacitated trees, even when all the demands are unit [77]. Recently the $O(\log^2 n)$-approximation of [34] was extended to the case where the objective function is submodular as opposed to linear [1].

Another interesting case is UFP on a cycle which is commonly called *the ring loading problem.* In fact this is the first UFP problem studied in the literature [53], before the term "unsplittable" was coined in [89]. The input is an undirected cycle with vertices numbered clockwise along the ring and demands $d_{ij} \geq 0$ for each pair of vertices $i < j$. The task is to route all demands unsplittably, that is demand d_{ij} needs to be routed either in clockwise or in counterclockwise direction. The objective is to minimize the maximum load on an edge of the ring. Let L^* be the optimal *split* (fractional) load. In a landmark paper Schrijver et al. [142] showed one can always achieve maximum load at most $L^* + \frac{3}{2}D$, where D is the maximum demand and conjectured that $L^* + D$ is achievable. Skutella [152] proved that any split routing can be turned into an unsplittable one while increasing the load on any edge by at most $\frac{19}{14}D$ and also disproved the above conjecture of [142].

Single-Source Unsplittable Flow. Constant approximation guarantees exist for the case where all commodities share the same source, the so-called *single-source* UFP (SUFP). In contrast to single-source EDP, SUFP is strongly NP-complete [89]. The version of SUFP with costs has also been studied. In the latter problem every edge $e \in E$, has a nonnegative cost c_e. The cost of an unsplittable flow solution is $\sum_{e \in E} c_e f_e$.

The first constant-factor approximations for all the three main objectives (minimizing congestion, maximizing demand and minimizing the number of rounds) were given by Kleinberg [89]. The factors were improved by Kolliopoulos and Stein in [98] where also the first approximation for minimizing congestion without NBA was given. The grouping-and-scaling technique of [98] partitions the original problem into a collection of independent subproblems, each of them with demands in a specified range. The fractional solution is then used to assign capacities to each subproblem. The technique is in general useful for translating within constant factors integrality gaps obtained for unit demand instances to arbitrary demand instances. It found further applications, e.g., in approximating CPIPs [99, 21], and weighted UFP on trees [42]. See also the problems treated in [29]. The currently best constant factors for SUFP were obtained by Dinitz, Garg and Goemans [58], though none of them is known to be best possible under some complexity-class separation assumption. Our understanding seems to be better for congestion. The 2-approximation in [58] is best possible if the fractional congestion is used as a lower bound. No ratio better than 3/2 is possible unless P = NP. The lower bound comes from minimizing makespan on parallel machines with allocation restriction [108] which reduces in an approximation-preserving manner to minimum-congestion SUFP. Interestingly, Svensson [160] has given for the latter problem a polynomial-time algorithm that estimates the optimal makespan within a factor $33/17 + \epsilon$. This scheduling problem is also a special case of the generalized assignment problem for which a simultaneous $(2, 1)$-approximation for makespan and assignment cost exists [150]. Naturally one wonders whether a simultaneous $(2, 1)$-approximation for congestion and cost is possible for SUFP. This is an outstanding open problem. More precisely, given a fractional solution f^* the conjecture is that there is there is an unsplittable solution with the same cost such that for every edge e the flow through it is at most $f^*_e + d_{\max}$. The currently best tradeoff is a $(3, 1)$-approximation algorithm due to Skutella [151] which cleverly builds on the earlier $(3, 2)$-approximation in [98]. Erlebach and Hall [63] show that it is NP-hard to obtain a $(2 - \varepsilon, 1)$-approximation, for any fixed $\varepsilon > 0$. Experimental evaluations of algorithms for congestion can be found in [97, 59]. For SUFP without NBA, the $O(n^{1-\varepsilon})$-hardness result of Azar and Regev [13] mentioned above for the cardinality objective holds on directed graphs. It was extended to undirected planar graphs by Shepherd and Vetta [147], always under the assumption that P $\neq$ NP. For a small demand regime, the lower bound for the cardinality objective becomes $\Omega(n^{1/2-\varepsilon}\sqrt{\log(d_{\max}/u_{\min})})$, with $d_{\max}/u_{\min} > 1$, for both directed and undirected graphs [147]. A lower bound of $\Omega(n^{1/2-\varepsilon})$ holds also for the maximum demand objective in both directed and undirected graphs [147]. Finally, we refer the interested reader to the excellent survey by Shepherd [146] on single-sink problems, including but not limited to unsplittable flow, which have been motivated by telecommunications networks.

1.4 Results specific to undirected graphs

The most interesting developments over the last decade have taken place mostly for undirected graphs. Chekuri, Khanna and Shepherd [41, 28, 29] introduced the framework of well-linked decompositions which brought into approximation algorithms some of the breakthrough insights from Topological Graph Theory and in particular the Graph Minors Project of Robertson and Seymour. This fertile exchange has produced a considerable body of work for problems such as All-or-Nothing Multicommodity Flow, Disjoint Paths with Congestion and Disjoint Paths on Planar Graphs, which we survey in this section. The story has recently come full-circle with some of the new algorithmic ideas contributing to a significant improvement [31, 47] of the bounds in the Grid Minor Theorem, whose first version was given by Robertson and Seymour in [134]. All graphs in this section are undirected unless mentioned otherwise. We denote by $\mathrm{tw}(G)$ the treewidth of G, a key parameter in structural graph theory. For informative yet accessible discussions of treewidth the reader is referred to the surveys in [131, 79]. Informally, a graph G with $\mathrm{tw}(G) \leq k$ can be recursively partitioned via "balanced" vertex separators of size at most $k+1$. Moreover, if k is a small constant then G is "tree-like".

Well-linked sets. The notion of a well-linked set was introduced by Reed [130] (see also [135] for a similar definition) in an attempt to capture the concept of a highly-connected graph in a global manner. The standard definition of vertex k-connectivity in a sense focuses on local properties: a $(k+1)$-cutset may disconnect the graph but not "shatter" it globally. Given a graph G, a set $X \subseteq V(G)$ is *well-linked* if for every pair $A, B \subseteq X$ such that $|A| = |B|$, there exists a set of $|A|$ vertex-disjoint paths from A to B in G. The *well-linked number* of G, denoted $\mathrm{wl}(G)$ is the size of the largest well-linked set in G. Reed [130] showed that $\mathrm{tw}(G) + 1 \leq \mathrm{wl}(G) \leq 4(\mathrm{tw}(G) + 1)$.

Chekuri et al. [29] generalized the definition of a well-linked set and connected it to the concurrent-flow/cut gap. Given $G = (V, E)$, let $X \subseteq V$ and $\pi : X \to [0, 1]$ a weight function on X. X is *π-edge-well-linked* in G if $|\delta_G(S)| \geq \pi(X \cap S)$ for all $S \subset V$ such that $\pi(X \cap S) \leq \pi\left(X \cap (V \setminus S)\right)$. If $\pi(u) = \alpha$ for all $u \in V$, we say that X is α-edge-well-linked; X is *edge-well-linked* if it is 1-edge-well-linked. For vertex-disjoint paths or more generally for vertex-capacitated problems π-vertex-well-linked sets are the natural choice; they are defined similarly with vertex cuts in the place of edge-cuts. Reed's definition above corresponds to 1-vertex-well-linked sets. We focus on edge-well-linkedness. The decomposition framework of Chekuri et al. [29] established that given an instance $(G, \mathcal{T})$ of MEDP, one can decompose it in polynomial-time into vertex-disjoint subinstances $(G_1, \mathcal{T}_1), \ldots (G_l, \mathcal{T}_l)$ where for every pair $s_j, t_j \in \mathcal{T}_i$, $s_j, t_j \in V(G_i)$ and if X_i is the set of terminals in $\mathcal{T}_i$ then X_i is π_i-well-linked in G_i for some appropriate weight function π_i. Moreover, this happens only at a polylogarithmic loss, i.e., $\sum_{i=1}^{l} \pi_i(X_i) = \Omega(OPT/\log^2 k)$. One can boost well-linkedness, i.e., Chekuri et al. [41] showed that given a π-well-linked set X, there is a subset $X' \subset X$ such that X' is α-well-linked for some constant $\alpha \leq 1$, and $|X'| = \Omega(\pi(X))$. The importance of well-linkedness stems from the fact that as long as one is interested in fractional routings, in an α-well-linked set X, any matching can be routed with congestion $O(\beta(G)/\alpha)$ where $\beta(G)$ is the concurrent-flow/cut gap for product multicommodity flow in G. The $\beta(G)$ factor in the congestion can be saved with the stronger notion of flow-well-linked sets, see [29] for details. This decomposition framework has been used as a black-box in many later papers, including several of those that we survey below. At a loss of a polylogarithmic factor in the total flow one can assume that the input is well-linked. Obtaining integral feasible routings or bringing congestion down to a small constant require numerous further powerful ideas that exceed the scope of our survey. This a complex body of work which reaps dividends from sustained interaction with Topological Graph Theory.

All-or-Nothing Multicommodity Flow. *All-or-Nothing Multicommodity Flow* (An-Mcf) is a relaxation of MEdp. A subset M' of the pairs $\{(s_1,t_1),\ldots,(s_k,t_k)\}$ is *routable* if there is a feasible (fractional) multicommodity flow that routes one unit of flow from s_i to t_i for every pair that belongs to M'. The objective is to maximize the cardinality of M. While the decision version of An-Mcf is in P via linear programming, An-Mcf is NP-hard and APX-hard to approximate, even in capacitated trees [77]. The best known approximations that are known are $O(\log^2 k)$ in edge-capacitated graphs [29] and $O(\log^4 k \log n)$ with congestion $(1+\varepsilon)$ in vertex-capacitated graphs [29]. These are improved to $O(\log k)$ and $O(\log^2 k \log n)$ for planar graphs [29]. The grouping-and-scaling technique of [99] (see also [42]) can be used to extend the results to arbitrary demands with nba. We note that the original paper of Chekuri et al. [41] on An-Mcf used the oblivious routing techniques of Räcke [126] to implement the (implicit) decomposition into well-linked sets. The general framework and its computation via near-optimal separators were made explicit in [29]. All bounds above hold also against the optimal LP solution without congestion. This is in contrast to the $\Omega(\sqrt{n})$ lower bound on the integrality gap for MEdp. The approximation ratios in [41] hold also for the weighted version of the problem. In [42] a 4-approximation was given for the weighted version on trees, while the cardinality version is 2-approximable [77]. The known hardness results for An-Mcf are the same as for MEdp, i.e., An-Mcf is hard to approximate within $\log^{1/2-\varepsilon}$. Even with congestion c, An-Mcf is $\log^{\Omega(1/c)} n$-hard to approximate for undirected graphs, and $\Omega(n^{1/c})$, for directed graphs. See Section 1.2 for details. For An-Mcf with congestion $2 \le c \le O(\log\log n/\log\log\log n)$ the integrality gap of the multicommodity flow relaxation is $\Omega(\frac{1}{c^2}(\frac{\log n}{(\log\log n)^2})^{1/(c+1)})$ [8].

Recently, there has been some progress for directed graphs when demand pairs are *symmetric:* for each routable pair (s_i,t_i) at least one unit of flow should be routed both from s_i to t_i and from t_i to s_i. For An-Mcf with symmetric demand pairs on vertex-capacitated directed graphs, an $O(\log^2 k)$-approximation with constant vertex congestion was given in [33].

Disjoint Paths with Congestion. In the *Edge-Disjoint Paths with Congestion* problem (EDPwC) we are given as an additional input an integer c. The objective is to route the maximum number of demand pairs so that the maximum edge congestion is c. For $c = 1$ we obtain MEdp. For $c > 1$, EDPwC is a special case of Ufp with small demands, which we examined in Section 1.3. We review some of the basic results for the latter problem. Allowing $c = \Omega(\log n/\log\log n)$ the classical randomized rounding of Raghavan and Thompson [128] gives a constant-factor approximation. For congestion $c \ge 1$, there is an $O(d^{1/\lfloor c \rfloor})$-approximation and a matching upper bound on the integrality gap, where d is the maximum length of a flow path in the optimal fractional solution [21]. For combinatorial algorithms see also [13] for a matching bound and [122] for an $O(\sqrt{n})$-approximation for the case $c = 2$ and all demands being unit. All these results hold also for directed graphs and in fact they are tight under complexity-theoretic assumptions, see Sections 1.2 and 1.3.

For the undirected case, the interest in routing with low congestion $c \ge 2$, is motivated by the large gap between the existing upper and lower bounds for the approximability of MEDP. In addition, for any congestion $2 \leq c \leq O(\log\log n/\log\log\log n)$ the integrality gap of the multicommodity flow relaxation of EDPwC is $\Omega(\frac{1}{c}(\frac{\log n}{(\log\log n)^2})^{1/(c+1)})$ [8]. Andrews [7] gave a $O(\log^{61} n)$-approximation with congestion $O((\log\log n)^6)$. Chuzhoy [46] improved this to an $O(\log^{22.5} k \log\log k)$ approximation with congestion 14. Chuzhoy and Li [52] gave a randomized algorithm that achieves $O(\text{polylog}\, k)$ approximation with congestion 2. All these algorithms round the solution to the multicommodity flow relaxation, therefore when congestion 2 is allowed the integrality gap improves from $O(\sqrt{n})$ to polylogarithmic. Before that, Kawarabayashi and Kobayashi [84] had given an $O(n^{3/7}\,\text{polylog}\, n)$-approximation for congestion $c = 2$. For planar graphs, improved bounds had been obtained earlier. Chekuri, Khanna and Shepherd had shown an $O(\log k)$-approximation with congestion 2 [29] and then an $O(1)$-approximation with congestion 4 [39]. Seguin-Charbonneau and Shepherd [143] showed a constant-factor approximation with congestion 2. For *Vertex-Disjoint Paths with Congestion,* Chekuri and Ene [32] gave an $O(\text{polylog}\, k)$-approximation with congestion 51. Chekuri and Chuzhoy have announced an $O(\text{polylog}\, k)$ approximation with congestion 2 [30]. Simple modifications to the reductions in [8] show that the $\log^{\Omega(1/c)} n$-hardness result for EDPwC holds also for Vertex-Disjoint Paths, assuming $\mathsf{NP} \not\subseteq \mathsf{ZPTIME}(n^{\text{polylog}(n)})$. In contrast to these intractability results, when $k \leq \delta(\log\log n)^{2/15}$, where δ is an appropriate constant, Kleinberg showed that for $c = 2$ the decision version is in P both for edge- and vertex-capacitated undirected graphs [90]. Finally, Chekuri et al. [35] gave a poly-logarithmic approximation with congestion 5 for MVDP with symmetric demands on planar digraphs.

Disjoint Paths without Congestion. An interesting result by Rao and Zhu [129] gives an $O(\text{polylog}\, n)$ approximation for MEDP (with congestion 1), if the global min cut has size $\Omega(\log^5 n)$. This result is not comparable to the above but the ideas have proved valuable for [7, 84, 46, 52].

For MVDP (with congestion 1) Chuzhoy, Kim and Li [51] gave recently an $O(n^{9/19}\,\text{polylog}\, n)$-approximation on planar graphs. Interestingly, this is the first result on planar graphs that improves upon the $O(\sqrt{n})$-approximation which the Greedy algorithm achieves on general graphs [99]. Recall that on the negative side, there is no $\log^{1/2-\varepsilon} n$ approximation algorithm on general undirected graphs unless $\mathsf{NP} \subseteq \mathsf{ZPTIME}(n^{\text{polylog}(n)})$ [8]. The only hardness result specific to planar graphs is APX-hardness on grid graphs [50].

Several other improved approximation ratios exist for special graph classes. See the references contained in [64, 51]. Notably, for graphs of treewidth r, improving upon [43] Ene et al. [61] obtained an $O(r^3)$-approximation for MEDP and a similar approximation for MVDP as a function of pathwidth. These guarantees hold with respect to the fractional solution. The standard gap example of [77] yields an $\Omega(r)$ lower bound on the integrality gap. Chekuri, Khanna and Shepherd [40] asked whether there is a matching upper bound. It should be noted that EDP remains NP-complete even in graphs of treewidth $r = 2$ [123]. Assuming $\mathsf{P} \neq \mathsf{NP}$, this rules out the existence of a fixed-parameter algorithm for EDP parameterized by treewidth. For a characterization of the fixed-parameter tractability of VDP see [116].

1.5 Further variants of the basic problems

Length-bounded flows are single- or multi-commodity flow problems in which every path used must obey a length constraint. See [16] and [17] for background on this topic, including results on the associated flow/cut gaps. In the *bounded-length* EDP (BLEDP), an additional input parameter M is specified. One seeks a set of disjoint s_i-t_i paths under the constraint that the length of each path is at most M. In (s,t)-BLEDP all the pairs share the same source s and sink t. Cases that used to be tractable become NP-complete with the length constraint. Both in the vertex and the edge-disjoint case, (s,t)-BLEDP is NP-complete on undirected graphs even for fixed $M \geq 4$ [81]. For variable M and fixed k, the problems remain NP-complete [109]. In the optimization version BLMEDP, one seeks a maximum-cardinality set of edge-disjoint paths that satisfy the length constraint. It is NP-hard to approximate (s,t)-BLMEDP within $O(n^{1/2-\varepsilon})$ on directed graphs and, unless NP = ZPP, BLMEDP cannot be approximated in polynomial time within $O(n^{1/2-\varepsilon})$ on undirected graphs, for any fixed $\varepsilon > 0$ [78]. On the positive side it is easy to obtain an $O(\sqrt{m})$-approximation for BLMEDP. For the paths in the optimal solution with length at most $M' := \min\{\sqrt{m}, M\}$ the BGA with parameter $L = M'$ achieves an $O(M')$-approximation. This is because, in the notation of the BGA analysis in Section 1.2, the set $\mathcal{O}_2$ is empty. On the other hand there are at most $\sqrt{m}$ edge-disjoint paths of length more than $\sqrt{m}$. See [78] for other algorithmic results. For the *bounded-length weighted* UFP with NBA, the results in [155, 99, 26] yield $O(M)$-approximations. Using the ellipsoid algorithm one can find an optimal fractional solution whose support contains only paths that satisfy the length constraint. Using their result on CPIPs without NBA Chekuri et al. [34] showed that the same $O(M)$ ratio holds for bounded-length weighted UFP without NBA. Interestingly, the integrality gap of the natural relaxation for the maximum integral length-bounded s-t flow is $\Omega(\sqrt{n})$ even for directed or undirected planar graphs [17].

In transportation logistics a commodity may be splittable in different containers, each of them to be routed along a single path. One wishes to bound the number of containers used. This motivates the *K-splittable flow problem*, a relaxed version of UFP where a commodity can be split along *at most* $K \geq 1$ paths, K an input parameter. This problem was introduced and first studied by Baier, Köhler and Skutella [18]. Clearly for $K = m$, it reduces to solving the fractional relaxation; it is NP-complete for $K = 2$. See [114, 94, 93] for a continuation of the work in [18]. The author observed in [96] that the single-source 2-splittable flow problem admits a simultaneous $(2,1)$-approximation for congestion and cost. This was improved and generalized by Salazar and Skutella to a $(1 + 1/K + 1/(2K-1), 1)$-approximation for K-splittable flows [137].

Finally, a problem in a sense complementary to K-splittable flow is the *multiroute flow* where for reliability purposes the flow of a commodity *has to* be split along a given number of edge-disjoint paths. Observe that in the K-splittable flow, there is no edge-disjointness requirement. Given a source s and a sink t, and an integer $K \geq 1$, an *elementary K-route flow* is a set of K edge-disjoint paths between s and t. A *K-route flow* is a nonnegative linear combination of elementary K-route flows. The multiroute flow problem consists of finding a maximum K-route flow that respects the edge capacities. It was introduced by Kishimoto and Tagauchi [85] (see [86]) and is solvable in polynomial time. See [2] for simplifications and extensions of the basic s-t multiroute flow theory. A dual object of interest to a K-route flow is a *K-route cut:* a subset of the edges whose removal leaves at most $K-1$ edge-disjoint paths between s and t. We note that in the early literature [86, 2] a different definition of cut was considered for the purpose of a flow/cut duality result. It is a natural question to study multiroute flows in the multicommodity setting in conjunction with the K-route flow/K-route cut gap. A sample of this work can be found in [14, 25, 37, 102, 103]. Martens [113] gives a greedy algorithm for a related multicommodity problem which he calls the *k-disjoint flow* problem. For every commodity i one seeks K_i edge-disjoint paths and the flow has to be perfectly balanced: on each path it must be equal to d_i/K_i. The objective function is to maximize the sum of routed demands subject to the capacity constraints and the approximation ratio achieved is $O(K_{\max}\sqrt{m}/K_{\min})$ [113]. A single-commodity (s,t)-flow of value F is *K-balanced* if every edge carries at most F/K units of flow. Interestingly, an acyclic (s,t)-flow is K-balanced if and only if it is a K-route flow [86, 2, 15].

Acknowledgements for the first edition. Thanks to Chandra Chekuri, Sanjeev Khanna, Maren Martens, Martin Skutella, and Cliff Stein for valuable comments and suggestions. Thanks to Naveen Garg for a clarification on [77], to Aris Pagourtzis for pointing out [159], and to Lex Schrijver for information on EDP on planar graphs.
Acknowledgements for the second edition. Thanks to Chandra Chekuri, Petr Kolman, Bruce Shepherd, Dimitrios Thilikos for valuable comments and suggestions.

References

[1] Anna Adamaszek, Parinya Chalermsook, Alina Ene, and Andreas Wiese. Submodular unsplittable flow on trees. In *Integer Programming and Combinatorial Optimization - 18th International Conference, IPCO 2016, Liège, Belgium, June 1-3, 2016, Proceedings*, pages 337–349, 2016.

[2] Charu C. Aggarwal and James B. Orlin. On multiroute maximum flows in networks. *Networks*, 39(1):43–52, 2002.

[3] N. Alon and J. Spencer. *The Probabilistic method, 2nd edition.* John Wiley and Sons, 2000.

[4] Aris Anagnostopoulos, Fabrizio Grandoni, Stefano Leonardi, and Andreas Wiese. A mazing $2+\epsilon$ approximation for unsplittable flow on a path. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 26–41, 2014.

[5] M. Andrews and L. Zhang. Hardness of the undirected congestion minimization problem. In *Proceedings of the 37th annual ACM Symposium on Theory of Computing*, pages 284–293, 2005.

[6] M. Andrews and L. Zhang. Hardness of the undirected edge-disjoint paths problem. In *Proceedings of the 37th annual ACM Symposium on Theory of Computing*, pages 276–283, 2005.

[7] Matthew Andrews. Approximation algorithms for the edge-disjoint paths problem via Räcke decompositions. In *51st Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 277–286, 2010.

[8] 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.

[9] Matthew Andrews and Lisa Zhang. Almost-tight hardness of directed congestion minimization. *J. ACM*, 55(6), 2008.

[10] Y. Aumann and Y. Rabani. Improved bounds for all-optical routing. In *Proceedings of the 6th ACM-SIAM Symposium on Discrete Algorithms*, pages 567–576, 1995.

[11] Y. Aumann and Y. Rabani. An $O(\log k)$ approximate Min-Cut Max-Flow theorem and approximation algorithm. *SIAM Journal on Computing*, 27:291–301, 1998.

[12] B. Awerbuch, Y. Azar, and S. Plotkin. Throughput-competitive online routing. In *Proceedings of the 34th Annual IEEE Symposium on Foundations of Computer Science*, pages 32–40, 1993.

[13] Yossi Azar and Oded Regev. Combinatorial algorithms for the unsplittable flow problem. *Algorithmica*, 44(1):49–66, 2006.

[14] A. Bagchi, A. Chaudhary, and P. Kolman. Short length Menger's Theorem and reliable optical routing. *Theoretical Computer Science*, 339:315–332, 2005. Prelim. version in SPAA 03 (revue paper).

[15] A. Bagchi, A. Chaudhary, P. Kolman, and J. Sgall. A simple combinatorial proof of duality of multiroute flows and cuts. Technical Report 2004-662, Charles University, Prague, 2004.

[16] Georg Baier. *Flows with path restrictions*. PhD thesis, TU Berlin, 2003.

[17] Georg Baier, Thomas Erlebach, Alexander Hall, Ekkehard Köhler, Petr Kolman, Ondrej Pangrác, Heiko Schilling, and Martin Skutella. Length-bounded cuts and flows. *ACM Trans. Algorithms*, 7(1):4, 2010.

[18] Georg Baier, Ekkehard Köhler, and Martin Skutella. The k-splittable flow problem. *Algorithmica*, 42(3-4):231–248, 2005.

[19] Nikhil Bansal, Amit Chakrabarti, Amir Epstein, and Baruch Schieber. A quasi-PTAS for unsplittable flow on line graphs. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006*, pages 721–729, 2006.

[20] Jatin Batra, Naveen Garg, Amit Kumar, Tobias Mömke, and Andreas Wiese. New approximation schemes for unsplittable flow on a path. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 47–58, 2015.

[21] A. Baveja and A. Srinivasan. Approximating low-congestion routing and column-restricted packing problems. *Information Processing Letters*, 74:19–25, 2000.

[22] A. Baveja and A. Srinivasan. Approximation algorithms for disjoint paths and related routing and packing problems. *Mathematics of Operations Research*, 25:255–280, 2000.

[23] D. Bienstock and M. A. Langston. Algorithmic implications of the Graph Minor Theorem. In M. O. Ball, T. L. Magnanti, C. L. Monma, and G. L. Nemhauser, editors, *Handbook in Operations Research and Management Science 7: Network models*. North-Holland, 1995.

[24] Paul S. Bonsma, Jens Schulz, and Andreas Wiese. A constant-factor approximation algorithm for unsplittable flow on paths. *SIAM J. Comput.*, 43(2):767–799, 2014.

[25] Henning Bruhn, Jakub Cerný, Alexander Hall, Petr Kolman, and Jirí Sgall. Single source multiroute flows and cuts on uniform capacity networks. *Theory of Computing*, 4(1):1–20, 2008.

[26] Amit Chakrabarti, Chandra Chekuri, Anupam Gupta, and Amit Kumar. Approximation algorithms for the unsplittable flow problem. *Algorithmica*, 47(1):53–78, 2007. Prelim. version in APPROX 02.

[27] Parinya Chalermsook, Bundit Laekhanukit, and Danupon Nanongkai. Pre-reduction graph products: Hardnesses of properly learning DFAs and approximating EDP on DAGs. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 444–453, 2014.

[28] 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*, pages 71–80, 2004.

[29] C. Chekuri, S. Khanna, and F. B. Shepherd. Multicommodity flow, well-linked terminals, and routing problems. In *Proceedings of the 37th annual ACM Symposium on Theory of Computing*, pages 183–192, 2005.

[30] Chandra Chekuri and Julia Chuzhoy. Half-integral all-or-nothing flow. Unpublished manuscript.

[31] Chandra Chekuri and Julia Chuzhoy. Polynomial bounds for the grid-minor theorem (extended abstract). In *Proc. Symposium on Theory of Computing (STOC), New York, NY*, pages 60–69, 2014.

[32] Chandra Chekuri and Alina Ene. Poly-logarithmic approximation for maximum node disjoint paths with constant congestion. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 326–341, 2013.

[33] Chandra Chekuri and Alina Ene. The all-or-nothing flow problem in directed graphs with symmetric demand pairs. *Math. Program.*, 154(1-2):249–272, 2015.

[34] Chandra Chekuri, Alina Ene, and Nitish Korula. Unsplittable flow in paths and trees and column-restricted packing integer programs. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 12th International Workshop, APPROX 2009, and 13th International Workshop, RANDOM 2009, Berkeley, CA, USA, August 21-23, 2009. Proceedings*, pages 42–55, 2009.

[35] Chandra Chekuri, Alina Ene, and Marcin Pilipczuk. Constant congestion routing of symmetric demands in planar directed graphs. To appear in the Proc. 43rd ICALP, 2016.

[36] Chandra Chekuri and Sanjeev Khanna. Edge-disjoint paths revisited. *ACM Trans. Algorithms*, 3(4), 2007. Prelim. version in SODA 2003.

[37] Chandra Chekuri and Sanjeev Khanna. Algorithms for 2-route cut problems. In *Automata, Languages and Programming, 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7-11, 2008, Proceedings, Part I: Tack A: Algorithms, Automata, Complexity, and Games*, pages 472–484, 2008.

[38] 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(7):137–146, 2006.

[39] Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd. Edge-disjoint paths in planar graphs with constant congestion. *SIAM J. Comput.*, 39(1):281–301, 2009.

[40] Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd. A note on multiflows and treewidth. *Algorithmica*, 54(3):400–412, 2009.

[41] Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd. The all-or-nothing multicommodity flow problem. *SIAM J. Comput.*, 42(4):1467–1493, 2013. Prelim. version in STOC 04.

[42] Chandra Chekuri, Marcelo Mydlarz, and F. Bruce Shepherd. Multicommodity demand flow in a tree and packing integer programs. *ACM Trans. Algorithms*, 3(3), 2007.

[43] Chandra Chekuri, Guyslain Naves, and F. Bruce Shepherd. Maximum edge-disjoint paths in k-sums of graphs. In *Automata, Languages, and Programming - 40th International Colloquium, ICALP 2013, Riga, Latvia, July 8-12, 2013, Proceedings, Part I*, pages 328–339, 2013.

[44] H. Chernoff. A measure of the asymptotic efficiency for tests of a hypothesis based on sum of observations. *Ann. Math. Stat.*, 23:493–509, 1952.

[45] J. Chuzhoy and S. Naor. New hardness results for congestion minimization and machine scheduling. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pages 28–34, 1994.

[46] Julia Chuzhoy. Routing in undirected graphs with constant congestion. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 855–874, 2012.

[47] Julia Chuzhoy. Excluded grid theorem: Improved and simplified. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 645–654, 2015.

[48] Julia Chuzhoy, Venkatesan Guruswami, Sanjeev Khanna, and Kunal Talwar. Hardness of routing with congestion in directed graphs. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007*, pages 165–178, 2007.

[49] Julia Chuzhoy and Sanjeev Khanna. Polynomial flow-cut gaps and hardness of directed cut problems. *J. ACM*, 56(2):6:1–6:28, 2009.

[50] Julia Chuzhoy and David H. K. Kim. On approximating node-disjoint paths in grids. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2015, August 24-26, 2015, Princeton, NJ, USA*, pages 187–211, 2015.

[51] Julia Chuzhoy, David H. K. Kim, and Shi Li. Improved approximation for node-disjoint paths in planar graphs. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 556–569, 2016.

[52] Julia Chuzhoy and Shi Li. A polylogarithmic approximation algorithm for edge-disjoint paths with congestion 2. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 233–242, 2012.

[53] S. Cosares and I. Saniee. An optimization problem related to balancing loads on SONET rings. *Telecommunications Systems*, 3:165–181, 1994. Prelim. version as Technical Memorandum. Bellcore, Morristown, NJ, 1992.

[54] M. C. Costa, L. Létocart, and F. Roupin. A greedy algorithm for multicut and integral multiflow in rooted trees. *Operations Research Letters*, 31:21–27, 2003.

[55] M. C. Costa, L. Létocart, and F. Roupin. Minimal multicut and maximum integer multiflow: a survey. *European Journal of Operational Research*, 162:55–69, 2005.

[56] Marek Cygan, Dániel Marx, Marcin Pilipczuk, and Michal Pilipczuk. The planar directed k-vertex-disjoint paths problem is fixed-parameter tractable. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 197–206, 2013.

[57] E. A. Dinitz. Algorithm for solution of a problem of maximum flow in networks with power estimation. *Soviet Math. Dokl.*, 11:1277–1280, 1970.

[58] Y. Dinitz, N. Garg, and M. X. Goemans. On the single-source unsplittable flow problem. *Combinatorica*, 19:1–25, 1999. Prelim. version in FOCS 98.

[59] Jingde Du and Stavros G. Kolliopoulos. Implementing approximation algorithms for the single-source unsplittable flow problem. *ACM Journal of Experimental Algorithmics*, 10:2.3:1–2.3:21, 2005.

[60] P. Elias, A. Feinstein, and C. E. Shannon. Note on maximum flow through a network. *IRE Transactions on Information Theory IT-2*, pages 117–199, 1956.

[61] Alina Ene, Matthias Mnich, Marcin Pilipczuk, and Andrej Risteski. On routing disjoint paths in bounded treewidth graphs. In *15th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2016, June 22-24, 2016, Reykjavik, Iceland*, pages 15:1–15:15, 2016.

[62] P. Erdős and J. L. Selfridge. On a combinatorial game. *Journal of Combinatorial Theory A*, 14:298–301, 1973.

[63] T. Erlebach and A. Hall. NP-hardness of broadcast sheduling and inapproximability of single-source unsplittable min-cost flow. *Journal of Scheduling*, 7:223–241, 2004. Prelim. version in SODA 02.

[64] Thomas Erlebach. Approximation algorithms for edge-disjoint paths and unsplittable flow. In Evripidis Bampis, Klaus Jansen, and Claire Kenyon, editors, *Efficient Approximation and Online Algorithms: Recent Progress on Classical Combinatorial Optimization Problems and New Applications*, pages 97–134. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.

[65] S. Even, A. Itai, and A. Shamir. On the complexity of timetable and multicommodity flow problems. *SIAM Journal on Computing*, 5:691–703, 1976.

[66] S. Even and R. E. Tarjan. Network flow and testing graph connectivity. *SIAM Journal on Computing*, 4:507–518, 1975.

[67] Uriel Feige, MohammadTaghi Hajiaghayi, and James R. Lee. Improved approximation algorithms for minimum weight vertex separators. *SIAM J. Comput.*, 38(2):629–657, 2008.

[68] L. R. Ford and D. R. Fulkerson. Maximal flow through a network. *Canad. J. Math.*, 8:399–404, 1956.

[69] S. Fortune, J. Hopcroft, and J. Wyllie. The directed subgraph homeomorphism problem. *Theoretical Computer Science*, 10:111–121, 1980.

[70] A. Frank. Packing paths, cuts and circuits - a survey. In B. Korte, L. Lovász, H. J. Prömel, and A. Schrijver, editors, *Paths, Flows and VLSI-Layout*, pages 49–100. Springer-Verlag, Berlin, 1990.

[71] A. Frank. Connectivity and network flows. In R. Graham, M. Grötschel, and L. Lovász, editors, *Handbook of Combinatorics*, pages 111–177. North-Holland, 1995.

[72] A. M. Frieze. Edge-disjoint paths in expander graphs. *SIAM Journal on Computing*, 30:1790–1801, 2001. Prelim. version in SODA 00.

[73] Zachary Friggstad and Zhihan Gao. On linear programming relaxations for unsplittable flow in trees. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2015, August 24-26, 2015, Princeton, NJ, USA*, pages 265–283, 2015.

[74] D. R. Fulkerson and G. B. Dantzig. Computation of maximum flow in networks. *Naval Research Logistics Quarterly*, 2:277–283, 1955.

[75] N. Garg and J. Könemann. Faster and simpler algorithms for multicommodity flow and other fractional packing problems. In *Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science*, pages 300–309, 1998.

[76] N. Garg, V. Vazirani, and M. Yannakakis. Approximate max-flow min-(multi)cut theorems and their applications. *SIAM Journal on Computing*, 25:235–251, 1996. Prelim. version in STOC 93.

[77] N. Garg, V. Vazirani, and M. Yannakakis. Primal-dual approximation algorithms for integral flow and multicut in trees. *Algorithmica*, 18:3–20, 1997. Prelim. version in ICALP 93.

[78] V. Guruswami, S. Khanna, R. Rajaraman, B. Shepherd, and M. Yannakakis. Near-optimal hardness results and approximation algorithms for edge-disjoint paths and related problems. *Journal of Computer and System Sciences*, 67:473–496, 2003. Prelim. version in STOC 99.

[79] Daniel J. Harvey and David R. Wood. Parameters tied to treewidth. *Journal of Graph Theory*, Published online, 2016.

[80] T. C. Hu. Multi-commodity network flows. *Operations Research*, 11:344–360, 1963.

[81] A. Itai, Y. Perl, and Y. Shiloach. The complexity of finding maximum disjoint paths with length constraints. *Networks*, 12:277–286, 1982.

[82] D. R. Karger and M. S. Levine. Finding maximum flows in simple undirected graphs seems easier than bipartite matching. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, 1998.

[83] R. M. Karp. On the computational complexity of combinatorial problems. *Networks*, 5:45–68, 1975.

[84] Ken-ichi Kawarabayashi and Yusuke Kobayashi. Breaking $O(n^{1/2})$-approximation algorithms for the edge-disjoint paths problem with congestion two. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 81–88, 2011.

[85] W. Kishimoto and M. Takeuchi. On m-route flows in a network. *IEICE Trans.*, J-76-A(8):1185–1200, 1993. In Japanese.

[86] Wataru Kishimoto. A method for obtaining the maximum multiroute flows in a network. *Networks*, 27(4):279–291, 1996.

[87] Philip N. Klein, Serge A. Plotkin, Satish Rao, and Éva Tardos. Approximation algorithms for Steiner and directed multicuts. *J. Algorithms*, 22(2):241–269, 1997.

[88] J. M. Kleinberg. *Approximation algorithms for disjoint paths problems.* PhD thesis, MIT, Cambridge, MA, May 1996.

[89] J. M. Kleinberg. Single-source unsplittable flow. In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science*, pages 68–77, October 1996.

[90] J. M. Kleinberg. Decision algorithms for unsplittable flow and the half-disjoint paths problem. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, pages 530–539, 1998.

[91] J. M. Kleinberg and R. Rubinfeld. Short paths in expander graphs. In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science*, pages 86–95, 1996.

[92] J. M. Kleinberg and É. Tardos. Disjoint paths in densely-embedded graphs. In *Proceedings of the 36th Annual IEEE Symposium on Foundations of Computer Science*, pages 52–61, 1995.

[93] Ronald Koch, Martin Skutella, and Ines Spenke. Maximum k -splittable s , t -flows. *Theory Comput. Syst.*, 43(1):56–66, 2008.

[94] Ronald Koch and Ines Spenke. Complexity and approximability of k-splittable flows. *Theor. Comput. Sci.*, 369(1-3):338–347, 2006.

[95] S. G. Kolliopoulos. *Exact and Approximation Algorithms for Network Flow and Disjoint-Path Problems.* PhD thesis, Dartmouth College, Hanover, NH, August 1998.

[96] S. G. Kolliopoulos. Minimum-cost single-source 2-splittable flow. *Information Processing Letters*, 94:15–18, 2005.

[97] S. G. Kolliopoulos and C. Stein. Experimental evaluation of approximation algorithms for single-source unsplittable flow. In G. Cornuéjols, R. E. Burkard, and G. J. Woeginger, editors, *Proceedings of the 7th Conference on Integer Programming and Combinatorial Optimization*, volume 1610 of *Lecture Notes in Computer Science*, pages 328–344. Springer-Verlag, June 1999.

[98] S. G. Kolliopoulos and C. Stein. Approximation algorithms for single-source unsplittable flow. *SIAM Journal on Computing*, 31:919–946, 2002. Prelim. version in FOCS 97.

[99] S. G. Kolliopoulos and C. Stein. Approximating disjoint-path problems using packing integer programs. *Mathematical Programming A*, 99:63–87, 2004. Prelim. version in IPCO 98.

[100] P. Kolman. A note on the greedy algorithm for the unsplittable flow problem. *Information Processing Letters*, 88:101–105, 2003.

[101] Petr Kolman and Christian Scheideler. Improved bounds for the unsplittable flow problem. *J. Algorithms*, 61(1):20–44, 2006. Prelim. version in SODA 02.

[102] Petr Kolman and Christian Scheideler. Approximate duality of multicommodity multiroute flows and cuts: single source case. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 800–810, 2012.

[103] Petr Kolman and Christian Scheideler. Towards duality of multicommodity multiroute cuts and flows: Multilevel ball-growing. *Theory Comput. Syst.*, 53(2):341–363, 2013.

[104] M. R. Kramer and J. van Leeuwen. The complexity of wire-routing and finding minimum-area layouts for arbitrary VLSI circuits. In F. P. Preparata, editor, *VLSI Theory*, volume 2 of *Advances in Computing Research*, pages 129–146. JAI Press, 1984.

[105] F. T. Leighton, S. Rao, and A. Srinivasan. Multicommodity flow and circuit switching. In *Hawaii International Conference on System Sciences*, pages 459–465, 1998.

[106] T. Leighton, F. Makedon, S. Plotkin, C. Stein, É. Tardos, and S. Tragoudas. Fast approximation algorithms for multicommodity flow problems. *Journal of Computer and System Sciences*, 50:228–243, 1995. Prelim. version in STOC 91.

[107] 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. Prelim. version in FOCS 88.

[108] J. K. Lenstra, D. B. Shmoys, and É. Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Mathematical Programming A*, 46:259–271, 1990.

[109] C. Li, T. McCormick, and D. Simchi-Levi. The complexity of finding two disjoint paths with min-max objective function. *Discrete Applied Mathematics*, 26:105–115, 1990.

[110] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15:215–246, 1995.

[111] J. F. Lynch. The equivalence of theorem proving and the interconnection problem. *ACM SIGDA Newsletter*, 5:31–36, 1975.

[112] B. Ma and L. Wang. On the inapproximability of disjoint paths and minimum Steiner forest with bandwidth constraints. *Journal of Computer and System Sciences*, 60:1–12, 2000.

[113] Maren Martens. A simple greedy algorithm for the k-disjoint flow problem. In *Theory and Applications of Models of Computation, 6th Annual Conference, TAMC 2009, Changsha, China, May 18-22, 2009. Proceedings*, pages 291–300, 2009.

[114] Maren Martens and Martin Skutella. Flows on few paths: Algorithms and lower bounds. *Networks*, 48(2):68–76, 2006.

[115] Dániel Marx. Eulerian disjoint paths problem in grid graphs is NP-complete. *Discrete Applied Mathematics*, 143(1-3):336–341, 2004.

[116] Dániel Marx and Paul Wollan. An exact characterization of tractable demand patterns for maximum disjoint path problems. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 642–661, 2015.

[117] K. Menger. Zur allgemeinen kurventheorie. *Fundamenta Mathematicae*, 10:96–115, 1927.

[118] Matthias Middendorf and Frank Pfeiffer. On the complexity of the disjoint paths problems. *Combinatorica*, 13(1):97–107, 1993.

[119] R. H. Möhring and D. Wagner. Combinatorial topics in VLSI design, annotated bibliography. In Mauro Dell'Amico, Francesco Maffioli, and Silvano Martello, editors, *Annotated Bibliographies in Combinatorial Optimization*, pages 429–444. Wiley, 1997.

[120] R. H. Möhring, D. Wagner, and F. Wagner. VLSI network design: a survey. In M.O. Ball, T.L. Magnanti, C.L. Monma, and G.L. Nemhauser, editors, *Handbooks in Operations Research/Management Science, Volume on Networks*, pages 625–712. North-Holland, 1995.

[121] Guyslain Naves and András Sebő. Multiflow feasibility: An annotated tableau. In William Cook, László Lovász, and Jens Vygen, editors, *Research Trends in Combinatorial Optimization: Bonn 2008*, pages 261–283. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.

[122] Thành Nguyen. On the disjoint paths problem. *Oper. Res. Lett.*, 35(1):10–16, 2007.

[123] Takao Nishizeki, Jens Vygen, and Xiao Zhou. The edge-disjoint paths problem is NP-complete for series-parallel graphs. *Discrete Applied Mathematics*, 115(1-3):177–186, 2001.

[124] Y. Perl and Y. Shiloach. Finding two disjoint paths between two pairs of vertices in a graph. *Journal of the ACM*, 25:1–9, 1978.

[125] S. Plotkin, D. B. Shmoys, and É. Tardos. Fast approximation algorithms for fractional packing and covering problems. *Mathematics of Operations Research*, 20:257–301, 1995.

[126] Harald Räcke. Minimizing congestion in general networks. In *43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings*, pages 43–52, 2002.

[127] P. Raghavan. Probabilistic construction of deterministic algorithms: approximating packing integer programs. *Journal of Computer and System Sciences*, 37:130–143, 1988.

[128] P. Raghavan and C. D. Thompson. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7:365–374, 1987.

[129] Satish Rao and Shuheng Zhou. Edge disjoint paths in moderately connected graphs. *SIAM J. Comput.*, 39(5):1856–1887, 2010.

[130] B. A. Reed. Tree width and tangles: a new connectivity measure and some applications. In R. A. Bailey, editor, *Surveys in Combinatorics*, volume 241 of *London Math. Soc. Lecture Note Ser.*, pages 87–162. Cambridge University Press, 1997.

[131] B. A. Reed. Algorithmic aspects of tree width. In Bruce A. Reed and Cláudia L. Sales, editors, *Recent Advances in Algorithms and Combinatorics*, pages 85–107. Springer New York, New York, NY, 2003.

[132] H. Ripphausen-Lipa, D. Wagner, and K. Weihe. Survey on efficient algorithms for disjoint paths problems in planar graphs. In W. Cook, L. Lovász, and P. D. Seymour, editors, *DIMACS-Series in Discrete Mathematics and Theoretical Computer Science, Volume 20 on the "Year of Combinatorial Optimization"*, pages 295–354. AMS, 1995.

[133] N. Robertson and P. D. Seymour. Graph Minors XIII. The disjoint paths problem. *Journal of Combinatorial Theory B*, 63:65–110, 1995.

[134] Neil Robertson and Paul D. Seymour. Graph minors. V. Excluding a planar graph. *J. Comb. Theory, Ser. B*, 41(1):92–114, 1986.

[135] Neil Robertson, Paul D. Seymour, and Robin Thomas. Quickly excluding a planar graph. *J. Comb. Theory, Ser. B*, 62(2):323–348, 1994.

[136] Michael E. Saks, Alex Samorodnitsky, and Leonid Zosin. A lower bound on the integrality gap for minimum multicut in directed networks. *Combinatorica*, 24(3):525–530, 2004.

[137] Fernanda Salazar and Martin Skutella. Single-source k-splittable min-cost flows. *Oper. Res. Lett.*, 37(2):71–74, 2009.

[138] C. Scheideler. *Universal routing strategies for interconnection networks*, volume 1390 of *LNCS*. Springer-Verlag, 1998.

[139] A. Schrijver. Homotopic routing methods. In B. Korte, L. Lovász, H. J. Prömel, and A. Schrijver, editors, *Paths, Flows and VLSI-Layout*. Springer-Verlag, Berlin, 1990.

[140] A. Schrijver. Finding k disjoint paths in a directed planar graph. *SIAM Journal on Computing*, 23:780–788, 1994.

[141] A. Schrijver. *Combinatorial Optimization: polyhedra and efficiency*. Springer-Verlag, Berlin, 2003.

[142] Alexander Schrijver, Paul D. Seymour, and Peter Winkler. The ring loading problem. *SIAM J. Discrete Math.*, 11(1):1–14, 1998.

[143] Loïc Seguin-Charbonneau and F. Bruce Shepherd. Maximum edge-disjoint paths in planar graphs with congestion 2. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 200–209, 2011.

[144] P. D. Seymour. Disjoint paths in graphs. *Discrete Mathematics*, 29:293–309, 1980.

[145] F. Shahrokhi and D. W. Matula. The maximum concurrent flow problem. *Journal of the ACM*, 37:318 – 334, 1990.

[146] F. Bruce Shepherd. Single-sink multicommodity flow with side constraints. In William Cook, László Lovász, and Jens Vygen, editors, *Research Trends in Combinatorial Optimization: Bonn 2008*, pages 261–283. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.

[147] F. Bruce Shepherd and Adrian Vetta. The inapproximability of maximum single-sink unsplittable, priority and confluent flow problems. *CoRR*, abs/1504.00627, 2015.

[148] Y. Shiloach. A polynomial solution to the undirected two paths problem. *Journal of the ACM*, 27:445–456, 1980.

[149] D. B. Shmoys. Cut problems and their applications to Divide and Conquer. In D. S. Hochbaum, editor, *Approximation algorithms for NP-hard problems*, pages 192 –231. PWS, Boston, 1997.

[150] D. B. Shmoys and É. Tardos. An approximation algorithm for the generalized assignment problem. *Mathematical Programming A*, 62:461–474, 1993.

[151] M. Skutella. Approximating the single-source unsplittable min-cost flow problem. *Mathematical Programming B*, 91:493–514, 2002. Prelim. version in FOCS 2000.

[152] Martin Skutella. A note on the ring loading problem. *SIAM J. Discrete Math.*, 30(1):327–342, 2016.

[153] J. Spencer. *Ten Lectures on the Probabilistic Method*. SIAM, Philadelphia, 1987.

[154] A. Srinivasan. An extension of the Lovász Local Lemma and its applications to integer programming. In *Proceedings of the 7th ACM-SIAM Symposium on Discrete Algorithms*, pages 6–15, 1996.

[155] A. Srinivasan. Improved approximations for edge-disjoint paths, unsplittable flow and related routing problems. In *Proceedings of the 38th Annual IEEE Symposium on Foundations of Computer Science*, pages 416–425, 1997.

[156] A. Srinivasan. Improved approximations guarantees for packing and covering integer programs. *SIAM Journal on Computing*, 29:648–670, 1999. Prelim. version in STOC 95.

[157] A. Srinivasan. New approaches to covering and packing problems. In *Proceedings of the 12th ACM-SIAM Symposium on Discrete Algorithms*, pages 567–576, 2001.

[158] Aravind Srinivasan. A survey of the role of multicommodity flow and randomization in network design and routing. In Panos M. Pardalos, Sanguthevar Rajasekaran, and José Rolim, editors, *Randomization Methods in Algorithm Design, Proceedings of a DIMACS Workshop, Princeton, New Jersey, USA, December 12-14, 1997*, volume 43 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 271–302. DIMACS/AMS, 1997.

[159] S. Stefanakos. *On the Design and Operation of High-Performance Optical Networks.* PhD thesis, ETH Zurich, No. 15691, 2004.

[160] Ola Svensson. Santa claus schedules jobs on unrelated machines. *SIAM J. Comput.*, 41(5):1318–1341, 2012.

[161] C. Thomassen. 2-linked graphs. *European Journal of Combinatorics*, 1:371–378, 1980.

[162] K. Varadarajan and G. Venkataraman. Graph decomposition and a greedy algorithm for edge-disjoint paths. In *Proceedings of the 15th ACM-SIAM Symposium on Discrete Algorithms*, pages 379–380, 2004.

[163] V. V. Vazirani. *Approximation Algorithms.* Springer-Verlag, Berlin, 2001.

[164] Jens Vygen. NP-completeness of some edge-disjoint paths problems. *Discrete Applied Mathematics*, 61(1):83–90, 1995.

[165] N. E. Young. Randomized rounding without solving the linear program. In *Proceedings of the 6th ACM-SIAM Symposium on Discrete Algorithms*, pages 170–178, 1995.