Approximation Algorithms for the Unsplittable Flow Problem

A. Chakrabarti[1] and C. Chekuri[2] and A. Gupta[2] and A. Kumar[3]

[1] Computer Science Dept., Princeton University. amitc@cs.princeton.edu
[2] Bell Labs, Lucent Tech. {chekuri,anupamg}@research.bell-labs.com
[3] Computer Science Dept., Cornell University. amitk@cs.cornell.edu

Part of this work was done while the first and the last author visited Bell Labs.

Abstract. We present approximation algorithms for the *unsplittable flow problem* (UFP) on undirected graphs. As is standard in this line of research, we assume that the maximum demand is at most the minimum capacity. We focus on the *non-uniform capacity* case in which the edge capacities can vary arbitrarily over the graph. Our results are:

- For undirected graphs we obtain a $O(\Delta\alpha^{-1}\log^2 n)$ approximation ratio, where n is the number of vertices, Δ the maximum degree, and α the expansion of the graph. Our ratio is capacity independent and improves upon the earlier $O(\Delta\alpha^{-1}(c_{\max}/c_{\min})\log n)$ bound [15] for large values of $c_{\max}/c_{\min}$. Furthermore, if we specialize to the case where all edges have the same capacity, our algorithm gives an $O(\Delta\alpha^{-1}\log n)$ approximation, which matches the performance of the best-known algorithm [15] for this special case.
- For certain strong constant-degree expanders considered by Frieze [10] we obtain an $O(\sqrt{\log n})$ approximation for the uniform capacity case, improving upon the current $O(\log n)$ approximation.
- For UFP on the line and the ring, we give the first constant-factor approximation algorithms. Previous results addressed only the uniform capacity case.
- All of the above results improve if the maximum demand is bounded away from the minimum capacity.

Our results are based on randomized rounding followed by greedy alteration and are inspired by the use of this idea in recent work [21, 9].

1 Introduction

In the *unsplittable flow problem* (UFP), we are given an n-vertex graph G and a set of k vertex pairs (terminals) $\mathcal{T} = \{(u_i, v_i) : i = 1, \ldots, k\}$; each pair (u_i, v_i) in $\mathcal{T}$ has a *demand* ρ_i and a *weight* (or profit) w_i, and furthermore, the graph is *capacitated* with edge capacities $\{c_e\}$. The goal is to find the maximum weight subset of pairs from $\mathcal{T}$, along with a path for each chosen pair, so that the entire demand for each such pair can be routed on its path while respecting the capacity constraints.

We note at the outset that even very special cases of UFP are NP-hard. For instance, when G is just a single edge, UFP specializes to the KNAPSACK problem. When each $c_e = 1$ and each $\rho_i = w_i = 1$, UFP specializes to the well-known *maximum edge-disjoint paths* problem (MEDP), where the goal is simply to find the largest number of pairs from $\mathcal{T}$ which can be simultaneously connected by edge-disjoint paths in G. MEDP is NP-hard even for planar graphs.

A substantial amount of research has focused on obtaining good approximation algorithms for MEDP and UFP due to their importance in network routing and design. For MEDP, the best known approximation ratio on general graphs is $O(\sqrt{m})$, where m is the number of edges in the graph. In *directed* graphs it is NP-hard to approximate it to a ratio better than $\Omega(m^{1/2-\varepsilon})$ [11]. However, in *undirected* graphs, which will be the focus of this paper, MEDP is only known to be hard to approximate to within constant factors. Improved approximation ratios have been obtained for several classes of graphs including trees, mesh-like planar graphs, and graphs with high expansion (see, e.g., [12, 15] for references).

Let $\rho_{\max} = \max_i \rho_i$ be the maximum demand among the pairs and $c_{\min} = \min_e c_e$ be the minimum capacity of an edge. We only consider instances with $\rho_{\max} \leq c_{\min}$; this is a standard assumption in the literature and is sometimes referred to as the *no-bottleneck* assumption. In its absence, we can embed an UFP instance on a given graph $G = (V, E)$ into any other graph $G' = (V, E')$ with $E \subseteq E'$, thus limiting our ability to study the role of the graph structure on the approximability. It is also known that UFP in directed graphs without the no-bottleneck assumption is provably harder than MEDP; in fact, it is hard to approximate to within $\Omega(n^{1-\varepsilon})$ [3]. Finally, this restriction is useful in many applications and still includes MEDP as a special case. Henceforth we will assume that $c_{\min} = 1$ and that $0 < \rho_i \leq 1$ for all i.

A special case of UFP is the *uniform capacity unsplittable flow* problem (UCUFP) in which all edges have the same capacity. UCUFP has received more attention and its approximability is usually related to the corresponding MEDP problem. In contrast, probably the only well understood case for UFP is when $c_{\min}/\rho_{\max}$ is $\Omega(\log n/\log\log n)$. In this case, an offline constant-factor approximation [18] (via randomized rounding) as well as an $O(\log n)$-competitive online algorithm [2] are known.

Our Results: In this paper we address UFP with non-uniform edge capacities on undirected graphs. For general graphs, we obtain algorithms with approximation ratios that depend on the *flow number* F of the graph G. This parameter was first defined by Kolman and Scheideler [15] who also established some of its basic properties. In particular, they showed that F is related to the expansion of the graph by $F = O(\Delta\alpha^{-1}\log n)$, where α is the edge expansion and Δ is the maximum degree. Thus, bounds stated in terms of F are typically most interesting when G is an expander, although, as noted in [15], there are other interesting cases such as meshes and hypercubes where a logarithmic bound on flow number holds even though the expansion is large.

Our results for general graphs (see Corollary 2) are:

- An $O(F \log n) = O(\Delta\alpha^{-1} \log^2 n)$ approximation for UFP. Our bound is capacity-independent, unlike the $O(\Delta\alpha^{-1}(c_{\max}/c_{\min}) \log n)$ bound of [15].
- An $O(F) = O(\Delta\alpha^{-1} \log n)$ approximation for UCUFP. This matches the bound of [15] and provides an alternative algorithm.
- When $\rho_{\max} \leq c_{\min}/B$, for integer B, the above approximation bounds improve to $O((\Delta\alpha^{-1} \log^2 n)^{1/B})$ for UFP, and $O((\Delta\alpha^{-1} \log n)^{1/B})$ for UCUFP,

Additionally, we obtain even better approximation ratios on special classes of graphs, by further exploiting some of the techniques used in proving the above. We obtain:

- An $O(\sqrt{\log n})$ approximation for UCUFP on "sufficiently strong" constant degree expanders, as defined by Frieze [10] (see Definition 1 and Theorem 3). This improves on the current $O(\log n)$ ratio [10, 15] and is the first sublogarithmic approximation for constant degree expanders that we are aware of.
- An $O(1)$ approximation for UFP on line and ring networks (see Theorems 5 and 7). Only UCUFP results were known earlier.

Our Techniques: Most approaches to approximating MEDP and UCUFP on expanders rely on proving the existence of near-optimal solutions to the multicommodity flow relaxation of the problem that use *short* (polylogarithmic in length) flow paths. Kolman and Scheideler [15] generalize this to UFP through their notion of flow number. However, their upper bound on the length of the flow paths depends on $c_{\max}/c_{\min}$ which can be quite large. We take a different approach by showing that there exist flow paths that use only a *few* (polylogarithmic number of) edges of *low* capacity, even though the overall length of the flow path might be large. High capacity edges (of capacity $\Omega(\log n)$) behave well under randomized rounding and this leaves us to worry only about the behavior of the low capacity edges under randomized rounding.

Our second idea, which proves useful for the case of the line and the ring as well, is to perform the randomized rounding step with more care. The naïve rounding involves scaling down the fractional solution before randomized rounding, where the scaling factor is chosen to be large enough to argue that none of the constraints are violated. Typically the events corresponding to the violation of these constraints are not independent and the union bound is too weak to estimate the failure probability of the randomized rounding. To overcome these problems, either the Lovász Local Lemma [19, 15] or a correlation type inequality like the FKG inequality [19] is used. These approaches are technically involved and add substantial complexity to the algorithm. The approach we will use, called the method of *alteration* [1] is applicable to monotone problems. Applications to approximation algorithms have recently been shown independently by Srinivasan [21] (who applied it to general packing *and* covering problems) and Calinescu et al. [9] (who applied it to a specific packing problem). In this approach, the first step is the same as above: scaling followed by randomized rounding. However, instead of desiring that all constraints are satisfied (feasible), we look at the random solution and alter it if it is not feasible: we simply change

1's in the random solution to 0's to ensure feasibility. This greedy (problem-dependent) alteration is designed to obtain a feasible solution, and hence the burden shifts to analyzing the expected loss in the alteration step. This turns out to be simple and effective for various problems and we believe that this idea will find more applications in the future.

Relation to Previous Work: We first discuss work connecting UFP and related problems with expansion of the underlying graph. Culminating a long line of work, Frieze [10] recently showed that for regular expanders with sufficiently strong expansion and sufficiently large (but constant) degree, there exists a constant c such that *any* $cn/\log n$ pairs can be connected via edge-disjoint paths provided no vertex appears in more than a constant (depending on and less than the degree) number of pairs. This result is optimal to within constant factors, and has been also has been extended to expander digraphs [7]. An immediate consequence of this is an $O(\log n)$ approximation for MEDP on such expanders. Kleinberg and Rubinfeld [13] in 1996 had used an earlier result of Broder, Frieze, and Upfal [8] to show that a deterministic *online* algorithm, the so-called bounded greedy algorithm (BGA), gave an $O(\log n \log\log n)$ approximation guarantee. (In fact Frieze's result mentioned above implies an $O(\log n)$ bound for BGA.) In the same paper, Kleinberg and Rubinfeld also showed the existence of a near-optimal fractional solution to any multicommodity flow instance on an expander that used only *short* paths of length $O(\log^3 n)$. This latter result formed the basis of an $O(\log^3 n)$ approximation for UCUFP on expanders by Srinivasan [19]. The above results do not explicitly specify the dependence of the approximation ratio on Δ and α. Kolman and Scheideler [15] suggest a dependence of the form $\Omega(\Delta^2\alpha^{-2})$.

The results on short flow paths (for UCUFP) have been recently improved by Kolman and Scheideler, who show the existence of near-optimal solutions to multicommodity flow instances that use paths of length $O(\Delta R)$ where R is the *routing number* of the graph [14], or solutions that use paths of length $O(F)$, where F is the *flow number* of the graph [15]. It is known that $R = O(\Delta\alpha^{-1}\log n)$ and, as we have mentioned earlier, $F = O(\Delta\alpha^{-1}\log n)$. In addition to improving the approximation ratio for UCUFP, these new results offer other advantages. The dependence on α and Δ is improved and is made explicit. The bound on the flow path lengths is stronger, and the proof is direct and simple and is based on the work of Leighton and Rao [16]. It is also shown that the flow number is more appropriate than the routing number in capturing the flow path lengths (see [15] for more details).

We now turn to the problem of UFP on the line network. In this case, the MEDP problem corresponds to the maximum independent set problem on interval graphs, which has a polynomial time algorithm. However, the UCUFP problem on the line is NP-hard (it generalizes KNAPSACK), and is equivalent to the task assignment problem on a single machine with fixed time windows. Generalizations of the task assignment problem to multiple machines and time windows have been studied in the recent past [4, 17], and most of these problems have $O(1)$-approximation algorithms as well as $O(1)$ integrality gaps. This is not necessarily the case with UFP, and in particular, if the demands are not constrained to be less than the minimum capacity, the integrality gap of the natural LP could be $\Omega(\log \rho_{\max}) = \Omega(n)$. Two of the techniques that have been used to develop $O(1)$ approximations for the task assignment problem, i.e., the local-ratio method [4] and LP rounding [17], do not seem to extend to the case of UFP. However, the recent techniques of Calinescu et al. [9] using a combination of randomized rounding, alteration and dynamic programming, can be extended, and we build upon these ideas to give an $O(1)$ approximation for UFP when $\rho_{\max} \le c_{\min}$. Without this restriction, we show that the LP integrality gap can be $\Omega(\log \rho_{\max})$ and provide a matching approximation ratio. We extend the results on the line to the ring via a simple reduction.

2 Preliminaries

UFP has a natural IP formulation based on multicommodity flow. Let $\mathcal{P}_i$ denote the set of all paths in G from u_i to v_i. The IP is:

$$\begin{array}{rll} \max & \sum_{i=1}^{k} w_i x_i, \quad \text{s.t.} & \\ & \sum_{\pi \in \mathcal{P}_i} f_\pi = x_i & i = 1, \dots, k \\ \sum_{i=1}^{k} & \sum_{\pi \in \mathcal{P}_i \,:\, \pi \ni e} \rho_i f_\pi \le c_e & e \in E(G) \\ & x_i \in \{0,1\} & i = 1, \dots, k \\ & f_\pi \in \{0,1\} & \pi \in \cup_{i=1}^{k} \mathcal{P}_i. \end{array}$$

The LP relaxation, which we shall call LPMAIN, is obtained by allowing x_i and f_π to lie in the real interval $[0,1]$. Let $(x_1, \dots, x_k, f_{\pi_1}, f_{\pi_2}, \dots)$ be a fractional solution to LPMAIN. We shall refer to $\sum_{i=1}^{k} w_i x_i$ as the *profit* or the *value* of the solution. We say that the solution *uses* a flow path π if $f_\pi > 0$. An optimal solution to LPMAIN can be obtained in polynomial time[4], and efficient combinatorial methods that obtain a $(1+\varepsilon)$-approximation are also known. We can also find efficient $(1+\varepsilon)$-approximations when we restrict the LP to use only flow paths of length bounded by a given parameter ℓ. This fact will be useful in the sequel.

Let us also make explicit the notions of expansion used in this paper:

Definition 1. *Let G be an uncapacitated n-vertex graph and for $U \subseteq V(G)$ let ∂U denote the set of edges of G with exactly one end point in U. G is said to have* expansion *α if for all $U \subseteq V(G)$ we have $|U| \le n/2 \Rightarrow |\partial U| \ge \alpha |U|$. If G is Δ-regular, it is called a* strong expander *if for sufficiently small α, arbitrary $\beta < 1 - \alpha$, $\gamma < \frac{1}{2}$, and $U \subseteq V(G)$, we have*
$|U| \le \gamma n \Rightarrow |\partial U| \ge (1-\alpha)\Delta |U|$ *and* $\gamma n < |U| \le n/2 \Rightarrow |\partial U| \ge \beta \Delta |U|$.

[4] Notice that the LP, as given here, is path-based and has exponential size. To solve it, we first solve a different, polynomial sized edge-based LP which uses flow-conservation constraints, and then perform a simple path-decomposition on the solution.

3 Expansion-based approximation bounds for UFP

Our approach is to show that the flows in any fractional solution to LPMAIN can be rerouted so as to yield a new fractional solution in which the flow paths use few edges of small capacity (we call such a solution *favorable*), where "few" is quantified using a parameter of the graph called its *flow number* [15]. The rerouting reduces the profit of the solution by a factor that can be made arbitrarily close to 1. Next, we show that a favorable fractional solution can be rounded efficiently to an integral solution with good profit.

Our rerouting is similar to that in [14, 15] and has been obtained independently following [14]; our contribution is to restrict the notion of "short" to edges of low capacity. Note that the rerouting need not be algorithmically efficient since it is just used to establish a lower bound on the profit of a favorable fractional solution; the solution itself can be obtained by solving a modified LP.

We shall denote by F the flow number of our input graph. For a definition of flow number we refer the reader to Kolman and Scheideler [15]. We shall use the facts, proven in [15], that F can be computed in polynomial time and that $F = O(\Delta\alpha^{-1}\log n)$. We remark that we shall only be concerned with the flow number of *uncapacitated* graphs. The paper [15] in fact defines flow number for capacitated graphs as well, and proves a similar bound for F in this case, with a "capacitated" definition of Δ. Recall that we normalize the capacities and demands such that $c_{\min} = 1$ and $\rho_{\max} \leq 1$.

Definition 2. *A fractional solution to* LPMAIN *is said to be* (c, d)-favorable *if every flow path used by the solution has at most d edges of capacity at most c.*

Theorem 1. *For $0 < \varepsilon \leq 1$ and $c \geq 1$, given a fractional solution to* LPMAIN *with profit W, there exists a $(c, 4cF/\varepsilon)$-favorable fractional solution with profit at least $W/(1+\varepsilon)$. Alternatively, there exists a $(c, 4F/\varepsilon)$-favorable fractional solution with profit at least $W/(c(1+\varepsilon))$.*

Corollary 1. *Suppose, as in UCUFP, that $c_e = 1$ for all edges e. For any $0 < \varepsilon \leq 1$, given a fractional solution to* LPMAIN *with profit W, there exists a $(1, 4F/\varepsilon)$-favorable fractional solution with profit at least $W/(1+\varepsilon)$.*

Theorem 2. *Given a $(\log n, d)$-favorable fractional solution to* LPMAIN *with profit W, we can efficiently compute a (random) integral solution with expected profit W' such that, for large enough d,*

1. *$W' = \Omega(W/d)$.*
2. *If, additionally, $\rho_{\max} \leq 1/B$, for integer $B \geq 2$, then $W' = \Omega\left(W/d^{1/(B-1)}\right)$.*
3. *If, additionally, each $\rho_i = 1/B$, for integer $B \geq 1$, then $W' = \Omega\left(W/d^{1/B}\right)$.*

Corollary 2. *For graphs with expansion α and maximum degree Δ, there is an $O(\Delta\alpha^{-1}\log^2 n)$ approximation for UFP, and an $O(\Delta\alpha^{-1}\log n)$ approximation for UCUFP. For integer B, if $\rho_{\max} \leq 1/B$, then the approximations improve to $O\left((\Delta\alpha^{-1}\log^2 n)^{1/B}\right)$ and $O\left((\Delta\alpha^{-1}\log n)^{1/B}\right)$ for UFP and UCUFP respectively.*

Remark. It is worth noting that the $O(\Delta\alpha^{-1}\log n)$-approximation for UFP given in [15] uses the aforementioned "capacitated" definition of Δ; their Δ can be as large as $c_{\max}/c_{\min}$ times the maximum degree. Clearly, the two definitions coincide for UCUFP.

Proof. For the first result on UFP, using $c = \log n$ and $\varepsilon = 1$ in Theorem 1, we have a $(\log n,\, O(F\log n))$-favorable fractional solution. Applying part 1 of Theorem 2 and using $F = O(\Delta\alpha^{-1}\log n)$ gives us the desired approximation guarantee. For the next result, by Corollary 1 we have a $(1,\, O(F))$-favorable solution which, for UCUFP, is the same as a $(\log n,\, O(F))$-favorable solution. Applying part 1 of Theorem 2 completes the proof.

Now suppose we have an instance $\mathcal{I}$ of UFP with $\rho_{\max} \leq \frac{1}{B}$, where $B \geq 2$. Let us create two new UFP instances, $\mathcal{I}_1$ and $\mathcal{I}_2$, with the same underlying graph as $\mathcal{I}$ but with source-sink pairs restricted to those with demands at most $\frac{1}{B+1}$ and those with demands greater than $\frac{1}{B+1}$, respectively. Part 2 of Theorem 2 gives us an $O(d^{1/B})$-approximation for $\mathcal{I}_1$ from a $(\log n,\, d)$-favorable fractional solution to the LP relaxation of $\mathcal{I}_1$. As for $\mathcal{I}_2$, note that each demand in $\mathcal{I}_2$ can be set equal to $\frac{1}{B}$ without affecting the space of feasible solutions. After making this change, part 3 of Theorem 2 gives us an $O(d^{1/B})$-approximation for $\mathcal{I}_2$. Either $\mathcal{I}_1$ or $\mathcal{I}_2$ has optimal profit at least half that of $\mathcal{I}$, so we can simply pick the better of the two approximate solutions. Finally, by Theorem 1, we can take $d = O(F\log n)$. This proves the third result. The last result, for UCUFP, follows similarly; we use Corollary 1 to take $d = O(F)$. □

For strong expanders we use a counting argument and the UCUFP approximation with $B = 2$ in Corollary 2 to obtain the following theorem.

Theorem 3. *For sufficiently large constant Δ, there is an approximation algorithm with ratio $O(\sqrt{\log n})$ for UCUFP on Δ-regular strong expanders.*

In the remainder of this section we prove the above theorems.

3.1 Producing a favorable solution

In this section we prove Theorem 1. We shall need the concept of a *balanced multicommodity flow problem* (BMFP), from Kolman and Scheideler [15]. For our purposes, a BMFP shall consist of an uncapacitated graph, a set of vertex pairs $\{(u_i, v_i)\}$, and demands $0 \leq \rho_i \leq 1$, one for each vertex pair. The total demand entering or leaving a vertex is required to equal its degree.

Suppose, as in the statement of Theorem 1, that we are given a fractional solution to LPMAIN with profit W. Let $\mathcal{P}$ be the set of all flow paths used by this solution. Set $L = 2cF/\varepsilon$. Let $\mathcal{P}'$ denote the subset of $\mathcal{P}$ consisting of paths with at least $2L$ edges of capacity at most c. We shall define a BMFP on the underlying uncapacitated graph G. For each flow path $\pi \in \mathcal{P}'$, orient it in the direction in which it carries flow. For a vertex u on π, let $\mathrm{pred}_\pi(u)$ denote the vertex which is the predecessor of u on π. We call u a *good* vertex if $\mathrm{pred}_\pi(u)$

exists and the edge $(\text{pred}_\pi(u), u)$ has capacity at most c. Let $u_1, u_2, \ldots, u_L$ be the first L good vertices on π and let $v_1, v_2, \ldots, v_L$ be the last L good vertices. We add the pairs $\{(u_i, v_i) : 1 \leq i \leq L\}$, each with demand f_π/c, to the BMFP. Once we do this for all flow paths in $\mathcal{P}'$, the total demand entering or leaving any vertex is clearly *at most* its degree; we then add dummy demands, if required, to satisfy the definition of a BMFP. We shall need the following proposition, which follows from the work in [15].

Proposition 1 (Kolman and Scheideler). *A $\frac{1}{2F}$ fraction of all the demands in this BMFP can be concurrently satisfied on the underlying uncapacitated graph G, using a family of flow paths of length at most $2F$ each.*

Let $\mathcal{Q}$ be a family of flow paths guaranteed by Prop. 1. We reroute flow going through paths in $\mathcal{P}'$, through $\mathcal{Q}$. Notice that a path $\pi \in \mathcal{P}'$ is associated with L paths in $\mathcal{Q}$ each of which "shortcuts" π. We send f_π/L flow through each of these shortcuts, adjusting the flow on edges in π appropriately. When we do this for all paths, we obtain a candidate fractional solution with profit W that uses paths with at most $\max(L + 2F, 2L)$ edges of capacity at most c. Notice that $L + 2F \leq 2L$, if $\varepsilon \leq 1$. Thus, the flow paths in this candidate solution have at most $2L = 4cF/\varepsilon$ edges with capacity at most c.

This candidate solution could violate some edge capacities. However, by Prop. 1, had we sent $f_\pi/(2cF)$ flow through each shortcut for $\pi \in \mathcal{P}'$ we would have had a total flow of at most 1 on each edge. Since $L = 2cF/\varepsilon$, we in fact have a total flow of at most ε on each edge due to the shortcuts. Thus, throwing in the flow paths in $\mathcal{P} \setminus \mathcal{P}'$ and scaling each flow value and each x_i by $1/(1+\varepsilon)$, we will have a feasible solution. The new profit after scaling is clearly $W/(1+\varepsilon)$, which proves the first part of Theorem 1.

We adapt the above argument in a simple way to obtain a $(c, 4F/\varepsilon)$–favorable solution of profit $W/(c(1+\varepsilon))$. Given a fractional solution with profit W we scale each x_i and f_π by $1/c$, reducing the profit to W/c. For each edge in G with capacity at most c, we reset the capacity to 1; the other capacities stay intact. This keeps the solution feasible, since $\rho_{\max} \leq 1$. We apply the above argument to this modified solution with $L = 2F/\varepsilon$ to obtain the claimed result.

3.2 Rounding a favorable solution

We now turn to the proof of Theorem 2.

Srinivasan [19] and Baveja and Srinivasan [5] show that in a UCUFP instance, randomized rounding yields an $O(d)$ approximation if all flow path lengths are bounded by d. However, the proof is complicated and is based on the FKG inequality. In our case we have a fractional solution to LPMAIN that is $(\log n,\, d)$-favorable for some large enough d, and has profit W. It is conceivable that the analysis in [19, 5] can be generalized to obtain an $O(d)$ approximation for our case but we believe it will be complex also. Instead, following the work of Srinivasan [21], we use randomized rounding followed by alteration to obtain an $O(d)$ approximation when rounding $(\log n,\, d)$-favorable fractional solutions. The proof is simple and transparent. Similar arguments yield an $O(1)$ approximation if $c_{\min}$ is $\Omega(\log F)$ where F is the flow number of the graph. For this latter case our proof is substantially simpler than the one in [15] which uses the Lovász Local Lemma.

Consider the following procedure: Fix an i and randomly select exactly one of the paths in $\mathcal{P}_i$; for a path $\pi \in \mathcal{P}_i$, we select it with probability $\frac{f_\pi}{16d}$. Since $\sum_{\pi \in \mathcal{P}_i} f_\pi = x_i$, we will have selected *some* path in $\mathcal{P}_i$ with probability $\frac{x_i}{16d}$. Repeat this procedure for all $i \in \{1, \ldots, k\}$. We shall refer to this as the *selection phase*.

Associate a path $\pi \in \mathcal{P}_i$ with the demand ρ_i; we shall write this demand as $\rho(\pi)$. Order all paths in $\bigcup_{i=1}^{k} \mathcal{P}_i$ in such a way that paths with $\rho(\pi) \geq \frac{1}{2}$ precede paths with $\rho(\pi) < \frac{1}{2}$. Starting from an empty set, add to it the paths picked in the selection phase, one by one, according to the above order; it is understood that when π is added, the demand $\rho(\pi)$ is to be routed along π. If a path can be added to the current set of paths without violating any edge capacity, add it. Otherwise, discard it. We shall refer to this as the *pruning phase* or the *alteration phase*. It is clear that at the end of this phase we have a feasible integral solution.

Lemma 1. *The resulting random integral solution has expected profit $\Omega(W/d)$.*

Before we prove this lemma, we need a proposition, whose proof, based on Chernoff-Hoeffding bounds, we omit.

Proposition 2. *Let $a_1, \ldots, a_m, y_1, \ldots, y_m \in [0, 1]$ be such that $\forall\, i < j : \left(a_i < \frac{1}{2} \Rightarrow a_j < \frac{1}{2}\right)$ and furthermore $\sum_{i=1}^{m} a_i y_i \leq 1$. Let $0 < \theta < 1$ and let 0-1 random variables $Y_1, \ldots, Y_m$ and $Z_1, \ldots, Z_m$ be defined as follows:*

$$Y_i = \begin{cases} 1, \text{with probability } \theta y_i \\ 0, \text{otherwise.} \end{cases} ; \quad Z_i = \begin{cases} Y_i, \text{if } \sum_{j<i} a_j Z_j \leq 1 - a_i \\ 0, \text{ otherwise.} \end{cases}$$

Then, for each i, $\Pr[Z_i = 0 \mid Y_i = 1] \;\leq\; (2 + 2e)\theta$.

Proof (of Lemma 1). Consider an edge $e \in E(G)$ and let $\pi_1, \ldots, \pi_m$ be all the paths in the fractional solution that pass through e, arranged in the order that they are considered in the pruning phase. Let X_i and $Z_i(e)$ be 0-1 random variables such that $X_i = 1$ iff π_i was picked in the selection phase and $Z_i(e) = 1$ iff the inclusion of π_i in the pruning phase did *not* violate the capacity of edge e. Then X_i and $Z_i(e)$ are equivalent to Y_i and Z_i (respectively) of Prop. 2 with $y_i = f_{\pi_i}, a_i = \frac{\rho(\pi_i)}{c_e}$, and $\theta = \frac{1}{16d}$.

Note that $a_i \leq 1$ since $\rho(\pi_i) \leq 1$ and $c_e \geq 1$. Further, $\sum_{i=1}^{m} a_i y_i \leq 1$ because of the constraints in LPMAIN. Therefore Prop. 2 applies and we have

$$\Pr[Z_i(e) = 0 \mid X_i = 1] \;\leq\; \tfrac{2+2e}{16d}. \tag{1}$$

For those edges that satisfy $c_e > \log n$, we can say something stronger. Since $0 \leq a_i \leq 1/c_e$, the random variables $\{c_e a_i X_i\}_{i=1}^{m}$ are distributed in $[0, 1]$ and their sum X satisfies $\mathrm{E}[X] \le \frac{c_e}{16d}$. Now, for any i,

$$\begin{aligned}\Pr[Z_i(e) = 0 \mid X_i = 1] \;\le\; \Pr[\textstyle\sum_{j=1}^{i-1} a_j X_j > 1 - a_i] \;\le \\ \Pr[X > c_e - a_i c_e] \;\le\; \Pr[X > c_e - 1].\end{aligned}$$

Let $\beta = \frac{c_e - 1}{\mathrm{E}[X]}$. The Chernoff-Hoeffding bound, after some routine algebra, gives

$$\begin{aligned}\Pr[Z_i(e) = 0 \mid X_i = 1] \;&\le\; \left(\frac{e^\beta}{\beta^\beta}\right)^{E[X]} \;\le\; \left(\frac{e}{\beta}\right)^{c_e - 1} \\ &\le\; \left(\frac{e c_e}{16d(c_e - 1)}\right)^{c_e - 1} \;\le\; 2^{-3c_e} \;\le\; \frac{1}{n^3},\end{aligned} \tag{2}$$

for large enough d. Finally, recall that we started with a fractional solution that was $(\log n, d)$-favorable. Let Z_i be a 0-1 random variable with $Z_i = 1$ iff π_i survives the pruning phase. Then (1) and (2) yield

$$\Pr[Z_i = 0 \mid X_i = 1] \;\le\; \textstyle\sum_{e \in \pi} \Pr[Z_i(e) = 0 \mid X_i = 1] \;\le\; d \cdot \frac{2+2e}{16d} + n^2 \cdot \frac{1}{n^3} \;\le\; \frac{1}{2},$$

for large enough n, because there at most n^2 edges on π and at most d of them have $c_e \le \log n$. This means that $\Pr[Z_i = 1] \;\ge\; (1 - \frac{1}{2}) \Pr[X_i = 1] = \frac{f_\pi}{32d}$. Therefore the expected profit of the solution after the pruning phase is at least $\frac{W}{32d} = \Omega(W/d)$. □

The first part of Theorem 2 follows immediately from Lemma 1. For the other two parts we need the following proposition which replaces Prop. 2 and is proven in a very similar manner.

Proposition 3. *Let $\{a_i\}$, $\{y_i\}$, θ, $\{Y_i\}$ and $\{Z_i\}$ be exactly as in Prop. 2. Let B be an integer. Then*

1. *If each $a_i \in [0, \frac{1}{B}]$ with $B \ge 2$, then $\forall i$, $\Pr[Z_i = 0 \mid Y_i = 1] \;\le\; e^B \theta^{B-1}$.*
2. *If each $a_i = \frac{1}{B}$ with $B \ge 1$, then $\forall i$, $\Pr[Z_i = 0 \mid Y_i = 1] \;\le\; e^B \theta^B$.*

3.3 An $O(\sqrt{\log n})$-approximation for strong expanders

We now sketch the proof of Theorem 3. We shall need a result due to Frieze [10].

Theorem 4 (Frieze). *For an n-vertex Δ-regular strong expander, with Δ sufficiently large, there exist constants k_1, k_2 such that any $(k_1 \Delta n / \log n)$ pairs of vertices, with no vertex appearing in more than $k_2 \Delta$ pairs, can be connected by disjoint paths of length $O(\log n)$ in polynomial time.*

Proof (of Theorem 3). Suppose we have an instance $\mathcal{I}$ of UCUFP. Fix an optimal integral solution $\mathcal{O}$ for $\mathcal{I}$ and partition the terminals pairs of $\mathcal{O}$ into three parts as follows:

– $\mathcal{O}_1$ includes exactly those pairs with demand at most $\frac{1}{2}$.

- $\mathcal{O}_2$ includes exactly those pairs not in $\mathcal{O}_1$, and routed by $\mathcal{O}$ on paths of length at most $\sqrt{\log n}$.
- $\mathcal{O}_3$ includes the rest of the pairs.

Our algorithm partitions $\mathcal{I}$ into two instances: $\mathcal{I}_1$, which is $\mathcal{I}$ restricted to demands $\rho_i \leq \frac{1}{2}$, and $\mathcal{I}_2$, which is $\mathcal{I} \setminus \mathcal{I}_1$. By Corollary 2, applied with $B = 2$, we can find a solution to $\mathcal{I}_1$ that $O(\Delta\beta^{-1}\log^{1/2} n)$-approximates the optimum (here β is as in Definition 1). Since $\mathcal{O}_1$ is a feasible solution for $\mathcal{I}_1$, our solution is within the same factor of $\mathcal{O}_1$. Now we solve an LP relaxation of $\mathcal{I}_2$ with the added restriction that flow path lengths are at most $\sqrt{\log n}$. By part 1 of Theorem 2, the fractional solution can be rounded to give an $O(\sqrt{\log n})$-approximation to the LP optimum. Since $\mathcal{O}_2$ is feasible for the LP relaxation, we obtain an $O(\sqrt{\log n})$-approximation to the value of $\mathcal{O}_2$.

Finally, we pick the $(k_1 \Delta n/\log n)$ most profitable demands in $\mathcal{I}$, where k_1 is the constant in Theorem 4. By that theorem, we can find disjoint paths for all of them. Since each demand in $\mathcal{O}_3$ is more than $\frac{1}{2}$, any two flow paths in a feasible solution must be edge-disjoint, and thus the number of pairs in $\mathcal{O}_3$ is at most $|E(G)|/\sqrt{\log n} = O(\Delta n/\sqrt{\log n})$. It follows that simply accepting the most profitable pairs gives us an $O(\sqrt{\log n})$-approximation to the optimum of $\mathcal{O}_3$. Since one of $\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3$ has at least a third of the profit of $\mathcal{O}$ and we approximated each within an $O(\sqrt{\log n})$ factor, we get the desired result. $\square$

4 UFP on Line and Ring Networks

In this section we consider UFP restricted to the line network. We handle the ring network in a very similar fashion; we give the relevant details at the end of the section. Before we proceed, let us fix some notation. The terminal pairs now form intervals $I_1, I_2, \ldots, I_m$ on the line $[1, n]$, with I_j having demand ρ_j and weight (or profit) w_j. Edge e on the line has capacity c_e, which we shall henceforth write as $c(e)$. For an edge e, let $\mathcal{I}(e)$ be the set of all demands (intervals) that contain e. Recall that we are working under the no-bottleneck assumption: $\rho_{\max} \leq 1 = \min_e c(e)$.

The UCUFP on the line is equivalent to a problem in resource allocation that has been studied recently [17, 4, 6, 9]; however, we will not use the resource allocation terminology. Constant factor approximation algorithms for the resource allocation problem, and consequently UCUFP on the line, have been obtained via several different techniques — LP rounding [17], the local-ratio method [4, 6], and primal-dual algorithms [4, 6]. Most of these techniques do not seem to extend to UFP on the line where capacities are non-uniform. There is, however, one exception: a recent algorithm of Calinescu et al. [9] which gives constant factor approximations for UCUFP on the line. We extend their algorithm and analysis to non-uniform capacities. Their algorithm is the following: the demands are divided into two sets, one set containing demands which are "large" compared to the (common) capacity, say 1, and the other containing the rest. Dynamic programming is then invoked to find the optimal solution on the set of large demands. For the "small" demands, the algorithm solves the LP and then randomly rounds the solution (after scaling it by a constant $\alpha < 1$). The resulting set of demands has the right weight in expectation, but it may not be feasible. The alteration step then looks at the randomly chosen demands in order of their left end points, accepting a demand in the final output if adding it maintains feasibility. Since all edges have capacity 1, a demand I_j is rejected in this step if demands sharing an edge with it and that have been inserted earlier add up to $1 - \rho_j$. However, these demands are small and their expected sum is at most α, so applying a Chernoff bound shows the probability that a demand is chosen randomly and later rejected is small.

Our algorithm for UFP is very similar to that in [9], but the analysis requires new ideas. One difficulty is the following: in the alteration step, a demand ρ_j which spans edges $e_1, e_2, \ldots, e_k$ in the left-to-right order is rejected if for some edge e_i, the demands already accepted that are using edge e_i sum up to more than $c(e_k) - \rho_j$. In the uniform capacity case it is sufficient to just look at the edge e_1 for the rejection probability. In the non-uniform case, taking a union bound for the rejection probability over edges $e_1, \ldots, e_k$ is too weak to give a constant factor approximation and we need a more careful analysis.

Another idea is needed in defining small and large demands so that dynamic programming is still feasible for the large demands, and the small demands are still small enough to allow us to make the concentration arguments. To this end, we define the *bottleneck capacity* b_j of a demand I_j to be the capacity of the lowest capacity edge on this demand. Now a demand I_j is *δ-small* if $\rho_j \leq \delta b_j$, else it is *δ-large*.

In the sequel, we show how to find the optimal solution for the δ-large demands, and a constant factor approximation for the set of δ-small demands, for some appropriate choice of δ. We then output the better of the two solutions.

Large Demands: The following lemma is key to invoking dynamic programming to find an optimal solution for the δ-large demands in $n^{O(1/\delta^2)}$ time ; the remaining details are routine and we omit them for lack of space.

Lemma 2. *The number of δ-large demands that cross an edge in any feasible solution is at most $2\lfloor 1/\delta^2 \rfloor$.*

Proof. Fix a feasible solution $\mathcal{S}$, and consider an edge e. Let S_e be the set of all δ-large demands in $\mathcal{S}$ that cross e. We partition S_e in to two sets S_ℓ and S_r as follows: a demand in S_e is in S_ℓ if it has a bottleneck capacity edge to the left of e (including e), otherwise the demand is in S_r. We show that $|S_\ell| \leq \lfloor 1/\delta^2 \rfloor$, and a similar argument works for $|S_r|$.

Let A be the set of bottleneck edges for demands in S_ℓ and let e' be the rightmost edge in A. Since e' is the bottleneck edge for some δ-large demand $I_j \in S_\ell$, by definition, $\rho_j \geq \delta c(e')$. Since $\rho_j \leq c_{\min}$, it follows that $c(e') \leq c_{\min}/\delta$. Because e' is the rightmost edge in A, all demands in S_ℓ pass through e'. But each demand I_k in S_ℓ is δ-large, which implies that $\rho_k \geq \delta b_k \geq \delta c_{\min}$. It follows that $|S_\ell| \leq \lfloor c(e')/(\delta c_{\min}) \rfloor \leq \lfloor 1/\delta^2 \rfloor$. □

Small Demands: When all demands are δ-small, we give a 12-approximation for some small enough δ. In the interest of clarity, we have not tried to optimize the constants too much. We first solve the LP LPMAIN for the problem. Let x_j be the fractional value assigned to demand I_j. We define two $\{0,1\}$-random variables X_j and Y_j thus:

1. Let $\alpha < 1$ be a constant, to be chosen later. Let X_j be set to 1 independently with probability αx_j.
2. Sort the demands corresponding to $X_j = 1$ in order of their left end points (breaking ties arbitrarily). Look at them in this order, adding the current demand to the output if adding it does not violate any edge capacity. Set $Y_j = 1$ if demand I_j is output.

By construction, this procedure produces a feasible solution. Clearly, $\mathrm{E}[X_j] = \Pr[X_j = 1] = \alpha x_j$, and the expected weight of the final solution is $\mathrm{E}[Y_j] = \Pr[Y_j = 1] = \alpha x_j \cdot \Pr[Y_j = 1 | X_j = 1]$. The rest of the argument shows that $\Pr[Y_j = 0 \mid X_j = 1]$, the chance of rejection, is at most a half (for $\alpha = \frac{1}{6}$). This in turn shows that the expected weight is at least $\sum w_j x_j / 12$, giving us the claimed 12-approximation.

Let us focus on a particular demand I_j with $X_j = 1$, and let $E_j = \langle e_1, \ldots, e_k \rangle$ be the edges on I_j from left to right. The crucial idea is the following: when considering I_j, its probability of rejection depends on whether there is "enough room" on all these edges. Instead of taking a union bound over all edges, we choose a subsequence of edges such that the capacity of each edge drops by half, and such that for a demand to be rejected, a "bad" event happens at one of these chosen edges. Now a union bound on these bad events at these edges suffices. We show that this union bound gives us a geometric sum, and thus the chance of rejection is a constant times the probability of rejection on some edge e_i. Finally, arguments similar to that in [9] complete the proof.

Formally, create a subsequence $E'_j = \langle e_{i_1}, e_{i_2}, \ldots, e_{i_h} \rangle$ of E_j as follows: set $i_1 = 1$, and hence $e_{i_1} = e_1$. For $\ell > 1$, set $i_\ell = \min\{t : t > i_{\ell-1} \text{ and } c(e_t) < c(e_{i_{\ell-1}})/2\}$. In other words e_{i_ℓ} is the first edge to the right of $e_{i_{\ell-1}}$ with capacity at most half the capacity of $e_{i_{\ell-1}}$. If there is no such edge we stop the construction of the sequence. For $1 \leq a \leq h$, let $\mathcal{E}_a$ denote the (bad) event that the random demands chosen in step 1 use at least $\frac{1}{2}c(e_{i_a}) - \delta b_j$ capacity in the edge e_{i_a}. The following lemma shows that it is enough to bound the chance that no bad event occurs on these chosen edges.

Lemma 3. $\Pr[Y_j = 0 \mid X_j = 1] \leq \sum_{a=1}^{h} \Pr[\mathcal{E}_a]$.

Proof. If $Y_j = 0$ and $X_j = 1$ then some edge $e_g \in E_j$ had a capacity violation when I_j was considered for insertion. Let e_{i_a} be the edge in E'_j to the left of e_g and closest to it. (Here, an edge is considered to be "to the left" of itself.) Note that such an edge always exists since $e_{i_1} = e_1$, and e_1 is the left most edge in I_j. By the construction of the subsequence, $c(e_g) \geq \frac{1}{2}c(e_{i_a})$. If the capacity of e_g was violated while trying to insert I_j, it must be that the capacity of demands already accepted that cross e_g is at least $c(e_g) - \rho_j$ which is lower bounded by

$\frac{1}{2}c(e_{i_a}) - \delta b_j$: we use the fact that I_j is small which implies that $\rho_j \leq \delta b_j$ and the fact that $c(e_g) \geq \frac{1}{2}c(e_{i_a})$. However, any interval that is accepted before I_j and crosses e_g, must also cross e_{i_a}, and thus event $\mathcal{E}_a$ occurs. Applying the trivial union bound, we have $\Pr[Y_j = 0 \mid X_j = 1] \leq \sum_a \Pr[\mathcal{E}_a]$. □

We now show that the sequence $\Pr[\mathcal{E}_a]$ is *geometric*, which makes the union bound work:

Lemma 4. *For* $\alpha = \frac{1}{6}$*and*$\delta = \frac{3}{100}$*, we have* $\Pr[\mathcal{E}_a] \leq \left(\frac{1}{3}\right)^{c(e_{i_a})}$*, and therefore,* $\sum_{a=1}^{h} \Pr[\mathcal{E}_a] \leq \frac{1}{2}$.

Proof. Let $Q_a = \sum_{I_s \in \mathcal{I}(e_{i_a})} \rho_s X_s$ be the random variable that gives the sum of demands that edge a intersects *and* that are chosen in step 1. Since each $\rho_s \leq 1$, the independent variables $\{\rho_s X_s\}$ are distributed in $[0,1]$. We have $\Pr[\mathcal{E}_a] = \Pr[Q_a \geq \frac{1}{2}c(e_{i_a}) - \delta b_j]$. Setting $\beta = (1/2 - \delta - \alpha)/\alpha$, and using the fact that $b_j \leq c(e_{i_a})$ gives $\Pr[\mathcal{E}_a] = \Pr[Q_a \geq \frac{1}{2}c(e_{i_a}) - \delta b_j] \leq \Pr[Q_a \geq (1+\beta)\alpha c(e_{i_a})]$. Also, $\mathrm{E}[Q_a] = \sum_{I_s \in \mathcal{I}(e_{i_a})} \rho_s \mathrm{E}[X_s] = \sum_{I_s \in \mathcal{I}(e_{i_a})} \alpha \rho_s x_s \leq \alpha c(e_{i_a})$, where the last inequality follows from the feasibility of the LP solution. Since Q_a is a sum of independent random variables distributed in $[0,1]$ we apply a Chernoff-Hoeffding bound to get the following: $\Pr[Q_a \geq (1+\beta)\alpha c(e_{i_a})] \leq \left(e^{\beta}/(1+\beta)^{1+\beta}\right)^{\alpha c(e_{i_a})} \leq (1/3)^{c(e_{i_a})}$, where the final inequality follows by plugging in the constants we chose for α and δ. Since $c(e_{i_a}) < c(e_{i_{a-1}})/2$ and each $c(e_{i_a}) \geq 1$, we now get: $\sum_{a=1}^{h} \Pr[\mathcal{E}_a] \leq \sum_a (1/3)^{c(e_{i_a})} \leq \sum_{i \geq 0} (1/3)^{2^i} \leq \sum_{j \geq 1} (1/3)^j = \frac{1}{2}$. □

Note that this implies that $\Pr[Y_j = 0 \mid X_j = 1] \leq \frac{1}{2}$, and hence gives us the claimed 12-approximation for the δ-small demands. Combining this with the fact that we can optimally solve for the set of δ-large demands, we get our main result of this section. We have not optimized the values of δ and α to obtain the best constant in the theorem below.

Theorem 5. *There is a* 13*-approximation for UFP on the line if* $\rho_{\max} \leq c_{\min}$.

Corollary 3. *There is a constant factor approximation for UFP on the line when* $\rho_{\max}/\rho_{\min}$ *is bounded even without the no-bottleneck assumption. Hence, for arbitrary demands we get an* $O(\log \rho_{\max})$ *approximation.*

Proof. Since the analysis for the δ-small demands does not use the fact that $\rho_{\max} \leq c_{\min}$, we need to only consider the large demands. For the δ-large demands, an argument similar to that in Lemma 2 works when $\rho_{\max}/\rho_{\min}$ is bounded. □

The performance of our algorithm matches the integrality gap of the LP to within a constant factor; we defer the proof to the full version.

Theorem 6. *The integrality gap of the natural LP is* $O(1)$ *when* $\rho_{\max} \leq c_{\min}$*. For arbitrary demands the integrality gap is* $\Theta(\log \rho_{\max})$ *which can be* $\Omega(n)$.

4.1 UFP on a Ring Network

Finally, we consider UFP on the ring network. Unlike the line network, this gives us a choice of one of two paths for each demand. However, we can reduce the problem on the ring to that on a line network with a slight loss in the approximation factor as follows. Let e be any edge on the ring with $c(e) = c_{\min}$. Consider any integral optimal solution $\mathcal{O}$ to the problem. The demands routed in $\mathcal{O}$ can be partitioned into two sets $\mathcal{O}_1$ and $\mathcal{O}_2$ where those in $\mathcal{O}_1$ use e and those in $\mathcal{O}_2$ do not. We remove e and solve the problem approximately on the resulting line network. This clearly approximates the value of $\mathcal{O}_2$. To approximate the solution for $\mathcal{O}_1$, for each demand we choose the path that uses e and solve a knapsack problem to find the maximum weight set of demands that can be routed with capacity bounded by c_e. Since $c(e) = c_{\min}$ any solution feasible at e will be feasible for the entire network. Thus we obtain:

Theorem 7. *For UFP on the ring there is a* $(1 + \alpha)$ *approximation where* α *is the approximation factor for the problem on the line.*

5 Concluding Remarks

We also note that an online $O(F \log n)$-approximation for UFP can be obtained by combining the bounded greedy algorithm [13], and the online algorithm of Awerbuch, Azar and Plotkin [2]; we defer the details to the full version of the paper.

Acknowledgments: We are grateful to Bruce Shepherd for suggesting the unsplittable flow problem on the line and for several discussions.

References

1. N. Alon and J. Spencer. *The Probabilistic Method*. Wiley Interscience, New York, 1992.
2. B. Awerbuch, Y. Azar, and S. Plotkin. Throughput-competitive online routing. In *Proceedings of the 34th Annual IEEE Symposium on Foundations of Computer Science*, pp. 32–40. 1993.
3. Y. Azar and O. Regev. Strongly polynomial algorithms for the unsplittable flow problem. In *Proceedings of the 8th Integer Programming and Combinatorial Optimization Conference*. 2001.
4. A. Bar-Noy, R. Bar-Yehuda, A. Freund, J. S. Naor, and B. Scheiber. A unified approach to approximating resource allocation and scheduling. In *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing*, pp. 735–744, 2000.
5. A. Baveja and A. Srinivasan. Approximation algorithms for disjoint paths and related routing and packing problems. *Math. Oper. Res.*, 25(2):255–280, 2000.
6. P. Berman and B. DasGupta. Improvements in throughput maximization for real-time scheduling. In *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing*, pp. 680–687, 2000.

7. T. Bohman and A. M. Frieze. Arc-disjoint paths in expander digraphs. In *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science*. 2001.
8. A. Z. Broder, A. M. Frieze, and E. Upfal. Existence and construction of edge-disjoint paths on expander graphs. *SIAM Journal on Computing*, 23(5):976–989, 1994.
9. G. Calinescu, A. Chakrabarti, H. Karloff, and Y. Rabani. Improved approximation algorithms for resource allocation. In *Proceedings of the 9th Integer Programming and Combinatorial Optimization Conference*, 2002.
10. A. M. Frieze. Edge-disjoint paths on expander graphs. *SIAM Journal on Computing*, 30(6):1790–1801, 2001.
11. V. Guruswami, S. Khanna, R. Rajaraman, F. B. Shepherd, and M. Yannakakis. Near-optimal hardness results and approximation algorithms for edge-disjoint paths and related problems. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing*, pp. 19–28. 1999.
12. J. M. Kleinberg. *Approximation Algorithms for Disjoint Paths Problems*. Ph.D. thesis, MIT, 1996.
13. J. M. Kleinberg and R. Rubinfeld. Short paths in expander graphs. In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science*, pp. 86–95. 1996.
14. P. Kolman and S. Scheideler. Simple on-line algorithms for the maximum disjoint paths problem. In *Proceedings of 13th ACM Symposium on Parallel Algorithms and Architectures*. 2001.
15. P. Kolman and S. Scheideler. Improved bounds for the unsplittable flow problem. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms*. 2002.
16. F. T. Leighton and S. B. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM*, 46(6):787–832, 1999. (Preliminary version in *29th Annual Symposium on Foundations of Computer Science*, pages 422–431, 1988).
17. C. A. Phillips, R. N. Uma, and J. Wein. Off-line admission control for general scheduling problems. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 879–888. 2000.
18. P. Raghavan and C. D. Thompson. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7(4):365–374, 1987.
19. 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*, pp. 416–425. 1997.
20. A. Srinivasan. Improved approximation guarantees for packing and covering integer programs. *SIAM J. Comput.*, 29(2):648–670, 1999.
21. A. Srinivasan. New approaches to covering and packing problems. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 567–576. 2001.