Near-Optimal Hardness Results and Approximation Algorithms for Edge-Disjoint Paths and Related Problems

Venkatesan Guruswami* Sanjeev Khanna† Rajmohan Rajaraman‡ Bruce Shepherd†
Mihalis Yannakakis†

Abstract

We study the approximability of two classes of network routing problems. The first class of problems in our study correspond to classical multicommodity flow problems of the following form: We are given a network G with integer capacities on its edges, together with source-sink pairs (s_i, t_i), $1 \leq i \leq k$, such that a positive integer demand d_i and a positive "profit" r_i is associated with each pair. A feasible solution is a subset $\mathcal{S}$ of the (s_i, t_i) pairs such that demands associated with pairs in S can be fully met through a routing which respects all capacity constraints, and the objective is to maximize the total profit associated with the satisfied pairs. We consider two natural variants: *unsplittable flow* (USF) where each pair must be satisfied by routing all its demand on a *single* path, and *integral splittable flow* (ISF) where the flow satisfying a demand can be *split* in several paths, each carrying an integral amount of flow. In the special case when all demands, profits, and capacities equal one, both these variants reduce to the classical NP-hard *edge-disjoint paths* problem (referred to as EDP). An $O(\sqrt{m})$ approximation is known for EDP [15] and has recently been obtained for the more general USF problem [1] under the assumption that $d_{\max} \leq u_{\min}$, where m denotes the number of edges, $d_{\max}$ denotes the maximum demand and $u_{\min}$ denotes the minimum edge capacity. While it has been generally believed that these problems are very hard to approximate, only MAX SNP-hardness was known thus far. We prove here the tight result that in *directed* graphs, for any $\epsilon > 0$, EDP (and hence also USF), is NP-hard to approximate within $m^{1/2-\epsilon}$. We also give a simple $O(\sqrt{m \log m} \log \log m)$-approximation algorithm for USF under the assumption that $d_{\max} \leq u_{\min}$, and an $O(\sqrt{m} \log^{3/2} m)$ approximation when the only assumption is that $d_{\max}$ is polynomially bounded. Our algorithms also turn out to be much simpler than the existing ones. The hardness result for EDP trivially implies an identical result for approximating ISF on directed networks. On the algorithmic side, we give a simple greedy algorithm that gives a $O(\sqrt{m d_{\max}} \log^2 m)$-approximation for ISF.

The second class of routing problems in our study is another well known one: the goal here is to find a maximum number of length bounded edge-disjoint paths between given source-sink pairs in a graph G. We refer to this problem as the *bounded length edge-disjoint paths* (BLEDP) problem. We show that, for any $\epsilon > 0$, BLEDP is hard to approximate within $m^{1/2-\epsilon}$ even on undirected graphs. On the algorithmic side, we give an $O(\sqrt{m})$-approximation algorithm for BLEDP. We also consider a special case of BLEDP, which we refer to as (s,t)-BLEDP, in which there is only one source-sink pair. Without the length restriction, this problem reduces to the polynomial-time solvable max-flow problem. However, with the length restriction, we show that this problem is hard to approximate within $m^{1/2-\epsilon}$, for any $\epsilon > 0$, on directed graphs; for undirected graphs we show MAX SNP-hardness even when the length bound is a (small) constant.

1 Introduction

In the *edge-disjoint paths* problem, denoted EDP, we are given a (possibly directed) graph G and a set $\mathcal{T} = \{(s_i, t_i) : 1 \leq i \leq k\}$ of k source-sink pairs, and the objective is to connect a maximum number of these pairs via edge-disjoint paths. EDP turns out to be a fundamental, extensively studied problem of combinatorial optimization, algorithmic graph theory and operations research, and figures in Karp's original list of NP-hard problems [13]. This paper investigates the approximability of EDP and of two related classes of network routing problems (on undirected as well as directed graphs) that are natural generalizations and variants of EDP. We now define these classes of problems for the directed graph case; we omit defining their standard restrictions to the undirected case.

Multicommodity Flow Problems: In multicommodity problems, we are given a directed graph $G = (V, A)$ and a set $\mathcal{T}$ of k source-sink pairs as above – we let $m = |A|$ and $n = |V|$ throughout. In addition, we are also given an integer capacity function $u : A \rightarrow \mathbf{Z}$ on the arcs, and a positive integer demand d_i for each (s_i, t_i) pair in $\mathcal{T}$, $1 \leq i \leq k$; d_i represents the bandwidth requested for flow from source s_i to sink t_i. We use the notation u_{min} and d_{max} to denote the minimum capacity and maximum demand value respectively. In all versions discussed, we assume that the

*MIT Laboratory for Computer Science, 545 Technology Square, Cambridge, MA 02139. Email: venkat@theory.lcs.mit.edu.

†Bell Labs, 700 Mountain Avenue, Murray Hill, NJ 07974. Email: {sanjeev,bshep,mihalis}@research.bell-labs.com.

‡College of Computer Science, Northeastern University, Boston MA 02115. Email: rraj@ccs.neu.edu. Part of this work was done when the author was visiting Bell Labs during his DIMACS postdoctoral fellowship 1997-1998.

bandwidth assigned (or reserved) for a pair (s_i, t_i) induces a standard network flow between s_i and t_i of value d_i. Depending on the version addressed, we may require extra conditions on these flows, e.g., for *unsplittable flow problems*, we require this flow to be sent on a single path.

We denote by $\mathcal{P}_i$ the set of all (simple) directed paths in D from s_i to t_i. A *routing* (of $\mathcal{T}$) in D is an assignment $x : \mathcal{P} \to \mathbf{R}_+$ of weights to directed paths in D, where $\mathcal{P} = \cup_{1 \le i \le k} \mathcal{P}_i$. A routing is said to *fulfill* the demand d_i for pair i, if $\sum(x(P) : P \in \mathcal{P}_i) = d_i$, i.e, if the following demand constraint is satisfied with equality:

$$\sum(x(P) : P \in \mathcal{P}_i) \le d_i. \quad (1)$$

A routing x *satisfies the arc capacities* if

$$\sum(x(P) : P \in \mathcal{P}, e \text{ is an arc of } P) \le u(e) \quad (2)$$

holds for every arc $e \in A$.

We consider several versions of these multicommodity routing problems. In each case, there is a profit $r_i \ge 0$ associated with each demand pair (s_i, t_i) and we wish to maximize our total profit. The common thread is that we only gain profit for the pair (s_i, t_i) if our routing has *fulfilled* its demand d_i.

We can consider two basic feasibility models: (i) *unsplittable flow* (USF), where we only gain profit for the demand pair (s_i, t_i) if our routing has assigned a flow of value d_i to a single path in $\mathcal{P}_i$, and, (ii) *splittable flow* (SF), where we gain the profit r_i as long as our routing has assigned a total weight of d_i to the paths in $\mathcal{P}_i$. In the splittable version, we further require our routings to be integral – this is referred to as *integral splittable flow* (ISF). Throughout, we let O_u, O_{is}, O_s respectively denote the maximum earnable profit from a feasible routing for the *unsplittable, integral-splittable* and *fractional-splittable* versions. One natural assumption is made throughout for all versions of the problem:

$$\text{The supply graph always has enough capacity to satisfy any single demand.} \quad (3)$$

Note that in the special case when all demands, capacities and profits equal one, both USF and ISF reduce to EDP. Hence these problems are NP-hard themselves, which motivates the investigation of their efficient approximability, which is one of the focuses of this paper.

Bounded Length Edge Disjoint Paths (BLEDP) Problems: A second set of problems that we study here is that of routing a maximum number of edge-disjoint paths between specified source-sink pairs in a network such that each of the paths used is of length at most L, for some length-bound L that is also part of the problem instance. We refer to this problem as BLEDP. As observed in Kleinberg [15], this length constraint, which arises quite naturally in practical routing problems, can transform tractable disjoint paths problems into NP-hard variants. For instance, even the single source-single sink case of this problem, referred to as (s,t)-BLEDP, stays NP-hard.

1.1 Overview of Our Results

EDP, USF **and** ISF: Approximation algorithms for USF and EDP have been studied in prior works [20, 14, 15, 25, 1, 18, 19, 5]; [15] provides a comprehensive background on these and related problems. Most of these algorithms begin with a linear programming relaxation of the problem and "round" the solutions suitably to get an approximate integer solution. The best known approximation factor for EDP is $O(\sqrt{m})$ [15] (an $O(\sqrt{m})$ approximation for weighted EDP, where profits are not necessarily all one, is presented in [25]). In recent work, an $O(\sqrt{m})$ approximation algorithm has been obtained for the more general USF problem [1], under the assumption that $d_{\max} \le u_{\min}$. The preceding result improves upon the $O(\sqrt{m}\log m)$ approximation algorithm of [19] for USF. All these approximation bounds are rather weak and reflect the generally appreciated hardness of these problems. Yet no better hardness than MAX SNP-hardness is known for any of these problems. We prove that even EDP on directed graphs is NP-hard to approximate within a factor of $m^{1/2-\epsilon}$ for any $\epsilon > 0$; our proof is surprisingly simple and does not even rely on the PCP theorem. Recently, Ma and Wang [21] have independently shown a weaker hardness result, namely, EDP on directed graphs is *Quasi*-NP-hard to approximate within $2^{O(\log^{1-\epsilon} m)}$ for any $\epsilon > 0$. Their proof uses the hardness of approximating LABEL COVER, and hence relies on the PCP theorem.

On the algorithmic side, we present a simple randomized $O(\sqrt{m}\log^{3/2} m)$ approximation algorithm for USF with polynomially bounded demands, without making the assumption that $d_{max} \le u_{min}$. With the assumption that $d_{max} \le u_{min}$, the approximation guarantee improves to $O(\sqrt{m \log m}\log\log m)$. While the preceding approximation guarantee is weaker than the $O(\sqrt{m})$ bound achieved recently by [1], the significance of our result lies in the fact that our randomized algorithm uses, perhaps, the most basic rounding scheme introduced by Raghavan and Thompson [23] and our analysis relies on elementary combinatorial arguments and straightforward Chernoff-type bounds. We also achieve an $O(\sqrt{m}\log^2 m)$ approximation for USF using a simple greedy algorithm; the main contribution here is extending the analysis of [19] for EDP to handle general capacities. While our strong inapproximability results above apply only for the directed case, we can also prove the (once again tight up to polylogarithmic factors) result that *undirected* USF is NP-hard to approximate within $n^{1/2-\epsilon}$, for any $\epsilon > 0$, if we consider the *node-capacitated* version [16] instead.

The integral splittable version of the problem was shown to be NP-hard on directed as well as undirected graphs (even with just two sources and sinks) in [6], and to our knowledge no explicit results on its approximability appear in the literature. Our hardness result for EDP trivially implies a similar hardness bound for approximating ISF on directed graphs, and an easy reduction from Independent Set shows that the same hardness bound of $m^{1/2-\epsilon}$ applies for the undirected case as well. On the algorithmic side, we present a simple greedy algorithm, again generalizing the one in [19], that achieves an approximation guarantee of $O(\sqrt{md_{max}}\log^2 m)$.

BLEDP: We show that BLEDP can be approximated in polynomial time within a factor of $O(\sqrt{m})$. We prove a matching hardness result of $m^{1/2-\epsilon}$ for any $\epsilon > 0$ that works for undirected graphs as well. For (s,t)-BLEDP we prove an inapproximability result of $m^{1/2-\epsilon}$ for directed graphs, and MAX SNP-hardness for undirected (and directed) graphs. The MAX SNP-hardness applies even when the length bound is a (small) constant. We also present a simple greedy algorithm for BLEDP that achieves an $O(\sqrt{m})$ approximation.

REMARK: In general, directed versions of these problems appear to be harder than their undirected counterparts. Accordingly, all our algorithms are described for the directed case, but they all work for the undirected case as well. Regarding hardness results, unless mentioned otherwise the result applies only to the directed case – but if a hardness result is stated or proved specifically for the undirected case, we stress that a similar result will hold for the directed case as well.

1.2 Organization

A portion of our algorithmic work follows a linear programming based approach, and hence we begin by describing the relaxations we use, bound their integrality gaps, and note a useful property about the structure of basic feasible solutions in Section 2. In Section 3, we present the hardness results for EDP, USF and ISF, and present an LP-based approximation for USF. We study the hardness of BLEDP problems in Section 4. Finally, in Section 5, we present simple greedy algorithms for all our problems that almost match the corresponding hardness bounds. The inapproximability result for undirected node-capacitated USF appears in Appendix A.

2 LP Formulations and Rounding

LP Formulations: A natural relaxation (in the sense of the objective function) of our network routing problems is the following linear program (LP) LP-Basic.

$$\max\{\sum_{i=1}^{k}\frac{r_i}{d_i}(\sum_{P\in\mathcal{P}_i}x(P)):\ x\geq 0 \text{ and satisfies the constraints (1), (2)}\} \tag{4}$$

Let LP_s denote the optimum value of this LP. Of course, in solving such an LP, one would resort to the well-known compact formulation which uses variables f_e^i for the flow for the demand pair (s_i,t_i) through edge e, for each $i\in[k]$ and $e\in A$. For the purposes of exposition, however, we view our solutions as vectors in $\mathbf{R}^{\mathcal{P}}$. For any such vector x, we denote by $supp(x)$ the set of paths P for which $x(P)>0$. One sees immediately that LP_s may be much more than our desired optimum.

Example 2.1 *Let G be a cycle with unit capacities on the arcs. Also between each pair of vertices consider a demand of size 2 and a profit of 1. Then clearly $LP_s = m/2$, however, $O_{is}=O_s=1$ (N.B. $O_u=0$ and would hence violate (3)).*

For many versions of our problems, we may easily amend the LP formulation to get a tighter relaxation. For instance, in the unsplittable version we define $\mathcal{P}_i^*$ to be those paths in $\mathcal{P}_i$ for which each arc has capacity at least d_i (and let $\mathcal{P}^*$ be the union of these sets). Clearly, any feasible solution to the unsplittable problem may only use the paths in $\mathcal{P}^*$ in its support. Thus we define the linear program LP-Unsplit:

$$\max\{\sum_{i=1}^{k}\frac{r_i}{d_i}(\sum_{P\in\mathcal{P}_i^*}x(P)):\ x\geq 0 \text{ and satisfies the constraints (1), (2)}\} \tag{5}$$

We use LP_u to denote the optimum value of this LP. Clearly we have $O_u\leq LP_u\leq LP_s$ and $O_u\leq O_{is}\leq O_s\leq LP_s$. Again, we view solutions as vectors in $\mathbf{R}^{\mathcal{P}^*}$. It turns out that this LP is a better approximation for the unsplittable flow problem than (4) was for the splittable flow problem. We will see that: $O_u=\Omega(LP_u/\sqrt{m}\,\mathrm{poly}(\log m))$.

For a solution x to any of our LP's, we use $G(x)$ to denote the set of (*good*) demand pairs (s_i,t_i) for which (1) is satisfied with equality. $B(x)$ denotes the set of demand pairs (s_i,t_i) which are not good but for which the left hand side of (1) is positive. Finally, let $S(x)=B(x)\cup G(x)$ be the set of demand pairs which are at least partially satisfied. There seems to be very little known about the structure of basic solutions for such multicommodity LP's. One elementary result we can prove is the following.

Proposition 2.1 *If x is a basic optimal solution to (4) or (5), then $|supp(x)|\leq|G(x)|+m$.*

Proof: We argue this for (4) but the same holds for (5). One easily sees that each unit vector in $\mathbf{R}^{\mathcal{P}}$ is feasible for (4) and hence the solution space is full-dimensional. It follows that any basic solution must satisfy some linearly independent subsystem of $|\mathcal{P}|$ constraints. The result now follows. □

Proposition 2.1 gives the following weak lower bound on the number of pairs which are satisfied by an unsplit flow.

Proposition 2.2 *If x is a basic optimal solution to (4), then at least $|S(x)|-m$ pairs are satisfied by an unsplit flow under x.*

Proof: Let x be a basic optimal solution to (4) and for each demand pair (s_i,t_i), let n_i denote $|supp(x)\cap\mathcal{P}_i|$. Hence (s_i,t_i) is satisfied by an unsplit flow under x if and only if $n_i=1$ and $(s_i,t_i)\in G(x)$. Let $U(x)$ denote the set of such pairs. Then, by Proposition 2.1, we have:

$$\begin{aligned}|G(x)|+m &\geq |supp(x)|=\sum_i n_i\\ &\geq |U(x)|+|B(x)|+2|G(x)-U(x)|\end{aligned}$$

and hence $|U(x)|\geq|B(x)|+|G(x)|-m=|S(x)|-m$, as required. □

If our demands set $\mathcal{T}$ consists of every pair of vertices in G so that for each pair of vertices $i\neq j$ there is a positive integral demand $d(ij)$, then we call this an *all-pairs* instance. The preceding immediately implies:

Corollary 2.1 *If all profits equal one, i.e if $r\equiv\bar{1}$, then $O_{is}\geq O_u\geq LP_s-m$, and in particular for any all-pairs instance, the LP (4) gives a factor 2 approximation algorithm as long as $d(ij)\leq u(ij)$ for each $(i,j)\in A$.* □

Another corollary, using Example 2.1, is a tightness result for the formulation (4) – the proof is easy and is omitted.

Corollary 2.2 $O_s\geq O_{is}\geq O_u\geq\frac{1}{2m}LP_s$. □

Standard Rounding and Deviation Bounds: We will later show how the LP formulation for USF can be rounded to obtain an approximate solution with performance ratios almost matching our hardness result. We now review a standard rounding technique and develop some bounds on its performance that we will use in our analysis later.

Consider a solution to our LP formulation for USF. Let z_i denote the fraction of the demand between the pair (s_i,t_i) that is satisfied by this solution. Decompose the flow of value z_id_i into a set of flow paths $\{\Gamma_{i,1},\Gamma_{i,2},...,\Gamma_{i,q_i}\}$ where flow on path $\Gamma_{i,j}$ is given by $f_{i,j}$. Now consider the following rounding procedure, introduced by Raghavan and Thompson [23] in a classic paper. Each (s_i,t_i) pair is routed independently with probability z_i and once a pair is chosen, we toss a q_i-sided dice with the property that the jth face shows up with probability $f_{i,j}/(z_i d_i)$. We choose to route the pair along the path $\Gamma_{i,j}$ if the jth face turns up. This rounding procedure is referred to as the *standard rounding* from here on. In what follows, we develop some properties of this rounding procedure. We start by stating some well-known deviation bounds [3, 9].

Proposition 2.3 *(Chernoff-Hoeffding Bounds) Let X_1, X_2, $\ldots$, X_ℓ be a set of k independent random variables in $[0,1]$ and let $X = \sum_{i=1}^{\ell} X_i$.*

1. *For any $\delta \geq 0$, we have:*

$$\Pr[X > (1+\delta)E(X)] \leq \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{E(X)}$$

2. *For $0 \leq \delta < 1$, we have:*

$$\Pr[X < (1-\delta)E(X)] \leq e^{-\delta^2 E(X)/2}.$$

Proposition 2.4 *Consider an LP solution to the* USF *problem with all demands between 1 and 2. Let S be a set of arcs and let f_e denote the flow through an arc $e \in S$ in an LP solution. Then if X is the random variable indicating the total number of paths in the standard rounding solution that use at least one arc in the set S, $\Pr[X > c \cdot \max\{\sum_{e\in S} f_e, \log m\}] < 1/m^2$ for some suitably large constant c.*

Proof: For each demand pair (s_i, t_i) with some non-zero flow in the LP solution, let $\{\Gamma_{i,j}\}_{j\geq 1}$ denote the set of paths in the flow decomposition that use at least one arc in S. Let $f_{i,j}$ be the flow on the path $\Gamma_{i,j}$. We assign every $\Gamma_{i,j}$ to some arc in S used by it. Let $\Gamma(e)$ denote the set of paths assigned to an arc $e \in S$. Let X_e denote the number of paths in $\Gamma(e)$ chosen to go through e using standard rounding. Thus, the random variable $X = \sum_{e\in S} X_e$.

Now consider the following (alternate) rounding scheme: Round each $\Gamma_{i,j}$ independently with probability $f_{i,j}$; thus multiple paths may be chosen for the same (s_i, t_i) pair. Let Y_e be a random variable indicating the number of paths in $\Gamma(e)$ that are chosen due to this rounding. It is easy to verify the following: (i) the random variable Y_e dominates the random variable X_e (i.e., for any $x \geq 0$, $\Pr[Y_e \geq x] \geq \Pr[X_e \geq x]$), and (ii) $E[Y_e] \leq f_e$. Hence it suffices to show that $\Pr[Y = \sum_{e\in S} Y_e > c \cdot \max\{\sum_{e\in S} f_e, \log m\}] < 1/m^2$. Since each $\Gamma_{i,j}$ is rounded independently, the bound follows from Proposition 2.3(1) if we choose c to be sufficiently large. □

Corollary 2.3 *If $u_{\min}/d_{\max} \geq c \log m$ for some suitably large constant c, then* USF *can be approximated to within a constant factor.*

Proof: Solve LP-UNSPLIT, scale the flow down by a factor of $1/c$, where c is a positive constant, and perform the standard rounding. If c is chosen sufficiently large, then it follows from Proposition 2.3(1) that w.h.p., no arc capacity is violated in the rounding procedure. On the other hand, the expected value of the solution obtained is $\Omega(LP_u)$. The result follows. □

3 EDP, USF and ISF

3.1 Hardness of Approximating EDP (USF and ISF)

Recall that in the *edge-disjoint paths problem* (EDP), we are given a graph G and a set $\mathcal{T}$ of k source-sink pairs $(s_1, t_1), \ldots, (s_k, t_k)$, and the goal is to find a subset $\mathcal{S} \subseteq \mathcal{T}$ of maximum cardinality such that all (s_i, t_i) pairs in $\mathcal{S}$ can be connected by edge-disjoint paths. Only an $O(\sqrt{m})$-approximation is known for EDP [15]. We now prove an essentially matching hardness result on directed graphs.

Theorem 1 *For any $\epsilon > 0$, it is NP-hard to distinguish given a directed instance $\left[G = (V, A), \mathcal{T} = \{(s_i, t_i) : i \in [k], s_i, t_i \in V\}\right]$ with $|A| = m$, whether all k pairs in $\mathcal{T}$ can be connected by edge-disjoint paths or at most a fraction $1/m^{1/2-\epsilon}$ of the k pairs can be connected.*

Remark: Note that the theorem says that the difficulty of the problem is not due to determining which subset of pairs to "route", but lies in determining the routing itself.

Proof of Theorem 1: The proof is by a reduction from the following well-known NP-hard problem [7]:
PROBLEM: 2DIRPATH:
INSTANCE: A *directed* graph $H = (V, A)$, *distinct* vertices $x_1, x_2, y_1, y_2 \in V$.
QUESTION: Are there two edge-disjoint directed paths, one from x_1 to y_1 and the other from x_2 to y_2, in H?

Given an $\epsilon > 0$, we construct a directed graph G (which will be the directed graph underlying our EDP instance) from H as follows. The skeleton of G will be the graph G' whose basic structure is as described below. G' will comprise of vertices s_i, t_i for $1 \leq i \leq N$, where $N = |A|^{\lceil 1/\epsilon \rceil}$, together with vertices $\{h_{ij}^{(1)}, h_{ij}^{(2)}, v_{ij}^{(1)}, v_{ij}^{(2)} : 1 \leq j < i \leq N\}$ and the "diagonal" vertices $\{d_{ii} : 1 \leq i \leq N\}$ connected in a grid-like fashion. Each s_i is connected by a directed path $\mathcal{P}_i$ to t_i where

$$\mathcal{P}_i = \left[s_i, v_{i1}^{(1)}, v_{i1}^{(2)}, v_{i2}^{(1)}, v_{i2}^{(2)}, \ldots, v_{i,i-1}^{(1)}, v_{i,i-1}^{(2)}, d_{ii}, h_{i+1,i}^{(1)}, h_{i+1,i}^{(2)}, \ldots, h_{n,i}^{(1)}, h_{n,i}^{(2)}, t_i\right]$$

The edge set of G' will be the (disjoint) union of the edges in the paths $\mathcal{P}_i$ for $1 \leq i \leq N$. (For geometric intuition, one can identify s_i, t_i with the points $(N-i+1, 0)$ and $(0, i)$ on the 2-dimensional grid, and visualize $\mathcal{P}_i$ as the path connecting $(N-i+1, 0)$ with $(0, i)$ on the 2-dimensional grid that first goes "up" the y-direction, makes a *single bend*, and then goes "left" along the x-direction.)

The graph G is constructed by making the following modification at each "ij intersection" for $1 \leq j < i \leq N$: (i) remove the edges $(h_{ij}^{(1)}, h_{ij}^{(2)})$ and $(v_{ij}^{(1)}, v_{ij}^{(2)})$, and (ii) place a copy of the graph H at the ij intersection while identifying the x_1, y_1, x_2, y_2 of the copy of H with the vertices $h_{ij}^{(1)}, h_{ij}^{(2)}, v_{ij}^{(1)}, v_{ij}^{(2)}$, respectively, of the ij intersection. The instance of EDP will now comprise of the graph G together with the N source-sinks pairs $\mathcal{T} = \{(s_i, t_i) : i \in [N]\}$.

Claim 1 *If there are edge-disjoint paths from x_1 to y_1 and x_2 to y_2 in H, then there are N edge-disjoint paths in G, one connecting s_i to t_i, for each $i \in [N]$.*

Proof: Suppose A_1 and A_2 are two edge-disjoint paths in H connecting x_1 to y_1 and x_2 to y_2 respectively. For $r \in [N]$, define the path $\mathcal{Q}_r$ in G to be the s_r-t_r path that is the same as the path $\mathcal{P}_r$ of the skeleton G' except that instead of using the edges $(h_{ij}^{(1)}, h_{ij}^{(2)})$ and $(v_{ij}^{(1)}, v_{ij}^{(2)})$ (for the relevant values of i, j for the path $\mathcal{P}_r$), it uses the edge-disjoint paths A_1 and A_2, respectively, of the local copy of H at the ij intersection. The paths $\mathcal{Q}_r$ thus defined are clearly edge-disjoint. □

Claim 2 *If there exist edge-disjoint paths $\mathcal{Q}_{i_1}$ and $\mathcal{Q}_{i_2}$ in G connecting s_{i_1} to t_{i_1} and s_{i_2} to t_{i_2}, respectively, for any $1 \leq i_1 \neq i_2 \leq N$, then there must be two edge-disjoint paths in H from x_1 to y_1 and x_2 to y_2.*

Proof: Identify G with its embedding in the plane. Clearly, one can extend $\mathcal{Q}_{i_1}$ to a closed contour where s_{i_2}, t_{i_2} are on the outside and inside respectively. It follows (cf. [22]) that $\mathcal{Q}_{i_2}$ must cross this contour, but this implies that there exists p, q such that $\mathcal{Q}_{i_1}, \mathcal{Q}_{i_2}$ each uses one of the edges associated with the pq intersection.

Now going back to the graph G, the two paths $\mathcal{Q}_{i_1}$ and $\mathcal{Q}_{i_2}$ will intersect at a point as guaranteed by the above fact, i.e $\mathcal{Q}_{i_1}$ enters at x_1 and leaves at y_1 in the copy of H at this intersection, while $\mathcal{Q}_{i_2}$ enters at x_2 and leaves at y_2, and the edge-disjointness of $\mathcal{Q}_{i_1}$ and $\mathcal{Q}_{i_2}$ implies that there must exist two edge-disjoint paths in H from x_1 to y_1 and x_2 to y_2. □

The above two claims imply that YES instances of 2DIRPATH are mapped to instances of EDP where all N pairs can be satisfied, while NO instances are mapped to instances of EDP where at most one pair is satisfied. This creates a gap of N, and since the number of arcs in G equals $m = O(N^2|A|) = O(N^{2+\epsilon})$ (recall that $N = |A|^{\lceil 1/\epsilon \rceil}$), the gap equals $\Omega(m^{1/(2+\epsilon)})$, and since $\epsilon > 0$ was arbitrary, we are done. □

Corollary 3.1 *The* USF *and* ISF *problems on directed graphs are NP-hard to approximate within $m^{1/2-\epsilon}$ for any $\epsilon > 0$.*

An easy approximation-preserving reduction from the independent set problem, together with Hastad's inapproximability result for independent set [10], gives

Fact 3.1 *Unless* NP = ZPP, ISF *on undirected graphs cannot be approximated to within $m^{1/2-\epsilon}$ of the optimum in polynomial time, for any $\epsilon > 0$.*

In the next section we give an (essentially tight) approximation algorithm for USF based on rounding an optimal solution to the LP (5). Standard rounding techniques are not as easily applied to ISF; we defer the positive results for the latter problem to Section 5.

3.2 An LP-Based Approximation for USF

We now give a simple algorithm to approximate USF to within $O(\sqrt{m \log m} \log\log m)$ when $u_{\min} \geq d_{\max}$, and an $O(\sqrt{m}(\log m)^{3/2})$-approximation without this assumption, but requiring that $d_{\max}$ is polynomially bounded.

Theorem 2 *If all demands are between 1 and 2, then* USF *can be approximated to within a factor of $O(\sqrt{m \log m})$.*

Proof: We assume w.l.o.g. that $\max_i r_i$ is 1; otherwise, we can scale the profits appropriately. We first solve LP-UNSPLIT. If the LP value LP_u is less than $\sqrt{m}$, then routing the demand with the highest profit yields a $\sqrt{m}$ approximation. Therefore, in the remainder of the proof, we assume that $LP_u \geq \sqrt{m}$. We scale down the LP flow by a factor of 1/6, and perform standard rounding. In the rounding procedure, each demand is selected independently with a probability equal to one-sixth of the fraction routed in the LP. Thus, the expected total profit of the rounded solution is $\Omega(LP_u)$. Since $LP_u \geq \sqrt{m}$ and the profit of each demand is in $[0, 1]$, it follows from Proposition 2.3(2) that w.h.p.the total profit of the rounded solution is $\Omega(LP_u)$. The rounded solution, however, may violate the capacity of many arcs. We now show that we can obtain a solution which w.h.p. has value $\Omega(LP_u/\sqrt{m \log m})$ and does not violate any arc capacities.

HEAVY AND LIGHT EDGES. Let c_0 denote the quantity $12 \ln m$. Call an arc e *heavy* if $u(e) > c_0$ and *light* otherwise. Our first claim is that with probability at least $(1 - 1/m)$, the rounded solution does not violate the capacity of an heavy arc. To see this, define X_e to be a random variable that indicates the *total number of demands* routed through an arc e. By Proposition 2.3(1),

$$\Pr[X_e > \frac{u(e)}{2}] \leq \Pr[X_e > (1+2)2\ln m] \leq \frac{1}{m^2}.$$

Thus with probability at least $1 - 1/m$, at most $u(e)/2$ paths are chosen to go through an arc of capacity at least $12 \ln m$. Since all demands are between 1 and 2, it follows that, with probability at least $1 - 1/m$, no heavy arc is violated.

PROCESSING THE ROUNDED SOLUTION. Now let $\mathcal{R}$ denote the set of paths in a solution obtained by standard rounding. A path $P \in \mathcal{R}$ is said to be *α-light* if the total capacity associated with the light arcs that lie on this path is exactly α. Partition $\mathcal{R}$ into two sets $\mathcal{R}_1$ and $\mathcal{R}_2$ such that $\mathcal{R}_1$ is the set of all α-light paths with $\alpha \geq \sqrt{mc_0}$, and $\mathcal{R}_2 = \mathcal{R} \setminus \mathcal{R}_1$. It is easy to see that $\mathcal{R}$ contains at most mc_0/α paths which are α-light. Thus, $|\mathcal{R}_1|$ is at most $\sqrt{mc_0}$.

Initialize S_1 to be the single-element set consisting of a path of largest profit in $\mathcal{R}_1$. For a given set S of paths, let $\text{PROFIT}(S)$ denote the sum of the profits of the source-sink pairs routed in S. Clearly, $\text{PROFIT}(S_1) \geq \text{PROFIT}(\mathcal{R}_1)/\sqrt{mc_0}$. Construct also another set S_2 of paths as follows: Repeatedly pick a path P of largest profit from $\mathcal{R}_2$, add P to S_2, and delete P along with all paths in $\mathcal{R}_2$ that share a light arc with P. Since the total capacity of light arcs is at most $\sqrt{mc_0}$, it follows from Proposition 2.4 that for any path P in $\mathcal{R}_2$, the total number of paths in $\mathcal{R}$ that share a light arc with P is $O(\sqrt{mc_0})$ with probability at least $1 - 1/m^2$. Therefore, with probability at least $1 - 1/m$, for every path added to S_2, only $O(\sqrt{mc_0})$ other paths in $\mathcal{R}_2$ are thrown away. Thus, upon termination (i.e. when $\mathcal{R}_2 = \emptyset$), $\text{PROFIT}(S_2)$ is $\Omega(\text{PROFIT}(\mathcal{R}_2)/\sqrt{mc_0})$ w.h.p. Output the better among the solutions S_1 and S_2.

APPROXIMATION RATIO AND CORRECTNESS. It is clear from the preceding description that we are guaranteed a profit of $\Omega(\text{PROFIT}(\mathcal{R})/\sqrt{mc_0})$; this yields the claimed approximation guarantee as $\text{PROFIT}(\mathcal{R})$ is $\Omega(LP_u)$ w.h.p. To see that with w.h.p., the set of paths in S_2 does not violate the capacity of any arc, observe that: (a) no two paths in S_2 share a light arc, and (b) with probability $1 - 1/m$, the set of paths in $\mathcal{R}$ does not violate the capacity of a heavy arc. □

Theorem 3 USF *with polynomially bounded demands can be approximated to within a factor of $O(\sqrt{m}\log^{3/2} m)$.*

Proof: Let $d_{\max} = O(m^{O(1)})$ denote the largest demand. Partition the demands into $O(\log m)$ classes, say $D_0, D_1, \ldots$ where the class D_i contains all demands in the interval $[2^i, 2^{i+1})$. Find a solution for each demand class using the approach of Theorem 2. Output the solution with the largest profit. By pigeonhole principle, one of these classes contains $\Omega(1/\log m)$-fraction of the optimum solution's total profit. The theorem follows. $\square$

Theorem 4 USF *with arbitrary demands can be approximated to within a factor of* $O(\sqrt{m \log m} \log \log m)$ *if minimum arc capacity is at least as large as the largest demand.*

Proof: Without loss of generality, we may assume that $d_{min} = 1$. As before, partition the demands into classes $D_0, D_1, ..., D_q$ where the class D_i contains all demands in the interval $[2^i, 2^{i+1})$. Set $p = q - (\log \log m + \log c)$ where c is the constant of Corollary 2.3, and define $\mathcal{D}_1 = \bigcup_{i=1}^{p} D_i$ and $\mathcal{D}_2 = \bigcup_{i=p}^{q} D_i$. Also, for some optimal solution, let OPT_1 and OPT_2 denote the profit generated by commodities in the demand sets $\mathcal{D}_1$ and $\mathcal{D}_2$ respectively.

Since the maximum demand in $\mathcal{D}_1$ is at most

$$(d_{max}/c \log m) \leq (u_{min}/c \log m),$$

by Corollary 2.3, it follows that OPT_1 can be approximated to within a constant factor. On the other hand, we can use Theorem 2 to find an approximate solution for each demand class D_i with $p \leq i \leq q$. By pigeonhole principle, one of these classes contains $\Omega(1/\log \log m)$-fraction of OPT_2. Choosing the better of the two solutions, we get the desired approximation ratio. $\square$

4 Bounded Length Edge-Disjoint Paths

We are given a (possibly directed) graph $G = (V, E)$, in which each edge e has a given nonnegative length. We are also given an integer L, which is referred to as the *length bound*. In BLEDP, we are given a set $\mathcal{T}$ of k source-sink pairs $(s_1, t_1), \ldots, (s_k, t_k)$, and the goal is to find a subset $\mathcal{S} \subseteq \mathcal{T}$ of maximum cardinality such that all (s_i, t_i) pairs in $\mathcal{S}$ can be connected by edge-disjoint paths, each path of "length" at most L. In the (s,t)-BLEDP problem, we are given a *single* pair (s, t), the goal is to find a maximum number of edge-disjoint s–t paths each of "length" at most L.

We consider (α, β)-approximations for the BLEDP problems. An (α, β)-approximation algorithm is one that obtains for each instance $\mathcal{I}$ with length bound L, $\alpha \cdot \text{OPT}(\mathcal{I})$ edge-disjoint paths each of length at most βL, where $\text{OPT}(\mathcal{I})$ is the size of the optimal solution. For convenience, we refer to an $(\alpha, 1)$-approximation as an α-approximation.

4.1 (α, β)-Approximating (s,t)-BLEDP

4.1.1 Inapproximability results

It is easy to see, by a simple reduction from 2DIRPATH, that the directed version of (s,t)-BLEDP does not have a polynomial time $(\frac{1}{2} + \epsilon, \frac{4}{3} - \epsilon)$-approximation algorithm unless P = NP. This reduction, however, is inexorably tied to the directed case; indeed the undirected version of 2DIRPATH is solvable in polynomial time by the work of Robertson and Seymour [24]. One could possibly attack the undirected case, however, by considering the related NP-hard problem (see [6]) INTEGER2COMMODITY where in we are given an *undirected* graph $G = (V, E)$ and *distinct* vertices $x_1, x_2, y_1, y_2 \in V$, and the objective is to find a maximum collection of edge-disjoint paths, each joining x_i to y_i for $i = 1$ or 2. We note that the half-integral version of INTEGER2COMMODITY flow was shown to be polynomially solvable in a classical paper of Hu [11]. We now give a reduction which shows MAX SNP-hardness of both undirected (s,t)-BLEDP and INTEGER2COMMODITY; a similar result for undirected (s,t)-BLEDP has recently been obtained independently in [2].

Theorem 5 *For both undirected as well as directed graphs,* (s,t)-BLEDP *with a length bound of six is MAX SNP-hard. Hence, there exist constants* α *and* β, $\alpha < 1$ *and* $\beta > 1$, *such that there is no polynomial-time* (α, β)*-approximation for* (s,t)-BLEDP *unless P=NP.*

Proof We prove the desired claim for undirected graphs only. The proof can be easily amended to apply to directed graphs.

The reduction is from Bounded Occurrence 3-DIMENSIONAL MATCHING (3DM) and follows the reduction presented in [18] which in turn is motivated by ideas from [8]. In an instance $\mathcal{I}$ of 3DM, we are given three disjoint sets $A = \{a_1, a_2, \ldots, a_n\}$, $B = \{b_1, b_2, \ldots, b_n\}$ and $C = \{c_1, c_2, \ldots, c_n\}$, and a set $\mathcal{T}$ of m triples $T_\mu \in A \times B \times C$, $\mu \in [m]$. It is shown in [12] that there exists an $\epsilon_0 > 0$ such it is NP-hard to distinguish between instances $\mathcal{I}$ where there exist n *disjoint* triples in $\mathcal{T}$ (call them "satisfiable" instances) and those where there are at most $(1 - \epsilon_0)n$ disjoint triples in $\mathcal{T}$ (call these "ϵ_0-unsatisfied" instances), even if we assume that each element of A, B, C is in the same constant number M of triples in $\mathcal{T}$.[1] We denote the μ^{th} triple T_μ as $(a_{p_\mu}, b_{q_\mu}, c_{r_\mu})$ for some $1 \leq p_\mu, q_\mu, r_\mu \leq n$.

Create an undirected graph $H = (V, E)$ as follows:

$$\begin{aligned} V &= \{b_i, c_i : i \in [n]\} \bigcup \{x_\mu, y_\mu : \mu \in [m]\} \\ &\quad \bigcup \{a_{il} : i \in [n], l \in [M-1]\} \bigcup \{s, t\} \\ E &= \{(s, b_i), (c_i, t), (a_{il}, t) : i \in [n], l \in [M-1]\} \\ &\quad \bigcup \{(s, x_\mu), (y_\mu, a_{p_\mu, l}) : \mu \in [m], l \in [M-1]\} \\ &\quad \bigcup \{(b_{q_\mu}, x_\mu), (x_\mu, y_\mu), (y_\mu, c_{r_\mu}) : \mu \in [m]\} \end{aligned}$$

We define positive integer lengths on every edge of H and consider the problem of finding the maximum number of edge-disjoint s-t paths of total length at most 6 in H. The edges (s, x_μ) for $\mu \in [m]$ have a "length" of 3 each, and the edges (y_μ, c_{r_μ}) have a length of 2, while all other edges have a length equal to 1.

We now claim that (a) [completeness] if the instance $\mathcal{I}$ of 3DM is *satisfiable*, then we can find Mn edge-disjoint paths from s to t of length at most 6, and (b) [soundness] if $\mathcal{I}$ is ϵ_0-unsatisfied, then there are at most $(M - \epsilon_0/2)n$ such s-t paths in H.

Let us first verify the completeness. Suppose T_{μ_1}, T_{μ_2}, $\ldots, T_{\mu_n}$ are n *disjoint* triples. Denote $f_i = q_{\mu_i}$ and $g_i = r_{\mu_i}$ for $i \in [n]$. Define the paths $P_i = [s, b_{f_i}, x_{\mu_i}, y_{\mu_i}, c_{g_i}, t]$ for $i \in [n]$, and the paths $Q_\mu = [s, x_\mu, y_\mu, a_{p_\mu, l_\mu}, t]$ for $\mu \in [m] - \{\mu_1, \mu_2, \ldots, \mu_n\}$ and an $l_\mu \in [M-1]$ defined such that no two μ's whose corresponding triples "share" an element of A, have the same value of l_μ. (This is possible to achieve since the fact that each element of A is present in exactly one triple among $T_{\mu_1}, \ldots, T_{\mu_n}$ implies that, for each element $a_i \in A$, there are $M - 1$ such μ's in the set $[m] - \{\mu_1, \ldots, \mu_n\}$.)

[1] Actually we only need the fact that each element in A occurs in some constant number of triples.

It is now easy to check that the paths P_i and Q_μ are all edge-disjoint and that each has length exactly 6.

To prove soundness, suppose there is a collection $\mathcal{C}$ of $(Mn-\delta n)$ edge-disjoint s-t paths of length at most 6. Clearly at most $(M-1)n$ of these paths use a final edge of the form $a_{p_\mu,l}t$ and hence at least $(1-\delta)n$ of the paths in $\mathcal{C}$ use a final edge of type c_it. Any such path of length at most 6 is easily seen to be of the form P_i above. The b_i's and c_j's used in these paths are all distinct and yield a set $\mathcal{S}$ of $(1-\delta)n$ triples which have distinct B and C "coordinates". If some a_i is present in say $t>1$ triples in $\mathcal{S}$, we can only take 1 of these in a 3-matching. But then we also "lost" at least $t-1$ potential paths of the form Q_μ for our collection $\mathcal{C}$. Thus a_i was used to get at most $(M-t)$ paths of the form Q_μ (as opposed to the $M-1$ paths in the completeness case). We can clearly lose at most δn such paths in all, and hence by retaining in $\mathcal{S}$ at most one triple containing any a_i, for each $i \in [n]$, we will be left with a set of at least $(1-2\delta)n$ **disjoint** triples.

We have thus proved that the undirected (s,t)-BLEDP problem (even with a length bound of 6) is MAX SNP-hard, and in fact for any $\epsilon > 0$, there is no polynomial-time $(1-\frac{\epsilon_0}{2M}, \frac{7}{6}-\epsilon)$-approximation for (s,t)-BLEDP unless P=NP. □
One can easily amend the previous reduction to obtain the following claim.

Corollary 4.1 INTEGER2COMMODITY *is MAX SNP-hard.*

4.1.2 An $(1-\epsilon, 1/\epsilon)$-approximation algorithm for (s,t)-BLEDP

We are given an undirected graph G, length bound L, a source s and a sink t in G. Let p be the optimal number of edge-disjoint paths from s to t, each path of length at most L. In this section, we present a simple algorithm that obtains, for any positive real ϵ, at least $(1-\epsilon)k$ paths, each of length at most L/ϵ.

We first define a minimum-cost flow problem in G from s to t, where the cost of an edge is defined to be its length and the value of the desired flow is p. Since there exists a flow of value p and cost at most pL, the minimum-cost flow algorithm returns a set S of edge-disjoint paths from s to t such that the sum of all of the path lengths is at most pL. By averaging, it follows that at least $(1-\epsilon)p$ s-t paths in S have length at most L/ϵ, thus yielding the desired approximation.

4.2 Hardness of α-approximating BLEDP

4.2.1 Hardness of (s,t)-BLEDP

Theorem 6 *For directed graphs, (s,t)-BLEDP is NP-hard to approximate within a factor of $m^{1/2-\epsilon}$ for any $\epsilon > 0$.*

Proof: The reduction is from 2DIRPATH. Given an instance $\left[H; x_1, x_2, y_1, y_2\right]$ of 2DIRPATH, we construct an instance $\left[G=(V,A'); s, t; L\right]$ where G is the same directed graph as the one constructed in the proof of Theorem 1 together with two new vertices s,t and arcs joining s to $s_1, s_2, \ldots, s_N$ and joining $t_1, t_2, \ldots, t_N$ to t, and L is a suitable length bound to be mentioned later in the proof.

We now define non-negative integer lengths on the arcs and consider the problem of finding edge-disjoint paths of total length at most L between s and t. In the directed graph G, the arcs at the bend (the ones incident to d_{ii} for $i \in [N]$) and the arcs which lie wholly within the various copies of H, get a length of zero. The arc (s, s_i) gets a length of i, and the arc (t_i, t) gets a length of $N-i$, for $i \in [N]$, and all other arcs have a length of 1.

We now claim that any s,t path P with total length at most $2N-1$ must use a "bend", i.e must go through d_{rr} for some $r \in [N]$. Indeed, let the path P go through s_i and t_j. Then, if it does not use a bend, then all arcs it uses in the underlying grid of (the skeleton of) G have a length 1, and at least j such "vertical" arcs and $N-i+1$ such "horizontal" arcs are required for any directed path through s_i and t_j. This, together with the lengths i and $N-j$ on (s,s_i) and (t_j,t) respectively, implies that P has total length at least $i+(N-j)+j+(N-i+1) = 2N+1$. Hence, using the same *geometric argument* as in the proof of Theorem 1, we can prove that in the case when there are no edge-disjoint paths between x_1, y_1 and x_2, y_2 in H, the maximum number of edge-disjoint (s,t)-paths in G of length at most $2N-1$ is one.

Also, when we start from a YES instance of 2DIRPATH, there are N edge-disjoint (s,t)-directed paths of length $2N-1$ in G, namely the N paths $\mathcal{P}_i$ where $\mathcal{P}_i$ is the (s,t) path that goes through s_i and t_i via the "bend" vertex d_{ii} (i.e it follows i vertical arcs from s_i up to d_{ii} and then $N-i+1$ horizontal arcs to t_i, for an overall length of $i+(i-1)+(N-i)+(N-i) = 2N-1$).

Thus the gap between the optimum values of the instances created from the YES and NO instances of 2DIRPATH is N, which as before will prove an inapproximability bound of $m^{1/2-\epsilon}$ by choosing N suitably. □

4.2.2 Hardness of BLEDP

We now employ an approximation preserving reduction from the independent set problem to establish that BLEDP, even on undirected graphs, is hard to approximate within $m^{1/2-\epsilon}$.

Theorem 7 *Unless* NP = ZPP, BLEDP *on undirected graphs cannot be approximated in polynomial time within a factor of $m^{1/2-\epsilon}$ for any $\epsilon > 0$.* □

Proof: The reduction is from the independent set problem. Assume we are given a graph $G=(V,E)$ with $V=\{1,2,\ldots,n\}$. It is known that unless NP = ZPP, it is not possible to distinguish in polynomial time between the cases when $\alpha(G) \geq n^{1-\epsilon_0}$ and when $\alpha(G) \leq n^{\epsilon_0}$ for any fixed $\epsilon_0 > 0$ [10].

Starting from G, we construct a "grid graph" instance of s_i-t_i-BLEDP $\left[H; \{(s_i,t_i) : i \in [n]\}\right]$ as follows. The n sources s_i correspond to the vertices $(i,0)$ respectively, while the n sinks t_i correspond to $(0, n-i+1)$ respectively. Each s_i is connected to t_i through a "canonical" path $\mathcal{Q}_i = \left[s_i, a[i,1], a[i,2], \ldots, a[i,n-i], b[i-1,n-i+1], b[i-2,n-i+1], \ldots, b[1,n-i+1], t_i\right]$ – here each $a[i,j]$ and $b[i,j]$ is to be thought of as *an edge* (with a designated entry and exit point) placed at vertex (i,j) of the plane grid (for $1 \leq i,j \leq n$ with $i+j \leq n$). The adjacency information of G is encoded in H by *identifying* the edges $a[i,n-j+1]$ and $b[i,n-j+1]$ whenever vertices i and j (with $i<j$) are adjacent in G – this serves the purpose that whenever i and j are adjacent in G, the canonical paths $\mathcal{Q}_i$ and $\mathcal{Q}_j$ will share an edge (viz. the "identified" copy of $a[i,n-j+1]$ and $b[i,n-j+1]$), and thus cannot both be simultaneously routed.

We assign a length of 2 to the following edges: (i) $(s_i, a[i,1])$ for $1 \leq i < n$, and $(b[1,n-i+1], t_i)$ for $1 < i \leq n$; (ii) $a[i,j]$ and $b[i,j]$ for i,j such that $i+j \leq n$; and (iii) edges joining the "exit" of $a[i,j]$ (respectively $b[i,j]$) to the "entry" of

```
GREEDY MULTICOMMODITY ROUTING (GREED)
{
/* Input is a directed network where some arcs have length 0
and others have length 1 */
    S = ∅
    While (There is an unsatisfied demand)
        For each demand s in D − S
            Find a shortest flow f for s in the residual network
        EndFor
        If no routings are found, then Output Current Solution
        Let s be such a demand with an overall minimum routing f
        Decrement the residual network by the flow f
        Add s to S
    EndWhile
    Output Current Solution
}
```

Figure 1: The "core" greedy algorithm

$a[i, j+1]$ (respectively $b[i-1, j]$) for i, j such that $i > 1$ and $i + j < n$. We assign a length of 1 to all other edges.

We now claim that the only s_i-t_i path in H of total length at most $4n-3$ is the canonical path $\mathcal{Q}_i$. Indeed let P_i be any shortest s_i-t_i path. Then, arguing informally, it can be shown that if P_i "leaves" a column (or a row) at a vertex v_1 and comes back to meet it again at v_2, then the path which for the v_1-v_2 leg of P_i uses the v_1-v_2 segment of that column (or row as the case may be), will have strictly smaller length than P_i, a contradiction. Thus the shortest s_i-t_i path lies wholly within the portion of the grid $\{(p, q) : p \leq i, q \leq n - i + 1\}$, and now it is easy to check that P_i is the unique s_i-t_i shortest path and that P_i has a total length of $4n - 3$.

Thus, starting from G, we can, in polynomial time, construct a graph H with n source-sink pairs (s_i, t_i) such that the maximum number of pairs which can be joined by edge-disjoint paths of length at most $4n-3$, equals $\alpha(G)$. The gap of $n^{1-2\epsilon_0}$ in $\alpha(G)$ thus clearly translates into a gap of $m^{1/2-\epsilon_0}$ for the s_i-t_i-BLEDP instance we create, and the proof is completed. □

5 Greedy Algorithms for USF, ISF and BLEDP

We now show that the hardness bounds for all problems considered can be matched (ignoring certain polylogarithmic factors) by suitable adaptations of a greedy algorithm. We give the "core" algorithm in Figure 1, although subroutines must be tailored for the different problems considered. In the code description, a network refers to a directed graph together with integer capacities on the arcs. The operation of *decrementation* of a network by a flow f, results in a new network for which the capacity on an arc e is precisely f_e units less than before.

The flows (or routings) f found by GREED depend (e.g., split or unsplit) on the problem version being addressed. In any case, "shortest" will always refer to minimizing the measure $\sum_e f_e l_e$ where l_e is either 0 or 1. In the cases of USF and ISF, the problem of finding such an f reduces to solving a shortest path or mincost flow problem respectively. For BLEDP with length bound L, the "shortest flow" f is a source-sink path with the fewest number of edges among all of the source-sink paths of length L. In the following, we briefly describe a simple subroutine that computes such a path. It is sufficient to provide an algorithm that given a directed graph $D = (V, A)$, a source s, a sink t, and a positive integer k, determines a path that has length at most L and has less than k edges, if such a path exists. We now define such an algorithm.

We construct a layered directed graph $D' = (V', A')$ in which the set V' consists of k copies of each vertex in V. We denote the ith copy, $1 \leq i \leq k$ of a vertex v in V by v_i. For every arc (u, v) in E, we have the following $k-1$ arcs in A', each of which has the same length as that of (u, v): (u_i, v_{i+1}) for $1 \leq i < k$. In addition, for each v in V, we have the following $k-1$ arcs in A', each of length zero: (u_i, u_{i+1}) for $1 \leq i < k$. It is now easy to see that there is a path in D from s to t that has length at most L and has at most k arcs if and only if there is a path in D' from s_1 to t_k that has length at most L. Moreover, given a path P' in D' from s_1 to t_k that has length at most L, we can construct in polynomial time a path in D from s to t that has length at most L and has at most k arcs. To complete the description of the algorithm, we note that we can determine the existence of path P' by finding the shortest path between s_1 and t_k.

In [15], [17], and [19], the greedy algorithm is analyzed to show that it satisfies at least $O_u/\sqrt{m}$ demands for an instance of EDP (also see [4]). (Recall that O_u denotes the optimum value for a given USF instance.) Their arguments may be extended easily to give $\sqrt{u(A)}$-approximations for unsplittable problems with general demands and capacitated networks (where $u(A) = \sum_e u(e)$). We improve this bound for capacitated networks by invoking the greedy algorithm several times.

Theorem 8 *Consider a unit-profit instance of* USF *in a directed graph D. If $d_{max} \leq d_{min}\Delta$, then the greedy algorithm may be used to find a solution which satisfies at least $\frac{O_u}{2\Delta\sqrt{m}}$ demands.*

Proof: First we render D as a directed graph D' having u_e parallel arcs, numbered from 1 to u_e, for each *topological arc* $e \in A(D)$. Any optimal solution corresponds to a collection of O_u arc-disjoint paths in D'. Moreover, we may assume that in any solution, each flow system is *contiguous*, that is, for any demand and any topological arc e, the flow uses contiguously numbered arcs parallel to e.

Next, consider a guess O for O_u. Let $u_e^O = \min\{u_e, \frac{d_{min} O}{\sqrt{m}}\}$. Call an arc *clipped* if $u_e^O \neq u_e$. An arc is assigned a length of 0 if it is clipped and 1 otherwise. Run the greedy algorithm on the graph D^O obtained by eliminating from D', any arc numbered greater than u_e^O. Let $P_1, P_2, \ldots, P_t$ be the (arc-disjoint) flow systems chosen by the greedy algorithm in D^O.

Let $\mathcal{Q}$ be those demands which are satisfied in some optimal solution but not by the greedy algorithm. Let Q be a path (contiguously allocated) in the optimal solution satisfying some demand in $\mathcal{Q}$. It follows that there was an original arc e in D such that (i) Q is using one of the u_e arcs in D' parallel to e, and (ii) the greedy algorithm is using at least the first $u_e^O - d_{max}$ copies of these arcs, otherwise the greedy algorithm would have found a path which is topologically equivalent to Q. Now if e was clipped, then the greedy algorithm must have routed $(\frac{O d_{min}}{\sqrt{m}} - d_{max})/d_{max} \geq \frac{O}{2\Delta\sqrt{m}}$ demands and we are done. So we may assume this is not the case and hence each such Q must contain an arc of length 1 from D^O that is in common with some P_i. We choose i to be minimal. Thus for each demand in $\mathcal{Q}$ we associate a unique such P_i which *blocks* that demand. Conversely, let k_i be the number of demands in $\mathcal{Q}$ which are blocked by P_i.

We now build on the ideas used in [17, 19]. Since we may assume that both the optimal and greedy solutions are contiguous, we deduce that P_i may block at most $n_i\Delta$ flow systems, where n_i is the number of arcs of length 1 with positive flow under P_i. Thus $k_i \leq n_i\Delta$. But also by greedy we have that each such blocked flow must have used at least $n_i d_{min} \geq k_i \frac{d_{min}}{\Delta}$ units of length 1 capacity. Hence the total length 1 capacity used by the unrouted paths from the optimal solution is at least $\frac{d_{min}}{\Delta}\sum_i k_i^2$. Combining this with the observation that all length 1 arcs are unclipped, the total available length 1 capacity is at most $\sum_{e\in A(D)} u_e^O$, and we obtain

$$\frac{d_{min}}{\Delta}\cdot\frac{(\sum_{i=1}^t k_i)^2}{t} \leq \frac{d_{min}}{\Delta}\sum_{i=1}^t k_i^2 \leq \sum_{e\in A(D)} u_e^O \leq \sqrt{m} O d_{min}. \tag{6}$$

Since $\sum_i k_i = |\mathcal{Q}| = O_u - t$, we obtain $\frac{(O_u - t)^2}{t} \leq \Delta\sqrt{m}O$. Running the above procedure for various possible values of O (guessed within a factor of 2 by doubling), we may deduce that $t \geq \frac{O_u}{2\Delta\sqrt{m}}$ as required. □

If $\Delta > 2$, we can divide the demands into $O(\log\Delta)$ classes and do the above analysis for the class that has the largest number of pairs satisfied by an optimal solution. This would give us a solution that satisfies at least $O(\frac{O_u}{\log\Delta\sqrt{m}})$ demands. The preceding argument applies equally well (with $\Delta = 1$ of course) to BLEDP and so we obtain the following result.

Theorem 9 *There is a polynomial-time $O(\sqrt{m})$ approximation algorithm for* BLEDP. □

At first glance, it may appear that the proof of Theorem 8 also gives a similar bound for the integral splittable problem. There is a snag, however; we may only deduce the following.

Theorem 10 *Consider a unit-profit instance of the* ISF *problem. Then by applying the greedy algorithm $O(\log n)$ times, we may find a solution which satisfies at least $\frac{O_{is}}{2\sqrt{m d_{max}}}$ demands.*

Proof: Mimic the proof of Theorem 8 directly, except clip each capacity at $\frac{O\sqrt{d_{max}}}{\sqrt{m}}$. The proof breaks down only because we cannot guarantee that the flow P_i blocks at most $n_i\Delta$ flow systems from an integral splittable optimum. Indeed, we can only deduce that P_i must use at least k_i units of capacity, and hence so do each of the flows which it blocks. The equation (6) becomes $\sum_i k_i^2 \leq \sqrt{m d_{max}}O$, and the rest then follows. □

Returning to the general versions with profits and arbitrary demands, the repeated greedy algorithm, together with techniques from the proof of Theorem 4, gives us an upper bound on the approximation factor that nearly matches the hardness result of Corollary 3.1.

Corollary 5.1 *The general* USF *problem with arbitrary polynomially bounded demands can be approximated to within a factor of $O(\sqrt{m}\log^2 m)$. The general* ISF *problem with arbitrary polynomially bounded demands can be approximated to within a factor of $O(\sqrt{m d_{max}}\log^2 m)$.* □

References

[1] A. Baveja and A. Srinivasan. Approximation algorithms for disjoint paths and related routing and packing problems. *Submitted*, January 1998.

[2] A. Bley. On the complexity of vertex-disjoint length-restricted path problems. *Konrad-Zuse-Zentrum für Informationstechnik Berlin Tech. Report SC-98-20*, 1998.

[3] H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Mathematical Statistics*, 23:493–507, 1952.

[4] E. A. Dinitz. Algorithm for solution of a problem of maximum flow in networks with power estimation. *Soviet Math. Dokl.*, Vol. 11 (1970), pp. 1277-1280.

[5] Y. Dinitz, N. Garg and M. Goemans. On the single source unsplittable flow problem. *Proc. of FOCS '98*, pp. 290-299.

[6] S. Even, A. Itai and A. Shamir. On the complexity of timetable and multicommodity flow problems. *SIAM Journal on Computing*, Vol. 5, No. 4 (1976), pp. 691-703.

[7] S. Fortune, J. Hopcroft and J. Wyllie. The directed subgraph homeomorphism problem. *Theoretical Computer Science*, Vol. 10, No. 2 (1980), pp. 111-121.

[8] N. Garg, V. Vazirani and M. Yannakakis. Primal-dual approximation algorithms for integral flow and multicut in trees. *Algorithmica*, 18 (1997), pp. 3-20.

[9] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:13–30, 1963.

[10] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *ECCC Technical Report TR97-038.* (Preliminary version in *Proc. of FOCS '96*, pp. 627-636).

[11] T. C. Hu. Multi-commodity Network Flows. *Operations Research*, 11(1963), pp. 344-360.

[12] V. Kann. Maximum bounded 3-dimensional matching is MAX SNP-complete. *Information Processing Letters*, 37(1991), pp. 27-35.

[13] R. M. KARP. Reducibility among combinatorial problems. *Complexity of Computer Computations*, R.E. Miller, J.W. Thatcher, Eds., New York: Plenum Press, 1972, pp. 85-103.

[14] J. M. KLEINBERG. Single-source unsplittable flow. *Proc. of FOCS '96*, pp. 68-77.

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

[16] J. M. KLEINBERG. Decision algorithms for unsplittable flow and the half-disjoint paths problem. *Proc. of STOC '98*, pp. 530-539.

[17] S. G. KOLLIOPOULOS. Exact and approximation algorithms for network flow and disjoint-path problems. *PhD Thesis*, Dartmouth College, Hanover, NH, August 1998.

[18] S. G. KOLLIOPOULOS AND C. STEIN. Improved approximation algorithms for unsplittable flow problems. *Proc. of FOCS '97*, pp. 426-435.

[19] S. G. KOLLIOPOULOS AND C. STEIN. Approximating disjoint-path problems using greedy algorithms and Packing Integer Programs. *IPCO 98.*

[20] F. T. LEIGHTON AND S. B. RAO. An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms. *Proc. of FOCS '88*, pp. 422-431.

[21] B. MA AND L. WANG. On the inapproximability of disjoint paths and minimum steiner forest with bandwidth constraints. *Journal of Computer and Systems Sciences*, to appear.

[22] W. S. MASSEY. *Algebraic Topology: An Introduction.* Graduate texts in mathematics 56, Springer-Verlag, 1967.

[23] P. RAGHAVAN AND C. D. THOMPSON. Randomized rounding: A technique for provably good algorithms and algorithmic proofs. *Combinatorica*, Vol. 7 (1987), pp. 365-374.

[24] N. ROBERTSON AND P. D. SEYMOUR. Outline of a disjoint paths algorithm. In B. Korte, L. Lovász, H. J. Prömel, and A. Schrijver, Eds., *Paths, Flows and VLSI-Layout.* Springer-Verlag, Berlin, 1990.

[25] A. SRINIVASAN. Improved approximations for edge-disjoint paths, unsplittable flow, and related routing problems. *Proc. of FOCS'97*, pp. 416-425.

A Near-tight inapproximability result for undirected node-capacitated USF

While our inapproximability bound of $m^{1/2-\epsilon}$ for USF applies only for directed graphs, we now prove that if we instead look at the node-capacitated version of USF, considered for instance by Kleinberg [16] under a *boundedness* assumption, then we can get an inapproximability bound of $n^{1/2-\epsilon}$, for any $\epsilon > 0$, where n is the number of nodes, even on **undirected** graphs. We only sketch the idea of the proof and defer the complete details to the full version of the paper.

We first define the undirected node-capacitated USF problem, denoted Undir-Node-USF, below: We are given an undirected graph $G = (V, E)$ with positive integral capacities on the **nodes**, and k source-sink pairs (s_i, t_i) with a positive integer demand d_i and a "profit" r_i for $i \in [k]$ — the objective is to find a subset of the pairs that can be routed feasibly (i.e all node capacities are obeyed) and that maximizes the total profit (as before we only get the profit for a pair (s_i, t_i) when its demand d_i is fully routed on a single path). When all profits are equal, we refer to the version as being "unweighted". Note that unweighted Undir-Node-USF is simply the classical vertex-disjoint paths problem when all demands and capacities are equal to 1.

Theorem 11 *It is NP-hard to approximate unweighted* Undir-Node-USF *within a factor of $n^{1/2-\epsilon}$ for any $\epsilon > 0$.*

We remark that the above theorem is nearly tight since LP-based or greedy algorithms like the ones we presented for the edge-capacity version can achieve an approximation ratio of $O(\sqrt{n}\ \mathrm{poly}(\log n))$ for Undir-Node-USF.

The proof of Theorem 11 uses the following theorem, whose proof we defer to the final version of the paper. We just remark that the proof is based on a reduction from SAT along the lines of the NP-hardness proof for two-commodity integral flow presented in [6].

Theorem 12 *Given an instance of* Undir-Node-USF *with two source-sink pairs, it is NP-hard to decide if both pairs can be feasibly routed. Moreover, the hardness holds even if all node capacities equal 1 or 2, and the two demands are 1 and 2.*

Now, to prove Theorem 11, we use a construction similar to the one used in Theorem 1 – the main difference is that at each "ij intersection" we now place a hard instance $\mathcal{I}$ of 2PAIR-Undir-Node-USF above, instead of the 2DIRPATH instance we used in Theorem 1. We set the demand of the pair (s_i, t_i) to $n - i$, for $i \in [N]$. For the nodes in the copy of $\mathcal{I}$ placed at the "ij intersection" (for $1 \leq j < i \leq N$) we change the capacities of nodes with capacity 1 (recall that all nodes in instance $\mathcal{I}$ have capacity 1 or 2) to $n - i$, and those of nodes with capacity 2 to $n - j$.

Omitting the details, we now claim that if the instance $\mathcal{I}$ was feasible, then the demands of all the N (s_i, t_i) pairs can be fully met by a feasible routing, while if $\mathcal{I}$ was not feasible, than at most one pair's demand can be met by any feasible routing in the Undir-Node-USF instance constructed above. As before, this gives us a gap of $n^{1/2-\epsilon}$ where n is the number of nodes in the Undir-Node-USF instance created. □

Whether Theorem 11 extends to edge-capacitated USF in undirected graphs, and indeed to EDP in undirected graphs, remains an interesting unresolved question.