

Max-Weight Integral Multicommodity Flow in Spiders and High-Capacity Trees

Jochen Könemann¹, Ojas Parekh², and David Pritchard¹

¹ University of Waterloo

jochen@math.uwaterloo.ca, dagpritchard@math.uwaterloo.ca
Dept. of Combinatorics & Optimization, 200 University Avenue West,
Waterloo, Ontario, Canada, N2L 3G1. Fax: 519-725-5441

² Emory University

ojas@mathcs.emory.edu
Math/CS Department, 400 Dowman Dr., Suite 401,
Atlanta, GA, USA, 30322. Fax: 404-727-5611

Abstract. We consider the max-weight integral multicommodity flow problem in trees. In this problem we are given an edge-capacitated tree and weighted pairs of terminals, and the objective is to find a max-weight integral flow between terminal pairs subject to the capacities. This problem was shown to be APX-hard by Garg, Vazirani and Yannakakis [Algorithmica, 1997], and a 4-approximation was given by Chekuri, Mydlarz and Shepherd [ACM Trans. Alg., 2007]. Some special cases are known to be exactly solvable in polynomial time, including when the graph is a path or a star.

First, when every edge has capacity at least $\mu \geq 2$, we use *iterated LP relaxation* to obtain an improved approximation ratio of $\min\{3, 1 + 4/\mu + 6/(\mu^2 - \mu)\}$. We show this ratio bounds the integrality gap of the natural LP relaxation. A complementary hardness result yields a $1 + \Theta(1/\mu)$ threshold of approximability (if $P \neq NP$). Second, we extend the range of instances for which exact solutions can be found efficiently. When the tree is a *spider* (i.e. if only one vertex has degree greater than 2) we give a polynomial-time algorithm to find an optimal solution, as well as a polyhedral description of the convex hull of all integral feasible solutions.

1 Introduction

In the max-weight integral multicommodity flow problem (WMCF), we are given an undirected *supply* graph $G = (V, E)$, terminal pairs $(s_1, t_1), \dots, (s_k, t_k)$ where $s_i, t_i \in V$, non-negative weights w_1, \dots, w_k , and non-negative integral edge-capacities c_e for all $e \in E$. The goal is to simultaneously route integral s_i - t_i flows of value y_i subject to the capacities, so as to maximize the weight $\sum w_i y_i$.

The single-commodity version ($k = 1$) of WMCF is well-known to be solvable in polynomial time. If we drop the integrality restriction the problem can be solved in polynomial time via linear programming for any k . However, when integrality is required, even the 2-commodity unit-capacity, unit-weight version

is NP-complete — see Even, Itai, and Shamir [12]. Moreover, recent results of Andrews et al. [1] show that WMCF is hard to approximate to within $\log^{\frac{1}{2}-\epsilon}(n)$ unless $\text{NP} \subset \text{ZPTIME}(n^{\text{polylog}(n)})$, even with unit capacities and weights. On the other hand, *randomized LP rounding* [22] yields a $O(\frac{\log n}{\log \log n})$ -approximation.

An easier and significant special case is where the supply graph G is a tree, which we denote by WMCFT. Garg, Vazirani and Yannakakis [15] considered the unit-weight case and showed APX-hardness even if G 's height is at most 3 and all capacities are 1 or 2; but on the positive side, they gave a 2-approximate polynomial-time primal-dual algorithm. Techniques of Garg et al. show that WMCFT can be solved in polynomial time when G has unit capacity (using dynamic programming and matching) or is a star (this problem is essentially equivalent to b -matching). The case where G is a path (so-called *interval packing*) is also polynomial-time solvable [4, 6, 17], e.g. by linear programming since the natural LP has a totally unimodular constraint matrix. For general WMCFT, without restrictions on capacities or weights, Chekuri, Mydlarz and Shepherd [6] gave a 4-approximation algorithm, and this remains the best ratio known to date.

Results. Throughout the paper, we use μ to denote the minimum capacity of any edge in the WMCFT instance. In the first part of the paper we use *iterated rounding/relaxation* to develop improved approximation ratios for WMCFT when μ is suitably large. Iterated relaxation yields an integral solution with optimal value or better but exceeding edge capacities by up to 2 (additively). This resolves a problem stated in Chekuri et al. [6]; in their words we prove “the c -relaxed integrality gap is 1” for $c = 2$ whereas they could not prove it for any constant c .³ We remark that our iterated rounding does not require *uncrossing*.

When the minimum capacity μ is $\Omega(\log |V|)$, randomized rounding gives a $1 + O(\frac{\log |V|}{\mu})$ -approximation for WMCFT. This suggests that the problem is easier to approximate as the minimum capacity increases. This was anticipated by Chekuri et al. [6], and indeed, by plugging our iterated rounding result into [6, Cor. 3.5] we get a $1 + O(1/\sqrt{\mu})$ -approximation for WMCFT when $\mu > 36$. In this paper, our first main result is an improvement on the best ratio for all $\mu \geq 2$.

Theorem 1 *For WMCFT, there are polynomial-time algorithms achieving (a) approximation ratio 3 for $\mu \geq 2$, and (b) approximation ratio $(1 + 4/\mu + 6/(\mu^2 - \mu))$ for general μ .*

A slight modification of Garg et al.’s hardness proof shows that for some $\epsilon > 0$, for all $\mu \geq 2$, it is NP-hard to approximate WMCFT within a ratio of $1 + \epsilon/\mu$; we detail this modification in Section 2.4. Thus (if $\text{P} \neq \text{NP}$) the ratio in Theorem 1(b) is tight up to the constant in the $\frac{1}{\mu}$ term.

³ To say that the c -relaxed integrality gap is 1 for an LP means that, where OPT is the LP’s optimal value, there is an integral solution of value at least OPT but violating each constraint by up to $+c$. E.g. when $c = 0$ this means the LP has an integral optimum.

Our methodology for Theorem 1 is to decrease the additively-violating solution towards feasibility, without losing too much weight. Part (a) uses an argument of Cheriyan, Jordán and Ravi [7]. Part (b) relies on an auxiliary *covering* problem; every feasible cover, when subtracted from the +2-violating WMCF solution, results in a feasible WMCF solution. An approach due to Jain [18] shows that iterated LP rounding, applied to the auxiliary problem, leads to a provably low-weight integral solution for the covering problem. A crucial fact in obtaining Theorem 1(b) is that the approximation guarantee of Jain’s approach is relative to the optimal *fractional* solution of the natural LP.

Our second result maps out more of the landscape of “easy” and “hard” WMCF instances. An *all-ror* instance is one in which, for some choice of root vertex, each commodity path either goes through the root or is *radial*. (A path is radial if one of its endpoints is an ancestor of the other with respect to the root.) For example, every instance of WMCF in which G is a spider is an all-ror instance.

Theorem 2 *All-ror WMCF instances can be solved in strongly polynomial time.*

One way to view this result is as an efficient solution for a common generalization of b -matching and interval packing. Our proof of Theorem 2 is via a combinatorial reduction to *bidirected flow* [9]. This reduction also yields a polyhedral characterization of the feasible solutions for all-ror instances.

We remark that our methodology can be generalized. Generalizing Theorem 1, for flows/covers of undirected/bidirected/vertex-capacitated trees, we can get constant-approximation algorithms, as well as $1 + O(1/\mu)$ -approximation algorithms where μ is the minimum capacity/requirement; for all these settings $1 + \Omega(1/\mu)$ -hardness can be shown. Similarly, Theorem 2 and its polyhedral counterpart generalize from flows to covers.

Related Work. WMCF appears in the literature under a variety of names including *cross-free-cut matching* [15] in the unit-capacity case and *packing of a laminar family* [7]. One generalization is the *demand* version [6] in which each commodity i is given a requirement r_i and we require $y_i \in \{0, r_i\}$ for each feasible solution.

The word *bidirected* has two meanings in the literature. We discuss *bidirected flows* later. In contrast, a *bidirected tree* is obtained from an undirected tree by replacing each edge by two antiparallel directed edges. WMCF extends naturally to directed graphs (with directed demands). A slight modification of Garg et al.’s hardness proof shows that WMCF on bidirected trees is APX-hard even for unit capacities and weights. On the other hand, WMCF on bidirected trees admits a $(\frac{5}{3} + \epsilon)$ -approximation for unit capacities [10] and a 4-approximation [6] for general capacities. On bidirected trees obtained from paths, stars, and spiders the problem can be solved in polynomial time, e.g., by reduction to a max-weight circulation problem, see also [11].

The WMCF problem on unit capacity trees is equivalent to the *weighted edge-disjoint paths* (WEDP) problem. The polylog-hardness result on general

graphs [1] applies to WEDP even with unit weights, however for fixed k , WEDP with at most k commodities is polynomial-time solvable. See e.g. [23, §70.5] for further discussion.

The extreme points of the natural LP for WMCFT arise frequently in the literature of LP-based network design [6, 7, 10, 13, 14, 16, 18, 20, 25]. From this perspective, WMCFT is a natural starting point for an investigation of how large capacities/requirements affect the difficulty of weighted network design problems. We remark that the *min-cardinality k -edge-connected spanning subgraph* problem has a similar history to WMCFT: a $1 + 2/(k + 1)$ -approximation was given by Cheriyan & Thurimella [8], $(1 + \epsilon/k)$ -hardness of approximation was shown by Gabow et al. [14], and the best current algorithm for the problem, due to Gabow & Gallagher [13], uses iterated rounding.

Organization of the Paper. Section 1.1 contains some basic definitions and notation. In Section 2 we give the proof of Theorem 1 and the complementary hardness result. In Section 3 we provide a proof of Theorem 2. In Section 3.1, we state our polyhedral results. Finally, we suggest some directions for future work in Section 4.

Some of the proofs which we did not have space to include are omitted, but can be obtained in the full version from the authors' webpages.

1.1 Formulation

We define the commodities by a set of *demand edges* $D = \{\{s_1, t_1\}, \dots, \{s_k, t_k\}\}$ on vertex set V with a weight w_d assigned to each demand edge $d \in D$; this is without loss of generality since the supply graph and demand edges are undirected. Since we discuss WMCF only on undirected trees, each commodity has a unique path along which flow is sent. For each demand edge d , let its *demand path* p_d be the unique path in G joining the endpoints of d . We thus may represent a multicommodity flow by a vector $\{y_d\}_{d \in D}$ where y_d is the amount of commodity d that is routed (along p_d). Then a flow y is *feasible* if it meets the two conditions

$$\forall d \in D : \quad y_d \geq 0 \quad (1)$$

$$\forall e \in E : \quad \sum_{d: e \in p_d} y_d \leq c_e \quad (2)$$

The objective of WMCFT is to find a feasible *integral* y that maximizes $w \cdot y$.

2 Improved Approximation

In this section we obtain a $\min\{3, 1 + 4/\mu + 6/(\mu^2 - \mu)\}$ -approximation algorithm for WMCFT, assuming $c_e \geq \mu \geq 2$ for each edge e . The algorithm uses the iterated rounding paradigm, which was used first by Jain [18] and more recently by Lau et al. [20] and others [3, 13, 25] for network design. The idea is that in each iteration, we use an extreme fractional optimum y^* of the natural LP to

develop an integral solution. If some demand edge d has value 1 or greater in y^* , we route the integer part and decrease the capacity of edges on p_d accordingly. If $0 < y < 1$ we perform a *relaxation* step; a counting argument (Section 2.3) guarantees that a particular relaxation can always be performed. At the end, we obtain an integral solution which exceeds the capacity of each edge by at most 2, and which has weight at least as large as that of an optimal feasible solution. In Sections 2.1 and 2.2 we show how to compute high-weight feasible solutions from this +2-violating solution. In Section 2.4 we show it is NP-hard to $(1 + \epsilon/\mu)$ -approximate WMCFT, for any fixed $\mu \geq 2$, and for some ϵ independent of μ .

The natural LP for WMCFT, which we denote by \mathcal{P}_0 , is as follows:

$$(\mathcal{P}_0) : \quad \text{maximize } w \cdot y \text{ subject to (1) and (2).}$$

This program has a linear number of variables and constraints, and thus can be solved in polynomial time. Any integral vector y is feasible for \mathcal{P}_0 iff it is feasible for the WMCFT instance. However, the linear program has fractional extreme points in general, and thus solving the LP does not give us the type of solution we seek. Nonetheless, optimal solutions to the LP have certain properties that permit an iterated rounding approach, such as the following.

Lemma 3 *Let y^* be an optimal solution to \mathcal{P}_0 , define $OPT = w \cdot y^*$, and suppose $y_d^* \geq t$ for some $d \in D$ and some integer $t \geq 1$. Reduce the capacity of each edge $e \in p_d$ by t and let OPT' denote the new optimal value of \mathcal{P}_0 . Then $OPT' = OPT - tw_d$.*

Proof. Let z denote the vector such that $z_d = t$ and $z_{d'} = 0$ for each $d' \neq d$. Then it is easy to see that $y^* - z$ is feasible for the new LP, and hence $OPT' \geq w \cdot (y^* - z) = OPT - tw_d$. On the other hand, where y' denotes the optimal solution to the new LP, it is easy to see that $y' + z$ is feasible for the original LP; so $OPT \geq OPT' + tw_d$. Combining these inequalities, we are done. \square

From now on, let OPT denote the optimal value to \mathcal{P}_0 .

In general terms, our iterated rounding approach works on the following principles. Define the following restricted version of \mathcal{P}_0 :

$$(\overline{\mathcal{P}}_0) : \quad \text{maximize } w \cdot y \text{ subject to (1) and (2) and } \forall d \in D : y_d \leq 1.$$

Assume for the moment that $\overline{\mathcal{P}}_0$ also has optimal value OPT . We iteratively build an integral solution to $\overline{\mathcal{P}}_0$ with value at least equal to OPT . The first step in each iteration is to solve $\overline{\mathcal{P}}_0$, obtaining solution y^* . If $y_d^* = 0$ for some demand edge d , then we can discard d without affecting the optimal value of $\overline{\mathcal{P}}_0$. If $y_d^* = 1$ for some d , then we can route one unit of flow along p_d and update capacities accordingly. Similar to Lemma 3, the optimal LP value will drop by an amount equal to the weight of the flow that was routed. If neither of these cases applies, we use the following lemma, whose proof appears in Section 2.3.

Lemma 4 *Suppose that y^* is an extreme point solution to \mathcal{P}_0 , and that $0 < y_d^* < 1$ for each demand edge $d \in D$. Then there is an edge $e \in E$ so that $|\{d \in D : e \in p_d\}| \leq 3$.*

Our algorithm discards the capacity constraint (2) for e from our LP. We call this *contracting* e because the effect is the same as if we had merged the two endpoints of e in the tree G . Pseudocode for our algorithm is given in Figure 1.

ITERATEDSOLVER

1. Set $\hat{y} = \mathbf{0}$
2. If $D = \emptyset$ terminate and return \hat{y}
3. Let y^* be an optimal extreme point solution to $\overline{\mathcal{P}_0}$
4. For each d such that $y_d^* = 0$, discard d
5. For each d such that $y_d^* = 1$, increase \hat{y}_d by 1, decrease c_e by 1 for each $e \in p_d$, and discard d
6. If neither step 4 nor 5 applied, find e as specified by Lemma 4 and contract e
7. Go to step 2

Fig. 1. The iterated rounding algorithm.

To justify our assumption that \mathcal{P}_0 and $\overline{\mathcal{P}_0}$ have the same optimal value, we preprocess the problem as follows. First, we compute any optimal solution y^* to \mathcal{P}_0 . Then we route the integer part $\lfloor y^* \rfloor$ of the solution (i.e., we initialize $\hat{y} = \lfloor y^* \rfloor$) and reduce each capacity c_e by $\sum_{d:e \in p_d} \lfloor y_d^* \rfloor$. The residual problem has $\langle y^* \rangle := y^* - \lfloor y^* \rfloor$ as an optimal solution, and since $0 \leq \langle y^* \rangle \leq 1$, our assumption is justified.

Assuming Lemma 4, we now prove the main properties of our iterated rounding algorithm: it runs in polynomial time, it exceeds each capacity by at most 2, and it produces a solution of value at least OPT .

Property 1 ITERATEDSOLVER runs in polynomial time.

Proof. Recall that \mathcal{P}_0 and $\overline{\mathcal{P}_0}$ can be solved in polynomial time. In each iteration we decrease $|D| + |E|$, so polynomially many iterations occur, and the result follows. \square

Property 2 The integral solution computed by ITERATEDSOLVER violates each capacity constraint (2) by at most 2.

Proof. Consider what happens to any given edge e during the execution of the algorithm. In the preprocessing and in each iteration, the flow routed through e equals the decrease in its residual capacity. If in some iteration, e 's residual capacity is decreased to 0, all demand paths through e will be discarded in the following iteration. Thus if e is not contracted, its capacity constraint (2) will be satisfied by the final solution.

The other case is that we contract e in step 6 of some iteration because e lies on at most 3 demand paths. The residual capacity of e is at least 1, and at most one unit of flow will be routed along each of these 3 demand paths in future iterations. Hence the final solution violates (2) for e by at most +2. \square

Property 3 *The integral solution computed by ITERATEDSOLVER has objective value at least equal to OPT .*

Proof. When we contract an edge e we just remove a constraint from $\overline{\mathcal{P}}_0$, which cannot decrease the optimal value of $\overline{\mathcal{P}}_0$ since it is a maximization LP. In every other iteration and in preprocessing, Lemma 3 implies that the LP optimal value drops by an amount equal to the increase in $w \cdot \hat{y}$. When termination occurs, the optimal value of $\overline{\mathcal{P}}_0$ is 0. Thus the overall weight of flow routed must be at least as large as the initial value of OPT . \square

2.1 Minimum Capacity $\mu = 2$

As per Property 2, our iterated solver may exceed some of the edge capacities. When $c_e \geq \mu \geq 2$ for each edge e we can invoke the following theorem, which appears as Thm. 6 in [7], to produce a high-weight feasible solution.

Theorem 5 (Cheriyán, Jordán, Ravi) *Suppose that \hat{y} is a nonnegative integral vector so that for each edge e , the constraint (2) is violated by at most a multiplicative factor of 2 by \hat{y} . Then in polynomial time, we can find an integral vector y' with $w \cdot y' \geq (w \cdot \hat{y})/3$, and $0 \leq y' \leq \hat{y}$, and such that y' satisfies all constraints (2).*

The algorithm as literally described in [7] is actually pseudo-polynomial, but it is straightforward to modify it to have polynomial running time; the proof appears in the full version. Assuming this fact, we now prove part (a) of Theorem 1.

Proof (of Theorem 1(a)). Let \hat{y} be the output of ITERATEDSOLVER. Since $c_e \geq 2$ for each edge e , and since by Property 2 each edge's capacity is additively violated by at most +2, Theorem 5 applies. Thus y' is a feasible solution to the WMCFT instance with objective value $w \cdot y' \geq w \cdot \hat{y}/3 \geq OPT/3$, using Property 3. Finally, since \mathcal{P}_0 is an LP-relaxation of the WMCFT problem, OPT is at least equal to the optimal WMCFT value, and so y' is a 3-approximate feasible integral solution. \square

2.2 Arbitrary Minimum Capacity

Given the infeasible solution \hat{y} produced by ITERATEDSOLVER, we want to reduce \hat{y} in a minimum-weight way so as to attain feasibility. For each edge e let $f_e = \max\{0, \sum_{d:e \in p_d} \hat{y}_d - c_e\}$, i.e. f_e is the amount by which \hat{y} violates the capacity of e . Note now that a reduction z with $0 \leq z \leq \hat{y}$ makes $\hat{y} - z$ a feasible (integral) WMCFT solution if and only if z is a feasible (integral) solution to the following linear program.

$$(\mathcal{P}_c) : \quad \text{minimize } w \cdot z \text{ subject to } 0 \leq z \leq \hat{y} \text{ and } \forall e \in E : \sum_{d:e \in p_d} z_d \geq f_e.$$

Notice that \mathcal{P}_c is a covering analogue of \mathcal{P}_0 (with added upper bounds). Furthermore, the integer program \mathcal{P}_c can be 2-approximately solved using Jain's iterated rounding framework [18].

Theorem 6 *There is a polynomial-time algorithm which returns an integral feasible solution \widehat{z} for \mathcal{P}_c such that $w \cdot \widehat{z}$ is at most twice the LP optimal value of \mathcal{P}_c .*

Proof. The idea is very similar to the main result of [18] but simpler in that no *uncrossing* is needed, because we already have a tree structure. Hence we only sketch the details. In each iteration, we obtain an extreme point optimal solution z^* to the linear program \mathcal{P}_c . We increase \widehat{z} by the integer part of z^* and accordingly decrease the requirements f . If $z_d^* = 0$, d is discarded. Finally if $0 < z^* < 1$ a token redistribution argument of Jain shows that some $d^* \in D$ has $z_{d^*}^* \geq 1/2$. In this case we increase \widehat{z}_{d^*} by 1 and update the requirements accordingly. Standard arguments then give the claimed bound on the cost of \widehat{z} and polynomial running time. \square

Here is how we use Theorem 6 to approximate WMCFT instances on trees.

Proof (of Theorem 1(b)). Notice that $z = \frac{2}{\mu+2}\widehat{y}$ is a feasible fractional solution to \mathcal{P}_c . Hence, the optimal value of \mathcal{P}_c is at most $\frac{2}{\mu+2}\widehat{y} \cdot w$. Thus the solution \widehat{z} produced by Theorem 6 satisfies $\widehat{z} \cdot w \leq \frac{4}{\mu+2}\widehat{y} \cdot w$, so $\widehat{y} - \widehat{z}$ is a feasible solution to the WMCFT problem, with $w \cdot (\widehat{y} - \widehat{z}) \geq (1 - \frac{4}{\mu+2})\widehat{y} \cdot w \geq (1 - \frac{4}{\mu+2})OPT$. This gives us a $1/(1 - \frac{4}{\mu+2}) = 1 + 4/\mu + O(1/\mu^2)$ approximation algorithm for WMCFT.

To obtain the exact bound claimed in Theorem 1(b), we refine this slightly by taking a two-round approach. In the first round we set f_e to be the characteristic vector of those edges which \widehat{y} violates by +2, obtaining $\widehat{y}' := \widehat{y} - \widehat{z}$. Then \widehat{y}' has only +1 additive violation, and the same reasoning as before shows $\widehat{y}' \cdot w \geq (1 - \frac{2}{\mu+2})OPT$. The second round analogously extracts from \widehat{y}' a feasible solution with weight at least $(1 - \frac{2}{\mu+1})\widehat{y}' \cdot w$. This gives approximation ratio $1/(1 - \frac{2}{\mu+2})(1 - \frac{2}{\mu+1}) = 1 + 4/\mu + 6/(\mu^2 - \mu)$, as desired. \square

2.3 Proof of Lemma 4

First, we need the following simple counting argument.

Lemma 7 *Let T be a tree with n vertices and let n_i denote the number of its vertices that have degree i . Then $n_1 > (n - n_2)/2$.*

Proof. Using the handshake lemma and the fact that T has $n - 1$ edges, we have $2(n - 1) = \sum_i i \cdot n_i$. But $\sum_i i \cdot n_i \geq n_1 + 2n_2 + 3(n - n_1 - n_2) = 3n - 2n_1 - n_2$ and hence $2n - 2 \geq 3n - 2n_1 - n_2$. Solving for n_1 gives $n_1 \geq (n - n_2 + 2)/2$ as needed. \square

Proof (of Lemma 4). Using basic facts from polyhedral combinatorics, it follows that there exists a set $E^* \subset E$ of edges with $|E^*| = |D|$ such that y^* is the unique solution to

$$\sum_{d \in D: e \in p_d} y_d = c_e \quad \forall e \in E^*. \quad (3)$$

In particular, the characteristic vectors of the sets $\{d : e \in p_d\}$ for $e \in E^*$ are linearly independent.

Contract each edge of $E \setminus E^*$ in (V, E) , resulting in the tree $T' = (V', E^*)$; call elements of V' *nodes*. We now use a counting argument to establish the existence of the desired edge e within E^* . We call the two ends of each $d \in D$ *endpoints* and say that node $v' \in V'$ *owns* k endpoints when the degree of v' in (V', D) is k .

First, consider any node $v' \in V'$ that has degree 2 in T' ; let e_1, e_2 be its incident edges in T' . If v' owns no endpoints then $\{d : e_1 \in p_d\} = \{d : e_2 \in p_d\}$, contradicting linear independence. If v' owns exactly one endpoint, the symmetric difference $\{d : e_1 \in p_d\} \Delta \{d : e_2 \in p_d\}$ consists of a single demand edge; but since y^* satisfies (3), $\mathbf{0} < y^* < \mathbf{1}$, and c is integral, this is a contradiction. Hence v' owns two or more endpoints.

If there exists a leaf node v' of T' that owns at most 3 endpoints then we are done, since this implies that the edge of E^* incident to v' , viewed in the original graph, lies on a most 3 demand paths. Otherwise, we apply a counting argument to T' , seeking a contradiction. Let n_i denote the number of nodes of T' of degree i . Then our previous arguments establish that the total number of endpoints is at least $4n_1 + 2n_2$. Lemma 7 then shows that the total number of endpoints is more than $2(|V'| - n_2) + 2n_2 = 2|V'| > 2|E^*| = 2|D|$. This is the desired contradiction, since there are only $2|D|$ endpoints in total. \square

We remark that Lemma 4 is tight in the following sense: if we replace the bound $|\{d \in D : e \in p_d\}| \leq 3$ with $|\{d \in D : e \in p_d\}| \leq 2$, the resulting statement is false. An example of an extreme point solution for which the modified version fails, due to Cheriyan et al. [7], is given in Figure 2.




Fig. 2. An extreme point solution to \mathcal{P}_0 . There are 9 edges in the supply graph, shown as thick lines; each has capacity 1. There are 9 demand edges, shown as thin lines; the solid ones have value $1/2$, and the dashed ones have value $1/4$. This is a tight example for Lemma 4 because each edge lies on at least three demand paths.

2.4 Hardness of MCFT with Constant Lower Bounds

We use MCFT to stand for the special case of WMCFT where all costs are 1.

Theorem 8 For some $\epsilon > 0$, for all integers $\mu \geq 2$, it is NP-hard to approximate MCFM to a factor of less than $1 + \epsilon/\mu$, even when restricted to instances where all capacities are at least μ .

Proof. The proof appears in the full version. □

3 Exact Solution for Spiders

In this section we show that WMCFT can be exactly solved in polynomial time when the supply graph is a spider. (A *spider* is a tree with exactly one vertex of degree greater than 2.) Call the vertex of degree ≥ 3 the *root* of the spider. Call each maximal path having the root as an endpoint a *leg* of the spider. Observe that in WMCFT when (V, E) is a spider, each demand path p_d either goes through the root, or else lies within a single leg. We generalize this observation into the following definition.

Definition 9 Consider an instance of WMCFT on graph (V, E) . With respect to a chosen root vertex $r \in V$, a demand edge d is said to be

- root-using, if r is an internal vertex of p_d ;
- radial, if one endpoint of d is a descendant of the other, with respect to the orientation of (V, E) induced by the root r .

The instance is *all-ror* (short for “all root-using or radial”) if there exists a choice of $r \in V$ for which every demand edge is either root-using or radial.

Instances with only radial demand paths can be exactly solved via the LP \mathcal{P}_0 since the constraint matrix is unimodular. Instances with only root-using demand paths can be solved using a matching approach, see e.g. the work of Nguyen [21]. To solve all-ror instances in general we use *bidirected flows*, which were introduced by Edmonds and Johnson [9]. Bidirected flow problems can be solved via a combinatorial reduction to *b*-matching (e.g., see [23]) which increases the instance size by a constant factor. We present in this section a reduction from all-ror WMCFT to bidirected flow.

A bidirected graph is an undirected graph together with, for each edge e and each of its endpoints $v \in e$, a sign $\sigma_{v,e} \in \{-1, +1\}$. Thus an edge can have two negative ends, two positive ends, or one end of each type; these are respectively called *negative* edges, *positive* edges, and *directed* edges. We will speak of directed edges as having the $+1$ end as their head and -1 end as their tail. An instance of capacitated max-weight bidirected flow is an integer program of the following form.

$$\text{maximize } \sum_{e \in E} \pi_e x_e \tag{4}$$

$$\forall v \in V : \quad a_v \leq \sum_{e \ni v} \sigma_{v,e} x_e \leq b_v \tag{5}$$

$$\forall e \in E : \quad \ell_e \leq x_e \leq u_e \tag{6}$$

$$x \text{ integral} \tag{7}$$

When $a = b = \mathbf{0}$ and all edges are directed, (4)–(7) becomes a max-weight circulation problem; when all edges are positive and $a = \mathbf{0}$, (4)–(7) becomes a b -matching problem. We now describe the reduction.

Proof (of Theorem 2). Let r denote the root vertex, i.e., assume every demand edge is either radial or root-using with respect to r . We construct a bidirected graph whose underlying undirected graph is $(V, E \cup D)$. Make each edge $e \in E$ directed, with head pointing towards r in the tree (V, E) . We make each root-using $d \in D$ a positive edge; we make each radial $d \in D$ a directed edge, with head pointing away from r . See Figure 3 for an illustration.




Fig. 3. An all-ror multicommodity flow instance. The tree graph (V, E) is depicted using thick lines, and the demand edges D are thin. Radial demand edges are dashed and root-using demand edges are solid. The root is r . An arrowhead denotes a positive endpoint, while the remaining endpoints are negative; these signs correspond to the reduction in the proof of Theorem 2.

Set $a_r = -\infty$, $b_r = +\infty$ and $a_v = b_v = 0$ for each $v \in V \setminus \{r\}$ in the bidirected flow problem (4)–(7). For each demand edge $d \in D$, call $C(d) := \{d\} \cup p_d$ the *demand cycle* of d . For a set F let χ^F denote the characteristic vector of F . Our choices of signs for the endpoints ensure that for each demand cycle $C(d)$, its characteristic vector $x = \chi^{C(d)}$ satisfies constraint (5). Moreover, any linear combination of these vectors is easily seen to satisfy (5), and the following converse holds.

Claim 10 *Any x satisfying (5) is a linear combination of characteristic vectors of demand cycles.*

Proof. Let $x' = x - \sum_d x_d \chi^{C(d)}$, and observe that x' also satisfies (5). Moreover, as each particular demand edge d^* occurs only in one demand cycle, namely $C(d^*)$, we have $x'_{d^*} = x_{d^*} - \sum_d x_d \chi_{d^*}^{C(d)} = x_{d^*} - x_{d^*} = 0$ for each $d^* \in D$. In other words, x' vanishes on D .

Now consider any leaf $v \neq r$ of G and its incident edge $uv \in E$. Since x' satisfies (5) at v and x' is zero on every edge incident to v except possibly uv , we deduce that $x_{uv} = 0$. By induction we can repeat this argument to show that x' also vanishes on all of E , so $x' = 0$. Then $x = x' + \sum_d x_d \chi^{C(d)} = \sum_d x_d \chi^{C(d)}$, which proves Claim 10. \square

By Claim 10, we may change the variables in the optimization problem from x to instead have one variable y_d for each $d \in D$; the variables are thus related by $x = \sum_d y_d \chi^{C(d)}$. In the bidirected optimization problem, set $\ell_e = 0, u_e = c_e$ for each $e \in E$, and $\ell_d = 0, u_d = +\infty$ for each $d \in D$. Rewriting (6) in terms of the new variables gives precisely the constraints (1) and (2). In other words, feasible integral flows x correspond bijectively to feasible integral solutions y for the WMCFT instance. Setting $\pi_d = w_d$ for $d \in D$ and $\pi_e = 0$ for $e \in E$, the objective function of (4) represents the weight for y , completing the reduction.

As mentioned earlier, this bidirected flow problem can in turn be reduced to a b -matching problem with a constant factor increase in the size of the problem. Using the strongly polynomial b -matching algorithm of Anstee [2], the proof of Theorem 2 is complete. \square

3.1 Polyhedral Results

The reduction used in the proof of Theorem 2 can also be used to derive the following polyhedral characterization; note that it is independent of which vertex is the root.

Theorem 11 *The convex hull of all integral feasible solutions in an all-ror WM-CFT problem has the following description:*

$$y_d \geq 0, \quad \forall d \in D \quad (8)$$

$$\sum_{e \in p_d} y_d \leq c_e, \quad \forall e \in E \quad (9)$$

$$\sum_{d \in D} y_d \lfloor |p_d \cap F|/2 \rfloor \leq \lfloor c(F)/2 \rfloor, \quad \forall F \subset E \text{ with } c(F) \text{ odd} \quad (10)$$

Proof. The proof appears in the full version. \square

Interestingly, results of Garg et al. [15, preliminary version] show that (8)–(10) is also integral in *unit-capacity* WMCFT. It is possible to synthesize our Theorems 2 and 11 with corresponding results of [15] for the unit-capacity case, by “gluing” all-ror instances at capacity-1 edges.

Caprara and Fischetti [5] gave a strongly polynomial-time algorithm to separate over the family (10) of inequalities. The ubiquity of the polyhedral formulation (8)–(10) suggests that it might be useful in designing a better approximation algorithm for WMCFT. One roadblock is that we do not know how to perform *uncrossing* (see, e.g., [18]) for that LP.

4 Closing Remarks

For the problem of finding a k -edge connected subgraph with the smallest number of edges, it is known [8, 14] that the best approximation ratio possible (if $P \neq NP$) is $1 + \Theta(1/k)$. We have proven a similar phenomenon for WMCFT (in terms of μ) and these results rely on similar techniques, notably iterated LP rounding. It would be nice to resolve the following outstanding question: for the problem of finding a *min-weight* k -edge connected subgraph, does the best possible approximation ratio decrease as k increases?

There is a close relation between WMCFT and its “demand” version where every flow variable is restricted according to $y_d \in \{0, r_d\}$ for some constants $\{r_d\}_{d \in D}$. E.g., combining ITERATEDSOLVER and Cor. 3.5 of [6], we obtain a $1 + O(\sqrt{r_{\max}/\mu})$ approximation for demand WMCFT where r_{\max} is the maximum demand. Shepherd and Vetta [24] showed that when the tree is a star, the $O(r_{\max})$ -relaxed integrality gap of the demand analogue of \mathcal{P}_0 is 1, and we can develop this result to get a $1 + O(r_{\max}/\mu)$ -approximation algorithm for demand WMCFT on stars. It would be interesting to demonstrate the same results on arbitrary trees.

When every edge has capacity one WMCFT is exactly solvable [15], and Theorem 1(a) gives a 3-approximation when there are no capacity-1 edges. Can we combine these results to improve upon the 4-approximation by Chekuri et al. [6] for general instances?

Acknowledgement

We would like to thank Joseph Cheriyan, Jim Geelen, and András Sebő for helpful discussions.

References

1. M. Andrews, J. Chuzhoy, S. Khanna, and L. Zhang. Hardness of the undirected edge-disjoint paths problem with congestion. In *Proc. 46th FOCS*, pages 226–244, 2005.
2. R. P. Anstee. A polynomial algorithm for b -matchings: An alternative approach. *Inf. Process. Lett.*, 24(3):153–157, 1987.
3. N. Bansal, R. Khandekar, and V. Nagarajan. Additive guarantees for degree bounded directed network design. In *Proc. 40th STOC*, pages 769–778, 2008.
4. G. Calinescu, A. Chakrabarti, H. J. Karloff, and Y. Rabani. Improved approximation algorithms for resource allocation. In *Proc. 9th IPCO*, pages 401–414, 2002.
5. A. Caprara and M. Fischetti. $\{0, \frac{1}{2}\}$ -Chvátal–Gomory cuts. *Math. Program.*, 74(3):221–235, 1996.
6. C. Chekuri, M. Mydlarz, and F. B. Shepherd. Multicommodity demand flow in a tree and packing integer programs. *ACM Trans. Algorithms*, 3(3):27, 2007. Extended abstract appeared in *Proc. 30th ICALP*, pages 410–425, 2003.
7. J. Cheriyan, T. Jordán, and R. Ravi. On 2-coverings and 2-packings of laminar families. In *Proc. 7th ESA*, pages 510–520, 1999.

8. J. Cheriyan and R. Thurimella. Approximating minimum-size k -connected spanning subgraphs via matching. *SIAM J. Comput.*, 30(2):528–560, 2000. Preliminary version appeared in *Proc. 37th FOCS*, pages 292–301, 1996.
9. J. Edmonds and E. Johnson. Matching: A well-solved class of integer linear programs. In R. Guy, H. Hanani, N. Sauer, and J. Schonheim, editors, *Proceedings, Calgary International Conference on Combinatorial Structures and their Applications*, pages 82–92. Gordon and Breach, 1970.
10. T. Erlebach and K. Jansen. Conversion of coloring algorithms into maximum weight independent set algorithms. In J. D. P. Rolim et al., editor, *Proc. ICALP Satellite Workshops*, pages 135–146, 2000.
11. T. Erlebach and K. Jansen. The maximum edge-disjoint paths problem in bidirected trees. *SIAM J. Discrete Math.*, 14(3):326–355, 2001. Preliminary version appeared as “Maximizing the Number of Connections in Optical Tree Networks” in *Proc. 9th ISAAC*, pages 179–188, 1998.
12. S. Even, A. Itai, and A. Shamir. On the complexity of timetable and multicommodity flow problems. *SIAM Journal on Computing*, 5(4):691–703, 1976. Preliminary version appeared in *Proc. 16th FOCS*, pages 184–193, 1975.
13. H. N. Gabow and S. Gallagher. Iterated rounding algorithms for the smallest k -edge-connected spanning subgraph. In *Proc. 19th SODA*, pages 550–559, 2008.
14. H. N. Gabow, M. X. Goemans, Éva Tardos, and D. P. Williamson. Approximating the smallest k -edge connected spanning subgraph by LP-rounding. In *Proc. 16th SODA*, pages 562–571, 2005.
15. N. Garg, V. V. Vazirani, and M. Yannakakis. Primal-dual approximation algorithms for integral flow and multicut in trees. *Algorithmica*, 18(1):3–20, May 1997. Preliminary version appeared in *Proc. 20th ICALP*, pages 64–75, 1993.
16. M. X. Goemans. Minimum bounded degree spanning trees. In *Proc. 47th FOCS*, pages 273–282, 2006.
17. I. B.-A. Hartman. Optimal k -colouring and k -nesting of intervals. In *Proc. 4th Israel Symposium on Theory of Computing and Systems*, pages 207–220, 1992.
18. K. Jain. A factor 2 approximation algorithm for the generalized Steiner network problem. *Combinatorica*, 21(1):39–60, 2001. Preliminary version appeared in *Proc. 39th FOCS*, pages 448–457, 1998.
19. V. Kann. *On the Approximability of NP-complete Optimization Problems*. PhD thesis, Royal Institute of Technology Stockholm, 1992.
20. L. C. Lau, J. S. Naor, M. R. Salavatipour, and M. Singh. Survivable network design with degree or order constraints. In *Proc. 39th STOC*, pages 651–660, 2007.
21. T. Nguyen. On the disjoint paths problem. *Oper. Res. Lett.*, 35(1):10–16, 2007.
22. P. Raghavan and C. Thompson. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7:365–374, 1987.
23. A. Schrijver. *Combinatorial optimization*. Springer, New York, 2003.
24. F. B. Shepherd and A. Vetta. The demand-matching problem. *Mathematics of Operations Research*, 32(3):563–578, 2007. Preliminary version appeared in *Proc. 9th IPCO*, pages 457–474, 2002.
25. M. Singh and L. C. Lau. Approximating minimum bounded degree spanning trees to within one of optimal. In *Proc. 39th STOC*, pages 661–670, 2007.