

Single-source k -splittable min-cost flows

Fernanda Salazar*

Martin Skutella^{†‡}

Abstract. Motivated by a famous open question on the single-source unsplittable minimum cost flow problem, we present a new approximation result for the relaxation of the problem where, for a given number k , each commodity must be routed along at most k paths.

Keywords: Approximation Algorithm; Multi-Commodity Flow; Network Flow; Routing; Unsplittable Flow; k -Splittable Flow

1 Introduction

We study a relaxation of the following network flow problem introduced by Kleinberg [5]:

Single-source unsplittable min-cost flow problem

Given: Digraph $G = (V, E)$ with capacities $u = (u_e)_{e \in E}$ and costs $c = (c_e)_{e \in E}$; source node $s \in V$ and p sink nodes $t_1, \dots, t_p \in V$ with demands $d = (d_1, \dots, d_p) \in \mathbb{R}_{\geq 0}^p$.

Task: Find a flow $(y_e)_{e \in E}$ with $y \leq u$ of minimum cost $c(y) = \sum_{e \in E} c_e y_e$ and with a path decomposition $(y_{P_i})_{i=1, \dots, p}$ such that P_i is an s - t_i -path and $y_{P_i} = d_i$ for $i = 1, \dots, p$.

The condition on the path decomposition of y simply says that the demand d_i of each commodity i must be routed along one single s - t_i -path. Any such flow y is called an *unsplittable flow satisfying demands d* . Already the problem of deciding whether an unsplittable flow satisfying demands d and obeying capacity constraints $y \leq u$ exists is NP-complete [5]. It contains several well-known NP-complete problems as special cases, such as, for example, Partition and Bin Packing. On the other hand, if we drop the constraint on y to be unsplittable, what remains is a classical minimum cost flow problem that can be solved efficiently.

Let $d_{\max} := \max_i d_i$ denote the maximum demand value. A popular assumption in the context of unsplittable flows is

*Escuela Politécnica Nacional, Departamento de Matemática, Quito, Ecuador

†Corresponding author.

‡TU Berlin, Institut für Mathematik, MA 5-2, Str. des 17. Juni 136, 10623 Berlin, skutella@math.tu-berlin.de

the *no bottleneck condition* which says that no demand may exceed the capacity of any arc, that is,

$$d_{\max} \leq u_e \quad \text{for all } e \in E. \quad (1)$$

Our work is motivated by the following conjecture.

Conjecture 1 (Goemans [4]). *For any flow x satisfying demands d , there is an unsplittable flow y satisfying demands d with*

$$y_e \leq x_e + d_{\max} \quad \text{for all } e \in E \quad (2)$$

and $c(y) \leq c(x)$.

Dinitz, Garg, and Goemans [2] prove that the conjecture without costs (i.e., removing the bound $c(y) \leq c(x)$) is true and provide an efficient algorithm that computes y .

The *congestion* of a given flow y is the minimum value $\alpha \geq 1$ with $y \leq \alpha u$. In particular, a flow of congestion 1 obeys the capacity constraints. The first approximation results for the min-congestion version of the single-source unsplittable flow problem (without costs) are given by Kleinberg [6]. Since a flow x satisfying demands d with minimum congestion can be computed with classical network flow techniques, the result of Dinitz et al. [2] implies the existence of a 2-approximation algorithm for the min-congestion problem without costs. For this and all further approximation results mentioned below we assume that the no-bottleneck condition (1) holds.

Kolliopoulos and Stein [8] prove the weaker version of Conjecture 1 where condition (2) is replaced by

$$y_e \leq 2x_e + d_{\max} \quad \text{for all } e \in E \quad (3)$$

and with the relaxed cost bound $c(y) \leq 2c(x)$. Their result implies the existence of a bicriteria (3, 2)-approximation algorithm for congestion and cost. Improving upon this result, Skutella [10] gives a (3, 1)-approximation algorithm. He proves Conjecture 1 with (2) replaced by (3) but with the original cost bound $c(y) \leq c(x)$. Notice that an efficient algorithm that computes an unsplittable flow y as in Conjecture 1 would yield a (2, 1)-approximation algorithm. On the negative side, Erlebach and Hall [3] prove that, for arbitrary $\varepsilon > 0$, there is no $(2 - \varepsilon, 1)$ -approximation algorithm, unless P=NP.

Kolliopoulos and Stein [8] and Skutella [10] both build upon the result that Conjecture 1 holds for the special case where all demand values are powers of 2. In [8] the case of

general demands is handled by *rounding up* demand values to the nearest power of 2. This yields an increase in cost by a factor of at most 2. In contrast to this, the improved result in [10] is achieved by *rounding down* demand values to the nearest power of 2 and carefully adjusting the given flow x .

We mention the following reformulation of Conjecture 1 stated in [9].

Conjecture 2 ([9]). *Any flow x satisfying demands d can be written as a convex combination of unsplittable flows y^ℓ , $\ell \in L$, with property (2) and satisfying demands d .*

It is not difficult to observe that Conjecture 2 is equivalent to Conjecture 1 (see [9] for details). Building upon [8] and [10], Martens, Salazar, and Skutella [9] prove the following result.

Theorem 1 ([9]). *Conjecture 2 holds if all demands are powers of 2. Moreover, the family of unsplittable flows y^ℓ , $\ell \in L$, can be obtained in polynomial time.*

In particular, the cardinality of L is polynomially bounded in the input size. More precisely, it is at most $|E| + 1$ (as a consequence of Carathéodory's Theorem).

Baier, Köhler, and Skutella [1] introduce the following relaxation of unsplittable flows. For a given $k \geq 1$, a k -splittable flow must route each commodity along at most k paths. In particular, 1-splittable flows are unsplittable flows. The resulting relaxation of our unsplittable flow problem is the single-source k -splittable min-cost flow problem. It follows from the classical flow decomposition theorem that k -splittability is not a meaningful restriction for $k \geq |E|$. We therefore assume in the remainder of the paper that $k < |E|$.

Kolliopoulos [7] presents an efficient algorithm that, given a flow x satisfying demands d , finds a 2-splittable flow y satisfying demands d with $y_e \leq \frac{4}{3}x_e + \frac{2}{3}d_{\max}$ for all $e \in E$ and $c(y) \leq c(x)$. This yields a $(2, 1)$ -approximation algorithm for the single-source 2-splittable min-cost flow problem. The main idea behind this result is to *round down* the demand values to the nearest sum of two powers of 2 and to carefully adjust the given flow x (as in [10]).

Our contribution. Inspired by the work of Kolliopoulos [7], we present the following improved and generalized result yielding a $(1 + \frac{1}{k} + \frac{1}{2k-1}, 1)$ -approximation algorithm for the single-source k -splittable min-cost flow problem. In particular, for $k = 2$ we get an $(\frac{11}{6}, 1)$ -approximation algorithm.

Theorem 2. *Let $k \in \mathbb{Z}_{>0}$. For any flow x satisfying demands d , there is a k -splittable flow y satisfying demands d with*

$$y_e \leq \frac{2k}{2k-1}x_e + \frac{d_{\max}}{k} \quad \text{for all } e \in E \quad (4)$$

and $c(y) \leq c(x)$. Moreover, such a flow y can be found in polynomial time.

In order to achieve this result, we build upon results mentioned above and introduce several new ideas and techniques. The history of bicriteria approximations for unsplittable flows outlined above suggests that rounding down demands to powers of 2 (as in [10]) leads to superior results compared to rounding up (as in [8]). Consequently, Kolliopoulos [7] also uses rounding-down in his algorithm. Surprisingly, the algorithm behind Theorem 2 is based on rounding-up. For the special case $k = 1$, our result coincides with the best known result for the unsplittable min-cost flow problem given in [10]. This implies, in particular, that the original idea of Kolliopoulos and Stein [8] to round up demands to powers of 2 can still lead to unsplittable flows y that are no more expensive than the given flow x . This insight also sheds new light on Conjecture 1.

Moreover, in contrast to earlier approximation results, in our result we use a more sophisticated technique based upon Theorem 1. That is, for the problem with rounded demands we compute an entire family of k -splittable flows which contain an accordingly rounded version of the given flow x in their convex hull. We emphasize that, in our approach, it is not sufficient to only compute a member of this family that has minimum cost. Only after going back to the original demands and rounding down all k -splittable flows we can check which members of the family do not violate the cost bound.

2 Proof of Theorem 2

The next lemma provides the basis for rounding up demand values to sums of k powers of 2.

Lemma 1. *For any $a \leq k$ there exist integers $q_1 \leq \dots \leq q_k \leq 0$ such that*

$$a \leq \sum_{j=1}^k 2^{q_j} < \frac{2k}{2k-1} a .$$

Moreover, given a , the numbers q_1, \dots, q_k can be obtained with $O(k)$ elementary operations.

Proof. Since $a \leq k$, there exist $q_1 \leq \dots \leq q_k \leq 0$ such that

$$a \leq \sum_{j=1}^k 2^{q_j} . \quad (5)$$

Among all possible choices of q_1, \dots, q_k , consider one that minimizes the right-hand side of (5). Decreasing q_1 by 1 yields a smaller right-hand side and thus

$$2^{q_1-1} + \sum_{j=2}^k 2^{q_j} < a . \quad (6)$$

Moreover, since $q_1 \leq q_j$ for $j = 2, \dots, k$,

$$\begin{aligned} \sum_{j=1}^k 2^{q_j} &= 2^{q_1-1} + \sum_{j=2}^k 2^{q_j} + \frac{(2k-1)2^{q_1-1}}{2k-1} \\ &\leq 2^{q_1-1} + \sum_{j=2}^k 2^{q_j} + \frac{2^{q_1-1} + 2 \sum_{j=2}^k 2^{q_j-1}}{2k-1} \\ &= \frac{2k}{2k-1} \left(2^{q_1-1} + \sum_{j=2}^k 2^{q_j} \right). \end{aligned}$$

This inequality together with (6) yields the desired result.

The following simple procedure can be used to compute q_1, \dots, q_k with $O(k)$ elementary operations: Set $a_k := a$. For $j = k$ down to 2 set $q_j := \min\{0, \lfloor \log_2 a_j \rfloor\}$ and $a_{j-1} := a_j - 2^{q_j}$. Finally, set $q_1 := \lceil \log_2 a_1 \rceil$. \square

Take an instance I of the single-source unsplittable min-cost flow problem. Without loss of generality we assume in the following that $d_{\max} = k$ (scaling). For each commodity $i = 1, \dots, p$, apply Lemma 1 to its demand $d_i \leq k$ in order to find $q_{i,1}, \dots, q_{i,k} \leq 0$ and $\bar{d}_i := \sum_{j=1}^k 2^{q_{i,j}}$ so that

$$1 \leq \alpha_i := \frac{\bar{d}_i}{d_i} < \frac{2k}{2k-1}. \quad (7)$$

Let $\bar{d} := (\bar{d}_1, \dots, \bar{d}_p)$ and let \bar{I} denote a modified instance with demand vector d replaced by the rounded demand vector \bar{d} . Moreover, let \bar{I}' denote the instance obtained from \bar{I} where each commodity $i = 1, \dots, p$ is replaced by k sub-commodities i_j with demands $2^{q_{i,j}}$ and sharing the same sink $t_{i_j} := t_i$, $j = 1, \dots, k$. We denote the corresponding demand vector by \bar{d}' .

Observation 1.

(i) A flow \bar{x} satisfies \bar{d} if and only if it satisfies \bar{d}' .

(ii) The maximum demand value \bar{d}'_{\max} is equal to d_{\max}/k .

Let I' denote the instance obtained from \bar{I}' by setting the demand of sub-commodity i_j to $2^{q_{i,j}}/\alpha_i$, $i = 1, \dots, p$, $j = 1, \dots, k$. The corresponding demand vector is denoted by d' . Notice that the maximum demand value d'_{\max} in d' is equal to $1 = d_{\max}/k$. Since $\sum_{j=1}^k 2^{q_{i,j}}/\alpha_i = \bar{d}_i/\alpha_i = d_i$, for $i = 1, \dots, p$, a flow satisfies d if and only if it satisfies d' . This yields the following observation.

Observation 2. Any unsplittable flow satisfying d' is a k -splittable flow satisfying d .

We can now state the algorithm used to prove Theorem 2.

Algorithm

Input: A flow x satisfying d .

Output: A k -splittable flow satisfying d .

1. Compute \bar{d} and \bar{d}' as discussed above.
2. Compute a minimum cost flow \bar{x} satisfying \bar{d} with

$$x \leq \bar{x} \leq \frac{2k}{2k-1} x. \quad (8)$$

3. Write \bar{x} as a convex combination of unsplittable flows \bar{y}^ℓ , $\ell \in L$, satisfying \bar{d}' with

$$\bar{y}_e^\ell \leq \bar{x}_e + \bar{d}'_{\max} \quad \text{for all } e \in E. \quad (9)$$

4. For each $\ell \in L$, construct y^ℓ from \bar{y}^ℓ by scaling flow of all sub-commodities i_j by $1/\alpha_i$ for $i = 1, \dots, p$, $j = 1, \dots, k$.

5. Determine $\ell^* \in L$ with $c(y^{\ell^*})$ minimal and output y^{ℓ^*} .

Lemma 2. Given a flow x satisfying d , the algorithm above computes in polynomial time a k -splittable flow satisfying d and with congestion bounded as in (4).

Proof. Step 1 of the algorithm can be done efficiently by Lemma 1. Notice that the flow \bar{x} in Step 2 exists due to (7). It can be obtained by a standard minimum cost flow computation. By Observation 1 (i), \bar{x} satisfies demands \bar{d}' . Due to Theorem 1, the convex combination in Step 3 exists and can be obtained efficiently. Since the cardinality of the index set L is bounded by $|E| + 1$ (remark after Theorem 1), also Steps 4 and 5 can be done efficiently.

By construction, y^ℓ is an unsplittable flow satisfying d' , for each $\ell \in L$. Thus, by Observation 2, each y^ℓ is a k -splittable flow satisfying d . Finally, we prove that each y^ℓ , $\ell \in L$, satisfies the bounds (4). For each arc $e \in E$

$$\begin{aligned} y_e^\ell &\leq \bar{y}_e^\ell && \text{since } \alpha_i \geq 1 \text{ for all } i, \\ &\leq \bar{x}_e + \bar{d}'_{\max} && \text{by (9),} \\ &\leq \frac{2k}{2k-1} x_e + \frac{d_{\max}}{k} && \text{by (8) and Observation 1 (ii).} \end{aligned}$$

This concludes the proof. \square

In the next lemma, we state the desired cost bound for the k -splittable flow computed by the algorithm.

Lemma 3. The cost of the k -splittable flow computed by the algorithm is at most the cost of the given flow x .

If x was contained in the convex hull of the y^ℓ , $\ell \in L$, the result would immediately follow. This is, however, in general not the case.

Proof. By construction, \bar{x} is a convex combination of the \bar{y}^ℓ , $\ell \in L$, that is

$$\bar{x} = \sum_{\ell \in L} \lambda_\ell \bar{y}^\ell \quad (10)$$

with $\lambda_\ell \geq 0$ for each $\ell \in L$ and $\sum_{\ell \in L} \lambda_\ell = 1$. Let

$$\tilde{x} := \sum_{\ell \in L} \lambda_\ell y^\ell. \quad (11)$$

Since each y^ℓ , $\ell \in L$, satisfies demands d , the same holds for the convex combination \tilde{x} .

By our choice of ℓ^* , the cost of y^{ℓ^*} is at most $c(\tilde{x})$. It therefore suffices to prove that $c(\tilde{x}) \leq c(x)$. To this end, we show that $\bar{x} - \tilde{x} + x$ is a feasible solution to the min-cost flow problem in Step 2 of the algorithm. By linearity of the cost function and optimality of \bar{x} , this implies that

$$0 \geq c(\bar{x}) - c(\bar{x} - \tilde{x} + x) = c(\tilde{x}) - c(x)$$

and thus the desired cost bound.

To check that flow $\bar{x} - \tilde{x} + x$ has the properties requested in Step 2, first notice that it satisfies the same demands \bar{d} as \bar{x} because both x and \tilde{x} satisfy d and thus cancel out. It remains to prove the required lower and upper bounds on $\bar{x} - \tilde{x} + x$. The bounds on the α_i 's in (7) yield

$$\frac{2k-1}{2k} \bar{y}^\ell < y^\ell \leq \bar{y}^\ell \quad \text{for each } \ell \in L.$$

Taking convex combinations (see (10) and (11)) we get

$$\frac{2k-1}{2k} \bar{x} \leq \tilde{x} \leq \bar{x}.$$

Together with (8) this finally yields

$$\begin{aligned} x &\leq \bar{x} - \tilde{x} + x \leq \bar{x} - \frac{2k-1}{2k} \bar{x} + x \\ &= \frac{1}{2k} \bar{x} + x \leq \frac{1}{2k} \frac{2k}{2k-1} x + x = \frac{2k}{2k-1} x. \end{aligned}$$

Thus, $\bar{x} - \tilde{x} + x$ is a feasible solution to the min-cost flow problem in Step 2 of the algorithm. \square

Lemmas 2 and 3 complete the proof of Theorem 2.

Acknowledgements. This work was supported by DFG Research Center MATHEON in Berlin and by Berlin Mathematical School.

References

- [1] G. Baier, E. Köhler, and M. Skutella, *On the k -splittable flow problem*, *Algorithmica* **42** (2005), 231–248.
- [2] Y. Dinitz, N. Garg, and M. X. Goemans, *On the single source unsplittable flow problem*, *Combinatorica* **19** (1999), 17–41.
- [3] T. Erlebach and A. Hall, *NP-hardness of broadcast scheduling and inapproximability of single-source unsplittable min-cost flow*, *Journal of Scheduling* **7** (2004), 223–241.
- [4] M. X. Goemans, January 2000, Personal communication.
- [5] J. M. Kleinberg, *Approximation algorithms for disjoint paths problems*, Ph.D. thesis, M.I.T., 1996.
- [6] J. M. Kleinberg, *Single-source unsplittable flow*, Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science (Burlington, VT), 1996, pp. 68–77.
- [7] S. G. Kolliopoulos, *Minimum-cost single-source 2-splittable flow*, *Information Processing Letters* **94** (2005), 15–18.
- [8] S. G. Kolliopoulos and C. Stein, *Approximation algorithms for single-source unsplittable flow*, *SIAM Journal on Computing* **31** (2002), 919–946.
- [9] M. Martens, F. Salazar, and M. Skutella, *Representing single source multicommodity flows as convex combinations of unsplittable flows*, Algorithms — ESA '07 (L. Arge, M. Hoffmann, and E. Welzl, eds.), Lecture Notes in Computer Science, vol. 4698, Springer, Berlin, 2007, pp. 395–406.
- [10] M. Skutella, *Approximating the single source unsplittable min-cost flow problem*, *Mathematical Programming* **91** (2002), 493–514.