Available online at www.sciencedirect.com
SCIENCE DIRECT®
Information Processing Letters 94 (2005) 15–18
Information Processing Letters
www.elsevier.com/locate/ipl

Minimum-cost single-source 2-splittable flow ☆

Stavros G. Kolliopoulos [1]

Department of Informatics and Telecommunications, University of Athens, Panepistimiopolis Ilissia, Athens 157 84, Greece

Received 24 June 2004; received in revised form 9 December 2004

Available online 21 January 2005

Communicated by F. Meyer auf der Heide

Abstract

In the *single-source unsplittable flow* problem, commodities must be routed simultaneously from a common source vertex to certain sinks in a given directed graph with edge capacities and costs. The demand of each commodity must be routed along a single path so that the total flow through any edge is at most its capacity. Moreover the cost of the solution should not exceed a given budget. An important open question is whether a simultaneous $(2, 1)$-approximation can be achieved for minimizing congestion and cost, i.e., the budget constraint should not be violated. In this note we show that this is possible for the case of 2-splittable flows, i.e., flows where the demand of each commodity is routed along at most two paths. Our result holds under the "no-bottleneck" assumption, i.e., the maximum demand does not exceed the minimum capacity.

Keywords: Analysis of algorithms; Approximation algorithms; Graph algorithms; Maximum flow; Unsplittable flow

1. Introduction

In the *single-source unsplittable flow* problem (UFP), we are given a directed graph $G = (V, E)$ with edge capacities $u : E \to \mathbb{R}^+$ and edge costs $c : E \to \mathbb{R}^+$, a budget $B > 0$, a designated source vertex $s \in V$, and k commodities each with a sink vertex $t_i \in V$ and associated demand $d_i > 0$, $i = 1, \ldots, k$. A feasible unsplittable flow routes for each i, d_i units of commodity i along a single path from s to t_i so that the total flow through an edge e is at most its capacity u_e. As is standard in the relevant literature we assume that no edge can be a bottleneck, i.e., the minimum edge capacity is assumed to have value at least $\max_i d_i$. The cost $c(f)$ of flow f is given by $c(f) = \sum_{e \in E} c_e f_e$. The cost $c(P_i)$ of a path P_i is defined as $c(P_i) = \sum_{e \in P_i} c_e$. We seek a feasible unsplittable flow whose total cost does not exceed the budget. The cost of an unsplittable flow f given by paths $P_1, \ldots, P_k$ can also be written as $c(f) = \sum_{i=1}^{k} d_i \cdot c(P_i)$. The feasibility question for UFP is strongly NP-complete [4] even without a

☆ A preliminary version of this work appeared in the Proceedings of the Cologne–Twente Workshop on Graph Theory and Combinatorial Optimization, 2004.

E-mail address: sgk@di.uoa.gr (S.G. Kolliopoulos).
URL: http://www.di.uoa.gr/~sgk.

[1] Research partially supported by NSERC Grant 227809-00.

0020-0190/$ – see front matter
doi:10.1016/j.ipl.2004.12.009

budget constraint. An optimization version which has attracted considerable attention is to *minimize congestion*: Find the smallest $\alpha \geqslant 1$ such that there exists a feasible unsplittable flow if all capacities are multiplied by α.

A relaxation of UFP is *b-splittable flow*, $b > 1$; the definition is the same as for UFP except we allow the demand of each commodity to be split along at most b paths. The paths for a single commodity do not have to be disjoint; their interaction is regulated by the capacity constraints in the same manner as with any two paths for different commodities. If there is no bound on b or if $b \geqslant |E|$, this yields a standard maximum flow problem. We will call a maximum flow solution for the original UFP instance, a *fractional flow*. In this note we study the simultaneous approximation of congestion and cost for the single-source 2-splittable flow problem. This corresponds to the strictest possible relaxation of UFP as far as the usage of paths is concerned. We slightly abuse terminology and view such flows as 2-splittable solutions to UFP.

UFP was introduced by Kleinberg [4]. It is an interesting special case of the general multisource unsplittable flow problem (for the latter problem see, e.g., [7]). UFP contains several well-known NP-complete problems as special cases: Partition, Bin Packing, scheduling on parallel machines to minimize makespan [4]. In addition UFP generalizes single-source edge-disjoint paths and models aspects of virtual circuit routing. The first constant-factor approximations were given in [5]. Kolliopoulos and Stein [6] gave a 3-approximation algorithm for congestion which also guarantees a flow cost of value at most 2 times the optimal cost of a fractional solution. A *bicriteria* (ρ_1, ρ_2)-approximation algorithm for congestion and cost is a polynomial-time algorithm which is guaranteed to output a solution which simultaneously has congestion at most ρ_1 times the optimal and cost at most ρ_2 times the given budget. In this notation, [6] gave a $(3, 2)$-approximation. Dinitz et al. [2] obtained a congestion bound of 2 but their algorithm cannot handle the budget constraint. To be more precise, their basic result is that any splittable flow satisfying all demands can be turned into an unsplittable flow while increasing the total flow through any edge by less than the maximum demand. This result is tight if the congestion achieved by the fractional flow is used as a lower bound. Skutella [9] obtained a $(3, 1)$-approximation algorithm and this is currently the best bicriteria bound. Various results on b-splittable flows were obtained by Baier et al. in [1]. The main focus of their paper is that of finding a maximum value b-splittable s-t flow. Finding a feasible solution to UFP reduces to solving optimally a k-splittable s-t flow instance on an auxiliary network where for each original commodity i there is an edge of capacity d_i connecting t_i to a new supersink vertex t. As Baier et al. observe [1] the analysis of any algorithm for UFP that uses as lower bounds the optima of a fractional flow applies for the single-source b-splittable flow problem. Therefore the result of [9] extends to a $(3, 1)$-approximation for the 2-splittable case. There does not seem to be a further connection between the results in [1] and the formulation we examine.

In terms of negative results, Erlebach and Hall [3] prove that for UFP and arbitrary $\varepsilon > 0$ there is no $(2 - \varepsilon, 1)$-approximation algorithm for congestion and cost unless P = NP. Matching this bicriteria lower bound is an important open question that has attracted a lot of attention. Such a $(2, 1)$-approximation is known only for the scheduling problem $R|p_{ij} = p_j$ or $\infty|C_{\max}$ with assignment costs [8]. This scheduling problem reduces to a UFP instance on a 2-level graph where minimizing congestion is equivalent to minimizing the makespan. We remark that this scheduling problem is also a special case of the generalized assignment problem. For the latter problem a simultaneous $(2, 1)$-approximation for makespan and cost is well known [8].

In this note we show that the simultaneous $(2, 1)$-approximation for congestion and cost can be obtained for the single-source 2-splittable flow problem. The precise bound we achieve is given in Theorem 2.4. Better bounds can be achieved for b-splittable solutions with $b > 2$; we omit the details.

2. The algorithm for 2-splittable flow

We set $d_{\max} = \max_{1 \leqslant i \leqslant k} d_i$, $d_{\min} = \min_{1 \leqslant i \leqslant k} d_i$ for the instance of interest. We assume without loss of generality that there is a feasible fractional solution f_0 for the given UFP instance I and that $d_{\max} \leqslant 1$. The “no-bottleneck” assumption implies that $u_{\min} \geqslant 1$. The cost of the final splittable solution will not exceed the cost of the initial fractional solution f_0. If the latter solution is a minimum-cost flow we obtain a best possible budget. The following result of Kolliopoulos and Stein [6] will be of use.

Theorem 2.1 [6]. *Given a* UFP *instance where all demands are powers of* $1/2$ *and an initial fractional flow solution, there is an algorithm, called* POWER-ALG, *which finds an unsplittable flow* f *that violates the capacity of any edge by at most* $d_{\max} - d_{\min}$ *and whose cost is bounded by the cost of the initial fractional flow.*

To solve the 2-splittable flow problem one can naively break a commodity into two equal demands and run an appropriate UFP algorithm. The Skutella algorithm guarantees that in the unsplittable solution the flow on edge e increases from f_e to at most $2f_e + d_{\max}$. Together with the fact that $d_{\max} \leqslant 0.5$ for the broken commodities, this yields a $(2.5, 1)$-approximation. Instead we will break every demand d_i, $i = 1, \ldots, k$, into two commodities with demands a_i and b_i s.t. $a_i + b_i \leqslant d_i$, and both a_i and b_i are powers of $1/2$. In a repairing stage at the end we will ensure that all of d_i is routed.

Let $\text{floor}_2(x)$ denote the largest number which is a power of $1/2$ and does not exceed x. We define the following operator $\lfloor \cdot \rfloor_2$

$$\lfloor x \rfloor_2 = \begin{cases} \text{floor}_2(x) & \text{if } x < 1, \\ 1/2 & \text{if } x = 1. \end{cases} \tag{1}$$

Set $a_i \doteq \lfloor d_i \rfloor_2$ and $b_i \doteq \lfloor d_i - a_i \rfloor_2$. Create a new UFP instance I^2 with $2k$ commodities where commodity i of I is mapped to two commodities with demands a_i, b_i. Run the POWER-ALG of Theorem 2.1 on I^2 to obtain a flow f. Observe that $x \leqslant y \Rightarrow \lfloor x \rfloor_2 \leqslant \lfloor y \rfloor_2$.

Lemma 2.2. *Given the* UFP *instance* I^2 *with initial fractional solution* f_0^2 *one can find an unsplittable flow* f *which* (i) *violates the capacity of any edge by at most* $\lfloor d_{\max} \rfloor_2$ *and* (ii) *whose cost is bounded by the cost of the initial fractional flow* f_0^2*. Flow* f *corresponds to a* 2*-splittable flow for instance* I *which routes* $a_i + b_i$ *units of flow for commodity* i, $i = 1, \ldots, k$.

The task that remains is to transform the flow f of Lemma 2.2 so that d_i units of flow are routed for commodity i. By the definition of a_i and b_i we have that $a_i \geqslant \frac{d_i}{2}$ and $b_i \geqslant \frac{d_i - a_i}{2}$. Therefore

$$a_i + b_i \geqslant \frac{3d_i}{4}, \quad i = 1, \ldots, k.$$

We obtain a flow f' from f by scaling the flow on each of the at most two s-t_i paths used in f *by the same amount* $\lambda_i \in (1, 4/3)$ so that

$$\lambda_i(a_i + b_i) = d_i.$$

This transformation yields a 2-splittable flow f' which (i) satisfies all demands d_i and (ii) satisfies $f'_e \leqslant (4/3)u_e + (4/3)\lfloor d_{\max} \rfloor_2$ for all edges $e \in E$. We distinguish two cases.

Case 1. $d_{\max} \geqslant 1/2$. Then $\lfloor d_{\max} \rfloor_2 = 1/2$. Therefore $f'_e \leqslant (4/3)u_e + 2/3$.

Case 2. $d_{\max} < 1/2$. Then $\lfloor d_{\max} \rfloor_2 \leqslant 1/4$. Therefore $f'_e \leqslant (4/3)u_e + 1/3$.

We have shown the following lemma.

Lemma 2.3. *Given a* UFP *instance with initial fractional solution* f_0 *one can find a* 2*-splittable flow* f' *such that*

(i) f' *satisfies all demands* d_i, $i = 1, \ldots, k$,
(ii) $f'_e \leqslant (4/3)u_e + 2/3$ *for all* $e \in E$, *and*
(iii) *the cost of* f' *is bounded by at most* $4/3$ *times the cost of the initial fractional flow* f_0.

To obtain the bound on the cost we define carefully the fractional solution for the instance I^2 using a method proposed by Skutella [9]. The fractional solution f_0^2 is obtained from f_0 as follows: for the commodities for which $d_i = a_i + b_i$, solution f_0^2 sends flow as f_0. For the remaining commodities we decrease the flow by $\bar{d}_i = d_i - (a_i + b_i)$ along the most expensive, i.e., higher cost, s-t_i paths as in [9]. For a given commodity i, this is implemented through the following iterative procedure. Find the most expensive s-t_i flow path P_1^i, and let f_1^i be its flow amount due to commodity i. Cancel $\gamma_i = \min\{f_1^i, \bar{d}_i\}$ units of flow from it and decrease $\bar{d}_i$ by γ_i. Remove any edge whose flow becomes zero and repeat on the remaining network by finding the next most expensive path P_2^i and so on until $\bar{d}_i = 0$. Computing each time a maximum-cost path can be implemented in polynomial time because we can assume without loss of generality that f_0 does not send flow along cycles [9]. After this preprocessing step $c(f_0^2) \leqslant c(f_0)$. By Lemma 2.2, $c(f) \leqslant c(f_0^2)$. Inspection of the POWER-ALG yields that the paths used in the unsplittable solution must be paths with nonzero flow in the initial fractional solution. Therefore all the paths from s to t_i that carry nonzero flow in f have cost less than or equal to the cost of the paths on which flow was decreased while obtaining f_0^2 from f_0. Routing $d_i - (a_i + b_i)$ additional units from s to t_i along the at most two paths used in f results in a solution f' for which $c(f') \leqslant c(f_0)$. The main result has been proved.

Theorem 2.4. *Given a* UFP *instance I with initial fractional solution f_0 one can find in polynomial time a* 2*-splittable flow f' such that*

(i) *f' satisfies all demands d_i, $i = 1, \ldots, k$,*
(ii) *$f'_e \leqslant (4/3)u_e + 2/3$ for all $e \in E$, and*
(iii) *the cost of f' is bounded by the cost of the initial fractional flow f_0.*

Corollary 2.5. *Given a* UFP *instance I one can find in polynomial time a* 2*-splittable flow solution that achieves a simultaneous* $(2, 1)$*-approximation for congestion and cost.*

Acknowledgement

Thanks to the two anonymous referees for corrections to the original proof.

References

[1] G. Baier, E. Köhler, M. Skutella, On the k-splittable flow problem, Algorithmica (special issue on ESA 2002), in press.

[2] Y. Dinitz, N. Garg, M.X. Goemans, On the single-source unsplittable flow problem, Combinatorica 19 (1999) 1–25.

[3] T. Erlebach, A. Hall, NP-hardness of broadcast sheduling and inapproximability of single-source unsplittable min-cost flow, in: Proc. 13th ACM-SIAM Symp. on Discrete Algorithms, 2002.

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

[5] J.M. Kleinberg, Single-source unsplittable flow, in: Proc. 37th Annual IEEE Symp. on Foundations of Computer Science, October 1996, pp. 68–77.

[6] S.G. Kolliopoulos, C. Stein, Approximation algorithms for single-source unsplittable flow, SIAM J. Comput. 31 (2002) 919–946.

[7] S.G. Kolliopoulos, C. Stein, Approximating disjoint-path problems using packing integer programs, Math. Programming A 99 (2004) 63–87.

[8] D.B. Shmoys, É. Tardos, An approximation algorithm for the generalized assignment problem, Math. Programming A 62 (1993) 461–474.

[9] M. Skutella, Approximating the single-source unsplittable min-cost flow problem, Math. Programming B 91 (2002) 493–514.