Doctopic: Algorithms, automata, complexity and games ARTICLE IN PRESS TCS:11206

Theoretical Computer Science ••• (••••) •••–•••

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Maximum flow under proportional delay constraint

Pierre Bonami [a,b], Dorian Mazauric [a,c,*], Yann Vaxès [a]

[a] *Aix-Marseille Université, CNRS, LIF UMR 7279, 13000, Marseille, France*
[b] *IBM ILOG CPLEX, Madrid, Spain*
[c] *Université Côte d'Azur, Inria, Sophia Antipolis, France*

A R T I C L E I N F O

Article history:
Received 27 May 2016
Received in revised form 5 April 2017
Accepted 30 May 2017
Available online xxxx
Communicated by T. Erlebach

Keywords:
Maximum flow
On–off delay constraints
Polynomial time approximation scheme (PTAS)
Dynamic programming
Bounded treewidth
Linear programming

A B S T R A C T

Given a network and a set of source destination pairs (connections), we consider the problem of maximizing the sum of the flow under proportional delay constraints. In this paper, the delay for crossing a link is proportional to the total flow crossing this link. If a connection supports non-zero flow, then the sum of the delays along any path corresponding to that connection must be lower than a given bound. The constraints of delay are on–off constraints because if a connection carries zero flow, then there is no constraint for that connection. The difficulty of the problem comes from the choice of the connections supporting non-zero flow. We first prove a general approximation ratio using linear programming for a variant of the problem. We then prove a linear time 2-approximation algorithm when the network is a path. We finally show a Polynomial Time Approximation Scheme when the graph of intersections of the paths has bounded treewidth.

1. Introduction

The multi-commodity flow problem is a classical network flow problem with multiple commodities (flow demands) between different source and sink nodes. This problem has been widely studied in the literature. Given a network, a set of capacities on edges, and a set of demands (commodities), the problem consists in finding a network flow satisfying all the demands and respecting capacities and flow conservation constraints. The integer version of the problem is NP-complete [6], even for only two commodities and unit capacities (making the problem Strongly NP-complete in this case). In that version, the problem consists of producing an integer flow satisfying all demands and respecting the previously mentioned constraints. However, if fractional flows are allowed, the problem can be solved in polynomial time through linear programming [1,4,8].

Network operators must satisfy some Quality of Service requirements for their clients. One of the most important parameters in telecommunications networks is the end-to-end delay of a unit of flow between a source node and a destination node. This requirement is not taken into account in the multi-commodity flow problem. The delay through a link depends on the amount of flow supported by this link; classically it is modeled by a convex function. The end-to-end delay for a demand and a path associated with this demand, is the sum of the delays through all links of this path. Some papers focus on minimizing the mean end-to-end delay. This problem consists of minimizing a convex function under linear con-

* Corresponding author.
E-mail address: dorian.mazauric@inria.fr (D. Mazauric).

http://dx.doi.org/10.1016/j.tcs.2017.05.034
0304-3975/

Please cite this article in press as: P. Bonami et al., Maximum flow under proportional delay constraint, Theoret. Comput. Sci. (2017), http://dx.doi.org/10.1016/j.tcs.2017.05.034

straints [3,9] and can be solved using semidefinite programming [12]. Other papers focus on finding a multi-commodity flow that satisfies the demands and that minimizes the maximum end-to-end delay [5].

As the authors of [2], we think that a more realistic problem consists of adding strict end-to-end delay constraints for all connections. Indeed, in communication networks, there are multiple classes of services, and for each of them, it is crucial to respect a certain level of Quality of Service, i.e. respecting a threshold for the end-to-end delay. It is why we study a multicommodity flow problem in which all demands have to respect an end-to-end delay constraint.

In this paper, we focus on multicommodity network flow problems in which each edge possesses a proportional latency function that describes the common delay experienced by the flow on that edge as a function of the flow rate. These problems model congestion effects that appear in a variety of applications such as communication networks, vehicular traffic, supply chain management, or evacuation planning. In such applications, the latency function need not be proportional to the flow rate. We investigate the problem of maximizing the sum of the flow under proportional delay constraints. We allow fractional flows for all connections. As mentioned before, we have an end-to-end delay constraint for each demand. But, if a path associated with one connection carries zero flow, then the constraint is not active. These on–off constraints make the problem more difficult to solve than just a simple linear program [7].

In [2], it is proved that this problem is NP-hard even if there are only two paths per connection. In their proof, delay and capacity constraints are used. The complexity of the problem is unknown when only considering delay constraints. We formally present our problem in Section 1.1. We then describe a simple example in Section 1.2 and preliminary results in Section 1.3. We present our contributions in Section 1.4.

1.1. Problem formulation and model

Let $G=(V,E)$ be a connected undirected graph (that represents a network) with a coefficient α_e for each edge $e \in E$. Let $\{(s_1,t_1),\dots,(s_m,t_m)\}$ be a set of m source destination pairs (connections). Let $\mathcal{P}=\{P_1,\dots,P_m\}$ be a set of m paths in G. For all $i=1,\dots,m$, P_i is the path, between the source s_i and the destination t_i, allocated to connection (s_i,t_i). Without loss of generality, we suppose that there is a unique path for each connection. Indeed, we can have several connections involving the same pair of nodes but with different associated paths. We denote by x_i the flow through the path P_i for all $i=1,\dots,m$. We suppose that the delay τ_e for crossing an edge $e \in E$ is proportional to the total flow $\sum_{i:e\in E(P_i)} x_i$ crossing the edge e, i.e. $\tau_e = \alpha_e \sum_{i:e\in E(P_i)} x_i$. Let $\lambda > 0$. For all $i=1,\dots,m$, we require for path P_i that, if $x_i > 0$, then the end-to-end delay $\sum_{e\in P_i} \tau_e$ is at most λ. By scaling the coefficients of the α_e, we can make this bound λ equal to one. Using the notation $\beta_{i,j} := \sum_{e\in E(P_i)\cap E(P_j)} \alpha_e$, these constraints can be written as follows: $\sum_{i=1}^m \beta_{i,j} x_i \le 1$. A multicommodity flow $x=(x_1,\dots,x_m)$ is *admissible* if the following latency requirement is satisfied: for all $j=1,\dots,m$ such that $x_j > 0$, then $\sum_{e\in E(P_j)} \tau_e \le 1$. Furthermore, for all $i=1,\dots,m$, if $x_i > 0$, then we say that connection (s_i,t_i) and path P_i are *active* (*inactive* otherwise). The Maximum Flow under Delay Constraint problem (FDC) consists in finding among the solutions satisfying these constraints a solution of maximum value $\sum_{i=1}^m x_i$, i.e.:

$$\begin{cases} \text{Max} & \sum_{i=1}^m x_i \\ & \sum_{i=1}^m \beta_{i,j} x_i \le 1 \quad j=1,\dots,m \quad x_j > 0 \\ & x_i \ge 0 \quad\quad i=1,\dots,m. \end{cases}$$

The hardness of this problem comes from the choice of the active paths. Indeed, if we are given a set of active paths $\mathcal{P}^* \subseteq \mathcal{P}$ in some optimal solution, then the problem becomes polynomial since it reduces to solving the linear program $LP(\mathcal{P}^*)$. Without loss of generality let $\mathcal{P}^* = \{P_1,\dots,P_{m^*}\}$ with $m^* \le m$.

$$LP(\mathcal{P}^*) \begin{cases} \text{Max} & \sum_{i=1}^{m^*} x_i \\ & \sum_{i=1}^{m^*} \beta_{i,j} x_i \le 1 \quad j=1,\dots,m^* \\ & x_i \ge 0 \quad\quad i=1,\dots,m^*. \end{cases}$$

The dual of this linear program is:

$$DLP(\mathcal{P}^*) \begin{cases} \text{Min} & \sum_{j=1}^{m^*} y_j \\ & \sum_{j=1}^{m^*} \beta_{i,j} y_j \ge 1 \quad i=1,\dots,m^* \\ & y_j \ge 0 \quad\quad j=1,\dots,m^*. \end{cases}$$

Note that $LP(\mathcal{P}^*)$ and its dual $DLP(\mathcal{P}^*)$ differ only by the sense of inequalities and the direction of optimization. In particular, if the system of equations $\sum_{i=1}^{m^*} \beta_{i,j} x_i = 1$, $j=1,\dots,m^*$, has a solution then this solution is optimal for the primal and for the dual since it satisfies both primal and dual complementary slackness conditions.

Please cite this article in press as: P. Bonami et al., Maximum flow under proportional delay constraint, Theoret. Comput. Sci. (2017), http://dx.doi.org/10.1016/j.tcs.2017.05.034