

Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

TOBIAS HARKS¹ STEFAN HEINZ² MARC E. PFETSCH

Competitive Online Multicommodity Routing

¹Technical University Berlin, Institute of Mathematics, Strasse des 17. Juni 136,
10623 Berlin, Germany.

²Supported by the DFG Research Center MATHEON *Mathematics for key technologies*
in Berlin.

COMPETITIVE ONLINE MULTICOMMODITY ROUTING

TOBIAS HARKS, STEFAN HEINZ*, AND MARC E. PFETSCH

ABSTRACT. In this paper we study online multicommodity routing problems in networks, in which commodities have to be routed sequentially. The flow of each commodity can be split on several paths. Arcs are equipped with load dependent price functions defining routing costs, which have to be minimized. We discuss a greedy online algorithm that routes each commodity by minimizing a convex cost function that only depends on the demands previously routed. We present a competitive analysis of this algorithm showing that for affine linear price functions this algorithm is $\frac{4K^2}{(1+K)^2}$ -competitive, where K is the number of commodities. For the single-source single-destination case, this algorithm is optimal. Without restrictions on the price functions and network, no algorithm is competitive. Finally, we investigate a variant in which the demands have to be routed unsplitably.

1. INTRODUCTION

In online multicommodity routing problems, commodities of a multicommodity flow have to be routed sequentially in a network. The cost of a flow is determined by dynamic load dependent price functions on links. The commodities arrive sequentially in time and have to be routed with lowest cost. We make four crucial assumptions: (i) demands for commodities are revealed in an *online* fashion; (ii) demands can be split along several paths; (iii) once a demand is routed, no rerouting is allowed; (iv) the routing cost on an arc is given by the integral over the arc flow with respect to the corresponding price function. Since at the time of routing a commodity, future demands are not known, this yields an online optimization problem that we call the *Online Multicommodity Routing Problem* (ONLINEMCRP).

We study an online greedy algorithm SEQ that routes a newly revealed commodity by solving a convex optimization problem that only depends on the previously routed demands. We investigate cases in which SEQ is *competitive*, that is, its cost is at most a constant factor larger than the cost of an offline solution for which all commodities are known. We see SEQ and the ONLINEMCRP as a first step towards an analytical methodology for the following practical application.

The problem under investigation arises in an inter-domain Quality of Service (QoS) market, where multiple service providers offer network resources (capacity) to enable Internet traffic with specific QoS constraints, see for example Yahaya and Suda [15] and Yahaya, Harks, and Suda [14]. In such a market, each service provider advertises prices for resources that he wants to

A preliminary version of this paper appeared in [11].

*Supported by the DFG Research Center MATHEON *Mathematics for key technologies* in Berlin.

sell. Buying providers reserve capacity along paths to route demand (coming from own customers) from source to destination via domains of other providers. The routing of a demand along paths is fixed by establishing a binding contract between the source domain and all domains along the paths. Prices in the market, however, are valid for a predefined bundle size, that is, routing flow of bundle size prompts an update of arc prices. In the limiting case, where the bundle size tends to zero, the routing cost on an arc is given by the integral over the arc flow with respect to the corresponding price function.

In [15, 14], a shortest path routing is introduced and investigated via simulations on real world networks and traffic demands. Their work provides empirical evidence for the efficiency of this working mechanism in such an inter-domain QoS market. The results in this paper provide analytical evidence. We show that for splittable demand and affine linear price functions the corresponding market is stable in the following sense: The cost of the greedy online algorithm SEQ does not deviate too much from the best possible outcome.

Multicommodity routing problems have been studied in the context of traffic engineering, see Fortz and Thorup [7, 8]. There, the goal is to route given demands subject to capacity constraints in order to minimize a convex load dependent penalty function. In this setting, a central planner has full knowledge of all demands, which is not the case in our approach.

Another related line of research is the investigation of efficient routing in decentralized noncooperative systems. This has been extensively studied using game theoretic concepts, cf. Roughgarden and Tardos [12], Correa, Schulz, and Stier Moses [4], Altman, Basar, Jimenez, and Shimkin [1], and the references therein. In these works the efficiency of Nash equilibria are studied. Hence, rerouting of demands is allowed in this context. In our model, once a routing decision has been made, it remains unchanged.

The main topic in online routing has been call admission control problems. An overview article about these problems is given by Leonardi in [6]. Perhaps closest to our work is the paper by Awerbuch, Azar, and Plotkin [2], where online routing algorithms are presented to maximize throughput under the assumption that routings are irrevocable. They, however, restrict the analysis to single path routing and present competitive bounds that depend on the number of nodes in the network.

In this paper, we first show that no online algorithm for the ONLINE-MCRP is competitive for general networks and price functions. If the price functions and network are restricted, however, one can obtain competitive results. The main result of this paper is that for affine linear price functions the greedy online algorithm SEQ is $\frac{4K^2}{(1+K)^2}$ -competitive, where K is the number of commodities; see Section 3.1. Furthermore, we prove in this case a lower bound of $\frac{4}{3}$ on the competitive ratio for any deterministic online algorithm. For SEQ, we prove a lower bound of $\frac{2K-1}{K}$. If we restrict the structure of the network to have a single-source and single-destination only, Section 3.2 shows that SEQ is optimal, i.e. is 1-competitive.

We also study the variant of the ONLINEMCRP in which the demands have to be routed unsplittably. In Section 4, we prove that the corresponding

offline problem is NP-hard, show that in general no competitive deterministic online algorithm exists, and present a lower bound of 2 on the competitive ratio for any deterministic online algorithm if the price functions are linear.

These results are preceded by a problem description together with optimality conditions for the convex problems that have to be solved by SEQ and to determine an optimal offline solution (Section 2). We close with further comments and open questions in Section 5.

2. PROBLEM DESCRIPTION

An instance of the *Online Multicommodity Routing Problem* (ONLINE-MCRP) consists of a directed network $D = (V, A)$ and nondecreasing continuous price functions $p_a : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ for each link $a \in A$. These functions define the price of reserving capacity on a link depending on the current load, as described below. Furthermore, a sequence $\sigma = 1, \dots, K$ of commodities must be routed one after the other. We assume that $K \geq 2$ and denote the set of commodities by $[K] := \{1, \dots, K\}$. Each commodity $k \in [K]$ has a demand $d_k > 0$ that has to be routed from a source $s_k \in V$ to a destination $t_k \in V$. To shorten notation we use the following convention throughout the paper: When we speak of a sequence $\sigma = 1, \dots, K$ of commodities, we refer to the full specification $(d_1, s_1, t_1), \dots, (d_K, s_K, t_K)$.

The routing decision for commodity k is *online*, that is, it only depends on the routings of commodities $1, \dots, k-1$. Once a commodity has been routed it remains unchanged.

A routing assignment, or *flow*, for commodity $k \in [K]$ is a nonnegative vector $\mathbf{f}^k \in \mathbb{R}_+^A$. This flow is *feasible* if for all $v \in V$

$$\sum_{a \in \delta^+(v)} f_a^k - \sum_{a \in \delta^-(v)} f_a^k = \gamma(v), \quad (1)$$

where $\delta^+(v)$ and $\delta^-(v)$ are the arcs leaving and entering v , respectively; furthermore, $\gamma(v) = d_k$ if $v = s_k$, $\gamma(v) = -d_k$ if $v = t_k$, and $\gamma(v) = 0$ otherwise. Note that splitting of demands is allowed. We say that $(\mathbf{f}^1, \dots, \mathbf{f}^K)$ form a *multicommodity flow*.

An alternative formulation uses a *path flow* for each commodity $k \in [K]$. Let \mathcal{P}_k be the set of all paths from s_k to t_k in D . A path flow is a nonnegative vector $(f_P^k)_{P \in \mathcal{P}_k}$. The corresponding flow on link $a \in A$ for commodity $k \in [K]$ is then

$$f_a^k := \sum_{P \ni a} f_P^k.$$

We define \mathcal{F}_k with $k \in [K]$ to be the set of vectors $(\mathbf{f}^1, \dots, \mathbf{f}^k)$ such that \mathbf{f}^i is a feasible flow for commodity $i = 1, \dots, k$. If $(\mathbf{f}^1, \dots, \mathbf{f}^k) \in \mathcal{F}_k$, we say that it is *feasible* for commodities $1, \dots, k$. The entire flow for the sequence $1, \dots, K$ of commodities is denoted by $\mathbf{f} = (\mathbf{f}^1, \dots, \mathbf{f}^K)$. Furthermore, the cost of a flow f_a^k on link $a \in A$ of commodity k is defined by

$$C_a^k(f_a^k; f_a^1, \dots, f_a^{k-1}) = \int_0^{f_a^k} p_a \left(\sum_{i=1}^{k-1} f_a^i + z \right) dz. \quad (2)$$

Note that $C_a^k(\cdot)$ is a convex function, because p_a is nondecreasing.

Remark 2.1. This cost function can be obtained as the limiting case of a single path routing: Assume that demand d_k is split into N equal pieces and the pieces are routed consecutively along a single path. The cost of this path is obtained by evaluating p_a at the flow on link a arising from the previous routings. Let z_a^ℓ be the flow on arc $a \in A$ arising from piece $\ell \in [N]$, i.e., $z_a^\ell = \frac{d_k}{N}$ if a is on the path and $z_a^\ell = 0$ otherwise. Then we have:

$$\lim_{N \rightarrow \infty} \sum_{n=1}^N p_a \left(\sum_{i=1}^{k-1} f_a^i + \sum_{\ell=1}^{n-1} z_a^\ell \right) z_a^n = C_a^k(f_a^k; f_a^1, \dots, f_a^{k-1}),$$

where $f_a^k := \sum_{n=1}^N z_a^n$ is a feasible flow for commodity k . Hence, the integral represents the fact that an infinitesimal amount of flow increases the price for each consecutive piece.

Given flows $\mathbf{f}^1, \dots, \mathbf{f}^{k-1}$, the cost for flow \mathbf{f}^k is

$$C^k(\mathbf{f}^k; \mathbf{f}^1, \dots, \mathbf{f}^{k-1}) := \sum_{a \in A} C_a^k(f_a^k; f_a^1, \dots, f_a^{k-1}),$$

and the total cost over all commodities is defined by

$$C(\mathbf{f}) = \sum_{k=1}^K C^k(\mathbf{f}^k; \mathbf{f}^1, \dots, \mathbf{f}^{k-1}).$$

In this paper, we study the greedy online algorithm SEQ that for a given sequence $\sigma = 1, \dots, K$ sequentially routes the requested demands with minimum cost. Thus, it solves for every $k \in [K]$ the following convex program

$$\begin{aligned} \min \quad & C^k(\mathbf{f}^k; \mathbf{f}^1, \dots, \mathbf{f}^{k-1}) \\ \text{s.t.} \quad & \sum_{a \in \delta^+(v)} f_a^k - \sum_{a \in \delta^-(v)} f_a^k = \gamma(v) & \forall v \in V \\ & f_a^k \geq 0 & \forall a \in A, \end{aligned} \quad (3)$$

where the vectors $\mathbf{f}^1, \dots, \mathbf{f}^{k-1}$ are fixed by solving the first $k-1$ problems. Optimal solutions to this problem are not necessarily unique. Problem (3) can be efficiently solved within arbitrary precision in polynomial time (see Grötschel, Lovász, and Schrijver [10]). Note that SEQ always produces a feasible flow.

Using the relation

$$\frac{\partial C^k}{\partial f_a^k}(\mathbf{f}^k) = p_a \left(\sum_{i=1}^k f_a^i \right),$$

we state in the following lemma necessary and sufficient optimality conditions of the above K problems.

Lemma 2.2. *A feasible flow $\mathbf{f} = (\mathbf{f}^1, \dots, \mathbf{f}^K) \in \mathcal{F}_K$ solves (3) if and only if for all $k \in [K]$ the following two equivalent conditions are satisfied:*

$$i) \quad \sum_{a \in A} p_a \left(\sum_{i=1}^k f_a^i \right) (f_a^k - x_a^k) \leq 0 \quad \begin{array}{l} \text{for all feasible flows } \mathbf{x}^k \\ \text{for commodity } k, \end{array} \quad (4)$$

$$ii) \quad \sum_{a \in P} p_a \left(\sum_{i=1}^k f_a^i \right) \leq \sum_{a \in Q} p_a \left(\sum_{i=1}^k f_a^i \right) \quad \begin{array}{l} \text{for all } P, Q \in \mathcal{P}_k, \\ P \text{ flow carrying w.r.t. } \mathbf{f}^k. \end{array} \quad (5)$$

The proof is based on the first order optimality conditions and the convexity of $C^k(\cdot)$, see Dafermos and Sparrow [5].

An *optimal offline flow* is given by a solution \mathbf{f}^* of the following convex optimization problem:

$$\begin{aligned} \min \quad & C(\mathbf{f}) \\ \text{s.t.} \quad & \sum_{a \in \delta^+(v)} f_a^k - \sum_{a \in \delta^-(v)} f_a^k = \gamma(v) \quad \forall v \in V, k \in K \\ & f_a^k \geq 0 \quad \forall a \in A, k \in K. \end{aligned} \quad (6)$$

We denote by $\text{OPT}(\sigma)$ the optimal value of this convex problem.

Using the relation

$$\frac{\partial C}{\partial f_a^k}(\mathbf{f}) = p_a \left(\sum_{i=1}^K f_a^i \right),$$

the necessary and sufficient optimality conditions of the above problem are given in the following lemma.

Lemma 2.3. *A feasible flow $\mathbf{f} = (\mathbf{f}^1, \dots, \mathbf{f}^K) \in \mathcal{F}_K$ solves (6) if and only if for all $k \in [K]$ the following two equivalent conditions are satisfied:*

$$i) \quad \sum_{a \in A} p_a \left(\sum_{i=1}^K f_a^i \right) (f_a^k - x_a^k) \leq 0 \quad \begin{array}{l} \text{for all feasible flows } \mathbf{x}^k \\ \text{for commodity } k, \end{array} \quad (7)$$

$$ii) \quad \sum_{a \in P} p_a \left(\sum_{i=1}^K f_a^i \right) \leq \sum_{a \in Q} p_a \left(\sum_{i=1}^K f_a^i \right) \quad \begin{array}{l} \text{for all } P, Q \in \mathcal{P}_k, \\ P \text{ flow carrying w.r.t. } \mathbf{f}^k. \end{array} \quad (8)$$

Note that the only difference to the optimality conditions in Lemma 2.2 is the summation in the price function up to commodity K instead of k . This reflects the offline aspect since all demands are known. For the proof see again Dafermos and Sparrow [5].

For a given sequence of commodities $\sigma = 1, \dots, K$ and a solution \mathbf{f} produced by an online algorithm ALG, we denote by $\text{ALG}(\sigma) = C(\mathbf{f})$ its cost. The online algorithm ALG is called (strictly) *c-competitive*, if the cost of ALG is never larger than c times the cost of an optimal offline solution. The *competitive ratio* of ALG is the infimum over all $c \geq 1$ such that ALG is c -competitive, see for instance Borodin and El-Yaniv [3] and Fiat and Woeginger [6].

Remark 2.4. If the price functions $p_a(z)$ are constant for every arc $a \in A$, the algorithm SEQ is optimal for the offline problem, i.e., its competitive ratio is 1. This holds because in this case the routing problems are independent




Figure 1: Graph construction for the proofs of Propositions 3.1, 4.4, and 4.5.

from each other. In fact, each routing decision is just a shortest path problem with respect to the constant costs. Furthermore, the offline problem is a min-cost flow problem without capacity constraints. Hence, both problems can be solved more efficiently than in the general case.

Clearly, also in the case $K = 1$, the competitive ratio of SEQ is 1.

3. COMPETITIVE ANALYSIS OF SEQ

First, we show that there exists no competitive deterministic online algorithm, if neither the network nor the price functions are restricted.

Proposition 3.1. *In general, there exists no competitive deterministic online algorithm for the ONLINEMCRP.*

Proof. Consider the network depicted in Figure 1. For all arcs a in the network, the price function is set to $p_a(z) = m \cdot z^{m-1}$ with $m > 2$. Let ALG be an arbitrary deterministic online algorithm. The first commodity of sequence σ has demand $d_1 = 1$ and has to be routed from node $s_1 = 1$ to node $t_1 = 4$. There are two possible paths for this commodity: path $P_1 = (1, 2, 4)$ and path $P_2 = (1, 3, 4)$. Because of symmetry, we can assume that ALG sends a flow of $\frac{1}{2} \leq \alpha \leq 1$ over path P_1 and $(1 - \alpha)$ along path P_2 . Now commodity 2 arises with demand $d_2 = 1$, source $s_2 = 1$, and target $t_2 = 2$. Algorithm ALG has to route this demand on the only possible path $P_3 = (1, 2)$. For this sequence σ , ALG produces a total cost of

$$\begin{aligned} \text{ALG}(\sigma) &= 2 \cdot \alpha^m + 2 \cdot (1 - \alpha)^m + \int_0^1 m(\alpha + z)^{m-1} dz \\ &= 2 \cdot \alpha^m + 2 \cdot (1 - \alpha)^m + (\alpha + 1)^m - \alpha^m. \end{aligned}$$

Routing the first commodity completely over path P_2 and the second over path P_3 leads to the total cost $2 \cdot 1^m + 1^m = 3 \geq \text{OPT}(\sigma)$. Letting m tend to infinity shows that in this case ALG is not competitive. \square

Despite the negative result of Proposition 3.1, we obtain competitive results in the following two sections. We first restrict the price functions to be affinely linear and then study networks with a single-source and single-destination.

3.1. AFFINELY LINEAR PRICE FUNCTIONS

In this section we will assume that the price functions are affinely linear and show that SEQ is $\frac{4K^2}{(1+K)^2}$ -competitive.

For affinely linear price functions $p_a(z) = q_a \cdot z + r_a$ with $q_a \geq 0$, $r_a \geq 0$ for $a \in A$, we have for a feasible flow $(\mathbf{f}^1, \dots, \mathbf{f}^k)$

$$C_a^k(\mathbf{f}^k; \mathbf{f}^1, \dots, \mathbf{f}^{k-1}) = q_a \left(\sum_{i=1}^{k-1} f_a^i + \frac{1}{2} f_a^k \right) f_a^k + r_a f_a^k.$$

It follows from the optimality conditions (4) that if $(\mathbf{f}^1, \dots, \mathbf{f}^k)$ is generated by SEQ, we have

$$\sum_{a \in A} \left(q_a \sum_{i=1}^k f_a^i + r_a \right) (f_a^k - x_a^k) \leq 0, \quad (9)$$

for all feasible flows \mathbf{x}^k .

Theorem 3.2. *If the price functions are affinely linear, SEQ is $\frac{4K^2}{(1+K)^2}$ -competitive for the ONLINEMCRP.*

Proof. We use the following useful relation at several places in the proof:

$$\sum_{k=1}^K \sum_{i=1}^K f_a^i f_a^k = 2 \sum_{k=1}^K \left(\sum_{i=1}^{k-1} f_a^i + \frac{1}{2} f_a^k \right) f_a^k. \quad (10)$$

Let $(\mathbf{x}_1, \dots, \mathbf{x}_K) \in \mathcal{F}_K$ be any feasible flow and let $(\mathbf{f}_1, \dots, \mathbf{f}_K) \in \mathcal{F}_K$ be the solution produced by SEQ. We start by considering the following inequality for arbitrary nonnegative real values α, β satisfying $1 \leq \alpha \leq \beta \leq 2$:

$$\begin{aligned} 0 &\leq \left(\alpha \sum_{k=1}^K f_a^k - \beta \sum_{k=1}^K x_a^k \right)^2 \\ &= \alpha^2 \sum_{k=1}^K \sum_{i=1}^K f_a^i f_a^k - 2\alpha\beta \sum_{k=1}^K \sum_{i=1}^K f_a^i x_a^k + \beta^2 \sum_{k=1}^K \sum_{i=1}^K x_a^i x_a^k. \end{aligned}$$

Using (10) for the first and last term, multiplying with q_a , and adding over all arcs yields:

$$\begin{aligned} 0 &\leq \sum_{a \in A} q_a \left(2\alpha^2 \sum_{k=1}^K \left(\sum_{i=1}^{k-1} f_a^i + \frac{1}{2} f_a^k \right) f_a^k - 2\alpha\beta \sum_{k=1}^K \sum_{i=1}^K f_a^i x_a^k + \right. \\ &\quad \left. + 2\beta^2 \sum_{k=1}^K \left(\sum_{i=1}^{k-1} x_a^i + \frac{1}{2} x_a^k \right) x_a^k \right). \quad (11) \end{aligned}$$

For the next step, consider the inequality

$$\begin{aligned} 0 &\leq \sum_{a \in A} \sum_{k=1}^K \left((2\alpha^2 - \frac{2\alpha\beta}{K}) r_a f_a^k + (2\beta^2 - 2\alpha\beta) r_a x_a^k \right) \\ &= \sum_{a \in A} \sum_{k=1}^K \left(2\alpha^2 r_a f_a^k - 2\alpha\beta r_a x_a^k + 2\beta^2 r_a x_a^k \right) - \frac{2\alpha\beta}{K} \sum_{a \in A} \sum_{k=1}^K r_a f_a^k. \quad (12) \end{aligned}$$

This inequality holds, because $K \geq 2$ and hence

$$2\alpha^2 - \frac{2\alpha\beta}{K} \geq 2\alpha^2 - \alpha\beta \geq 0,$$

since $1 \leq \alpha \leq \beta \leq 2$ and hence $2\alpha - \beta \geq 0$. Furthermore, $2\beta^2 - 2\alpha\beta \geq 2\beta^2 - 2\beta^2 = 0$. Adding Inequality (12) to (11) leads to:

$$\begin{aligned} 0 &\leq 2\alpha^2 C(\mathbf{f}) - 2\alpha\beta \sum_{a \in A} \sum_{k=1}^K (q_a \sum_{i=1}^k f_a^i + r_a) x_a^k + 2\beta^2 C(\mathbf{x}) \\ &\quad - \frac{2\alpha\beta}{K} \sum_{a \in A} \sum_{k=1}^K r_a f_a^k. \end{aligned}$$

We drop part of the second term and apply (9):

$$\begin{aligned} 0 &\leq 2\alpha^2 C(\mathbf{f}) - 2\alpha\beta \sum_{a \in A} \sum_{k=1}^K (q_a \sum_{i=1}^k f_a^i + r_a) f_a^k + 2\beta^2 C(\mathbf{x}) \\ &\quad - \frac{2\alpha\beta}{K} \sum_{a \in A} \sum_{k=1}^K r_a f_a^k \\ &= (2\alpha^2 - 2\alpha\beta) C(\mathbf{f}) - \alpha\beta \sum_{a \in A} q_a \sum_{k=1}^K f_a^k f_a^k + 2\beta^2 C(\mathbf{x}) \\ &\quad - \frac{2\alpha\beta}{K} \sum_{a \in A} \sum_{k=1}^K r_a f_a^k. \end{aligned}$$

Using the inequality of Cauchy-Schwarz and (10) yields:

$$\begin{aligned} 0 &\leq (2\alpha^2 - 2\alpha\beta) C(\mathbf{f}) + 2\beta^2 C(\mathbf{x}) - \frac{\alpha\beta}{K} \sum_{a \in A} q_a \left(\sum_{k=1}^K f_a^k \right)^2 - \frac{2\alpha\beta}{K} \sum_{a \in A} \sum_{k=1}^K r_a f_a^k \\ &= (2\alpha^2 - 2\alpha\beta) C(\mathbf{f}) + 2\beta^2 C(\mathbf{x}) - \frac{2\alpha\beta}{K} C(\mathbf{f}). \end{aligned}$$

This is equivalent to:

$$C(\mathbf{f}) \leq \frac{\beta^2}{-\alpha^2 + \alpha\beta + \frac{\alpha\beta}{K}} C(\mathbf{x}).$$

We now take $\alpha = (1 + \frac{1}{K})$ and $\beta = 2$ to get $C(\mathbf{f}) \leq \frac{4K^2}{(1+K)^2} C(\mathbf{x})$, i.e. the desired bound (if we let \mathbf{x} an optimal offline solution). \square

Remark 3.3. The parameters α and β in the previous proof are optimal solutions to the following minimization problem:

$$\min_{1 \leq \alpha \leq \beta \leq 2} \frac{\beta^2}{-\alpha^2 + \alpha\beta + \frac{\alpha\beta}{K}}.$$

We do not know whether the result in Theorem 3.2 is tight. The best known lower bound for *any* deterministic online algorithm is the following.

Proposition 3.4. *In case of linear cost functions, no deterministic online algorithm for the ONLINEMCRP is c -competitive for any $c < \frac{4}{3}$.*

Proof. Consider the network displayed in Figure 2. Each arc a leaving node 1 has the same price function $p_a(z) = 4z$. All the other arcs (leading to node 5) have price function $p_a(z) = 0$. Let ALG be an arbitrary deterministic online




Figure 2: Graph construction for the proof of Proposition 3.4.

algorithm. The first commodity with demand 1 has to be routed from $s_1 = 1$ to $t_1 = 5$.

Assume the algorithm behaves like SEQ. This means that the demand gets evenly divided into three pieces: one third is routed over path $P_1 = (1, 2, 5)$, another over path $P_2 = (1, 3, 5)$, and the final third over path $P_3 = (1, 4, 5)$ (compare Lemma 2.2). We then reveal commodity 2 with demand 1 between nodes 1 and 2. The algorithm ALG has to route this demand on the only possible path $P_4 = (1, 2)$. Therefore, the cost of ALG for this sequence σ is:

$$\text{ALG}(\sigma) = \text{SEQ}(\sigma) = 3 \cdot 4 \cdot \left(\frac{1}{2} \cdot \frac{1}{3}\right) \cdot \frac{1}{3} + 4 \cdot \left(\frac{1}{3} + \frac{1}{2} \cdot 1\right) \cdot 1 = 4.$$

An optimal offline solution is to route half of commodity 1 over path P_2 , the other half over path P_3 , and commodity 2 along P_4 (compare Lemma 2.3). Therefore,

$$\text{OPT}(\sigma) = 2 \cdot 4 \cdot \left(\frac{1}{2} \cdot \frac{1}{2}\right) \cdot \frac{1}{2} + 4 \cdot \left(\frac{1}{2} \cdot 1\right) \cdot 1 = 3.$$

This leads to

$$\frac{\text{ALG}(\sigma)}{\text{OPT}(\sigma)} = \frac{4}{3}.$$

If ALG does not behave like SEQ for the first commodity, we can assume by symmetry that ALG routes a flow of $\alpha > \frac{1}{3}$ over path P_1 . Hence, a demand of $1 - \alpha$ is routed over path P_2 and P_3 . The best way to do this is to route $(1 - \alpha)/2$ over each path. Then commodity 2 is released as above, again leaving no routing choice. The cost of ALG for this sequence σ is

$$\text{ALG}(\sigma) \geq 4 \cdot \left(\frac{1}{2} \cdot \alpha\right) \cdot \alpha + 2 \cdot 4 \cdot \left(\frac{1}{2} \cdot \frac{(1-\alpha)}{2}\right) \cdot \frac{(1-\alpha)}{2} + 4 \cdot \left(\alpha + \frac{1}{2} \cdot 1\right) \cdot 1 > 4.$$

since $\alpha > \frac{1}{3}$. Because $\text{OPT}(\sigma) = 3$, we have

$$\frac{\text{ALG}(\sigma)}{\text{OPT}(\sigma)} > \frac{4}{3}.$$

Therefore, ALG cannot have a competitive ratio less than $\frac{4}{3}$. \square

The following proposition provides an improved lower bound for SEQ.

Proposition 3.5. *In case of affine linear cost functions the online algorithm SEQ for ONLINEMCRP has a competitive ratio greater or equal to $\frac{2K-1}{K}$.*

Proof. Consider the network shown in Figure 3 with the following price functions: $p_{(s_i, s)}(z) = 0$, $p_{(t, t_i)}(z) = 0$, $p_{(s_i, t_i)}(z) = i$, and $p_{(s, t)}(z) = z$, for $i = 1, \dots, n$. For $i = 1, \dots, n$, we consecutively release a demand of size 1




Figure 3: Graph construction for the proof of Proposition 3.5.

from s_i to t_i . Using Lemma 2.2, we see that SEQ routes every demand over arc (s, t) . The cost for these n demands is:

$$\frac{1}{2} \cdot 1 + (1 + \frac{1}{2}) \cdot 1 + \dots + (n - 1 + \frac{1}{2}) \cdot 1 = \sum_{i=1}^n \frac{2i-1}{2} = \frac{1}{2}n^2.$$

The $(n + 1)$ -st demand of size $d \geq 1$ is released from s to t and incurs the following cost:

$$(n + \frac{1}{2}d)d = nd + \frac{1}{2}d^2.$$

Thus, the total cost for SEQ is given by:

$$\text{SEQ}(\sigma) = \frac{1}{2}(n^2 + 2nd + d^2).$$

In an optimal offline solution the first n demands are routed along the arcs (s_i, t_i) and the last demand is routed on (s, t) . Hence, the total cost is:

$$\text{OPT}(\sigma) = \sum_{i=1}^n i + \frac{1}{2}d^2 = \frac{n(n+1)}{2} + \frac{1}{2}d^2.$$

Setting $d = n + 1$ and substituting $n = K - 1$ yields

$$\frac{\text{SEQ}(\sigma)}{\text{OPT}(\sigma)} = \frac{n^2 + 2nd + d^2}{n^2 + n + d^2} = \frac{1 + 2n}{1 + n} = \frac{2K - 1}{K},$$

which proves the theorem. \square

Remark 3.6. The value $d = n + 1$ solves the following optimization problem with respect to d :

$$\max_{d \geq 1} \frac{n^2 + 2nd + d^2}{n^2 + n + d^2} = \frac{1 + 2n}{1 + n}.$$

This yields the best lower bound for the network in the proof of Proposition 3.5.

Figure 4 illustrates the lower and upper bounds on the competitive ratio of the online algorithm SEQ for affine linear price functions. The bounds asymptotically converge to 2 and 4, respectively, for K tending to infinity.

A characteristic of SEQ is that it splits demand and distributes it onto several paths. We now show that only algorithms that split demand can be competitive.

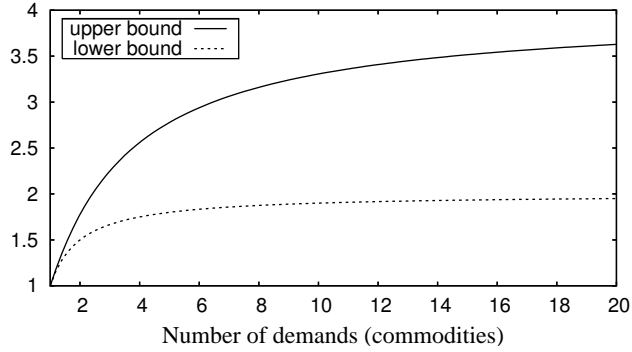


Figure 4: Upper bound $\frac{4K^2}{(K+1)^2}$ versus lower bound $\frac{2K-1}{K}$ on the competitive ratio of SEQ for affine linear price functions.

Proposition 3.7. *Every deterministic online algorithm for the ONLINE-MCRP that routes all demands unsplittably is not competitive, even for linear cost functions.*

Proof. Consider the network shown in Figure 5. This network contains $n+2$ nodes and n paths from node s to node t . The price functions are $p_a(z) = 2z$ for all arcs a . Let ALG be an arbitrary deterministic online algorithm that does not split demands. We consider a single commodity with demand 1 between nodes s and t . Since ALG does not split, the cost of its routing is independent from the chosen path:

$$\text{ALG}(\sigma) = 2 \cdot \left(\frac{1}{2} \cdot 1\right) \cdot 1 + 2 \cdot \left(\frac{1}{2} \cdot 1\right) \cdot 1 = 2.$$

An optimal solution splits the demand into n evenly divided pieces and sends each piece over a different path. This leads to an optimal cost of

$$\text{OPT}(\sigma) = n \left(2 \cdot \left(\frac{1}{2} \cdot \frac{1}{n}\right) \cdot \frac{1}{n} + 2 \cdot \left(\frac{1}{2} \cdot \frac{1}{n}\right) \cdot \frac{1}{n} \right) = n \cdot 2 \cdot \left(\frac{1}{n}\right)^2 = \frac{2}{n}.$$

Therefore, the competitive ratio of ALG is not smaller than n . Since this holds for all $n \in \mathbb{N}$, ALG is not competitive. \square

In Section 4, we further investigate the problem variant in which splitting demand is not allowed.




Figure 5: Graph construction for the proof of Proposition 3.7.

3.2. SINGLE-SOURCE SINGLE-DESTINATION NETWORKS

Restricting the considered networks to only contain a single source and single destination, i.e., $s_1 = \dots = s_k$ and $t_1 = \dots = t_k$, we can show that SEQ is optimal.

Theorem 3.8. *Consider an instance of ONLINEMCRP, in which each commodity has the same source and the same destination. Then, SEQ computes an offline optimal routing.*

Proof. The proof uses induction on the number of commodities K . For the case $K = 1$, the claim follows since by definition SEQ routes one commodity with minimum cost. Therefore, assume that the claim holds for any sequence containing $K - 1$ commodities.

For the sake of contradiction, assume that the flow \mathbf{f} that is generated by SEQ for a given sequence with K commodities is not offline optimal. Hence, this flow does not satisfy the conditions of Lemma 2.3. Therefore, there exist paths $P, Q \in \mathcal{P}$, where P is flow carrying, with

$$\sum_{a \in P} p_a \left(\sum_{i=1}^K f_a^i \right) > \sum_{a \in Q} p_a \left(\sum_{i=1}^K f_a^i \right). \quad (13)$$

By the induction hypothesis the routing computed by SEQ for the first $K - 1$ commodities is optimal and for this routing Lemma 2.3 holds. Inequality (13) is only valid if $f_P^K > 0$. To see this, assume $f_P^K = 0$. Then, it follows that P is flow carrying with respect to the first $K - 1$ commodities. Invoking the optimality conditions in Lemma 2.3 for the first $K - 1$ commodities (induction hypothesis) leads to:

$$\sum_{a \in P} p_a \left(\sum_{i=1}^K f_a^i \right) = \sum_{a \in P} p_a \left(\sum_{i=1}^{K-1} f_a^i \right) \leq \sum_{a \in Q} p_a \left(\sum_{i=1}^{K-1} f_a^i \right) \leq \sum_{a \in Q} p_a \left(\sum_{i=1}^K f_a^i \right),$$

where the last inequality follows because the p_a are nondecreasing. Since this contradicts (13), we have $f_P^K > 0$. In particular, $f_a^K > 0$ for all $a \in P$. For small enough $\varepsilon > 0$, we define the *nonnegative* flow $\mathbf{x} := (\mathbf{f}^1, \dots, \mathbf{f}^{K-1}, \mathbf{x}^K)$ with

$$x_a^K := \begin{cases} f_a^K - \varepsilon & a \in P \setminus Q \\ f_a^K + \varepsilon & a \in Q \setminus P \\ f_a^K & \text{otherwise.} \end{cases}$$

By construction this flow is feasible.

We obtain for the difference of costs:

$$\begin{aligned} C(\mathbf{x}) - C(\mathbf{f}) &= \\ &= \sum_{a \in P \setminus Q} \left[\int_0^{f_a^K - \varepsilon} p_a \left(\sum_{i=1}^{K-1} f_a^i + z \right) dz - \int_0^{f_a^K} p_a \left(\sum_{i=1}^{K-1} f_a^i + z \right) dz \right] + \\ &\quad + \sum_{a \in Q \setminus P} \left[\int_0^{f_a^K + \varepsilon} p_a \left(\sum_{i=1}^{K-1} f_a^i + z \right) dz - \int_0^{f_a^K} p_a \left(\sum_{i=1}^{K-1} f_a^i + z \right) dz \right] \end{aligned}$$




Figure 6: Construction for the proof of Proposition 4.1.

$$\begin{aligned}
&= - \sum_{a \in P \setminus Q} \int_{f_a^K - \varepsilon}^{f_a^K} p_a \left(\sum_{i=1}^{K-1} f_a^i + z \right) dz + \sum_{a \in Q \setminus P} \int_{f_a^K}^{f_a^K + \varepsilon} p_a \left(\sum_{i=1}^{K-1} f_a^i + z \right) dz \\
&= \int_0^\varepsilon \left(- \sum_{a \in P \setminus Q} p_a \left(\sum_{i=1}^K f_a^i + z - \varepsilon \right) + \sum_{a \in Q \setminus P} p_a \left(\sum_{i=1}^K f_a^i + z \right) \right) dz. \quad (14)
\end{aligned}$$

We now define

$$g(z, \varepsilon) := - \sum_{a \in P \setminus Q} p_a \left(\sum_{i=1}^K f_a^i + z - \varepsilon \right) + \sum_{a \in Q \setminus P} p_a \left(\sum_{i=1}^K f_a^i + z \right).$$

By (13) we have $g(0, 0) < 0$. Since $p_a(\cdot)$ is continuous, g is continuous, too. Hence, $g(z, \varepsilon) < 0$ for all z and δ with $0 \leq z, \varepsilon < \delta$, if δ is small enough. Therefore, the right hand side of (14) is strictly smaller than 0. It follows that $C(\mathbf{x}) < C(\mathbf{f})$. This is a contradiction since \mathbf{x} and \mathbf{f} only differ with respect to the last commodity K and SEQ solves problem (3). Hence, SEQ computes an offline optimal solution. \square

4. UNSPLITTABLE ROUTINGS

In this section we study the variant of the ONLINEMCRP with unsplittable routings, i.e., the demand of each commodity has to be routed on a single path. Such a restriction often occurs in practice, for instance in single path routing problems in telecommunication networks. It is possible to formulate a mixed integer convex program for this setting. In contrast to the splittable case, however, the offline problem is NP-hard in this case.

Proposition 4.1. *The offline problem for the ONLINEMCRP with unsplittable routings is NP-hard, even when the price functions are linear.*

Proof. Consider an instance of the *minimum sum of squares problem*, which is NP-complete in the strong sense (see Garey and Johnson [9]). Here, one is given nonnegative integers d_1, \dots, d_K and positive integers $N \leq K$ and J . The question is whether there exists a partition of $[K]$ into N sets A_1, \dots, A_N

such that

$$\sum_{i=1}^N \left(\sum_{k \in A_i} d_k \right)^2 \leq J?$$

For the reduction to the offline problem, we construct a network D with node set $\{s_1, \dots, s_K, u_1, \dots, u_N, t\}$ and the following arcs: For each $k \in [K]$ and $i \in [N]$ we have an arc (s_k, u_i) with price function 0. For each $i \in [N]$ we add an arc $a = (u_i, t)$ with price function $p_a(z) = 2z$; see Figure 6. Furthermore, for $k \in [K]$ there are demands d_k between s_k and t .

We claim that there exists an unsplittable solution to the offline problem of value at most J if and only if the answer to the minimum sum of squares problem is positive. To see this, first assume that A_1, \dots, A_N is the wanted partition. Then if $k \in A_i$, we route commodity k along u_i to t . Using (10), we obtain the following costs:

$$2 \sum_{i=1}^N \sum_{k \in A_i} \left(\sum_{\substack{j \in A_i \\ j < k}} d_j + \frac{1}{2} d_k \right) d_k = \sum_{i=1}^N \sum_{k \in A_i} \sum_{j \in A_i} d_k d_j = \sum_{i=1}^N \left(\sum_{k \in A_i} d_k \right)^2.$$

This proves the forward direction of the claim. Conversely, assume that there exists an unsplittable flow of value J . For $i = 1, \dots, N$, let A_i be the set of indices k whose corresponding demands are routed over the arc (u_i, t) . Again the cost is given as above, which shows that there exists a solution to the minimum sum of squares problem. \square

Remark 4.2. The unsplittable variant of (3) can be computed in polynomial time since it amounts to solving a shortest path problem.

Remark 4.3. When the price functions are constant, both the unsplittable variants of (3) and (6) can be written as (integer) min-cost flow problems. Hence, they can be solved in polynomial time, see e.g. Schrijver [13, Ch. 12].

The following two results show that the additional requirement of unsplittable routings does not improve competitiveness properties of the ONLINE-MCRP. The first is the unsplittable version of Proposition 3.1.

Proposition 4.4. *In general there exists no competitive deterministic online algorithm for the unsplittable variant of the ONLINEMCRP.*

Proof. Consider again the network shown in Figure 1, where each arc a has the price function $p_a(z) = m \cdot z^{m-1}$ for some $m > 2$. Let ALG be an arbitrary deterministic online algorithm. We first reveal a commodity with demand $d_1 = 1$, source $s_1 = 1$, and target $t_1 = 4$. Without loss of generality, we can assume that ALG uses path $P_1 = (1, 2, 4)$ to route this demand. We then release commodity 2 with demand $d_2 = 1$, source $s_2 = 1$, and target $t_2 = 2$. The algorithm ALG has to route this commodity on the single path $P_2 = (1, 2)$. Hence, for this sequence σ , ALG yields the cost

$$\text{ALG}(\sigma) = 2 \cdot 1^m + \int_0^1 m(1+z)^{m-1} dz = 2 + (1+1)^m - 1^m = 1 + 2^m.$$

The optimal cost is $\text{OPT}(\sigma) = 3$, which is achieved by routing commodity 1 over path $P_3 = (1, 3, 4)$ and commodity 2 along path P_2 . Therefore, for m going to infinity it follows that ALG is not competitive. \square

We can also improve the lower bound of Proposition 3.4 from $\frac{4}{3}$ to 2.

Proposition 4.5. *If we consider only linear price functions, no deterministic online algorithm has a competitive ratio less than 2 for the unsplittable variant of the ONLINEMCRP.*

Proof. Consider the network shown in Figure 1, where each link a is equipped with the same price function $p_a(z) = 2z$. Let ALG be an arbitrary deterministic online algorithm. We first reveal commodity 1 with demand $d_1 = 1$, source $s_1 = 1$, and target $t_1 = 4$. Without loss of generality this commodity is routed over path $P_1 = (1, 2, 4)$. Then we release one commodity from node 1 to 2 and one commodity from node 2 to 4. Both have a demand of 1. Since for each of the last two commodities there exists only a single path, the assignment by ALG for this sequence σ leads to a cost of

$$\text{ALG}(\sigma) = 2 \cdot 2 \cdot \left(\frac{1}{2} \cdot 1\right) \cdot 1 + 2 \cdot \left(1 + \frac{1}{2} \cdot 1\right) \cdot 1 + 2 \cdot \left(1 + \frac{1}{2} \cdot 1\right) \cdot 1 = 8.$$

An optimal routing is achieved by routing commodity 1 along path $P_2 = (1, 3, 4)$ and commodity 2 and 3 over their single paths. Since the optimal cost for σ is $\text{OPT}(\sigma) = 4$, the competitive ratio of ALG is at least 2. \square

5. FINAL COMMENTS AND FUTURE RESEARCH

We see the framework introduced in this paper as a first step towards modeling of real world online multicommodity routing problems. In practice, however, there are many more additional requirements. For instance, routings have to consider capacities, which we ignored in our approach. With capacities, however, one can easily construct examples in which any online algorithm does not even produce a feasible solution. Further requirements in practice include path length restrictions and survivability issues. Another important point is that in practice routings are only valid until a given time, after which they disappear. This has effects on the cost for future routings. We plan to study this problem in the future.

Open problems include the question whether other price functions, e.g. quadratic price functions, also allow competitive results. Further restrictions on the network structure might be interesting as well. It is also an open issue whether the competitiveness bound in Theorem 3.2 is tight, and whether there exists a competitive online algorithm for the unsplittable variant of the ONLINEMCRP.

REFERENCES

- [1] E. ALTMAN, T. BASAR, T. JIMENEZ, AND N. SHIMKIN, *Competitive routing in networks with polynomial costs*, IEEE Trans. Automat. Control, 47 (2002), pp. 92–96.
- [2] B. AWERBUCH, Y. AZAR, AND S. PLOTKIN, *Throughput-competitive on-line routing*, in Proc. 34th Annual IEEE Symposium on Foundations of Computer Science (FOCS) 1993, Palo Alto, 1993, pp. 32–40.
- [3] A. BORODIN AND R. EL-YANIV, *Online Computation and Competitive Analysis*, Cambridge University Press, 1998.
- [4] J. R. CORREA, A. S. SCHULZ, AND N. E. STIER MOSES, *Selfish routing in capacitated networks*, Math. Oper. Res., 29 (2004), pp. 961–976.
- [5] S. DAFERMOS AND F. SPARROW, *The traffic assignment problem for a general network*, J. Res. Natl. Bur. Stand., Sect. B, 73 (1969), pp. 91–118.

- [6] A. FIAT AND G. J. WOEGINGER, eds., *Online Algorithms: The State of the Art*, vol. 1442 of LNCS, Springer, 1998.
- [7] B. FORTZ AND M. THORUP, *Optimizing OSPF/IS-IS weights in a changing world*, IEEE J. Sel. Area Comm., 20 (2002), pp. 756–767.
- [8] ———, *Increasing internet capacity using local search*, Computational Optimization and Applications, 29 (2004), pp. 13–48.
- [9] M. R. GAREY AND D. S. JOHNSON, *Computers and Intractability. A Guide to the Theory of NP-Completeness*, W. H. Freeman and Company, New York, 1979.
- [10] M. GRÖTSCHEL, L. LOVÁSZ, AND A. SCHRIJVER, *Geometric Algorithms and Combinatorial Optimization*, vol. 2 of Algorithms and Combinatorics, Springer, Heidelberg, 2nd ed., 1993.
- [11] T. HARKS, S. HEINZ, AND M. E. PFETSCH, *Competitive online multicommodity routing*, in Proc. Workshop on Approximation and Online Algorithms (WAOA) 2006, T. Erlebach and C. Kaklamanis, eds., vol. 4368 of LNCS, Springer, Berlin Heidelberg, 2006, pp. 240–252.
- [12] T. ROUGHGARDEN AND E. TARDOS, *How bad is selfish routing?*, Journal of the ACM, 49 (2002), pp. 236–259.
- [13] A. SCHRIJVER, *Combinatorial Optimization: Polyhedra and Efficiency*, vol. 24 of Algorithms and Combinatorics, Springer, Berlin Heidelberg, 2003.
- [14] A. YAHAYA, T. HARKS, AND T. SUDA, *iREX: Efficient inter-domain QoS automation using economics*, in Proceedings of IEEE Globecom, November 2006.
- [15] A. YAHAYA AND T. SUDA, *iREX: Inter-domain QoS automation using economics*, in Proceedings of IEEE Consumer Communications and Networking Conference, Jan 2006, pp. 96–101.

TECHNICAL UNIVERSITY BERLIN, INSTITUTE OF MATHEMATICS, STRASSE DES 17. JUNI 136,
10623 BERLIN, GERMANY

E-mail address: `harks@math.tu-berlin.de`

ZUSE INSTITUTE BERLIN, TAKUSTR. 7, 14195 BERLIN, GERMANY

E-mail address: `{heinz,pfetsch}@zib.de`