Competitive Online Multicommodity Routing

Tobias Harks · Stefan Heinz · Marc E. Pfetsch

Abstract 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 depends on the previously routed flow. We present a competitive analysis of this algorithm showing that for affine price functions this algorithm is $\frac{4K^2}{(1+K)^2}$-competitive, where K is the number of commodities. For networks with two nodes and parallel arcs, this algorithm is optimal. Without restrictions on the price functions and network, no algorithm is competitive.

We then investigate a variant in which the demands have to be routed unsplittably. In this case, it is NP-hard to compute the offline optimum. The variant of the greedy algorithm that produces unsplittable flows is $(3 + 2\sqrt{2})$-competitive, and we prove a lower bound of 2 for the competitive ratio of any deterministic online algorithm.

T. Harks
Institute of Mathematics, Technische Universität Berlin, Straße des 17. Juni 136, 10623 Berlin, Germany
e-mail: harks@math.tu-berlin.de

S. Heinz (✉)
Zuse Institute Berlin, Takustr. 7, 14195 Berlin, Germany
e-mail: heinz@zib.de

M.E. Pfetsch
Institute for Mathematical Optimization, Technische Universität Braunschweig, Pockelsstr. 14, 38106 Braunschweig, Germany
e-mail: m.pfetsch@tu-bs.de

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 sequen-
tially in time and have to be routed with lowest cost. We make four crucial assump-
tions: (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, which depends only on the previously
routed flows. 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, for instance, in an inter-domain Quality of
Service (QoS) market, in which multiple service providers offer network resources
(capacity) to enable Internet traffic with specific QoS constraints, see for example
Yahaya and Suda [28] and Yahaya, Harks, and Suda [29]. In such a market, each
service provider advertises prices for resources that he wants to sell. Buying providers
reserve capacity along paths to route demand (coming from their 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 are valid for a predefined bundle
size, that is, routing flow of this 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 [28, 29], a shortest path routing is introduced and investigated via simulations
on a real world network and stochastically generated traffic demands. The results
show that the efficiency loss for affine price functions is bounded by approximately
30%. Hence, this provides empirical evidence for the efficiency of the above work-
ing mechanism in such an inter-domain QoS market. In this paper, we provide an-
alytical evidence. We show that for splittable demand and affine 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 engi-
neering in telecommunication networks, see, for intance, Fortz and Thorup [15, 16].
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 planer
has full knowledge of all demands, which is not the case in our approach.

Another related line of research is the investigation of efficient routings in decentralized noncooperative systems. Using game theoretic concepts, Roughgarden and Tardos [25], Correa, Schulz, and Stier Moses [11] and Altman et al. [1] studied the efficiency of Nash equilibria for nonatomic players. For atomic routing games, where players may control a significant part of the entire demand, Roughgarden and Tardos examined the price of anarchy for an unsplittable variant [25]. Awerbuch, Azar, and Epstein [4] and Christodoulou and Koutsoupias [9] derived bounds on the price of anarchy for affine latency functions. Atomic splittable routing games have been studied by Cominetti et al. [10] and Roughgarden [24].

In these works, the efficiency of Nash equilibria is studied. In this context, every player has full information about the demands and flows of the other players. In our model, however, a routing decision has to be made without knowing future demands. Furthermore, once a demand has been routed, this routing remains unchanged.

The closest topic in the online field has been online load balancing in the context of machine scheduling. Awerbuch et al. [3] considered a greedy online load balancing algorithm, where the goal is to minimize the L_2 norm of the aggregated server loads. They proved an upper bound on the competitive ratio of $3 + 2\sqrt{2}$. In the same context, Suri, Toth, and Zhou [26] and Caragiannis et al. [8] studied Nash solutions for every released job and showed that the resulting online algorithm outperforms the greedy strategy of [3]. Their setting, however, is restricted to m parallel arcs and all released jobs have to be assigned to exactly one machine (arc). Furthermore, Awerbuch, Azar, and Plotkin [2] present online routing algorithms that maximize throughput under the assumption that routings are irrevocable. They 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 ONLINEMCRP is competitive for general networks and price functions. If the price functions and networks are restricted, however, one can obtain competitive results. The main result of this paper is that for affine price functions the greedy online algorithm SEQ is $\frac{4K^2}{(1+K)^2}$-competitive, where K is the number of commodities; see Sect. 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 be a two node directed graph with parallel arcs, Sect. 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 Sect. 4, we prove that the corresponding offline problem is NP-hard and show that, in general, no competitive deterministic online algorithm exists. For the unsplittable variant of SEQ and affine price functions, we prove an upper bound of $3 + 2\sqrt{2}$ on the competitive ratio. This result can be seen as a generalization of the result of Awerbuch et al. [3] to multicommodity flow networks. Finally, we present a lower bound of 2 on the competitive ratio for any deterministic online algorithm if the price functions are linear.

In a follow up paper, Harks et al. [21] generalize the current setting: commodities are released in rounds and come with a lifetime; the commodities of a given round can be routed in an arbitrary order. In this context, upper and lower bounds for polynomial price functions with nonnegative coefficients are proved.

As noted by one of the reviewers, Olver [22] and Farzad et al. [13] considered a similar model to ours under a different viewpoint, independently from our work. They proved strengthened results for the lower bound in the splittable and unsplittable case. They showed that the asymptotic upper bound of 4 in the splittable case and $3 + 2\sqrt{2}$ in the unsplittable case are tight. Furthermore, they proved upper bounds for polynomial price functions (similar to [21]).

2 Problem Description

Instances of the *Online Multicommodity Routing Problem* (ONLINEMCRP) consist of a directed network $D = (V, A)$ and nondecreasing continuous price functions p_a : $\mathbb{R}_+ \to \mathbb{R}_+$ 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. There are no capacity bounds in this setting.

Furthermore, a sequence $\sigma = 1, \ldots, K$ of commodities must be routed one after the other. We assume that $K \geq 2$ and denote the set of commodities by $[K] := \{1, \ldots, K\}$. 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 $\sigma = 1, \ldots, K$ of commodities, we refer to the full specification $(d_1, s_1, t_1), \ldots, (d_K, s_K, t_K)$.

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

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

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

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 define $\mathcal{F}_k$ with $k \in [K]$ to be the set of vectors $(\boldsymbol{f}^1, \ldots, \boldsymbol{f}^k)$ such that $\boldsymbol{f}^i$ is a feasible flow for commodity $i = 1, \ldots, k$. If $(\boldsymbol{f}^1, \ldots, \boldsymbol{f}^k) \in \mathcal{F}_k$, we say that it is *feasible* for commodities $1, \ldots, k$. The entire flow for the sequence $1, \ldots, K$ of commodities is denoted by $\boldsymbol{f} = (\boldsymbol{f}^1, \ldots, \boldsymbol{f}^K)$; it is also called a *multicommodity flow*.

The cost of a flow f_a^k on link $a \in A$ of commodity k is defined by

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

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 single paths. The cost of each 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 corresponding path and $z_a^\ell = 0$ otherwise. Then we have:

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

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 subsequent piece.

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

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

and the total cost over all commodities is defined by

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

2.1 Greedy Online Algorithm SEQ

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

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

where the vectors $f^1, \ldots, f^{k-1}$ are fixed by solving the first $k-1$ problems. As a convex program, Problem (3) can be efficiently solved within arbitrary precision in polynomial time (see, e.g., Grötschel, Lovász, and Schrijver [18]). Note that SEQ always produces a feasible flow.

Remark 2.2 If all price functions are strictly increasing, the cost function is strictly convex, hence the optimal solution of (3) is unique. In general, optimal solutions are not necessarily unique, e.g., if some price functions are constant. In this case, SEQ produces an arbitrary optimal solution. In fact, different optimal solutions to (3) can lead to different total costs as shown in the following example.

Fig. 1 Graph for Example 2.3
and the proofs of
Propositions 3.1, 4.3, and 4.4
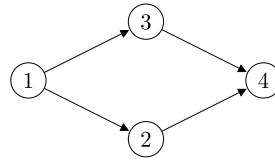
Example 2.3 Consider the network displayed in Fig. 1. All arcs entering node 4 have
a constant price function equal to 0. Arc $(1, 3)$ has a constant price function 1 and arc
$(1, 2)$ is equipped with the price function

$$p(z) = \begin{cases} 1, & 0 \le z \le 1 \\ z, & z > 1. \end{cases}$$

Note that $p(z)$ is continuous and nondecreasing. We are given a sequence of two
commodities of unit size, where the first commodity has to be routed from node 1 to
node 4 and the second from node 1 to node 2. Problem (3) for the first commodity
admits a variety of optimal solutions. More precisely, let $\alpha \in [0, 1]$ be the fraction of
the first demand routed over path $P_1 = (1, 2, 4)$ and $(1 - \alpha)$ be the demand routed
over $P_2 = (1, 3, 4)$. Since the demand of the first commodity is 1, it follows that the
cost is

$$\int_0^{1-\alpha} 1 \, dz + \int_0^{\alpha} 1 \, dz = (1 - \alpha) + \alpha = 1,$$

which is independent of α. The choice of α, however, influences the cost for the
second commodity:

$$\int_{\alpha}^{1} 1 \, dz + \int_1^{1+\alpha} z \, dz = (1 - \alpha) + \frac{1}{2}(1 + \alpha)^2 - \frac{1}{2}.$$

Hence, the choice of different optimal solutions computed by SEQ can lead to differ-
ent total costs.

   Using the relation

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

we state in the following lemma necessary and sufficient optimality conditions of
problems (3) for $k \in [K]$.

Lemma 2.4 *Given flows* $\boldsymbol{f}^1, \ldots, \boldsymbol{f}^{k-1} \in \mathcal{F}_{k-1}$, *a feasible flow* $\boldsymbol{f}^k$ *solves problem* (3)
if and only if the following two equivalent conditions are satisfied:

(i) $\displaystyle\sum_{a \in A} p_a \left(\sum_{i=1}^{k} f_a^i \right)(f_a^k - x_a^k) \le 0$ *for all feasible flows* $\boldsymbol{x}^k$,   (4)
for commodity k,

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

Here, $\mathcal{P}_k$ *is the set of all paths from* s_k *to* t_k *in the network.*

The proof is based on the first order optimality conditions and the convexity of $C^k(\cdot)$, see Dafermos and Sparrow [12] or the book by Boyd and Vandenberghe [6]. The relation of the above characterization to Wardrop equilibria is discussed in the next section.

2.2 Optimal Offline Solutions

An *optimal offline flow*, for a given sequence $\sigma = 1, \ldots, K$, is an optimal solution $\boldsymbol{f}^\star$ of the following convex optimization problem:

$$\min C(\boldsymbol{f})$$

$$\text{s.t.} \quad \sum_{a \in \delta^+(v)} f_a^k - \sum_{a \in \delta^-(v)} f_a^k = \gamma(v), \quad \forall v \in V, \ k \in K,$$ (6)

$$f_a^k \geq 0, \quad \forall a \in A, \ k \in K.$$

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

Using the relation

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

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

Lemma 2.5 *A feasible flow* $\boldsymbol{f} = (\boldsymbol{f}^1, \ldots, \boldsymbol{f}^K) \in \mathcal{F}_K$ *solves* (6) *if and only if for all* $k \in [K]$ *the following two equivalent conditions are satisfied*:

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

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

Note that differently from the optimality conditions in Lemma 2.4, the summation in the price function is 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 [12].

Remark 2.6 An easy calculation shows that the total cost can be written as

$$C(\boldsymbol{f}) = \sum_{a \in A} \int_0^{\sum_{k=1}^{K} f_a^k} p_a(z) \, dz.$$

Therefore, a solution to (6) can be seen as a Wardrop equilibrium [27] in the selfish routing model, since $C(f)$ is a potential function, cf. [23]. Similarly, a solution to (3) for some k is a Wardrop equilibrium with respect to the fixed previously routed flow and without considering the future flows of commodities $k + 1, \ldots, K$.

For a given sequence of commodities $\sigma = 1, \ldots, K$ and a solution f produced by an online algorithm ALG, we denote by $\text{ALG}(\sigma) = C(f)$ its cost. The online algorithm ALG is called (strictly) *c-competitive*, if $\text{ALG}(\sigma)$ is never larger than c times the cost of an optimal offline solution over all networks, price functions, and sequences σ. 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 [5] and Fiat and Woeginger [14].

Remark 2.7 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 since in this case the routing problems are independent of each other. In fact, each routing decision is just a shortest path problem with respect to the constant prices. Hence, a solution can be computed more efficiently than in the general case.

Clearly, the competitive ratio of SEQ is 1, if $K = 1$.

3 Competitive Analysis

To analyze the ONLINEMCRP, we first 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 Fig. 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 $\alpha \in [\frac{1}{2}, 1]$ over path P_1 and $(1 - \alpha)$ along path P_2. Now commodity 2 arrives 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

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

$$= 2 \cdot \alpha^m + 2 \cdot (1 - \alpha)^m + (\alpha + 1)^m - \alpha^m$$

$$\geq (\alpha + 1)^m.$$

Routing the first commodity completely over path P_2 and the second over path P_3 leads to the constant 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 affine and then study networks with two nodes and parallel arcs.

3.1 Affine Price Functions

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

For affine 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 $(f^1, \ldots, f^k)$ that

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

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

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

for all feasible flows x^k. The main idea to obtain a bound on the competitive ratio for SEQ is to bound the cost of the flow produced by SEQ in terms of the cost of an offline optimal flow by applying (9) for every $k \in [K]$.

Theorem 3.2 *If the price functions are affine*, 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:

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

Let $(x_1, \ldots, x_K) \in \mathcal{F}_K$ be any feasible flow and let $(f_1, \ldots, f_K) \in \mathcal{F}_K$ be the solution produced by SEQ. We start by considering the following inequality for arbitrary nonnegative real values α, β satisfying $1 \leq \alpha \leq \beta \leq 2$:

$$0 \leq \left(\alpha \sum_{k=1}^{K} f_a^k - \beta \sum_{k=1}^{K} x_a^k \right)^2$$

$$= \alpha^2 \sum_{k=1}^{K} \sum_{i=1}^{K} f_a^i f_a^k - 2\alpha\beta \sum_{k=1}^{K} \sum_{i=1}^{K} f_a^i x_a^k + \beta^2 \sum_{k=1}^{K} \sum_{i=1}^{K} x_a^i x_a^k.$$

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

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

$$\left. + 2\beta^2 \sum_{k=1}^{K} \left(\sum_{i=1}^{k-1} x_a^i + \frac{1}{2} x_a^k \right) x_a^k \right). \tag{11}$$

For the next step, consider the inequality

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

$$= \sum_{a \in A} \sum_{k=1}^{K} \left(2\alpha^2 r_a f_a^k - 2\alpha\beta r_a x_a^k + 2\beta^2 r_a x_a^k \right) - \frac{2\alpha\beta}{K} \sum_{a \in A} \sum_{k=1}^{K} r_a f_a^k. \tag{12}$$

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

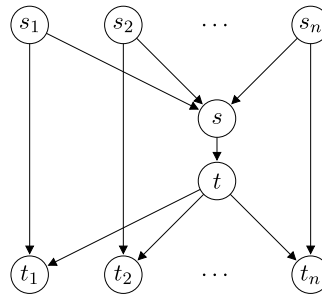$$2\alpha^2 - \frac{2\alpha\beta}{K} \geq 2\alpha^2 - \alpha\beta \geq 0,$$

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

$$0 \leq 2\alpha^2 C(f) - 2\alpha\beta \sum_{a \in A} \sum_{k=1}^{K} \left(q_a \sum_{i=1}^{K} f_a^i + r_a \right) x_a^k + 2\beta^2 C(x)$$

$$- \frac{2\alpha\beta}{K} \sum_{a \in A} \sum_{k=1}^{K} r_a f_a^k.$$

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

$$0 \leq 2\alpha^2 C(f) - 2\alpha\beta \sum_{a \in A} \sum_{k=1}^{K} \left(q_a \sum_{i=1}^{k} f_a^i + r_a \right) f_a^k + 2\beta^2 C(x)$$

$$- \frac{2\alpha\beta}{K} \sum_{a \in A} \sum_{k=1}^{K} r_a f_a^k$$

$$= (2\alpha^2 - 2\alpha\beta) C(f) - \alpha\beta \sum_{a \in A} q_a \sum_{k=1}^{K} f_a^k f_a^k + 2\beta^2 C(x)$$

$$- \frac{2\alpha\beta}{K} \sum_{a \in A} \sum_{k=1}^{K} r_a f_a^k.$$

Fig. 2 Network for the proof of Proposition 3.5

We now use the inequality of Cauchy-Schwarz as follows:

$$\langle \boldsymbol{f}_a, \mathbb{1} \rangle^2 \leq \|\boldsymbol{f}_a\|^2 \cdot \|\mathbb{1}\|^2 \quad \Leftrightarrow \quad \frac{1}{K}\left(\sum_{k=1}^{K} f_a^k\right)^2 \leq \sum_{k=1}^{K} (f_a^k)^2,$$

where $\mathbb{1}$ is the vector of all ones and $\boldsymbol{f}_a := (f_a^1, \ldots, f_a^K)$. Together with (10) this yields:

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

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

This is equivalent to:

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

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

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

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

Remark 3.4 For polynomials with nonnegative coefficients one can prove constant upper bounds on the competitive ratio, which depend on the maximum degree [19, 21, 22]. These bounds, however, do not depend on the number of commodities. Currently, we do not know how to generalize the proof of Theorem 3.2 to this case.

In the following, we derive lower bounds on the competitive ratio. We start with a lower bound for SEQ.

Proposition 3.5 *In case of affine price 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 Fig. 2 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 $i = 1, \ldots, n$. For $i = 1, \ldots, n$, we consecutively release a demand of unit size from s_i to t_i. Using Lemma 2.4, we see that SEQ routes every demand over arc (s, t). The cost for these n demands is:

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

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

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

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 over arc (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 proposition.   □

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.

Independently of our work, Olver [22] and Farzad et al. [13] proved an asymptotic lower bound of 4, improving the asymptotic 2 of Proposition 3.5. Therefore, the result in Theorem 3.2 is *asymptotically* tight. For fixed K, however, we do not know whether the bound is tight.

The best known lower bound for *any* deterministic online algorithm is the following.

Fig. 3 Graph for the proof of
Proposition 3.7
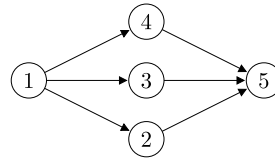
Proposition 3.7 *In case of linear cost functions*, *no deterministic online algorithm for the* ONLINEMCRP *is c-competitive for any $c < \frac{4}{3}$.*

Proof Consider the network displayed in Fig. 3. 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 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 2.4). 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 2.5). 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
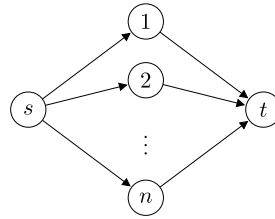
$$\text{ALG}(\sigma) \geq 4 \cdot \left(\frac{1}{2} \cdot \alpha \right) \cdot \alpha + 2 \cdot 4 \cdot \left(\frac{1}{2} \cdot \frac{(1 - \alpha)}{2} \right) \cdot \frac{(1 - \alpha)}{2}$$
$$+ 4 \cdot \left(\alpha + \frac{1}{2} \cdot 1 \right) \cdot 1 > 4,$$

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 $\frac{4}{3}$. □

Fig. 4 Network for the proof of
Proposition 3.8

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. This restricts the class of efficient algorithms to analyze.

Proposition 3.8 *Every deterministic online algorithm for the* ONLINEMCRP *that routes all demands unsplittably is not competitive, even for linear cost functions.*

Proof Consider the network shown in Fig. 4. 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. □

In Sect. 4, we further investigate the problem variant in which splitting demand is also not allowed for the offline optimum.
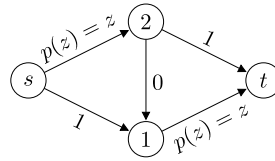
3.2 Parallel Arc Case

We now consider the parallel arc case, that is, D consists of two nodes s and t and parallel arcs from s to t only. We allow for arbitrary nondecreasing continuous price functions. We will show that SEQ is 1-competitive, i.e., optimal, in this case.

Recall from Lemma 2.4(ii) that if f is generated by SEQ for the sequence $\sigma = 1, \ldots, K$, then the following holds for all $a \in A$, $k \in [K]$ with $f_a^k > 0$:

$$p_a \left(\sum_{i=1}^{k} f_a^i \right) \leq p_{\hat{a}} \left(\sum_{i=1}^{k} f_{\hat{a}}^i \right) \quad \text{for all } \hat{a} \in A. \tag{13}$$

Fig. 5 Network and price functions for Example 3.11

By Lemma 2.5(ii), a flow x solves the offline problem (6) if and only if we have for all $k \in [K]$:

$$p_a\left(\sum_{i=1}^{K} x_a^i\right) \le p_{\hat{a}}\left(\sum_{i=1}^{K} x_{\hat{a}}^i\right) \quad \text{for all } a, \hat{a} \in A \text{ with } x_a^k > 0. \tag{14}$$

Theorem 3.9 *Algorithm* SEQ *computes an offline optimum in the parallel arc case.*

Proof For some sequence $\sigma = 1, \ldots, K$ in the parallel arc case, let f be the flow generated by SEQ for σ. Consider $a \in A$ with $\sum_{k=1}^{K} f_a^k > 0$. Let $k \in [K]$ be the largest index such that $f_a^k > 0$. Then for all $\hat{a} \in A$:

$$p_a\left(\sum_{i=1}^{K} f_a^i\right) = p_a\left(\sum_{i=1}^{k} f_a^i\right) \le p_{\hat{a}}\left(\sum_{i=1}^{k} f_{\hat{a}}^i\right) \le p_{\hat{a}}\left(\sum_{k=1}^{K} f_{\hat{a}}^i\right),$$

where the first inequality follows from the optimality condition (13) for the flow $(f^1, \ldots, f^k)$ and the second follows from the assumption that the price functions are nondecreasing. This shows (14) and hence the claim. $\square$

Remark 3.10 Theorem 3.9 also holds for the following slightly more general networks: There are two demand nodes s and t, demand is only sent from s to t, and all directed $(s - t)$-paths are induced paths, i.e., all nodes except s and t on the path have degree two. Note that paths from t to s do not make a difference.

A natural extension is the case, in which each commodity has the same source and same destination, but otherwise the network is arbitrary. The following example proves that SEQ is not optimal for this generalization, even if price functions are affine.

Example 3.11 Consider the network depicted in Fig. 5 with the specified price functions, see also Braess [7]. All commodities have node s as source and node t as destination. The first commodity has a demand of 1. SEQ assigns this commodity completely to path $(s, 2, 1, t)$ and yields a cost of

$$\int_0^1 z \, dz + 0 + \int_0^1 z \, dz = 1.$$

Then a second commodity of unit size is released. SEQ evenly splits this demand into two pieces. A flow of $\frac{1}{2}$ is routed over path $(s, 1, t)$ and $\frac{1}{2}$ along path $(s, 2, t)$. This

leads to a cost of

$$\int_0^{\frac{1}{2}} (1+z)\, dz + \frac{1}{2} + \int_0^{\frac{1}{2}} (1+z)\, dz + \frac{1}{2} = 2.25$$

for the second commodity. Hence, for these two commodities, SEQ produces a solution with total cost of 3.25. However, splitting the first and the second commodity evenly into two pieces and assigning one piece to the lower path $(s, 1, t)$ and the other to the upper path $(s, 2, t)$ gives a solution with cost

$$\int_0^1 z\, dz + 1 + \int_0^1 z\, dz + 1 = 3.$$

Hence, SEQ does not compute an optimal solution.

4 Unsplittable Routings

In this section, we study the *unsplittable* ONLINEMCRP: the variant of ONLINE-MCRP in which the demand of each commodity has to be routed along a single path. Such a restriction frequently 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 unsplittable* ONLINEMCRP *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 [17]). Here, one is given nonnegative integers $d_1, \ldots, d_K$ and positive integers $N \leq K$ and J. The question is whether there exists a partition of $[K]$ into N sets $A_1, \ldots, A_N$ such that
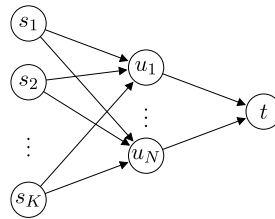
$$\sum_{i=1}^N \left(\sum_{k \in A_i} d_k \right)^2 \leq J?$$

For the reduction to the offline problem, we construct a network D with node set $\{s_1, \ldots, s_K, u_1, \ldots, u_N, t\}$ 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 Fig. 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 $A_1, \ldots, A_N$ is the wanted partition. Then if $k \in A_i$, we

Fig. 6 Construction for the
proof of Proposition 4.1

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 at most J. For $i = 1, \ldots, N$, 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 exits a solution to the minimum sum of squares problem. □

In the following we study the online algorithm U-Seq that, given unsplittable flows $f_1, \ldots, f_{k-1}$, solves the variant of (3) in which the flow has to be unsplittable.

Remark 4.2 The unsplittable variant of problem (3) amounts to solving a shortest path problem and hence can be solved in polynomial time.

When the price functions are constant, also the offline optimum of the unsplittable OnlineMCRP can be computed by solving a sequence of shortest path problems. Hence, the unsplittable variant of (6) can be solved in polynomial time, and U-Seq produces an optimal solution in this case.

The following result shows that the additional requirement of unsplittable routings does not allow for a competitive deterministic online algorithm for the Online-MCRP, in general. This result follows already from the proof of Proposition 3.1, since the offline optimum for the used instance is unsplittable. Hence, in this case, lower bounds for the splittable OnlineMCRP carry over to the unsplittable variant. We present a proof for the following proposition, since it contains an improvement of the implicit lower bound for polynomial price functions compared to the proof of Proposition 3.1.

Proposition 4.3 *In general, there exists no competitive deterministic online algorithm for the unsplittable* OnlineMCRP.

Proof Consider again the network shown in Fig. 1, where each arc a has 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 1, source node 1, and target node 4. Because of symmetry, we can assume that Alg uses path $P_1 = (1, 2, 4)$ to

route this demand. We then release commodity 2 with demand 1, source 1, and target 2. 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 tending to infinity it follows that ALG is not competitive. $\qquad\square$

We can also improve the lower bound of Proposition 3.7 from $\frac{4}{3}$ to 2.

Proposition 4.4 *For linear price functions, no deterministic online algorithm has a competitive ratio less than 2 for the unsplittable* ONLINEMCRP.

Proof Consider the network shown in Fig. 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 1, source 1, and target 4. Without loss of generality this commodity is routed over path $P_1 = (1, 2, 4)$. Then we release a commodity from node 1 to 2 and a commodity from node 2 to 4, both with 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. This leads to an optimal cost of $\text{OPT}(\sigma) = 4$, and hence the competitive ratio of ALG is at least 2. $\qquad\square$

In the following, we will derive an upper bound on the competitive ratio of U-SEQ for affine price functions.

The next lemma provides a straight-forward formulation of optimality in the unsplittable variant of (3).

Lemma 4.5 *Given unsplittable flows $f^1, \ldots, f^{k-1} \in \mathcal{F}_{k-1}$, the feasible unsplittable flow f^k solves the unsplittable variant of* (3) *if and only if the following inequality is satisfied*:

$$\sum_{a \in A} \int_0^{f_a^k} p_a \left(\sum_{i=1}^{k-1} f_a^i + z \right) dz \leq \sum_{a \in A} \int_0^{x_a^k} p_a \left(\sum_{i=1}^{k-1} f_a^i + z \right) dz,$$

for all feasible unsplittable flows x_a^k for commodity k.

Note that Lemma 4.5 also holds for the splittable case, i.e., for problem (3).

Theorem 4.6 *In the case of affine price functions,* U-SEQ *is* $(3 + 2\sqrt{2})$*-competitive for the unsplittable* ONLINEMCRP.

Proof We start with the cost of the flow f generated by U-SEQ for affine price functions $p_a(z) = q_a \cdot z + r_a$, with $q_a \geq 0$ and $r_a \geq 0$ for $a \in A$:

$$C(f) = \sum_{a \in A} \sum_{k=1}^{K} q_a \left(\sum_{i=1}^{k-1} f_a^i f_a^k + \frac{1}{2} f_a^k f_a^k \right) + r_a f_a^k.$$

Using Lemma 4.5 and $q_a \geq 0$, we get:

$$C(f) \leq \sum_{a \in A} \sum_{k=1}^{K} q_a \left(\sum_{i=1}^{k-1} f_a^i x_a^k + \frac{1}{2} x_a^k x_a^k \right) + r_a x_a^k.$$

Using $\sum x_a^k x_a^k \leq (\sum x_a^k)^2$, increasing the limits of the inner sum, and applying (10), we obtain:

$$\leq \sum_{a \in A} q_a \left(\sum_{k=1}^{K} \sum_{i=1}^{K} f_a^i x_a^k + \frac{1}{2} \left(\sum_{k=1}^{K} x_a^k \right)^2 \right) + r_a x_a^k$$

$$= C(x) + \sum_{a \in A} q_a \sum_{k=1}^{K} \sum_{i=1}^{K} f_a^i x_a^k.$$

For $\lambda > 1$ (to be specified later), we add and subtract $\lambda\, C(x)$, which yields
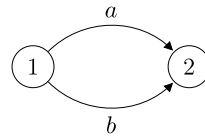
$$C(f) \leq (1 + \lambda)\, C(x) + \sum_{a \in A} q_a \left(\sum_{k=1}^{K} \sum_{i=1}^{K} f_a^i x_a^k - \frac{\lambda}{2} \left(\sum_{k=1}^{K} x_a^k \right)^2 \right) - \lambda\, r_a \sum_{k=1}^{K} x_a^k.$$

Dropping the last term and adding

$$0 \leq \left(\frac{1}{\sqrt{2\lambda}} \sum_{k=1}^{K} f_a^k - \frac{\sqrt{\lambda}}{\sqrt{2}} \sum_{k=1}^{K} x_a^k \right)^2$$

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

(multiplied by q_a and summed over a), we obtain:

$$C(f) \leq (1 + \lambda)\, C(x) + \frac{1}{2\lambda} \sum_{a \in A} q_a \left(\sum_{k=1}^{K} f_a^k \right)^2$$

$$\leq (1 + \lambda)\, C(x) + \frac{1}{\lambda}\, C(f),$$

Fig. 7 Graph for Example 4.10

where for the last inequality we added

$$\sum_{a \in A} r_a \sum_{k=1}^{K} f_a^k \geq 0.$$

Finally, defining $\lambda = 1 + \sqrt{2}$ and rewriting yields the desired result.   □

Remark 4.7 The parameter λ in the previous proof is the optimal solution to the following minimization problem:

$$\min_{\lambda > 1} \frac{(1 + \lambda)\lambda}{\lambda - 1}.$$

Note that the above proof also holds for the splittable case, but yields a weaker bound $3 + 2\sqrt{2} \approx 5.828$ compared to $\frac{4K^2}{(1+K)^2}$ in Theorem 3.2.

With the bound of Theorem 4.6 and Remark 4.2, we obtain a constant approximation algorithm for the NP-hard unsplittable version of (6).

Corollary 4.8 U-SEQ *is a* $(3 + 2\sqrt{2})$-*approximation algorithm for the unsplittable variant of* (6).

Remark 4.9 For polynomials with nonnegative coefficients, one can prove constant upper bounds on the competitive ratio, which depend on the maximum degree [21, 22]. Furthermore, in [22] it is shown that these upper bounds are tight.

In Sect. 3.2, we proved that SEQ is optimal in the special setting of parallel arcs in the splittable case. In the unsplittable case, however, the following example shows that U-SEQ does not always return an optimal solution, even if we restrict the network to parallel arcs.

Example 4.10 Consider the network depicted in Fig. 7 with the following price functions: $p_a(z) = z$ and $p_b(z) = 2z$. Each commodity has node 1 as source and node 2 as destination.

The first commodity is released with a demand of 2. Algorithm U-SEQ routes this commodity over arc a and produces a cost of 2. A second commodity arises with demand 3. This request is send by U-SEQ over arc b, which leads to a cost of 9. The total cost is 11.

Routing the first request over arc b and the second one over arc a yields a cost of 8.5. Hence, U-SEQ is not optimal for this setting.

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. First steps in this direction have been taken in [21].

References

1. Altman, E., Basar, T., Jimenez, T., Shimkin, N.: Competitive routing in networks with polynomial costs. IEEE Trans. Autom. Control **47**, 92–96 (2002)
2. Awerbuch, B., Azar, Y., Plotkin, S.: Throughput-competitive on-line routing. In: Proc. 34th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 32–40 (1993)
3. Awerbuch, B., Azar, Y., Grove, E.F., Kao, M.-Y., Krishnan, P., Vitter, J.S.: Load balancing in the L_p norm. In: Proc. 36th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 383–391 (1995)
4. Awerbuch, B., Azar, Y., Epstein, A.: The price of routing unsplittable flow. In: Proc. of the Thirty-seventh Annual ACM Symposium on Theory of Computing (STOC), pp. 57–66. ACM Press, New York (2005)
5. Borodin, A., El-Yaniv, R.: Online Computation and Competitive Analysis. Cambridge University Press, Cambridge (1998)
6. Boyd, S., Vandenberghe, L.: Convex Optimization. Cambridge University Press, Cambridge (2004)
7. Braess, D.: Über ein Paradoxon aus der Verkehrsplanung. Unternehmensforschung **12**, 258–268 (1969)
8. Caragiannis, I., Flammini, M., Kaklamanis, C., Kanellopoulos, P., Moscardelli, L.: Tight bounds for selfish and greedy load balancing. In: Proc. of the 33rd International Colloquium of Automata, Languages and Programming (ICALP). LNCS, vol. 4051, pp. 311–322. Springer, Berlin (2006)
9. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: Proc. of the Thirty-Seventh Annual ACM Symposium on Theory of Computing (STOC), pp. 67–73. ACM Press, New York (2005)
10. Cominetti, R., Correa, J.R., Stier-Moses, N.E.: The impact of oligopolistic competition in networks. Oper. Res. (2009, in press)
11. Correa, J.R., Schulz, A.S., Stier Moses, N.E.: Selfish routing in capacitated networks. Math. Oper. Res. **29**, 961–976 (2004)
12. Dafermos, S.C., Sparrow, F.T.: The traffic assignment problem for a general network. J. Res. Natl. Bur. Stand., Sect. B **73**, 91–118 (1969)
13. Farzad, B., Olver, N., Vetta, A.: A priority-based model of routing. Chic. J. Theor. Comput. Sci. Article 1 (2008)
14. Fiat, A., Woeginger, G.J. (eds.): Online Algorithms: The State of the Art. LNCS, vol. 1442. Springer, Berlin (2008)
15. Fortz, B., Thorup, M.: Optimizing OSPF/IS-IS weights in a changing world. IEEE J. Sel. Area Commun. **20**, 756–767 (2002)
16. Fortz, B., Thorup, M.: Increasing internet capacity using local search. Comput. Optim. Appl. **29**, 13–48 (2004)
17. Garey, M.R., Johnson, D.S.: Computers and Intractability. A Guide to the Theory of NP-Completeness. Freeman, New York (1979)

18. Grötschel, M., Lovász, L., Schrijver, A.: Geometric Algorithms and Combinatorial Optimization. Algorithms and Combinatorics, vol. 2, 2nd edn. Springer, Heidelberg (1993)
19. Harks, T.: Multicommodity routing problems—selfish behavior and online aspects. PhD thesis, TU Berlin (2007)
20. Harks, T., Heinz, S., Pfetsch, M.E.: Competitive online multicommodity routing. In: Erlebach, T., Kaklamanis, C. (eds.) Proc. 4th International Workshop on Approximation and Online Algorithms (WAOA). LNCS, vol. 4368, pp. 240–252. Springer, Berlin (2006)
21. Harks, T., Heinz, S., Pfetsch, M.E., Vredeveld, T.: Online multicommodity routing with time windows. ZIB Report 07-22, Zuse Institute Berlin (2007)
22. Olver, N.: The price of anarchy and a priority-based model of routing. Master's thesis, McGill University, Montréal (2006)
23. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. Int. J. Game Theory **2**, 65–67 (1973)
24. Roughgarden, T.: Selfish routing with atomic players. In: Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 973–974 (2005)
25. Roughgarden, T., Tardos, E.: How bad is selfish routing? J. ACM **49**, 236–259 (2002)
26. Suri, S., Toth, C., Zhou, Y.: Selfish load balancing and atomic congestion games. Algorithmica **47**, 79–96 (2007)
27. Wardrop, J.G.: Some theoretical aspects of road traffic research. Proc. of the Institute of Civil Engineers **1**, 325–378 (1952)
28. Yahaya, A., Suda, T.: iREX: Inter-domain QoS automation using economics. In: Proc. of IEEE Consumer Communications and Networking Conference, pp. 96–101 (2006)
29. Yahaya, A., Harks, T., Suda, T.: iREX: Efficient automation architecture for the deployment of inter-domain QoS policy. Netw. Service Manag. **5**, 50–64 (2008)