Operations Research Letters 46 (2018) 93–98

Contents lists available at ScienceDirect

Operations Research Letters

journal homepage: www.elsevier.com/locate/orl

Bottleneck routing with elastic demands

Tobias Harks [a,*], Max Klimm [b], Manuel Schneider [c]

[a] Universität Augsburg, Institut für Mathematik, 86135 Augsburg, Germany
[b] Humboldt-Universität zu Berlin, School of Business and Economics, 10178 Berlin, Germany
[c] Technische Universität Berlin, Institut für Mathematik, 10623 Berlin, Germany

ARTICLE INFO

Article history:
Received 21 February 2017
Received in revised form 9 November 2017
Accepted 9 November 2017
Available online 2 December 2017

Keywords:
Bottleneck routing game
Elastic demands
Pure Nash equilibrium

ABSTRACT

Bottleneck routing games are a well-studied model to investigate the impact of selfish behavior in communication networks. In this model, each user selects a path in a network for routing her *fixed* demand. The disutility of a user only depends on the most congested link visited. We extend this model by allowing users to continuously vary the demand rate at which data is sent along the chosen path. As our main result we establish tight conditions for the existence of pure strategy Nash equilibria.

1. Introduction

Bottleneck routing games are a theoretical model to study the effects of resource allocation in distributed communication networks [1,4]. Every user of the network is associated with a non-negative demand that she wants to send from her source to the respective destination, and her goal is to find a path that minimizes the congestion of the most congested link. It has been argued (cf. [5,28]) that in the context of packet-switched communication networks, the performance of a path is more closely related to the most congested link than the classical sum-aggregation of costs (as in [18,29,34]), and there are several proposals (cf. [26,36]) for replacing the sum-aggregation of congestion costs with the max-aggregation, primary, because the max-aggregation leads to favorable properties of protocols in terms of their stability in presence of communication delays [36].

While bottleneck routing games are an important step in terms of integrating routing decisions with bottleneck objectives, they lack one fundamental tradeoff inherent in packet-switched communication networks: once a path is selected, a user increases the sending rate in case of low congestion and decreases it in case of high congestion. In this paper, we address this tradeoff by introducing bottleneck congestion games with *elastic* demands, where users can continuously vary their demands. Formally, there is a finite set of resources and a strategy of a player is a tuple consisting of a subset of resources and a demand. Resources have player-specific cost functions that are non-decreasing and strictly convex. Every user is associated with a non-decreasing strictly concave utility function measuring the received utility from sending at a certain demand rate (cf. [18,31]). The goal of a user is to select both a subset of resources and a demand rate that maximizes the utility (from the demand rate) minus the congestion cost on the most expensive resource contained in the chosen resource set. Our model thus integrates as a special case (*i*) single-path routing (which is up to date standard as splitting packets over several routes leads to different packet inter-arrival times and synchronization problems) and (*ii*) congestion control via data rate adaption based on the maximum congestion experienced.

1.1. Our results

We derive conditions for the cost functions that ensure that the resulting bottleneck congestion games with elastic demands admit a pure Nash equilibrium (PNE). The existence of pure Nash equilibria is favorable for large communication network as they provide a deterministic steady state from which no player has an incentive to deviate. Mixed Nash equilibria, on the other hand are less desirable as they may lead to oscillating route choices which degrade the performance of the system [21]. Our condition requires that for every player the player-specific resource cost functions are non-decreasing, strictly convex and equal up to resource specific shifts in their argument. While monotonicity and convexity are natural conditions, the last assumption seems limiting. We can show, however, that without it there are examples without any PNE. Our proof is constructive, i.e., we devise an algorithm that computes a PNE.

Our results thus give further indication that the max-aggregation of congestion costs has desirable properties in terms of equilibrium existence. Specifically, our results imply that networks

* Corresponding author.
E-mail addresses: tobias.harks@math.uni-augsburg.de (T. Harks), max.klimm@hu-berlin.de (M. Klimm).

https://doi.org/10.1016/j.orl.2017.11.007
0167-6377/

with $M/M/1$ functions (cf. [23]) possess a PNE, if congestion costs are aggregated with the maximum. This stands in contrast to the classical sum-aggregation of congestion costs, where PNE need not exist [12].

An extended abstract of the results presented in this paper appeared in Harks et al. [14].

1.2. Related work

Bottleneck routing games with fixed demands admit strong equilibria [13,24], a strengthening of PNE that are resilient to coordinated deviations of groups of players. The complexity of computing PNE and strong equilibria in these games was investigated in [9]. For works on the price of anarchy of PNE and the worst-case quality of strong equilibria we refer to [2,3,5,7,16,17]. Further related is the model of [25] who studied generalizations of congestion games in which the sum-aggregation is replaced by an arbitrary aggregation function.

In previous work [11], we established the existence of an equilibrium for a class of aggregative location games. This result implies the existence of a PNE for the present model when the allowable sets of resource of players contain singletons only. Congestion games with variable demands coincide with the present model except that the traditional sum-aggregation of costs is used. For these games only affine and certain exponential cost functions lead to the existence of PNE [12]. These results imply that for $M/M/1$ delay functions a PNE does not always exist.

Integrated routing and congestion control has been studied in [8,19,20,32,33], where the existence of a PNE is proved by relating it to optimal solutions of an associated convex utility maximization problem. These models require that every user possibly splits the flow among a number of paths that may even be exponential in the size of the underlying graph. This issue has been addressed in [6,27], where controllable route splitting at routers is assumed which can effectively limit the resulting number of used routes. For all the above models, however, the end-to-end applications may suffer in service quality due to packet jitter caused by different path delays. Partly because of this issue, the standard TCP/IP protocol suite still uses single path routing. Also in contrast to our model, all these models assume that congestion feedback is aggregated via the sum instead of the max operator.

Another model for resource allocation in telecommunication networks is the class of *MAXBAR*-games. Here, players select a single path in a network with fixed edge capacities. Then, all players synchronously increase their rate until the capacity of an edge is reached. After such an event all rates of users using this tight edge are fixed. *MAXBAR* games possess a PNE [35], and a strong equilibrium [10] even when the rate increase is non-homogeneous.

2. The model

Let $R = \{1, \dots, m\}$ be a nonempty and finite set of $m \in \mathbb{N}$ resources, and let $N = \{1, \dots, n\}$ be a nonempty and finite set of $n \in \mathbb{N}$ players. For every $i \in N$, let $X_i \subset 2^R \setminus \{\emptyset\}$ be a nonempty set of nonempty subsets of resources available to player i and let $X = \times_{i\in N} X_i$ denote their Cartesian product. We call $x_i \in X_i$ an allocation of player i and we denote by $x = (x_i)_{i\in N} \in X$ the overall allocation vector. For every player i and every resource $r \in R$ we are given a player-specific cost function $c_{i,r} : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. Every player $i \in N$ has a utility function $U_i : [\sigma_i, \tau_i] \to \mathbb{R}_{\geq 0}$ where $[\sigma_i, \tau_i] \subseteq \mathbb{R}_{\geq 0}$ with $\sigma_i \in \mathbb{R}_{\geq 0}$, $\tau_i \in \mathbb{R}_{\geq 0} \cup \{\infty\}$, $\sigma_i \leq \tau_i$ is the interval of feasible demands of player i. A vector $d = (d_i)_{i\in N}$ with $d_i \in [\sigma_i, \tau_i]$ is called a demand vector. A bottleneck congestion game with elastic demands is a maximization game $G = (N, \mathcal{S}, \pi)$ with $\mathcal{S}_i = X_i \times [\sigma_i, \tau_i]$ for all $i \in N$. We set $\mathcal{S} = \times_{i\in N} \mathcal{S}_i$ and for $s = (x, d) \in \mathcal{S}$, we define the private payoff function of player i as $\pi_i(s) := U_i(d_i) - \max_{r\in x_i}\left\{c_{i,r}\left(\ell_r(x, d)\right)\right\}$. Here, $\ell_r(s) = \ell_r(x, d) = \sum_{j\in N: r\in x_j} d_j$ is the *load* of resource r under strategy profile $s = (x, d)$. We impose the following assumptions on the utility and cost functions, respectively.

Assumption 2.1. For every player i, the following properties are satisfied:

(A1) the utility function $U_i : [\sigma_i, \tau_i] \to \mathbb{R}_{\geq 0}$ is non-decreasing, differentiable and strictly concave;
(A2) for every resource, the cost function $c_{i,r} : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is non-decreasing, differentiable and strictly convex;
(A3) there is a function $c_i : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ and player-independent offsets $\upsilon_r \in \mathbb{R}_{\geq 0}$ such that $c_{i,r}(t) = c_i(t + \upsilon_r)$ for all $t \geq 0$ and $r \in R$.

Note that the monotonicity together with the strict concavity and convexity of the utility and cost functions required in (A1) and (A2) implies that these functions are strictly monotonic.

Strict concavity of the utility function as required in (A1) is justified by application-specific characteristics such as the rate-control algorithm used in common congestion control protocols, cf. [18,31]. Also, in many applications the considered cost functions are strictly convex as required in (A2), e.g., the polynomial delay functions considered in transportation networks (cf. [30]) and $M/M/1$ functions modeling queuing delays in telecommunication networks (cf. [32]). Assumption (A3) essentially requires that the maximum load (including the offsets υ_r) experienced on the chosen subset of resources determines the bottleneck. Since this assumption is certainly restrictive, a few remarks are in order. First, in games violating assumption (A3) a player's private payoff function need not be differentiable in the demand of the player, e.g., consider a game with a single player using both a resource with cost function $c_1(t) = 1+t$ and a resource with cost function $c(t) = 2t$. It is easy to verify that the player's private payoff is not differentiable at $d_1 = 1$. On the other hand, payoff functions are differentiable when (A3) is satisfied as the limit $\lim_{\epsilon\to 0} \frac{\pi_i(x,(d_i+\epsilon,d_{-i}))-\pi_i(x,(d_i,d_{-i}))}{\epsilon} = U_i'(d_i) - c_i'(\max_{r\in x_i}\{\ell_r(x, d) + \upsilon_r\})$ exists. Second, we will see in Section 6 that games violating assumption (A3) may actually fail to admit a PNE. Lastly, we observe that conditions (A1)–(A3) are still expressive enough to include relevant functions, for instance, $M/M/1$ delay functions $c_{i,r} : [0, z_r) \to \mathbb{R}_{\geq 0}$ of the form $c_{i,r}(t) = t_i/(z_r - t)$ with $z_r > 0$ and $t_i \geq 0$.

3. A characterization of pure Nash equilibria

In this section, we give a complete characterization of PNE in bottleneck congestion games. Our characterization relies on the notion of a *demand equilibrium* which we define as a strategy profile with the property that no player can increase her payoff by unilaterally changing her demand only.

Definition 3.1 (*Demand Equilibrium*). A strategy profile (x, d) is called a *demand equilibrium* if $\pi_i(x, d) \geq \pi_i(x, \tilde{d})$ for all $i \in N$ and $\tilde{d} = (d_{-i}, \tilde{d}_i)$ with $\tilde{d}_i \in [\sigma_i, \tau_i]$.

Every PNE is a demand equilibrium, but not vice versa. Assumption (A3) implies that the private payoffs are differentiable in the demand. Thus, the Karush–Kuhn–Tucker conditions impose the following necessary conditions for demand equilibria. For $s = (x, d)$, let $b_i(x, d) = \max_{r\in x_i}\{\ell_r(x, d) + \upsilon_r\}$ denote the maximal load (or *bottleneck*) that player i experiences and denote by $b_i^{-1}(x, d)$ the set of resources where the bottleneck is attained, i.e., $b_i^{-1}(x, d) = \arg\max_{r\in x_i}\{\ell_r(x, d) + \upsilon_r\}$.

Corollary 3.2. *Let (x, d) be a demand equilibrium. Then, for all $i \in N$ the following conditions are satisfied:*

(i) $d_i < \tau_i \Rightarrow U_i'(d_i) - c_i'\big(\ell_{r^*}(x, d) + \upsilon_{r^*}\big) \leq 0$ *for all* $r^* \in b_i^{-1}(x, d)$;
(ii) $d_i > \sigma_i \Rightarrow U_i'(d_i) - c_i'\big(\ell_{r^*}(x, d) + \upsilon_{r^*}\big) \geq 0$ *for all* $r^* \in b_i^{-1}(x, d)$.

Let $s = (x, d)$ be a demand equilibrium and let $\ell_r^{-i}(x, d) := \sum_{j \in N\setminus\{i\}: r \in x_j} d_j$ denote the *residual load* of player i on some $r \in R$. We then obtain

$$d_i \in \arg\max_{d_i' \in [\sigma_i, \tau_i]} \{U_i(d_i') - c_i(\ell_{r^*}^{-i}(x, d) + \upsilon_{r^*} + d_i')\}.$$

for all $r^* \in b_i^{-1}(s)$. We proceed to investigate how a demand equilibrium is adapted if the residual load on a resource changes. Formally, for a fixed residual load α on a resource r, we are interested in analyzing how the *best-reply demand function* $d_i(\alpha) := \arg\max_{d_i \in [\sigma_i, \tau_i]} \{U_i(d_i) - c_i(\alpha + d_i)\}$ depends on α. As for all $\alpha \in \mathbb{R}_{\geq 0}$ the function $f(y) := U_i(y) - c_i(\alpha + y)$ is continuous and strictly concave in y and the maximum is taken over a compact set, the above optimization problem has a unique solution, hence, the best reply demand function is well defined.

Lemma 3.3. *Let $\alpha, \beta \in \mathbb{R}_{\geq 0}$. Then, $\alpha < \beta$ if and only if $d_i(\alpha) + \alpha < d_i(\beta) + \beta$.*

Proof. "⇒": Assume $\alpha < \beta$ and $d_i(\alpha) + \alpha \geq d_i(\beta) + \beta$. This implies $d_i(\alpha) > d_i(\beta)$. We obtain $0 \leq U_i'(d_i(\alpha)) - c_i'(\alpha + d_i(\alpha)) < U_i'(d_i(\beta)) - c_i'(\beta + d_i(\beta)) \leq 0$. The first inequality follows from $d_i(\alpha) > \sigma_i$ and Corollary 3.2. The second inequality follows from $d_i(\alpha) > d_i(\beta)$ and U_i' being monotonically decreasing together with $d_i(\alpha) + \alpha \geq d_i(\beta) + \beta$ and c_i being convex.

"⇐": Assume $d_i(\alpha) + \alpha < d_i(\beta) + \beta$ and $\alpha \geq \beta$. This implies $\sigma_i \leq d_i(\alpha) < d_i(\beta)$ and we obtain the same contradiction as before by exchanging the roles of α and β. □

Given a strategy profile $s = (x, d)$, we call a strategy $s_i' = (x_i', d_i')$ a *better reply* of player i if $\pi_i(s_i', s_{-i}) > \pi_i(s)$. The following lemma characterizes better replies to demand equilibria.

Lemma 3.4. *Let $s = (x, d)$ be a demand equilibrium. Player i has a better reply $s_i' = (x_i', d_i') \in \mathcal{S}_i$ if and only if $\ell_r^{-i}(s) + \upsilon_r > \max_{t \in x_i'} \ell_t^{-i}(s) + \upsilon_t$ for some $r \in b_i^{-1}(s)$.*

Proof. "⇐": Since $\ell_r^{-i}(s) + \upsilon_r > \max_{t \in x_i'} \ell_t^{-i}(s) + \upsilon_t$ for some $x_i' \in X_i$ and cost functions are strictly increasing (cf. Assumption 2.1), player i improves her payoff deviating from $s_i = (x_i, d_i)$ to $s_i' = (x_i', d_i)$.

"⇒": Suppose $\ell_r^{-i}(s) + \upsilon_r \leq \max_{t \in x_i'} \ell_t^{-i}(s) + \upsilon_t$ for all $r \in b_i^{-1}(s)$. For $s' = (s_i', s_{-i})$ we obtain

$$\begin{aligned}\pi_i(s') &= U_i(d_i') - c_i\big(\max_{t \in x_i'}\{\ell_t^{-i}(s') + \upsilon_t\} + d_i'\big)\\ &\leq U_i(d_i') - c_i\big(\ell_r^{-i}(s) + \upsilon_r + d_i'\big)\\ &\leq U_i(d_i) - c_i\big(\ell_r^{-i}(s) + \upsilon_r + d_i\big) = \pi_i(s),\end{aligned}$$

where we use for the first inequality that $\ell_r^{-i}(s) + \upsilon_r \leq \max_{t \in x_i'} \ell_t^{-i}(s)$ $+ \upsilon_t = \max_{t \in x_i'} \ell_t^{-1}(s') + \upsilon_t$ and c_i is non-decreasing. The second inequality holds because s is a demand equilibrium and therefore d_i is player i's best reply demand. We conclude that s_i' is not a better reply. □

We obtain the following characterization of PNE as a direct corollary.

Corollary 3.5. *A strategy profile $s = (x, d)$ is a PNE if and only if s is a demand equilibrium and $\ell_r^{-i}(s) + \upsilon_r \leq \max_{t \in x_i'} \ell_t^{-i}(s) + \upsilon_t$ for all $i \in N$, $x_i' \in X_i$ and $r \in b_i^{-1}(s)$.*

4. Computing demand equilibria

Corollary 3.5 suggests that for computing a PNE we should be able to compute a demand equilibrium. In this section, we describe an algorithm for this purpose. First, we need the notion of a *distributed equilibrium*. Let $G = (N, \mathcal{S}, \pi)$ be a bottleneck congestion game, $M \subseteq N$, and let r be a resource. We define the restriction of G on M and r as the bottleneck congestion game $G|_{(M,r)} = (M, \mathcal{S}', \pi')$ with $\mathcal{S}_i' = \big\{\{r\}\big\} \times [\sigma_i, \tau_i]$ for all $i \in M$ and $\pi_i'(x, d) = U_i(d_i) - c_i(\ell_r(x, d) + \upsilon_r)$.

Definition 4.1 (*Distributed Equilibrium*). Let $x \in X$ and $N_r(x) := \{i \in N : r \in x_i\}$. A non-negative vector $\tilde{d} = (\tilde{d}_{i,r})_{i \in N_r(x), r \in R}$ is called a *distributed equilibrium* for $x \in X$ if for all $r \in R$ the strategy profile $(\tilde{d}_{i,r})_{i \in N_r(x)}$ is a PNE of $G|_{(N_r(x),r)}$.

Every restricted game $G|_{(N_r(x),r)}$ is a concave game on a compact action space, thus, by Kakutani's fixed point theorem [15] for every $x \in X$ the existence of a distributed equilibrium is guaranteed. For the computation of a distributed equilibrium, one can cast the problem as a nonlinear complementarity problem, which can be solved by sequential linearization methods similar to Lemke's algorithm (cf. [22]). For a distributed equilibrium $\tilde{d}$ with respect to $x \in X$, we define $\tilde{\ell}_r(x, \tilde{d}) := \sum_{i \in N_r(x)} \tilde{d}_{i,r}$. As $\tilde{d}$ is a PNE for the bottleneck games $G|_{(N_{r(x)},r)}$, $r \in R$, the optimality conditions of Corollary 3.2 imply $U_i'(\tilde{d}_{i,r}) - c_i'(\tilde{\ell}_r(x, d) + \upsilon_r) \leq 0$ for all $i \in N$, $r \in R$ with $d_{i,r} < \tau_i$, and $U_i'(\tilde{d}_{i,r}) - c_i'(\tilde{\ell}_r(x, d) + \upsilon_r) \geq 0$ for all $i \in N$ and $r \in R$ with $d_{i,r} > \sigma_i$.

We use these optimality conditions to prove the following lemma.

Lemma 4.2. *Let $x, x' \in X$ and let $\tilde{d}$ and $\tilde{d}'$ be two respective distributed equilibria. Then, $\tilde{d}_{i,r} \leq \tilde{d}_{i,r'}'$ for all $r, r' \in R$ with $\tilde{\ell}_r(x, \tilde{d}) + \upsilon_r \geq \tilde{\ell}_{r'}(x', \tilde{d}') + \upsilon_{r'}$ and all $i \in N_r(x) \cap N_{r'}(x')$.*

Proof. Assume by contradiction that there are two resources $r, r' \in R$ with $\tilde{\ell}_r(x, \tilde{d}) + \upsilon_r \geq \tilde{\ell}_{r'}(x', \tilde{d}') + \upsilon_{r'}$ and a player $i \in N_r(x) \cap N_{r'}(x')$ with $\tau_i \geq \tilde{d}_{i,r} > \tilde{d}_{i,r'}' \geq \sigma_i$. Thus,

$$\begin{aligned}0 &\geq U_i'(\tilde{d}_{i,r'}') - c_i'(\tilde{\ell}_{r'}(x', \tilde{d}') + \upsilon_{r'})\\ &> U_i'(\tilde{d}_{i,r}) - c_i'(\tilde{\ell}_r(x, \tilde{d}) + \upsilon_r) \geq 0,\end{aligned}$$

where the first and the third inequality are due to the fact that $\tilde{d}'$ and $\tilde{d}$ are distributed equilibria for x' and x, respectively. For the second inequality, we use that utilities are strictly concave and that costs are non-decreasing and convex. □

Using this lemma, we obtain the following immediate corollary.

Corollary 4.3 (*Uniqueness*). *Let $x, x' \in X$ and let $\tilde{d}$ and $\tilde{d}'$ be two respective distributed equilibria. Then, the following two statements hold:*

(i) $\tilde{\ell}_r(x, \tilde{d}) \leq \tilde{\ell}_r(x', \tilde{d}')$ *for all* $r \in R$ *with* $N_r(x) \subseteq N_r(x')$.
(ii) $\tilde{\ell}_r(x, \tilde{d}) = \tilde{\ell}_r(x', \tilde{d}')$ *for all* $r \in R$ *with* $N_r(x) = N_r(x')$.

Proof. To prove (i), let us assume for a contradiction that there is $r \in R$ with $N_r(x) \subseteq N_r(x')$ and $\tilde{\ell}_r(x, \tilde{d}) > \tilde{\ell}_r(x', \tilde{d}')$. We derive the existence of a player $i \in N_r(x)$ with $\tilde{d}_{i,r} > \tilde{d}_{i,r}'$. This gives a contradiction to Lemma 4.2. The second statement follows directly from the first one. □

Algorithm 1: Computation of a demand equilibrium

```
   Input: Bottleneck congestion game with elastic demands G and
          x ∈ X
   Output: Demand equilibrium (x, d) of G
1  initialize R' ← R;
2  while R' ≠ ∅ do
3  |   compute a distributed equilibrium d̃ for x;
4  |   choose index-minimal r ∈ arg max_{r∈R'} ℓ̃_r(x, d̃) + υ_r ;
5  |   d_i ← d̃_{i,r}; σ_i ← d_i; τ_i ← d_i for all i ∈ N_r(x);
6  |   R' ← R' \ {r};
```

Algorithm 2: Computation of a PNE

```
Input: Bottleneck congestion game with elastic demands G
Output: PNE (x, d) of G
(x, d') ← arbitrary strategy profile;
while true do
|   Compute a demand equilibrium (x, d) by Algorithm 1;      Phase I
|   if there is a player i with a better reply to (x, d) then Phase II
|   |   (x'_i, d'_i) ← best reply of player i to (x, d);
|   |   (x, d') ← (x'_i, x_{-i}, d'_i, d_{-i});
|   else
|   |   return (x, d);
```

Given the above lemma, we derive that for every $x \in X$, the corresponding distributed equilibrium d is unique. We proceed to give a lemma that will be useful to compute *demand* equilibria. Fix $x \in X$, a distributed equilibrium $\tilde{d}$, and a resource r^* with maximal load. The lemma states that when the demand of the players $i \in N_{r^*}(x)$ is fixed to $\tilde{d}_{i,r^*}$, and a new distributed equilibrium $\tilde{d}'$ is recomputed, then the load of all resources does not increase.

Lemma 4.4. *Let $\tilde{d}$ be a distributed equilibrium for $x \in X$ and $r^* \in \arg\max_{r \in R'(x)} \tilde{\ell}_r(x, \tilde{d}) + \upsilon_r$ where $R'(x) = \{r \in R : \exists i \in N_r(x)$ with $\sigma_i < \tau_i\}$. Define a new game G' that differs only in the fact that $\tau'_i = \sigma'_i = \tilde{d}_{i,r^*}$ for all $i \in N_{r^*}(x)$. Then, $\tilde{\ell}_r(x, \tilde{d}') \le \tilde{\ell}(x, \tilde{d})$ for all $r \in R$ and for all distributed equilibria $\tilde{d}'$ for x in G.*

Proof. For a contradiction, let us assume there is $r \in R$ with $\tilde{\ell}_r(x, \tilde{d}') > \tilde{\ell}_r(x, \tilde{d})$. This implies the existence of a player $i \in N_r(x)$ with $\tau_i \ge \tilde{d}'_{i,r} > \tilde{d}_{i,r} \ge \sigma_i$. We distinguish two cases. If $i \notin N_{r^*}(x)$, we obtain a contradiction similar to the proof of Lemma 4.2 by observing

$$0 \ge U'_i(\tilde{d}_{i,r}) - c'_i(\tilde{\ell}_r(x, \tilde{d}) + \upsilon_r) > U'_i(\tilde{d}'_{i,r}) - c'_i(\tilde{\ell}_r(x, \tilde{d}') + \upsilon_r) \ge 0,$$

where for the first and the last inequality we used the local optimality conditions of the distributed equilibria $\tilde{d}$ and $\tilde{d}'$, and the strict concavity and convexity of the utility and cost functions. If, on the other hand, $i \in N_{r^*}(x)$, we have $\tilde{d}_{i,r^*} > \tilde{d}_{i,r}$ as well as $\tilde{\ell}_{r^*}(x, \tilde{d}) + \upsilon_{r^*} \ge \tilde{\ell}_r(x, \tilde{d}) + \upsilon_r$ by the choice of r^*. This is a contradiction to Lemma 4.2 (applied for $x = x'$). □

We are now ready to propose an algorithm that takes as input an allocation $x \in X$ and computes a corresponding demand equilibrium $(x, d) \in \mathcal{S}$. The algorithm first computes a distributed equilibrium $\tilde{d}$. Then, a resource r with maximum load is chosen and the demand of each player $i \in N_r(x)$ is fixed to the demand $\tilde{d}_{i,r}$. We recompute a distributed equilibrium and reiterate. The formal description is given in Algorithm 1.

Theorem 4.5. *Algorithm 1 computes a demand equilibrium.*

Proof. Let $R = \{1, \dots, m\}$ be such that for each $k \in \{1, \dots, m\}$ in the kth iteration of the algorithm resource k is chosen in line 4. We claim that $\ell_1(x, d) + \upsilon_1 \ge \dots \ge \ell_m(x, d) + \upsilon_m$. To see this, note that in each iteration a resource r with maximum load is chosen and the demand of its users is fixed. If r is such that $\sigma_i = \tau_i$ for all $i \in N_r(x)$, fixing the demands does not change the game when recomputing the distributed equilibrium and distributed equilibria are unique by Corollary 4.3. Otherwise, by Lemma 4.4, fixing the demands does not increase the load of the other resources when recomputing an equilibrium. We derive that for each player i the bottleneck is attained at the resource $r \in x_i$ with minimal index. As the algorithm fixes the demand of each player i the first time one of the resources used by player i is considered (line 5), the demand vector d computed by Algorithm 1 is a demand equilibrium. □

Remark 4.6. For a given input G and x, Algorithm 1 computes a unique demand equilibrium (x, d). This follows since the distributed equilibria are unique (Corollary 4.3) and there is a fixed tie-breaking rule employed that determines the order in which resources are fixed (line 4).

5. An algorithm for computing PNE

In this section, we give an algorithm that computes a PNE. The algorithm starts with an arbitrary strategy profile and computes a demand equilibrium. Then, if a player can improve, we let this player play a best reply (where necessarily the player's resource set changes), and recompute a demand equilibrium. The technically involved part is to show that the algorithm terminates. To prove this, we first derive several properties of intermediate strategy profiles during the execution of the algorithm. We will prove that $b(s) = (b_i(s))_{i \in N}$ strictly decreases with respect to the sorted lexicographical order $\prec_{\text{lex}}$ that is defined as follows: For two vectors $u, v \in \mathbb{R}^n_{\ge 0}$ we say that u is *sorted lexicographically* smaller than v, written $u \prec_{\text{lex}} v$, if there is an index $k \in \{1, \dots, n\}$ such that $u_{\pi(i)} = v_{\psi(i)}$ for all $i < k$ and $u_{\pi(k)} < v_{\psi(k)}$ where π and ψ are permutations that sort u and v non-increasingly, i.e., $u_{\pi(1)} \ge u_{\pi(2)} \ge \dots \ge u_{\pi(n)}$ and $v_{\psi(1)} \ge v_{\psi(2)} \ge \dots \ge v_{\psi(n)}$.

Lemma 5.1. *Let $s = (x, d)$ be a demand equilibrium computed (by Algorithm 1) in Phase I of Algorithm 2. Let $s'_i = (x'_i, d'_i)$ be a better and best reply of player i and let $s' = (x', d')$, where $x' = (x'_i, x_{-i})$ and $d' = (d'_i, d_{-i})$. Denote by $\bar{s} = (x'_i, x_{-i}, \bar{d})$ the demand equilibrium that is computed in Algorithm 1 in Phase I in the following iteration. Then, $b(\bar{s}) \prec b(s)$.*

Proof. We first prove

$$b_i(s') < b_i(s). \tag{1}$$

Let $r_i \in b_i^{-1}(s)$ and $r'_i \in b_i^{-1}(s')$. We show $b_i(s) = \ell_{r_i}(s) + \upsilon_{r_i} > \ell_{r'_i}(s') + \upsilon_{r'_i} = b_i(s')$. Since player i strictly improves playing s'_i, Lemma 3.4 implies $\ell^{-i}_{r_i}(s) + \upsilon_{r_i} > \ell^{-i}_{r'_i}(s') + \upsilon_{r'_i}$. Applying Lemma 3.3 with $\beta := \ell^{-i}_{r_i}(s) + \upsilon_{r_i}$ and $\alpha := \ell^{-i}_{r'_i}(s') + \upsilon_{r'_i}$ yields

$$\begin{aligned} b_i(s) &= \ell_{r_i}(s) + \upsilon_{r_i} \\ &= \ell^{-i}_{r_i}(s) + d_i(\beta) + \upsilon_{r_i} \\ &> \ell^{-i}_{r'_i}(s') + d_i(\alpha) + \upsilon_{r'_i} \\ &= \ell_{r'_i}(s') + \upsilon_{r'_i} = b_i(s'). \end{aligned}$$

We show next that if $b_i(\bar{s}) \ge b_i(s)$, then there is $p \in N$ with $b_p(\bar{s}) < b_p(s)$ and $b_i(\bar{s}) \le b_p(\bar{s})$. Let $\bar{r}_i \in b_i^{-1}(\bar{s})$. By (1), we have $b_i(s') < b_i(s) \le b_i(\bar{s})$. Thus, there is $p \in N_{\bar{r}_i}(x')$ with $\bar{d}_p > d'_p$. If $p = i$, we obtain a contradiction observing that $0 \ge U'_i(d'_i) - c'_i(b_i(s')) > U'_i(\bar{d}_i) - c'_i(b_i(\bar{s})) \ge 0$, where for the first and the third inequality we used that d'_i and $\bar{d}_i$ are best replies to s' and $\bar{s}$, respectively.

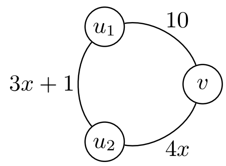

(a) Underlying graph.

Player 1 \ Player 2	$(\{(u_2, v)\}, 1)$	$(\{(u_2, u_1), (u_1, v)\}, 2)$
$(\{(u_1, v)\}, 2)$	2, 9	2, 6
$(\{(u_1, u_2), (u_2, v)\}, 1)$	3, 5	1, 6

(b) Private utility functions of four selected strategy profiles.

Fig. 1. Bottleneck congestion game with elastic demands that lacks a PNE. (a) The underlying graph. (b) The private utility functions for four selected strategy profiles as discussed in the text.

If $p \neq i$, we have $\bar{d}_p > d'_p = d_p$ which gives the inequalities $0 \geq U'_p(d_p) - c'_p(b_p(s))$ and $0 \leq U'_p(\bar{d}_p) - c'_p(b_p(\bar{s}))$ as d_p and $\bar{d}_p$ are best replies to s and $\bar{s}$, respectively. With $U'_p(d_p) > U'_p(\bar{d}_p)$ this implies $b_p(\bar{s}) < b_p(s)$, and since $p \in N_{\bar{r}_i}(x')$, we also have $b_p(\bar{s}) \geq b_i(\bar{s})$.

Next, we show that for every $j \in N \setminus \{i\}$ with $b_j(\bar{s}) > b_j(s)$, there is $p \in N$ with $b_p(\bar{s}) < b_p(s)$ and $b_j(\bar{s}) \leq b_p(\bar{s})$. Suppose there is such j with $b_j(\bar{s}) > b_j(s)$ and denote $\bar{r}_j \in b_j^{-1}(\bar{s})$. Then, one of the following two cases holds:

(a) there is $p \in N_{\bar{r}_j}(x)$ with $\bar{d}_p > d_p$;
(b) $\bar{d}_p \leq d_p$ for all $p \in N_{\bar{r}_j}(x)$ and $i \in N_{\bar{r}_j}(x')$ with $\bar{d}_i \geq \ell_{\bar{r}_j}(\bar{s}) - \ell_{\bar{r}_j}(s) > 0$.

For case (a), we again obtain $0 \geq U'_p(d_p) - c'_p(b_p(s))$ and $0 \leq U'_p(\bar{d}_p) - c'_p(b_p(\bar{s}))$ as d_p and $\bar{d}_p$ are best replies to s and $\bar{s}$, respectively. With $U'_p(d_p) > U'_p(\bar{d}_p)$ this implies $b_p(s) > b_p(\bar{s})$. Since $p \in N_{\bar{r}_j}(x')$, we also get $b_p(\bar{s}) \geq b_j(\bar{s})$.

For case (b), we obtain $b_i(\bar{s}) \geq b_j(\bar{s})$ since $i \in N_{\bar{r}_j}(x')$. For the case that $b_i(\bar{s}) \geq b_i(s)$, we have already shown the existence of another $j' \in N$ with $b_{j'}(\bar{s}) < b_{j'}(s)$ and $b_{j'}(\bar{s}) \geq b_i(\bar{s}) \geq b_j(\bar{s})$. Hence, we may assume $b_i(\bar{s}) < b_i(s)$ which implies the result with $i = p$.

In conclusion, we have shown that whenever $b_j(\bar{s}) > b_j(s)$ for a player j, then there is another player p with $b_p(s) > b_p(\bar{s}) \geq b_j(\bar{s}) > b_j(s)$. If $j = i$, this holds even if $b_i(\bar{s}) \geq b_i(s)$, and we may conclude that $b(\bar{s}) \prec b(s)$. □

We are now in a position to prove correctness of Algorithm 2.

Theorem 5.2. *Algorithm 2 terminates and computes a PNE.*

Proof of Theorem 5.2. Lemma 5.1 shows that the vector $b(s)$ strictly lexicographically decreases during the execution of Algorithm 2. Since for a fixed $x \in X$ the demand equilibrium computed by Algorithm 1 in Phase I of Algorithm 2 is always the same (Remark 4.6), no vector $x \in X$ is visited twice during the execution of Algorithm 2. Using that X contains only finitely many elements, we conclude that the algorithm terminates. □

6. A counterexample for general convex costs

In this section, we show that the assumption that the cost function of each player does not depend on the resource (up to a resource-specific shift) is a necessary assumption that cannot be dispensed with for the existence of a PNE. This is true, even if we impose the additional assumption that for each of the resources the cost functions of all players coincide, and that the strategies of each player i correspond to the set of (u_i, v)-paths in a network.

Theorem 6.1. *There are bottleneck congestion games with elastic demands and affine cost functions (which do not satisfy Assumption (A3)) without a PNE.*

Proof. Consider the graph shown in Fig. 1(a). There are two players associated with the source–sink pairs (u_1, v) and (u_2, v) and demands $\sigma_1 = \sigma_2 = 1$ and $\tau_1 = \tau_2 = 2$. Their utility functions are $U_1(d_1) = d_1 + 10$ for all $d_1 \in [1, 2]$ and $U_2(d_2) = 3d_2 + 10$ for all $d_2 \in [1, 2]$. When going clockwise, player 1 will use the maximal demand $d_1 = 2$ since the cost of (u_1, v) is constant. In contrast, when player 1 goes counterclockwise, for both edges $(u_1, u_2), (u_2, v)$, the congestion increases at least with rate 3 and the utility of player 1 increases only with rate 1. Thus, player 1 will use the minimal demand $d_1 = 1$. Equivalently, player 2 uses the minimal demand when going counterclockwise and the maximal demand going clockwise as in the latter case the bottleneck is always attained at edge (u_1, v). It can be checked that none of these four strategy profiles is a PNE; see Fig. 1(b). The game does not satisfy Assumption (A1) or Assumption (A2), but this can be achieved by adding a small strictly concave or convex function to the utility and cost functions, respectively. □

7. Conclusions

We studied bottleneck routing games where players vary both their rate and their path. As our main result, we derived an algorithm computing a PNE provided that cost functions are non-decreasing, convex and equal up to resource-specific shifts in the argument. This condition is met by the practically relevant $M/M/1$ delay functions with heterogeneous service rates, thus, our result implies the existence of PNE for this model. Our algorithm is centralized but mimics a decentralized dynamic in the following sense. In TCP/IP networks, rate adaption takes place on a shorter time scale than route adaption. In a similar vein, our algorithm first restores a demand equilibrium (where no player has an incentive to vary the demand) before changing a path of a player. Since the choice of the player that adapts the path is arbitrary as long as the path adaption increases the player's utility, this step can be implemented in a distributed manner. The question whether there is a completely decentralized algorithm that allows for asynchronous and concurrent updates is left for future research.

Acknowledgments

We wish to thank two anonymous referees for their intensive work to improve the paper.

The second author's research was carried out in the framework of Matheon supported by Einstein Foundation Berlin (MI8).

References

[1] R. Banner, A. Orda, Bottleneck routing games in communication networks, IEEE J. Sel. Areas Commun. 25 (6) (2007) 1173–1179.
[2] C. Busch, R. Kannan, A. Samman, Bottleneck routing games on grids, in: Proc. 2nd International ICST Conf. on Game Theory for Networks, 2011, pp. 294–307.
[3] C. Busch, M. Magdon-Ismail, Atomic routing games on maximum congestion, Theoret. Comput. Sci. 410 (36) (2009) 3337–3347.
[4] I. Caragiannis, C. Galdi, C. Kaklamanis, Network load games, in: Proc. 16th International Symposium on Algorithms and Computation, 2005, pp. 809–818.

[5] R. Cole, Y. Dodis, T. Roughgarden, Bottleneck links, variable demand, and the tragedy of the commons, in: Proc. 17th Annual ACM-SIAM Sympos. on Discrete Algorithms, 2006, pp. 668–677.
[6] R. Cominetti, C. Guzman, Network congestion control with Markovian multipath routing, Math. Program. A 147 (1–2) (2014) 231–251.
[7] B. de Keijzer, G. Schäfer, O. Telelis, On the inefficiency of equilibria in linear bottleneck congestion games, in: S. Kontogiannis, E. Koutsoupias, P. Spirakis (Eds.), Proc. 3rd Internat. Sympos. Algorithmic Game Theory, in: LNCS, vol. 6386, 2010, pp. 335–346.
[8] H. Han, S. Shakkottai, C.V. Hollot, R. Srikant, D.F. Towsley, Multi-path TCP: a joint congestion control and routing scheme to exploit path diversity in the Internet, IEEE/ACM Trans. Netw. 16 (6) (2006) 1260–1271.
[9] T. Harks, M. Hoefer, M. Klimm, A. Skopalik, Computing pure Nash and strong equilibria in bottleneck congestion games, Math. Program. A 141 (2013) 193–215.
[10] T. Harks, M. Hoefer, K. Schewior, A. Skopalik, Routing games with progressive filling, IEEE/ACM Trans. Netw. 24 (4) (2016) 2553–2562.
[11] T. Harks, M. Klimm, Equilibria in a class of aggregative location games, J. Math. Econom. 61 (2015) 211–220.
[12] T. Harks, M. Klimm, Congestion games with variable demands, Math. Oper. Res. 41 (1) (2016) 255–277.
[13] T. Harks, M. Klimm, R. Möhring, Strong equilibria in games with the lexicographical improvement property, Internat. J. Game Theory 42 (2) (2012) 461–482.
[14] T. Harks, M. Klimm, M. Schneider, Bottleneck routing with elastic demands, in: Proc. 11th Internat. Conference on Web and Internet Econom., 2015, pp. 384–397.
[15] S. Kakutani, A generalization of Brouwer's fixed point theorem, Duke Math. J. 8 (3) (1941) 457–458.
[16] R. Kannan, C. Busch, Bottleneck congestion games with logarithmic price of anarchy, in: S. Kontogiannis, E. Koutsoupias, and P. Spirakis, editors, Proc. 3rd Internat. Sympos. Algorithmic Game Theory, Vol. 6386, 2010, pp. 222–233.
[17] R. Kannan, C. Busch, A.V. Vasilakos, Optimal price of anarchy of polynomial and super-polynomial bottleneck congestion games, in: Proc. 2nd International ICST Conf. on Game Theory for Networks, 2011, pp. 308–320.
[18] F. Kelly, A. Maulloo, D. Tan, Rate control in communication networks: Shadow prices, proportional fairness, and stability, J. Oper. Res. Soc. 49 (1998) 237–252.
[19] P.B. Key, L. Massoulié, D.F. Towsley, Path selection and multipath congestion control, Commun. ACM 54 (1) (2011) 109–116.
[20] P.B. Key, A. Proutiere, Routing games with elastic traffic, SIGMETRICS Perform. Eval. Rev. 37 (2) (2009) 63–64.
[21] A. Khanna, J. Zinky, The revised ARPANET routing metric, SIGCOMM Comput. Commun. Rev. 19 (4) (1989) 45–56.
[22] C.D. Kolstad, L. Mathiesen, Computing Cournot-Nash equilibria, Oper. Res. 39 (5) (1991) 739–748.
[23] Y. Korilis, A. Lazar, On the existence of equilibria in noncooperative optimal flow control, J. ACM 42 (3) (1995) 584–613.
[24] N. Kukushkin, Acyclicity of improvements in games with common intermediate objectives. Russian Academy of Sciences, Dorodnicyn Computing Center, Moscow, 2004.
[25] N. Kukushkin, Congestion games revisited, Internat. J. Game Theory 36 (2007) 57–83.
[26] K. Miller, T. Harks, Utility max–min fair congestion control with time-varying delays, in: Proc. 27th IEEE Internat. Conf. on Computer Comm., 2008, pp. 331–335.
[27] F. Paganini, E. Mallada, A unified approach to congestion control and node-based multipath routing, IEEE/ACM Trans. Netw. 17 (5) (2009) 1413–1426.
[28] L. Qiu, Y. Yang, Y. Zhang, S. Shenker, On selfish routing in Internet-like environments, IEEE/ACM Trans. Netw. 14 (4) (2006) 725–738.
[29] R. Rosenthal, A class of games possessing pure-strategy Nash equilibria, Internat. J. Game Theory 2 (1) (1973) 65–67.
[30] Y. Sheffi, Urban Transportation Networks, Prentice-Hall, Upper Saddle River, NJ, USA, 1985.
[31] S. Shenker, Fundamental design issues for the future Internet, IEEE J. Sel. Areas Commun. 13 (1995) 1176–1188.
[32] R. Srikant, The Mathematics of Internet Congestion Control, Birkhäuser, Basel, Switzerland, 2003.
[33] T. Voice, Stability of multi-path dual congestion control algorithms, IEEE/ACM Trans. Netw. 15 (6) (2007) 1231–1239.
[34] J. Wardrop, Some theoretical aspects of road traffic research, Proc. Inst. Civ. Eng. 1 (Part II) (1952) 325–378.
[35] D. Yang, G. Xue, X. Fang, S. Misra, J. Zhang, A game-theoretic approach to stable routing in max–min fair networks, IEEE/ACM Trans. Netw. 21 (6) (2013) 1947–1959.
[36] Y. Zhang, D. Loguinov, On delay-independent diagonal stability of max–min congestion control, IEEE Trans. Automat. Control 54 (5) (2009) 1111–1116.