Wardrop Equilibria and Price of Stability for Bottleneck Games with Splittable Traffic*

Vladimir Mazalov[1], Burkhard Monien[2], Florian Schoppmann[2], and Karsten Tiemann[2]

[1] Institute of Applied Mathematical Research, Karelian Research Center, Russian Academy of Sciences, 11, Pushkinskaya st., Petrozavodsk, Russia, 185910.
vmazalov@krc.karelia.ru

[2] Faculty of Computer Science, Electrical Engineering and Mathematics, University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany.
{bm,fschopp,tiemann}@uni-paderborn.de

Abstract. We look at the scenario of having to route a continuous rate of traffic from a source node to a sink node in a network, where the objective is to maximize throughput. This is of interest, e.g., for providers of streaming content in communication networks. The overall path latency, which was relevant in other non-cooperative network routing games such as the classic Wardrop model, is of lesser concern here.

To that end, we define bottleneck games with splittable traffic where the throughput on a path is inversely proportional to the maximum latency of an edge on that very path—the bottleneck latency. Therefore, we define a Wardrop equilibrium as a traffic distribution where this bottleneck latency is at minimum on all used paths. As a measure for the overall system well-being—called social cost—we take the weighted sum of the bottleneck latencies of all paths.

Our main findings are as follows: First, we prove social cost of Wardrop equilibria on series parallel graphs to be unique. Even more, for any graph whose subgraph induced by all simple start-destination paths is not series parallel, there exist games having equilibria with different social cost. For the price of stability, we give an independence result with regard to the network topology. Finally, our main result is giving a new exact price of stability for Wardrop/bottleneck games on parallel links with M/M/1 latency functions. This result is at the same time the exact price of stability for bottleneck games on general graphs.

1 Introduction

Motivation and Framework. In recent years, the Wardrop model—which was already introduced in the 1950's (see, e.g., [4,25])—received a lot of attention with regard to analyzing the price of anarchy, i.e., quantifying the worst-case system loss due to selfish behavior of its participants. A Wardrop game can be understood as a game with infinitely many players each in control of a negligible fraction of the total traffic in a

* This work was partially supported by the IST Program of the European Union under contract number IST-15964 (AEOLUS). The research of Florian Schoppmann and Karsten Tiemann was supported by the International Graduate School of Dynamic Intelligent Systems (Paderborn, Germany).

network. Players choose a path according to their respective start and destination node, with the assumption that they act purely selfishly and therefore each takes the fastest path under current network conditions. A situation in which none of the selfish players has an incentive to switch to another path is called a *Wardrop equilibrium*.

While the Wardrop model has been successfully applied for researching road traffic, its basic assumption of drivers minimizing just their own travel time (more generally called path latency) is not appropriate in all networks. For instance, in communication networks such as the Internet, providers of streaming content would strive to maximize the *throughput* to their clients whereas the transmission time is of lesser concern. More generally, for a routing path in this network, one would be interested in the maximum latency of all its edges—in other words, the *latency of the bottleneck*—as it is inversely proportional to the achievable throughput on that very path. From a purely mathematical perspective, the bottleneck latency of a path corresponds to the ∞-norm of the (finite) vector of edge latencies whereas the sum of edge latencies—that was of interest in the Wardrop model—equals the 1-norm. For a broader discussion of when the ∞-norm should be used, confer also Banner and Orda [3]. They mention, for instance, that the ∞-norm is appropriate to model wireless networks where each node has a limited transmission energy. As another motivation, a celebrated result by Leighton et al. [15] implies that the bottleneck latency is also of interest in settings with individual traffic: Their result states that individual packets can be routed in time $\mathcal{O}(\text{congestion} + \text{dilation})$ when the paths for the packets are given in advance. Here dilation denotes the maximum length of a path and congestion denotes the maximum number of paths sharing a common edge.

We address the scenario described in the last paragraphs by studying what we call *bottleneck games with splittable traffic*. Similar to Wardrop games, one could think of infinitely many selfish players each controlling a negligible amount of traffic. However, their objective is now to choose a path such that their experienced bottleneck latency is at minimum. Likewise, we define a Wardrop equilibrium in a game of our new model as a traffic distribution where the bottleneck latency is at minimum on all used paths. As a measure for the overall system well-being—called *social cost*—we take the weighted sum of the bottleneck latencies of all paths. A similar weighted sum over the path latencies was used as social cost for Wardrop games (see e.g. [22]).

In this work, we will consider the degradation of social welfare due to selfish behavior of the players. To that end, we will determine the so called *price of stability*, a term that was coined by Anshelevich et al. [2] and denotes the worst-case ratio, over all instances, between the social cost of the *best* equilibrium and optimum social cost. Roughly speaking, it describes the worst-case inefficiency of optimum stable states, in which no player wants to unilaterally deviate, compared to an overall optimum solution. In contrast, the *price of anarchy*, which was introduced by Koutsoupias and Papadimitriou [14], denotes the worst-case ratio, over all instances, between the social cost of *any* equilibrium and that of social optima.

Related Work. Wardrop Games: Inspired by the arisen interest in the price of anarchy, Roughgarden and Tardos [22] re-investigated the Wardrop games and used the *total latency* as social cost. In this context the price of anarchy was shown to be $\frac{4}{3}$ for linear latency functions [22] and $\Theta(\frac{d}{\ln d})$ for polynomials of degree at most d with non-negative coefficients [21]. Roughgarden [19] proved that the price of anarchy is independent of the network topology if a class $\mathcal{F}$ of latency functions is considered that only fulfills relatively weak assumptions. Instead, it only depends on the so called "anarchy value" $\alpha(\mathcal{F})$ of $\mathcal{F}$, and the worst-case ratio is already achieved on parallel links. Roughgarden [21] also considered networks with *M/M/1 latency functions*. When r is the amount of traffic and $c_{\min} > r$ is the minimum capacity among all edge capacities in the network, an upper bound on the price of anarchy is given by $\frac{1}{2} \cdot (1 + \sqrt{c_{\min}/(c_{\min} - r)})$. Observe that this expression approaches ∞ as the amount of traffic r approaches $c_{\min}$. The upper bound is asymptotically tight even for games on so-called union of paths graphs, i.e., on graphs that consist of many disjoint paths from s to t only having the two nodes s and t in common. Note that the results on the price of stability for games with M/M/1 latency functions that we give in this paper even apply if $c_{\min} \leq r$.

In a recent paper, Cole et al. [8] studied Wardrop-like games where the latency of a path is defined as the p-norm, $1 < p \leq \infty$, of the vector of its edge latencies. In this context, they also looked at "elastic traffic", i.e., some share of the participants might be better off by not traveling at all. When $p = \infty$ and in the case of inelastic traffic, their games are equal to our bottleneck games with splittable traffic. However, they looked at a subclass of Wardrop equilibria that they define as "subpath-optimal", with the reason for their restricting being that otherwise the price of anarchy is infinite even if latency functions are just linear. They showed that the anarchy value is an upper bound on the price of anarchy for subpath optimal equilibria and hence also an upper bound on the price of stability.

There is also some work that focused on the original Wardrop games but did not use total latency to measure the social cost [10,9,20].

Finite Splittable Routing Games: In this setting, a finite number of players with non-negligible effect on each other is given who have to split their traffics over the available paths with the objective to minimize their private costs. Two papers [13,16] studied such games with certain player-specific private cost functions that are based on M/M/1 latency functions. Korilis et al. [13] studied what happens to the private costs of the players if new capacity is added to the network or if existing capacity is reallocated. Orda et al. [16] considered the (non-) uniqueness of Nash equilibria. Banner and Orda [3] studied finite splittable routing games where the private cost of a player is defined as the maximum among all latencies of edges to which this player assigns a non-zero flow, whereas social cost is given by the maximum edge latency in the network. Banner and Orda proved the existence and non-uniqueness of equilibria. They were also able to show that the price of anarchy is unbounded.

Finite Unsplittable Routing Games: Again, there are finitely many players with non-negligible effect on each other and each having to route all its traffic on the same path (see [11] for a survey). Two recent papers [6,7] studied the routing of unsplittable traffics where the private cost of a player is defined as the maximum latency of any edge on its path, i.e., the bottleneck latency in our words. Caragiannis et al. [7] allowed different amounts of traffic for the players, whereas in the setting of Busch and Magdon-Ismail [6] all players control traffic of unit size. Both studied the price of anarchy with respect to social cost defined as the maximum latency of any edge in the network.

M/M/1 Latency Functions: M/M/1 latency functions arise in queuing theory as the expected latency of queues with a Poisson arrival process and an exponentially distributed service time [12,18]. They are used in networking theory to model packet-switched networks. Here, a packet that starts at its entry node in the network or arrives at an intermediate node on its way to the destination is stored in a queue. It can leave the queue as soon as the next link on the path of the packet becomes available [5,24].

Contribution. In this work, we define and study *bottleneck games with splittable traffic*. The ingredients of such a game are a graph $G = (V, E)$, whose edges $e \in E$ are each endowed with a latency function $f_e : \mathbb{R}_0^+ \to \mathbb{R}_0^+ \cup \{\infty\}$, and distinct source and sink nodes $s, t \in V$ between which an arbitrarily splittable traffic $r > 0$ has to be routed. We will give special attention to instances having M/M/1 latency functions, i.e., functions of the form $f_e(u) = \frac{1}{c_e - u}$ where $c_e > 0$ denotes the *capacity* of edge $e \in E$.

Our investigations are two-fold: First, we study general properties of bottleneck games with splittable traffic such as existence and uniqueness of Wardrop equilibria and dependence of both the price of anarchy and stability of the network topology. Most of our results here are based on properties of maximum flows and minimum cuts. In the second part we prove an exact expression for the price of stability for bottleneck games on parallel links with splittable traffic and M/M/1 latency functions. We view this result as the main result of our paper. Especially the proof of the upper bound requires a very careful analysis. In detail, our main findings are:

- General results for bottleneck games with splittable traffic:
  - We define the notion of *capacity* of a network and then show that a bottleneck game with splittable traffic has a Wardrop equilibrium with finite social cost if the traffic is smaller than the network capacity.
  - For games on *series parallel graphs* with arbitrary latency functions we prove the social cost of Wardrop equilibria to be unique. On the other hand, we also show that for any graph whose subgraph induced by all simple start-destination paths is not series parallel, there exist games having equilibria with different social cost.
  - We show that the price of stability for bottleneck games with splittable traffic is independent of the network topology, i.e., the worst-case ratio, over all instances, between the best Nash equilibrium and an optimum is attained on *parallel links*. (See Section 3.3 for a comparison with a similar result by Cole et al. [8].)
- Bottleneck games with splittable traffic and M/M/1 latency functions:
  We prove that the expression

$$\frac{m \cdot \frac{r}{c_{\min}}}{\frac{r}{c_{\min}} + 2 \cdot (m-1) \cdot \left(\sqrt{\frac{r}{c_{\min}} + 1} - 1 \right)} \tag{1}$$

  describes the *exact* price of stability for games on m parallel links with M/M/1 latency functions, minimum edge capacity $c_{\min}$, and traffic r. Furthermore, the expression is increasing in both m and r and it converges to m for large r.

Interestingly, series parallel graphs form exactly the class of graphs where the prices of anarchy and stability conside for every class of latency functions. Parallel links are special series parallel graphs. Furthermore, on parallel link graphs bottleneck games and Wardrop games coincide. Our results imply that for every class of latency functions bounds for the price of stability for parallel link graphs also hold for the price of stability of bottleneck games on arbitrary graphs. This can be used when latency functions are restricted to polynomials where results of Roughgarden [19] can be used, and also for the class of M/M/1 latency functions where the expression (1) describes the price of stability for bottleneck games on arbitrary graphs.

Road Map. The rest of the paper is organized as follows. In Section 2 we give exact definitions for our bottleneck games with splittable traffic. We study general games in Section 3, whereas we restrict ourselves to M/M/1 latency functions in Section 4. Due to lack of space we have to omit most of the proofs.

2 Notation

For all $k \in \mathbb{N}$ denote $[k] = \{1, \dots, k\}$.

Latency Function, Network, Instance. A *latency function* $f : \mathbb{R}_0^+ \to \mathbb{R}_0^+ \cup \{\infty\}$ is a nonnegative, continuous, and nondecreasing function. Here, $\mathbb{R}_0^+ \cup \{\infty\}$ is meant to be endowed with the topology of the one-point compactification of $\mathbb{R}_0^+$. (This basically means that f has no jump discontinuities, not even to ∞.) A *network* is a tuple $(G, s, t, (f_e)_{e\in E})$, where $G = (V, E)$ is a directed *multigraph*, $s, t \in V$ are distinct *source* and *sink* (target) nodes, and the f_e are latency functions. A *bottleneck game with splittable traffic* with *general latency functions* is a tuple $\Gamma = (G, s, t, (f_e)_{e\in E}, r)$ where $(G, s, t, (f_e)_{e\in E})$ is a network in which a *traffic* of $r \in \mathbb{R}^+$ has to be routed from s to t. When obvious, we will usually refer to our bottlenecks games with splittable traffic only as bottleneck games or even just games.

M/M/1 latency function. For ease of notation, we write $\Gamma = (G, s, t, (c_e)_{e\in E}, r)$ for a bottleneck game with splittable traffic and *M/M/1 latency functions* where $c_e > 0$ is the *capacity* for edge $e \in E$. The M/M/1 latency functions f_e, $e \in E$, are implicitly defined by

$$f_e(u) = \begin{cases} \frac{1}{c_e - u} & \text{if } u < c_e \\ \infty & \text{otherwise.} \end{cases}$$

Observe that the latency $f_e(u)$ approaches ∞ as the load u approaches c_e. We denote by $\mathcal{M}$ the set of all M/M/1 latency functions and by $\mathcal{M}_{\geq c} \subset \mathcal{M}$ the functions with a capacity of at least c where $c > 0$.

Strategy Profiles, Wardrop Equilibria, and Social Cost. The traffic r can split arbitrarily over the set $\mathcal{P}_{st}$ of all possible *simple* paths from s to t. A *strategy profile* is a vector $\mathbf{x} = (x_P)_{P\in\mathcal{P}_{st}}$ where $\sum_{P\in\mathcal{P}_{st}} x_P = r$ and $x_P \geq 0$ for all $P \in \mathcal{P}_{st}$. The *load* δ_e on an edge $e \in E$ is given by $\delta_e(\mathbf{x}) = \sum_{P\in\mathcal{P}_{st}, P\ni e} x_P$. A strategy profile is a *Wardrop equilibrium* if the latency of each used path is not larger than the latency of any other path, i.e., if for all $P, R \in \mathcal{P}_{st}$

$$x_P > 0 \quad \Rightarrow \quad \max_{e\in P} f_e(\delta_e(\mathbf{x})) \leq \max_{e\in R} f_e(\delta_e(\mathbf{x})).$$

The *social cost* of a strategy profile $\mathbf{x}$ is defined as the “canonically” weighted sum of all path latencies, i.e.,

$$\mathsf{SC}(\Gamma, \mathbf{x}) = \sum_{P \in \mathcal{P}_{st}} x_P \cdot \max_{e \in P} f_e(\delta_e(\mathbf{x})).$$

If $\mathbf{x}$ is a Wardrop equilibrium, $l(\mathbf{x}) = \frac{\mathsf{SC}(\Gamma,\mathbf{x})}{r}$ denotes the unique latency of all paths with non-zero flow. The *optimum* associated with a bottleneck game with splittable traffic Γ is the minimum social cost of any strategy profile: $\mathsf{OPT}(\Gamma) = \min_{\mathbf{x}} \mathsf{SC}(\Gamma, \mathbf{x})$. The *price of anarchy* (PoA) and *price of stability* (PoS) for a set $\mathscr{G}$ of games are defined as

$$\mathsf{PoA}(\mathscr{G}) := \sup_{\substack{\Gamma \in \mathscr{G} \\ \mathbf{x} \text{ Wardr. Equ. in } \Gamma}} \frac{\mathsf{SC}(\Gamma, \mathbf{x})}{\mathsf{OPT}(\Gamma)} \quad \text{and} \quad \mathsf{PoS}(\mathscr{G}) := \sup_{\Gamma \in \mathscr{G}} \inf_{\mathbf{x} \text{ Wardr. Equ. in } \Gamma} \frac{\mathsf{SC}(\Gamma, \mathbf{x})}{\mathsf{OPT}(\Gamma)}$$

where by definition $\infty/\infty := 1$ and $0/0 := 1$. Furthermore, $u/0 := \infty$ if $u > 0$. For a given network $(G, s, t, (f_e)_{e \in E})$ its *capacity* is given by

$$C(G, s, t, (f_e)_{e \in E}) = \sup \left\{ r \in \mathbb{R}_0^+ \,\middle|\, \begin{array}{l} \exists \text{ strategy profile } \mathbf{x} \text{ with } \mathsf{SC}(\Gamma, \mathbf{x}) < \infty \\ \text{for } (G, s, t, (f_e)_{e \in E}, r) \end{array} \right\} \cup \{0\}.$$

Series Parallel Graphs. (Sometimes also called two terminal series parallel.) *Series parallel* is a recursively defined property: As the base case, the graph that only consists of two nodes s, t and a single edge (s, t) is series parallel with terminals (s, t). An arbitrary multigraph G is series parallel with terminals (s, t) if it can be constructed from two series parallel graphs with terminals (s_1, t_1) and (s_2, t_2) connected either in series or in parallel. In a series connection, $t_1 = s_2$, $s = s_1$, and $t = t_2$. In a parallel connection, $s = s_1 = s_2$ and $t = t_1 = t_2$.

Parallel Links. A graph of *parallel links* is a multigraph $G = (V, E)$ consisting of two nodes $V = \{s, t\}$ and m edges $E = \{1, 2, \ldots, m\}$ from s to t. Whenever a bottleneck game with splittable traffic and M/M/1 latency functions on parallel links is considered we assume that $c_1 \geq \ldots \geq c_m$ and denote $C := \sum_{i=1}^{m} c_i$ and $C^{\leq i} := \sum_{k=1}^{i} c_k$. Clearly, C is just the capacity of the network.

For a non-empty set $\mathcal{F}$ of latency functions we define $\mathscr{G}(\mathcal{F})$ as the set of all bottleneck games with latency functions drawn from $\mathcal{F}$. The subset $\mathscr{P}(\mathcal{F}) \subset \mathscr{G}(\mathcal{F})$ consists of all games in $\mathscr{G}(\mathcal{F})$ that are defined on a graph of parallel links. To further differentiate we denote by $\mathscr{G}(\mathcal{F}, m, r) \subset \mathscr{G}(\mathcal{F})$ the set of games with at most m edges and a traffic of at most r. Likewise, $\mathscr{P}(\mathcal{F}, m, r) := \mathscr{G}(\mathcal{F}, m, r) \cap \mathscr{P}(\mathcal{F})$.

3 General Results for Bottleneck Games with Splittable Traffic

For general bottleneck games with splittable traffic we will prove the existence of Wardrop equilibria (Section 3.1), study the (non-)uniqueness of equilibria social cost (Section 3.2), and show that the price of stability is independent of the the network topology (Section 3.3).

3.1 Existence of Wardrop Equilibria

Existence of Wardrop equilibria in bottleneck games with splittable traffic can be established by employing the general result of [23] (for a proof using more elementary maths, see [17]). To illustrate the connection to maximum flows, however, we start by giving a proof that makes use of the max-flow min-cut theorem (see, e.g., [1]). The construction described in the proof will also be used in the proof of Theorem 5. Obviously, the only interesting case is the traffic to be routed being smaller than the capacity of the network.

Theorem 1. *Let $\Gamma = (G, s, t, (f_e)_{e\in E}, r)$ be a bottleneck game with splittable traffic where $r < C(G, s, t, (f_e)_{e\in E})$. Then Γ possesses a Wardrop equilibrium of finite social cost.*

3.2 Uniqueness Results about Social Cost of Equilibria

We will show in this section that different equilibria for a bottleneck game with splittable traffic on a series parallel graph have the same social cost. The proof for this result employs a technique based on what we define as *strong cuts*.

Definition 1. *Let Γ be a bottleneck game with splittable traffic on a series parallel graph $G = (V, E)$ and let $\mathbf{x}$ be a Wardrop equilibrium for Γ. Then $D \subseteq E$ is called* strong cut *with respect to Γ and $\mathbf{x}$ if*

1. *each path $P \in \mathcal{P}_{st}$ contains exactly one edge that belongs to D, and*
2. *$f_e(\delta_e(\mathbf{x})) \geq l(\mathbf{x})$ for all edges $e \in D$.*

Observe that, given a strong cut D with respect to Γ and an equilibrium $\mathbf{x}$, all edges $e \in D$ with $\delta_e(\mathbf{x}) > 0$ have latency $l(\mathbf{x})$ whereas all other edges $e \in D$ with $\delta_e(\mathbf{x}) = 0$ have latency at least $l(\mathbf{x})$. Before making use of the crucial properties of strong cuts, we need to ensure their existence.

Theorem 2. *Let Γ be a bottleneck game with splittable traffic on a series parallel graph and let $\mathbf{x}$ be a Wardrop equilibrium for Γ. Then a strong cut with respect to Γ and $\mathbf{x}$ exists.*

Proof. The proof is by structural induction over all series parallel graphs. Our *induction hypothesis* is that every series parallel graph G with terminals (s, t) has the following property: For any bottleneck game Γ on G and all Wardrop equilibria of Γ, there is a strong cut.

The only *base case* to verify consists of the graph with two nodes s, t solely connected by the edge e. Obviously, for any game Γ on this graph, $\{e\}$ is a strong cut with respect to Γ and its trivial equilibrium. For the *induction step*, consider any arbitrary graph $G = (V, E)$ with terminals (s, t). Furthermore, assume that G is a series parallel connection of two series parallel graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ with terminals (s_1, t_1) and (s_2, t_2), respectively, and both G_1 and G_2 fulfill the induction hypothesis. To prove the induction step, we then have to show that G fulfills the induction hypothesis, too. Thus, let $\Gamma = (G, s, t, (f_e)_{e\in E}, r)$ be an arbitrary game on G and $\mathbf{x}$ be an arbitrary Wardrop equilibrium for Γ and consider the two cases:

Parallel Connection: Set $r_1 = \sum_{P \in \mathcal{P}_{s_1 t_1}} x_P$ and $r_2 = \sum_{P \in \mathcal{P}_{s_2 t_2}} x_P$ where $\mathcal{P}_{s_1 t_1}$ and $\mathcal{P}_{s_2 t_2}$ are meant to only contain paths from G_1 and G_2, respectively. Obviously, $r_1 + r_2 = r$ and the games $\Gamma_1 = (G_1, s_1, t_1, (f_e)_{e \in E_1}, r_1)$ and $\Gamma_2 = (G_2, s_2, t_2, (f_e)_{e \in E_2}, r_2)$ have Wardrop equilibria $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ where $x_P^{(1)} = x_P$ for all paths P with edges in E_1 and $x_P^{(2)} = x_P$ for all paths P with edges in E_2. It follows by the induction hypothesis that there are strong cuts D_1 and D_2 with respect to $G_1, \mathbf{x}^{(1)}$ and $G_2, \mathbf{x}^{(2)}$ that we can use to get a strong cut $D := D_1 \cup D_2$ for Γ and its equilibrium $\mathbf{x}$.
Series Connection: Consider the games $\Gamma_1 = (G_1, s_1, t_1, (f_e)_{e \in E_1}, r)$ and $\Gamma_2 = (G_2, s_2, t_2, (f_e)_{e \in E_2}, r)$. Obviously, $\mathbf{x}$ induces strategy profiles $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ for Γ_1 and Γ_2, respectively. At least one of $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ is a Wardrop equilibrium. Otherwise, there would be a path $P \in \mathcal{P}_{st}$ with non-zero flow on which the latency is larger than on another path $R \in \mathcal{P}_{st}$, and $\mathbf{x}$ cannot be a Wardrop equilibrium. If $\mathbf{x}^{(1)}$ is an equilibrium we set $D := D_1$ and $D := D_2$ otherwise. In either case, D is a strong cut for Γ and its equilibrium $\mathbf{x}$. □

We now use strong cuts in the proof of the next theorem to show that all Wardrop equilibria of a bottleneck game on a series parallel graph have the same social cost. Obviously, this implies that the price of stability does not differ from the price of anarchy for this class of games, i.e., $\mathsf{PoA}(\mathscr{S}) = \mathsf{PoS}(\mathscr{S})$ for any set $\mathscr{S}$ of bottleneck games on series parallel graphs.

Theorem 3. *Let Γ be a bottleneck game with splittable traffic on a series parallel graph, let $\hat{\mathbf{x}}$ and $\mathbf{x}$ be two Wardrop equilibria for Γ. Then* $\mathsf{SC}(\Gamma, \hat{\mathbf{x}}) = \mathsf{SC}(\Gamma, \mathbf{x})$.

Proof. The proof is by contradiction. Assume that two different Wardrop equilibria $\hat{\mathbf{x}}$ and $\mathbf{x}$ for $\Gamma = (G, s, t, (f_e)_{e \in E}, r)$ are given such that $\mathsf{SC}(\Gamma, \hat{\mathbf{x}}) < \mathsf{SC}(\Gamma, \mathbf{x})$. Clearly, $l(\hat{\mathbf{x}}) < l(\mathbf{x})$. Let D be a strong cut with respect to Γ and $\mathbf{x}$. Consider first of all an edge $e \in D$ with $\delta_e(\hat{\mathbf{x}}) > 0$. Since $\hat{\mathbf{x}}$ is a Wardrop equilibrium and e is an edge of the strong cut D with respect to Γ and $\mathbf{x}$, we get that $f_e(\delta_e(\hat{\mathbf{x}})) \leq l(\hat{\mathbf{x}}) < l(\mathbf{x}) \leq f_e(\delta_e(\mathbf{x}))$, which implies $\delta_e(\hat{\mathbf{x}}) < \delta_e(\mathbf{x})$ because f_e is nondecreasing. If instead an edge $e \in D$ with $\delta_e(\hat{\mathbf{x}}) = 0$ is considered we trivially obtain that $\delta_e(\hat{\mathbf{x}}) \leq \delta_e(\mathbf{x})$. Together we get that

$$r = \sum_{e \in D} \delta_e(\hat{\mathbf{x}}) \ < \ \sum_{e \in D} \delta_e(\mathbf{x}) = r,$$

which is a contradiction. □

We will now consider general graphs. If traffic is sent through a graph G then only edges that are on a simple path from s to t can be used. So the same equilibria are obtained when playing the game not on G but on the maximum subgraph of G containing only edges that are on a simple path from s to t, i.e., the subgraph induced by all paths from s to t. This idea is captured by the following definition.

Definition 2. *A directed multigraph $G = (V, E)$ without isolated vertices where $s, t \in V$, $s \neq t$, is called* strongly (s,t)-connected *if every edge $e \in E$ is contained in a simple path from s to t.*

Theorem 4. *Let G be a strongly (s,t)-connected graph that is not series parallel. Then there exists a bottleneck game $\Gamma = (G, s, t, (f_e)_{e \in E}, r)$ possessing Wardrop equilibria of different social cost.*

3.3 Price of Stability

In this section we will show that the price of stability for bottleneck games with splittable traffic and latency functions from an arbitrary non-empty set of nonnegative, continuous, and nondecreasing functions $\mathcal{F}$ is the same on general graphs as on parallel links, i.e., $\mathsf{PoS}(\mathscr{G}(\mathcal{F})) = \mathsf{PoS}(\mathscr{P}(\mathcal{F}))$. To do so, we will show that given a game Γ on a general graph with latency functions from $\mathcal{F}$ there exists a game Γ' on parallel links with latency functions from $\mathcal{F}$ and Wardrop equilibria $\mathbf{x}$ for Γ and $\hat{\mathbf{x}}$ for Γ' such that $\frac{\mathsf{SC}(\Gamma,\mathbf{x})}{\mathsf{OPT}(\Gamma)} \leq \frac{\mathsf{SC}(\Gamma',\hat{\mathbf{x}})}{\mathsf{OPT}(\Gamma')}$.

We assume that Cole et al. [8] proved a very similar result to establish their Theorem 4.6 (whose proof they had to omit due to lack of space). Since we need a rather technical formulation for our main result on the price of stability for games with M/M/1 latency functions, we give the following Theorem 5.

Theorem 5. *Let $\Gamma = (G, s, t, (f_e)_{e\in E}, r)$, $G = (V, E)$, be a bottleneck game with splittable traffic where $r < C(G, s, t, (f_e)_{e\in E})$. Then there exist*

- *a bottleneck game with splittable traffic on parallel links $\Gamma' = (G', s', t', (f'_e)_{e'\in E'}, r)$, $G' = (V', E')$, where $|E'| \leq |E|$ and for each $e' \in E'$ there is an edge $e \in E$ such that $f'_{e'} = f_e$ and*
- *Wardrop equilibria $\mathbf{x}$ for Γ and $\hat{\mathbf{x}}$ for Γ',*

such that $\frac{\mathsf{SC}(\Gamma,\mathbf{x})}{\mathsf{OPT}(\Gamma)} \leq \frac{\mathsf{SC}(\Gamma',\hat{\mathbf{x}})}{\mathsf{OPT}(\Gamma')}$.

Recall that in the case of parallel links bottleneck games do not differ from Wardrop games and hence the prices of stability coincide. This, together with Theorem 5 implies that the price of stability for bottleneck games on arbitrary graphs corresponds to the price of stability (or anarchy) for Wardrop games on parallel links. Consequently, the results by Roughgarden [19] on the price of anarchy for Wardrop games lead to the following corollary.

Corollary 1. *Let $\Gamma = (G, s, t, (f_e)_{e\in E}, r)$ be a bottleneck game with splittable traffic where all functions f_e, $e \in E$, are polynomials of degree at most d with non-negative coefficients. Then there exists a Wardrop equilibrium $\mathbf{x}$ where*

$$\frac{\mathsf{SC}(\Gamma,\mathbf{x})}{\mathsf{OPT}(\Gamma)} \leq \frac{(d+1)\cdot\sqrt[d]{d+1}}{(d+1)\cdot\sqrt[d]{d+1}-d}\,.$$

For readers who are familiar with the anarchy value defined in [19] we would like to mention that it is possible to draw a more general conclusion that is also the result of Cole et al. [8, Theorem 4.6]: If the anarchy value $\alpha(\mathcal{F})$ exists for a set of functions $\mathcal{F}$ this value $\alpha(\mathcal{F})$ is an upper bound on the price of stability for bottleneck games on general graphs and with latency functions from $\mathcal{F}$, i.e., $\mathsf{PoS}(\mathscr{G}(\mathcal{F})) \leq \alpha(\mathcal{F})$. Under some moderate assumptions being made on $\mathcal{F}$ even equality holds. This result, however, cannot be used to prove our main finding, i.e., the price of stability for bottleneck games with M/M/1 latency functions, since we will include other game properties in the sets of games under consideration. Therefore, Theorem 5 is essential for the generalization from parallel links to arbitrary graphs in the M/M/1 case.

4 Bottleneck Games with Splittable Traffic and M/M/1 Functions

In the rest of this paper we focus on bottleneck games with splittable traffic and M/M/1 latency functions. We want to remark that in this setting there are instances for which social cost of Wardrop equilibria may be arbitrarily worse than those of an optimum. This justifies looking at the price of stability instead. Unfortunately, also $\mathsf{PoS}(\mathscr{G}(\mathcal{M})) = \mathsf{PoS}(\mathscr{P}(\mathcal{M})) = \infty$, which will be a trivial consequence of Theorem 8. Hence, we need to consider other game properties, too, in order to get a meaningful result for the price of stability. To achieve this goal we will derive the exact value for $\mathsf{PoS}(\mathscr{P}(\mathcal{M}_{\geq c}, m, r))$ where $m \in \mathbb{N}$, $c > 0$, $r > 0$, and then argue that it is the same as $\mathsf{PoS}(\mathscr{G}(\mathcal{M}_{\geq c}, m, r))$. By our notation, m is meant here to denote the maximum number of edges, c the minimum edge capacity, and r the maximum amount of traffic.

4.1 Social Cost of Equilibria and Optimum Solutions in the Parallel Link Case

For our later proofs on the price of stability we need some insight into the social cost of Wardrop equilibria and optimum solutions. Thus we now give the exact social cost of Wardrop equilibria.

Theorem 6. *Let $\Gamma = (G, s, t, (c_e)_{e \in E}, r)$ be a bottleneck game with splittable traffic and M/M/1 latency functions on m parallel links where $r < C$, and let $\mathbf{x}$ be a Wardrop equilibrium. Furthermore, let $s = |\{i \in [m] \mid x_i > 0\}|$ denote the number of links used in $\mathbf{x}$. Then*

$$s = \max\left\{i \in [m] \mid r + i \cdot c_i > C^{\leq i}\right\} \quad \textit{and} \quad \mathsf{SC}(\Gamma, \mathbf{x}) = \frac{s \cdot r}{C^{\leq s} - r}.$$

We will now derive an expression that describes the social cost of an optimum solution.

Theorem 7. *Let $\Gamma = (G, s, t, (c_e)_{e \in E}, r)$ be a bottleneck game with splittable traffic and M/M/1 latency functions on m parallel links where $r < C$, and let $\mathbf{x}$ be a strategy profile with optimal social cost. Furthermore, let $t = |\{i \in [m] \mid x_i > 0\}|$ denote the number of links used in $\mathbf{x}$. Then:*

$$t = \max\left\{i \in [m] \mid r + \sqrt{c_i} \cdot \sum_{k=1}^{i} \sqrt{c_k} > C^{\leq i}\right\} \textit{ and } \mathsf{OPT}(\Gamma) = \frac{\left(\sum_{i=1}^{t} \sqrt{c_i}\right)^2}{C^{\leq t} - r} - t.$$

4.2 Price of Stability for Games on Parallel Links

Combining our knowledge about the social cost of Wardrop equilibria and optimum solutions, we will now give the exact price of stability for games on $m \in \mathbb{N}$ parallel links routing a traffic of $r > 0$. To that end, we require the capacities $c_1, \ldots, c_m$ and the traffic r to be normalized such that $c_m = 1$, i.e., we will derive an exact expression for $\mathsf{PoS}(\mathscr{P}(\mathcal{M}_{\geq 1}, m, r))$. This is not a restriction as for an $\alpha > 0$ the bijective mapping $\Gamma = (G, s, t, (c_e)_{e \in E}, r) \mapsto \Gamma_\alpha := (G, s, t, (\alpha \cdot c_e)_{e \in E}, \alpha \cdot r)$ associates both the Wardrop equilibrium and optimum in Γ with the respective equilibrium and optimum in Γ_α. Note that, if $(x_P)_{P \in \mathcal{P}_{st}}$ is a strategy profile in Γ, $(\alpha \cdot x_P)_{P \in \mathcal{P}_{st}}$ is a strategy profile in Γ_α with social costs α times as much as that of $(x_P)_{P \in \mathcal{P}_{st}}$ in Γ. Hence, $\mathsf{PoS}(\mathscr{P}(\mathcal{M}_{\geq c}, m, r)) = \mathsf{PoS}(\mathscr{P}(\mathcal{M}_{\geq 1}, m, \frac{r}{c}))$, where again $m \in \mathbb{N}$ denotes the maximum number of edges, $c > 0$ the minimum edge capacity, and $r > 0$ the maximum amount of traffic.

Theorem 8. *For bottleneck games $\Gamma = (G, s, t, (c_e)_{e \in E}, r)$ with splittable traffic and M/M/1 latency functions on m parallel links where $c_m = 1$ the price of stability is exactly*

$$\mathsf{PoS}(\mathscr{P}(\mathcal{M}_{\geq 1}, m, r)) = \frac{m \cdot r}{r + 2 \cdot (m-1) \cdot (\sqrt{r+1} - 1)} =: \Psi(m, r).$$

Note the following properties of $\Psi(m, r)$:

- $\Psi(m, r)$ is strictly increasing in both $m \in \mathbb{N}$ and $r > 0$. To see the latter it can similarly be shown with standard methods that $\frac{\partial}{\partial r}\Psi(m, r) > 0$.
- $\lim_{m \to \infty} \Psi(m, r) = \frac{r}{2\sqrt{r+1}+2}$ and $\lim_{r \to \infty} \Psi(m, r) = m$, hence we always have the bound $\Psi(m, r) \leq m$.

We conclude this section by the following corollary which is a direct consequence of the preceding result together with Theorem 5.

Corollary 2. *The price of stability for general bottleneck games with M/M/1 latency functions on a graph with no more than $m \in \mathbb{N}$ edges each with a minimum capacity of at least $c > 0$ and traffic at most $r > 0$ is the same as in the parallel links case, i.e.,* $\mathsf{PoS}(\mathscr{G}(\mathcal{M}_{\geq c}, m, r)) = \mathsf{PoS}(\mathscr{P}(\mathcal{M}_{\geq c}, m, r))$.

Acknowledgment. We would like to thank Marios Mavronicolas for many helpful comments.

References

1. R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.
2. E. Anshelevich, A. Dasgupta, J. Kleinberg, T. Roughgarden, É. Tardos, and T. Wexler. The Price of Stability for Network Design with Fair Cost Allocation. In *Proc. of the 45th Annual Symposium on Foundations of Computer Science (FOCS'04)*, 2004.
3. R. Banner and A. Orda. Bottleneck Routing Games in Communication Networks. In *Proc. of the 25th Conference on Computer Communications (INFOCOM'06)*, to appear 2006.
4. M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
5. D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, 1987.
6. C. Busch and M. Magdon-Ismail. Atomic Routing Games on Maximum Congestion. In *Proc. of the 2nd International Conference on Algorithmic Aspects in Information and Management (AAIM'06)*, to appear 2006.
7. I. Caragiannis, C. Galdi, and C. Kaklamanis. Network Load Games. In X. Deng and D. Du, editors, *Proc. of the 16th Annual International Symposium on Algorithms and Computation (ISAAC'05)*, LNCS Vol. 3827, Springer Verlag, pages 809–818, 2005.

8. R. Cole, Y. Dodis, and T. Roughgarden. Bottleneck Links, Variable Demand, and the Tragedy of the Commons. In *Proc. of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'06)*, pages 668–677, 2006.
9. J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. Computational Complexity, Fairness, and the Price of Anarchy of the Maximum Latency Problem. In D. Bienstock and G. Nemhauser, editors, *Proc. of the 10th Conference on Integer Programming and Combinatorial Optimization (IPCO'04)*, LNCS Vol. 3064, Springer Verlag, pages 59–73, 2004. An extended version will appear in *Operations Research*.
10. A. Czumaj, P. Krysta, and B. Vöcking. Selfish Traffic Allocation for Server Farms. In *Proc. of the 34th Annual ACM Symposium on Theory of Computing (STOC'02)*, pages 287–296, 2002.
11. M. Gairing, T. Lücking, B. Monien, and K. Tiemann. Nash Equilibria, the Price of Anarchy and the Fully Mixed Nash Equilibrium Conjecture. In *Proc. of the 32nd International Colloquium on Automata, Languages, and Programming (ICALP'05)*, LNCS Vol. 3580, Springer Verlag, pages 51–65, 2005.
12. L. Kleinrock. *Queueing Systems, Volume 1: Theory*. Wiley-Interscience, 1975.
13. Y. A. Korilis, A. A. Lazar, and A. Orda. Capacity Allocation Under Noncooperative Routing. *IEEE Transactions on Automatic Control*, 42(3):309–325, 1997.
14. E. Koutsoupias and C. H. Papadimitriou. Worst-Case Equilibria. In C. Meinel and S. Tison, editors, *Proc. of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS'99)*, LNCS Vol. 1563, Springer Verlag, pages 404–413, 1999.
15. F. T. Leighton, B. M. Maggs, and S. B. Rao. Packet Routing and Job-Scheduling in O(Congestion + Dilation) steps. *Combinatroica*, 14(2):167–186, 1994.
16. A. Orda, R. Rom, and N. Shimkin. Competitive Routing in Multiuser Communication Networks. *IEEE/ACM Transactions on Networking*, 1(5):510–521, 1993.
17. K. P. Rath. A direct proof of the existence of pure strategy equilibria in games with a continuum of players. *Economic Theory*, 2(3):427–433, 1992.
18. T. G. Robertazzi. *Computer Networks and Systems: Queuing Theory and Performance Evaluation*. Springer-Verlag, 2000.
19. T. Roughgarden. The Price of Anarchy is Independent of the Network Topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.
20. T. Roughgarden. The Maximum Latency of Selfish Routing. In *Proc. of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'04)*, pages 980–981, 2004.
21. T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
22. T. Roughgarden and É. Tardos. How Bad Is Selfish Routing? *Journal of the ACM*, 49(2):236–259, 2002.
23. D. Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7(4):295–300, 1973.
24. M. Schwarz. *Telecommunication Networks: Protocols, Modeling and Analysis*. Addison-Wesley, 1987.
25. J. G. Wardrop. Some Theoretical Aspects of Road Traffic Research. In *Proc. of the Institute of Civil Engineers, Pt. II, Vol. 1*, pages 325–378, 1952.