

Nonadaptive Selfish Routing with Online Demands

Tobias Harks¹ and László A. Végh^{2,*}

¹ Institute of Mathematics, Technical University Berlin, 10623 Berlin, Germany
`harks@math.tu-berlin.de`

² Department of Operations Research, Eötvös University, Budapest, Hungary, H-1117
`veghal@cs.elte.hu`

Abstract. We study the efficiency of selfish routing problems in which traffic demands are revealed *online*. We go beyond the common Nash equilibrium concept in which possibly all players reroute their flow and form a new equilibrium upon arrival of a new demand.

In our model, demands arrive in n sequential games. In each game, the new demands form a Nash equilibrium and their routings remain unchanged afterwards. We study the problem both with nonatomic and atomic player types and with continuous and nondecreasing latency functions on the edges. For polynomial latency functions, we give constant upper and lower bounds on the competitive ratio of the resulting online routing in terms of the maximum degree, the number of games and in the atomic setting the number of players. In particular, for nonatomic players and affine latency functions we show that the competitive ratio is at most $\frac{4n}{n+2}$. Finally, we present improved upper bounds for the special case of two nodes connected by parallel arcs.

1 Introduction

Recent contributions in the field of algorithmic game theory provided much insight into the structure and efficiency of Nash equilibria in networks that lack a central coordination. Among others, a prominent result in this field states that the *price of anarchy* for a nonatomic selfish routing game, is bounded by a small constant depending on the class of feasible latency functions, see Roughgarden and Tardos [30], Roughgarden [29], and Correa Schulz, and Stier-Moses [11]. It is well known that this kind of games applies to the source routing concept in telecommunication networks, see Qiu, Yang, Zhang, and Shenker [25] and Friedman [19] for an engineering perspective and Roughgarden [28] and Altman, Basar, Jimenez, and Shimkin [1] for a theoretical perspective on this topic. In the source routing model, sources are responsible for selecting paths to route data to the corresponding sink.

* Supported by the Hungarian National Foundation for Scientific Research, OTKA K60802 and NK 67867 and by European MCRTN ADONET, Grant Number 504438. Work was done while visiting the Konrad Zuse Institute in Berlin.

The main focus of the research done so far regarding the source routing concept is to quantify the efficiency loss of a Nash equilibrium compared to the system optimum. Here, one assumption is crucial: if the traffic matrix changes, all sources may possibly change their routes and converge to a new equilibrium, see Even-Dar and Mansour [16] for a further discussion about the convergence behavior. This assumption, however, has some important implications: Each source would have to *continuously* maintain the current state of all available routes, which in turn introduces additional traffic overhead by signaling these needed informations. Furthermore, frequent rerouting attempts during data transmission may not only produce transient load oscillations as observed by Fischer and Vöcking [18], but may also interfere with the widely used congestion control protocol TCP that determines the data rate, as reported by La, Walrand, and Anantharam in [24]. For these reasons, rerouting attempts in reaction to traffic changes in the network are not necessarily beneficial and efficient.

In this paper, we study a different model in which demands of players are released in n sequential games in an online fashion. In each game, the new demands form a Nash equilibrium, and their routing remains unchanged afterwards, that is, the routing becomes *nonadaptive*.

We can interpret this model as follows. Let us introduce a cost for each player quantifying the *cost of rerouting* after some initial time frame. Within the standard equilibrium concept, rerouting comes at no cost. On the other hand, if this rerouting cost is sufficiently large for each player, then, fixing the initial equilibrium routing is the best response strategy.

If rerouting is not allowed in general, then, the problem of finding efficient routings becomes an *online* optimization problem. In this regard, nonadaptive selfish routing constitutes an online algorithm, where the goal is to minimize average congestion cost for all commodities. We present two *distributed* online algorithms, called NSEQNASH and ASEQNASH for this setting. Upon release of a set of commodities (network game), the online algorithm NSEQNASH routes the commodity such that the flow is at Nash equilibrium provided *nonatomic* agents are carrying the flow. The *atomic* splittable variant is given by ASEQNASH.

1.1 Related Work

The fact that the cost of a Nash equilibrium may strictly exceed that of a system optimum is well known in the transportation literature, see Braess [6] and Dubey [15]. A first successful attempt to exactly quantify this so called “price of anarchy” is given by Papadimitriou and Koutsoupias [23] in the context of a load balancing game in communication networks. Roughgarden and Tardos [30] studied the price of anarchy in nonatomic selfish routing games. In nonatomic games, a large number of players is assumed, each consuming an infinitesimal part of the resources. In particular, they proved for affine latency functions a bound of $\frac{4}{3}$ on the price of anarchy. A series of several other follow-up papers analyzed the price of anarchy for more general cost functions and model features; see for example Czumaj and Vcking [13], Correa Schulz, and Stier-Moses [11], and Roughgarden [28].

For atomic routing games, that is, some players may control a significant part of the entire demand, Roughgarden and Tardos [30] examined the price of anarchy for unsplittable flow. Awerbuch, Azar, and Epstein [2] and Christodoulou and Koutsoupias [9] studied the price of anarchy for linear atomic congestion games. Cominetti, Correa, and Stier-Moses [10] presented new bounds on the price of anarchy for splittable atomic routing games that revised previous work of Roughgarden [29] and Correa, Schulz, and Stier-Moses [12]. Hayrapetyan, Tardos, and Wexler [22] improved these bounds for special network topologies.

In the online routing field, several papers considered online load balancing in the context of machine scheduling. Awerbuch et al. [3] considered a greedy online load balancing strategy, where the goal is to minimize the L_2 norm of the aggregated server loads. Similar to this paper, Suri et al. [31] and Caragiannis et al. [7] studied Nash solutions for every released job and showed that the resulting online algorithm outperforms the greedy strategy of [3]. These results, however, are restricted to m parallel arcs and all jobs have to be assigned to exactly one machine. In the paper by Awerbuch, Azar, and Plotkin [4], online routing algorithms are presented to maximize throughput under the assumption that routings are irrevocable. They presented online algorithms whose competitive bounds depend on the number of nodes in the network.

Our work is motivated by the paper by Harks, Heinz, and Pfetsch [21], where online multicommodity routing problems are considered. They considered affine latency functions and presented a greedy online algorithm for a different convex cost function that is $\frac{4K^2}{(1+K)^2}$ competitive, where K is the number of commodities. In their framework, only single demands are released consecutively.

1.2 Our Results and Techniques

We introduce the framework *Online Network Games* (ONLINENG) to analyze *nonadaptive selfish routing* under the assumption that demands (network games) are released *online*. For the online algorithm NSEQNASH that is characterized by selfish routing of *nonatomic* players for a sequence of network games, we obtain the following results. The online algorithm NSEQNASH that produces a flow that is at Nash equilibrium for every game is $\frac{4n}{2+n}$ -competitive for affine latency functions, where n is the number of games within a given sequence. This result contains the bound on the price of anarchy of $\frac{4}{3}$ for affine latency functions of Roughgarden and Tardos [30] as a special case of our result, where $n = 1$. We prove a lower bound of $\frac{3n-2}{n}$ of NSEQNASH showing that for $n = 2$, the upper bound is tight. For linear latency functions, we further improve this bound to $\frac{4n^2}{(1+n)^2}$. For polynomial latency functions with nonnegative coefficients, we prove lower and upper bounds on the competitive ratio of NSEQNASH that grow both exponentially in the degree of the considered polynomials. We further show that for parallel arcs, the competitive ratio is significantly lower. In particular, we show that in this case, the competitive ratio of the online algorithm NSEQNASH does not exceed the price of anarchy of a related nonatomic network game in which all games of a given sequence are considered at the same time.

Furthermore, we consider online network games in which atomic players route their demand selfishly. Note that the atomic players may split their flow along different paths. The online algorithm ASEQNASH, which produces a flow that is at Nash equilibrium for every game is $\min\{\frac{2(3\mathcal{K}+1)n}{n\mathcal{K}+3n+3\mathcal{K}+1}, \frac{5\mathcal{K}+1}{\mathcal{K}+5}, 4.92\}$ -competitive for affine latency functions. Here, \mathcal{K} denotes the total number of players and n is the number of games within a given sequence. For general polynomial latency functions, we prove lower and upper bounds on the competitive ratio of ASEQNASH that grow both exponentially in the degree of the considered polynomials. Finally, we prove better bounds for the parallel arc case for ASEQNASH, relating the cost of ASEQNASH to the cost of a nonatomic game, generalizing a result of Hayrapetyan, Tardos, and, Wexler [22].

To prove our main results (Theorem 1 and 2) we generalize the variational inequality approach previously used by Correa, Schulz, and Stier-Moses [11], Roughgarden [27], Cominetti, Correa, and Stier-Moses [10], and Harks [20]. The techniques used to prove upper bounds for ASEQNASH in the parallel arc case (Theorem 4) are based on ideas of Hayrapetyan, Tardos, and, Wexler [22]. Our extended approach incorporates the known price of anarchy results as a special case with $n = 1$.

Note that the online algorithms NSEQNASH and ASEQNASH are fully distributed, hence, no coordination mechanism is needed to implement these algorithms. Furthermore, all results for the parallel arc case directly carry over to the online load balancing problem with parallel (splittable) jobs, where the objective is to minimize the L_2 norm of the server loads.

2 Online Network Games

An instance of the *Online Network Game* (ONLINENG) consists of a directed network $D = (V, A)$ together with nondecreasing continuous and convex latency functions $\ell_a : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ for each arc $a \in A$. Furthermore, a sequence $\sigma = 1, \dots, n$ of network games are given. A network game i is characterized by a set of commodities $[K_i] := \{i_1, \dots, i_{n_i}\}$. For each commodity $ij \in [K_i]$, a flow of rate $d_{ij} > 0$ must be routed from the origin s_{ij} to the destination t_{ij} . The routing decision for game i is *online*, that is, it only depends on the routings of previous games $1, \dots, i-1$. Once the commodities of a game have been routed, they remain unchanged. Let $[K] = \bigcup_{i=1}^n [K_i]$ denote the union of the sets $[K_1], \dots, [K_n]$. The total number of commodities is given by $\mathcal{K} = \sum_{i=1}^n n_i$.

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

$$\sum_{a \in \delta^+(v)} f_a^{ij} - \sum_{a \in \delta^-(v)} f_a^{ij} = \gamma_{ij}(v),$$

where $\delta^+(v)$ and $\delta^-(v)$ are the arcs leaving and entering v , respectively; furthermore, $\gamma_{ij}(v) = d_{ij}$, if $v = s_{ij}$, $-d_{ij}$ if $v = t_{ij}$, and 0, otherwise. Alternatively, one can consider a *path flow* for a commodity $ij \in [K_i]$. Let \mathcal{P}_{ij} be the set of all paths

from s_{ij} to t_{ij} in D . A path flow is a nonnegative vector $(f_P^{ij})_{P \in \mathcal{P}_{ij}}$. The corresponding flow on link $a \in A$ for commodity $ij \in [K_i]$ is then $f_a^{ij} := \sum_{P \ni a} f_P^{ij}$. We denote by $f_a^i = \sum_{ij \in [K_i]} f_a^{ij}$ the aggregated flow of game i on link a . The total aggregate flow on link a is given by $f_a = \sum_{i=1}^n f_a^i$. We define \mathcal{F}_i with $i \in [n]$ to be the set of vectors $(\mathbf{f}^1, \dots, \mathbf{f}^i)$ such that \mathbf{f}^j is a feasible flow for games $j = 1, \dots, i$. If $(\mathbf{f}^1, \dots, \mathbf{f}^i) \in \mathcal{F}_i$, we say that it is *feasible* for the sequence of network games $1, \dots, i$. The entire flow for the sequence $1, \dots, n$ of games is denoted by $\mathbf{f} = (\mathbf{f}^1, \dots, \mathbf{f}^n)$.

The *current cost* of a feasible flow for game i on link $a \in A$ is given by $\ell_a(\sum_{j=1}^i f_a^j) f_a^i$. This expression can be obtained as the routing cost on arc a for a feasible flow for game i , *given* the flows $(\mathbf{f}^1, \dots, \mathbf{f}^{i-1})$ of previous games $1, \dots, i-1$ and *without* knowing about future games $j = i+1, \dots, n$. The individual current cost for commodity $ij \in [K_i]$ on arc a is given by $\ell_a(\sum_{j=1}^i f_a^j) f_a^{ij}$. Note that this individual current cost on arc a may increase if later commodities are routed on a . The *total cost* of all sequentially played games is given by:

$$C(\mathbf{f}) = \sum_{a \in A} \ell_a(f_a) f_a = \sum_{a \in A} \ell_a \left(\sum_{i=1}^n f_a^i \right) \left(\sum_{i=1}^n f_a^i \right). \quad (1)$$

This cost function reflects the routing cost provided all commodities of the entire sequence of games have been routed. Thus, the cost of routing commodities of a sequence of games is not separable with respect to the games. That is, if an online algorithm routes flow for the games $i+1, \dots, n$ along arcs that are used by commodities of games $1, \dots, i$, the latter commodities may face higher individual cost on these arcs compared to their initial routing costs.

2.1 Player Types

Motivated by the source routing model in telecommunication networks, we focus on selfish behavior of players routing the demands d_{ij} , $ij \in [K]$. In the following, we use the word commodity ij interchangeably with player ij to indicate that this player decides on the routing assignment \mathbf{f}^{ij} for the demand d_{ij} .

In the nonatomic routing variant, we assume infinitely many agents carrying the flow of a player, where each agent controls only an infinitesimal fraction of the flow. This is in contrast to the atomic routing variant, where it is assumed that each player ij controls and coordinates the entire flow for his demand d_{ij} . For a sequence of games, we investigate the online algorithms NSEQNASH and ASEQNASH (a formal definition follows) that produce a feasible flow $\mathbf{f}^1, \dots, \mathbf{f}^n \in \mathcal{F}_n$, where each \mathbf{f}^i is at Nash equilibrium for the corresponding network game i .

2.2 Nash Equilibria for Nonatomic and Atomic Players

A flow for game i is at Nash equilibrium, if no player has an incentive to unilaterally change her strategy. We assume that players of game i decide on their strategies without taking future games $j = i+1, \dots, n$ into account. It is straightforward to check that a Nash flow \mathbf{f}^i for nonatomic players minimizes the

potential function $\Phi^i(\mathbf{f}) = \sum_{a \in A} \int_0^{f_a^i} \ell_a(\sum_{k=1}^{i-1} f_a^k + z) dz$, see for example Roughgarden and Tardos [30]. Furthermore, using convexity of the potential function two different Nash equilibria incur the same cost. The following conditions are necessary and sufficient to characterize a Nash equilibrium for game i .

Lemma 1. *A feasible flow \mathbf{f}^i for the nonatomic game i is at Nash equilibrium if and only if it satisfies:*

$$\sum_{a \in A} \ell_a \left(\sum_{k=1}^i f_a^k \right) (f_a^i - x_a^i) \leq 0 \text{ for all feasible flows } \mathbf{x}^i \text{ for game } i. \quad (2)$$

The proof is based on the first order optimality conditions and the convexity of the potential function $\Phi^i(\mathbf{f})$, see Dafermos and Sparrow [14].

Definition 1 (NSEQNASH for the OnlineNG). *Consider an instance of the ONLINENG with a given sequence σ of n network games. The deterministic online algorithm NSEQNASH produces a feasible flow $\mathbf{f} = (\mathbf{f}^1, \dots, \mathbf{f}^n) \in \mathcal{F}_n$, such that each flow \mathbf{f}^k minimizes $\Phi^k(\mathbf{f})$, that is, each \mathbf{f}^k is at Nash equilibrium for the corresponding games $k \in [n]$.*

Note that the problem of minimizing $\Phi^k(\mathbf{f})$ is well defined and admits an optimal solution with a unique objective value. Hence, NSEQNASH is also well defined by this property. Since this convex program may have several different solutions (with the same objective value), the flow that NSEQNASH produces is not necessarily unique. As this might contradict the notion of a *deterministic* online algorithm, we can advise a selection rule to make the flow unique. We omit this issue in the following, since our results hold for *every* sequence of Nash flows for the games $1, \dots, n$.

In network games with atomic players, some players may control a significant part of the entire demand. In the following, we characterize the strategy of an atomic player. It is straightforward to see that a best reply strategy for player ij of game i is to minimize its individual current cost $C^{ij}(\mathbf{f}) := \sum_{a \in A} \ell_a(\sum_{k=1}^i f_a^k) f_a^{ij}$.

The following conditions are necessary and sufficient to characterize a Nash equilibrium for game i .

Lemma 2. *A feasible flow \mathbf{f}^i for the atomic game i is at Nash equilibrium if and only if for every player $ij \in [K_i]$ the following inequality is satisfied:*

$$\sum_{a \in A} \left(\ell_a \left(\sum_{k=1}^i f_a^k \right) + \ell'_a \left(\sum_{k=1}^i f_a^k \right) f_a^{ij} \right) (f_a^{ij} - x_a^{ij}) \leq 0, \quad (3)$$

for all feasible flows \mathbf{x}^{ij} for game i .

The proof relies on the convexity of $\ell_a(z)z$.

Definition 2 (ASeqnash for the OnlineNG). *Consider an instance of the ONLINENG with a given sequence σ of n network games. The deterministic*

online algorithm `ASEQNASH` produces a feasible flow $\mathbf{f} = (\mathbf{f}^1, \dots, \mathbf{f}^n) \in \mathcal{F}_n$, such that each flow \mathbf{f}^{ij} , $ij \in [K_i]$, $i \in [n]$ minimizes $C^{ij}(\mathbf{f})$, that is, each \mathbf{f}^i is at Nash equilibrium for the corresponding games $i \in [n]$.

Since we assume convex latency functions, the minimization problem is well defined and admits an optimal solution with a unique objective value. Then, the existence of a flow at Nash equilibrium is guaranteed by the result of Rosen [26]. Hence, `ASEQNASH` is also well defined by this property.

Finally, the *total offline optimum* minimizes the total cost $C(\mathbf{f})$ among all feasible flows. For a given sequence σ , we denote by $\text{OPT}(\sigma)$ the optimal value of this convex problem.

3 Competitive Analysis

For a solution \mathbf{f} produced by an online algorithm `ALG` for a given sequence of games σ , we denote by $\text{ALG}(\sigma) = C(\mathbf{f})$ its cost. An online algorithm `ALG` is called (strictly) *c-competitive*, if the cost of `ALG` is never larger than c times the cost of an optimal offline solution. The *competitive ratio* of `ALG` is the infimum over all $c \geq 1$ such that `ALG` is c -competitive, see for instance Borodin and El-Yaniv [5] and Fiat and Woeginger [17].

3.1 Competitive Analysis for `NSEQnash`

In order to derive competitive results for `NSEQNASH` for a sequence of games, we make use of the variational inequality (2). Using the notation $\vartheta_a^n(\ell_a, \mathbf{f}_a) := \ell_a(\mathbf{f}_a)f_a - \sum_{i=1}^n \ell_a(\sum_{k=1}^i \mathbf{f}_a^k) f_a^i$, we define for every $a \in A$, nonnegative vectors $\mathbf{f}_a, \mathbf{x}_a \in \mathbb{R}_+^K$, and a nonnegative real number $\lambda \geq 0$, the following value (we assume by convention $0/0 = 0$):

$$\omega(\ell_a; n, \lambda) := \sup_{\mathbf{x}_a, \mathbf{f}_a \geq 0} \frac{(\ell_a(\mathbf{f}_a) - \lambda \ell_a(\mathbf{x}_a))x_a + \vartheta_a^n(\ell_a, \mathbf{f}_a)}{\ell_a(\mathbf{f}_a)f_a}. \quad (4)$$

Figure 1 illustrates the value $\omega(\ell_a; n, \lambda)$ for $n = 3$. For a given class \mathcal{L} of nondecreasing latency functions we further define $\omega(\mathcal{L}; n, \lambda) := \sup_{\ell_a \in \mathcal{L}} \omega(\ell_a; n, \lambda)$. Furthermore, we define the following feasible set for the parameter λ .

Definition 3. *The feasible scaling set for λ is defined as*

$$\Lambda(\mathcal{L}, n) := \{\lambda \in \mathbb{R}^+ \mid (1 - \omega(\mathcal{L}; n, \lambda)) > 0\}.$$

Theorem 1. *Consider an instance of the `ONLINE` involving a sequence of n games and latency functions in \mathcal{L} . Then, the competitive ratio of `NSEQNASH` is at most*

$$\inf_{\lambda \in \Lambda(\mathcal{L}, n)} \left[\frac{\lambda}{1 - \omega(\mathcal{L}; n, \lambda)} \right].$$




Fig. 1. Illustration of the value $\omega(\ell_a; \lambda, n)$ for $n = 3$. The entire shaded area corresponds to the value $\vartheta_a^n(\ell_a, \mathbf{f})$. For some $\lambda > 1$, the dark-gray shaded rectangle corresponds to the first term $(\ell_a(f_a) - \lambda \ell_a(x_a))x_a$.

Proof. Let \mathbf{f} be the flow generated by NSEQNASH and let \mathbf{x} be any feasible flow for a given sequence of games $\sigma = 1, \dots, n$. Then, we obtain:

$$C(\mathbf{f}) \leq \sum_{a \in A} (\ell_a(f_a) f_a + \sum_{i=1}^n \ell_a(\sum_{j=1}^i f_a^j) (x_a^i - f_a^i)) \quad (5)$$

$$\begin{aligned} &\leq \sum_{a \in A} (\vartheta_a^n(\ell_a, \mathbf{f}_a) + \ell_a(f_a) x_a) \\ &= \lambda C(\mathbf{x}) + \sum_{a \in A} (\vartheta_a^n(\ell_a, \mathbf{f}_a) + (\ell_a(f_a) - \lambda \ell_a(x_a)) x_a) \\ &\leq \lambda C(\mathbf{x}) + \omega(\mathcal{L}; n, \lambda) C(\mathbf{f}). \end{aligned} \quad (6)$$

Here, (5) follows by applying the variational inequality in Lemma 1. The last inequality (6) follows from the definition of $\omega(\mathcal{L}; n, \lambda)$ and since $\lambda \in \Lambda(\mathcal{L}, n)$. Taking \mathbf{x} as the optimal offline solution yields the claim.

In the following we relate the value $\omega(\mathcal{L}; n, \lambda)$ to the anarchy value $\alpha(\mathcal{L})$ introduced by Roughgarden in [27], the parameter $\beta(\mathcal{L})$ introduced by Correa, Schulz, and Stier-Moses in [11], and the value $\omega(\mathcal{L}; \lambda)$ introduced in Harks [20]. Our definition of $\omega(\mathcal{L}, n, \lambda)$ is equal to $\omega(\mathcal{L}; 1) = \beta(\mathcal{L}) = 1 - \frac{1}{\alpha(\mathcal{L})}$ if we have $\lambda = 1$ and $n = 1$. For arbitrary $\lambda \geq 0$ and $n = 1$, we have $\omega(\mathcal{L}, n, \lambda) = \omega(\mathcal{L}, \lambda)$ as defined in Harks [20]. The difference between these two values is the nonnegative value $\vartheta_a^n(\ell_a, \mathbf{f}_a)$, which accounts for the online setting. It increases for $n \geq 1$ making the value $\omega(\mathcal{L}, n, \lambda)$ larger and, hence, increases the competitive ratio.

Upper Bounds for Linear Latency Functions. In the following, we bound the value $\omega(\mathcal{L}; n, \lambda)$ for affine linear latency functions. We start with some useful prerequisites.

Lemma 3. *For affine functions $\ell(z) = c_1 z + c_0$, $c_1 \geq 0, c_0 \geq 0$, the value $\omega(\mathcal{L}; n, 1)$ is at most $\frac{3n-2}{4n}$.*

The proof of the lemma follows from the Cauchy-Schwarz inequality and the inequality $(f - x)x \leq \frac{1}{4}f^2$.

Equipped with the above lemma, we can apply Theorem 1 to derive an upper bound on the competitive ratio of NSEQNASH for affine latency functions.

Corollary 1. *If the latency functions of the ONLINENG are affine, the online algorithm NSEQNASH is $\frac{4n}{n+2}$ -competitive, where n is the number of games.*

For $n = 1$, we obtain the bound of $\frac{4}{3}$ for nonatomic network games involving affine latency functions first proved in Roughgarden and Tardos [30].

For purely linear latency functions we can improve the upper bound by defining $\lambda := \frac{n}{n+1}$ below 1.

Corollary 2. *If the latency functions of the ONLINENG are linear, the online algorithm NSEQNASH is $\frac{4n^2}{(n+1)^2}$ -competitive, where n is the number of games.*

Using a geometric proof as illustrated in Figure 1 the upper bound of 4 also holds for general continuous, nondecreasing, and concave latency functions.

Corollary 3. *If the latency functions of the ONLINENG are concave, the online algorithm NSEQNASH is 4-competitive.*

Upper Bounds for Polynomial Latency Functions. Now we consider the class \mathcal{L}_d of polynomials with nonnegative coefficients and degree at most $d \in \mathbb{N}$:

$$\mathcal{L}_d := \{c_d x^d + \dots + c_1 x + c_0 : c_s \geq 0, s = 0, \dots, d\}.$$

Note that polynomials in \mathcal{L}_d are nonnegative for nonnegative arguments, nondecreasing, and convex. We can easily see that $\sup_{f_a \geq 0} \vartheta_a^n(\ell_a, f_a) \leq \frac{d}{d+1} \ell_a(f_a) f_a$ for $\ell_a \in \mathcal{L}_d$. Observe that the cost function $C(\mathbf{f})$ is linear in each of the latency functions $\ell_a(\cdot)$. Therefore, we can reduce the analysis to monomial price functions by subdividing each arc a into d arcs a_1, \dots, a_d with monomial latency functions $\ell_{a_s}(x) = c_s x^s$ for every $s = 1, \dots, d$.

Lemma 4. *For the class \mathcal{M}_d of monomials $c_s x^s$ of degree $1 \leq s \leq d$ and $\lambda \geq 1$, we have*

$$\omega(\mathcal{M}_d; n, \lambda) \leq \max_{0 \leq \mu} \mu - \lambda \mu^{d+1} + \frac{d}{d+1}.$$

Proof. For $\ell_a(\cdot) \in \mathcal{M}_d$, we can assume that $\ell_a(f_a) f_a > 0$, since otherwise $\omega(\mathcal{M}_d; n, \lambda) = 0$ and the claim is trivially true. By definition, we have

$$\omega(\ell_a; n, \lambda) = \sup_{x_a, f_a \geq 0} \frac{(\ell_a(f_a) - \lambda \ell_a(x_a)) x_a + \vartheta_a^n(\ell_a, f_a)}{\ell_a(f_a) f_a}.$$

Defining $\mu := \frac{x_a}{f_a}$ (recall that $f_a > 0$), we obtain

$$\omega(\ell_a; n, \lambda) \leq \sup_{0 \leq \mu} \frac{(\ell_a(f_a) - \lambda \ell_a(\mu f_a)) \mu f_a}{\ell_a(f_a) f_a} + \frac{d}{d+1}.$$

Consider now the monomial price function $\ell_a(x_a) = c_s x_a^s$ of degree $s \in [d]$. To bound the value $\omega(\ell_a; n, \lambda)$ from above, we have to consider:

$$\sup_{0 \leq \mu} \frac{(c_s f_a^s - \lambda c_s \mu^s f_a^s) \mu f_a}{c_s f_a^{s+1}} = \max_{0 \leq \mu} \mu - \lambda \mu^{s+1}. \quad (7)$$

Because of the assumption $\lambda \geq 1$ the maximum is attained at a point with $\mu \leq 1$. Thus, it follows that $\max_{0 \leq \mu} \mu - \lambda \mu^{s+1} \leq \max_{0 \leq \mu} \mu - \lambda \mu^{d+1}$. This shows the claim.

Proposition 1. *For polynomial latency functions $\ell \in \mathcal{L}_d$ and $\lambda := (d+1)^{(d-1)}$, the value $\omega(\mathcal{L}; n, \lambda)$ is at most $\frac{d^2+2d}{(d+1)^2}$.*

Proof. The unique solution of the maximization problem in Lemma 4 is given by $\mu^* = \frac{1}{d+1}$. Evaluating the objective with $\lambda := (d+1)^{(d-1)}$ proves the claim:

$$\omega(\ell_a, n; \lambda) \leq \frac{1}{d+1} - (d+1)^{(d-1)} \left(\frac{1}{d+1}\right)^{d+1} + \frac{d}{d+1} = \frac{d^2+2d}{(d+1)^2}.$$

Corollary 4. *Consider the ONLINEENG with latency functions in \mathcal{L}_d . Then, the competitive ratio of the online algorithm NSEQNASH is at most $(d+1)^{d+1}$.*

Proof. Let the flow \mathbf{f} be produced by the online algorithm NSEQNASH and let \mathbf{x} be an arbitrary feasible flow for the ONLINEENG. We define $\lambda := (d+1)^{(d-1)}$ and apply Proposition 1, which yields $\omega(\mathcal{L}; n, \lambda) \leq \frac{d^2+2d}{(d+1)^2}$. In order to apply Theorem 1, we have to verify that $\lambda \in \Lambda(\mathcal{L}, n)$. What remains to be shown is that $1 - \frac{d^2+2d}{(d+1)^2} > 0$ holds. This inequality is equivalent to $\frac{1}{d+1} > 0$. Then, applying Theorem 1 yields

$$C(\mathbf{f}) \leq \frac{(d+1)^{d-1}}{\left(1 - \frac{d^2+2d}{(d+1)^2}\right)} C(\mathbf{x}) = (d+1)^{d+1} C(\mathbf{x}).$$

Taking \mathbf{x} as the optimal offline solution proves the claim.

3.2 Competitive Analysis for ASEQNASH

In this section, we analyze the efficiency of the online algorithm ASEQNASH, which produces a flow \mathbf{f}^i that is at Nash equilibrium for every game i provided that we also allow for atomic players. Recently, Cominetti, Correa, and Stier-Moses [10] discovered that the price of anarchy may be quite large in network games with atomic players. Based on the work of Catoni and Pallotino [8], they presented an example, where the price of anarchy in a network game with atomic players is larger than that of the corresponding nonatomic game. As we show in this section, our upper bounds on the competitive ratio of the online algorithm ASEQNASH also exceed that of NSEQNASH. In the following we only present the main ideas. Complete proofs are left for the full version of this paper.

We define for every $a \in A$, for any nonnegative vectors $\mathbf{f}_a, \mathbf{x}_a \in \mathbb{R}_+^{\mathcal{K}}$ the value

$$\theta_a(\ell_a; \mathbf{f}_a, \mathbf{x}_a) := \sum_{i=1}^n \left(\ell'_a \left(\sum_{k=1}^i f_a^k \right) \sum_{ij \in [K_i]} (f_a^{ij} x_a^{ij} - f_a^{ij} f_a^{ij}) \right).$$

By assuming $0/0 = 0$, we further define

$$\omega(\ell_a; n, \mathcal{K}, \lambda) := \sup_{\mathbf{f}_a, \mathbf{x}_a \geq 0} \frac{(\ell_a(f_a) - \lambda \ell_a(x_a))x_a + \vartheta_a^n(\ell_a, \mathbf{f}_a) + \theta_a(\ell_a; \mathbf{f}_a, \mathbf{x}_a)}{\ell_a(f_a)f_a}. \quad (8)$$

For a given class \mathcal{L} of nondecreasing latency functions and a nonnegative real number $\lambda \geq 0$, we further define $\omega(\mathcal{L}; n, \mathcal{K}, \lambda) := \sup_{\ell_a \in \mathcal{L}} \omega(\ell_a; n, \mathcal{K}, \lambda)$. We define the following feasible set for the parameter λ .

Definition 4. *The feasible scaling set for λ is defined as*

$$\Lambda(\mathcal{L}, n, \mathcal{K}) := \{ \lambda \in \mathbb{R}^+ \mid (1 - \omega(\mathcal{L}; n, \mathcal{K}, \lambda)) > 0 \}.$$

Theorem 2. *Consider an instance of the ONLINENG involving a sequence of n games with \mathcal{K} players and latency functions in \mathcal{L} . Then, the competitive ratio of ASEQNASH is at most*

$$\inf_{\lambda \in \Lambda(\mathcal{L}, n, \mathcal{K})} \left[\frac{\lambda}{1 - \omega(\mathcal{L}; n, \mathcal{K}, \lambda)} \right].$$

The proof proceeds along the same lines as the proof of Theorem 1 except that the value $\omega(\ell_a; n, \mathcal{K}, \lambda)$ contains derivatives ℓ' , which account for the ability of atomic players to coordinate their flow.

Linear and Polynomial Latency Functions. To facilitate the result of Theorem 2, we bound $\omega(\mathcal{L}; n, \mathcal{K}, \lambda)$ for linear latency functions.

Lemma 5. *For affine latency functions $\ell(z) = c_1 z + c_0$, $c_1 \geq 0, c_0 \geq 0$, and $\lambda \geq 1$ the value $\omega(\mathcal{L}; n, \mathcal{K}, \lambda)$ is less than or equal to $\frac{4(\mathcal{K}-1)}{5\mathcal{K}+1}$.*

Applying Theorem 2 with $\lambda = 1$ yields the following result.

Corollary 5. *If the latency functions of the ONLINENG are affine, the online algorithm ASEQNASH is $\frac{5\mathcal{K}+1}{\mathcal{K}+5}$ -competitive, where \mathcal{K} is the total number of players.*

Corollary 5 gives a bound that only depends on the total number of players in the sequence σ of games. This bound states that ASEQNASH is asymptotically 5-competitive for online atomic network games. By choosing $\lambda = 1.13$ and applying Theorem 2 it is possible to improve the upper bound to 4.92.

In the following, we derive a bound that depends on the number of games.

Corollary 6. *If the latency functions of the ONLINENG are affine, the online algorithm ASEQNASH is $\frac{2(3\mathcal{K}+1)n}{n\mathcal{K}+3n+3\mathcal{K}+1}$ -competitive, where n is the number of games and \mathcal{K} is the total number of players.*




Fig. 2. Graph construction for the proof of the lower bound in Proposition 2 and Corollary 9

This bound is asymptotically 6-competitive. It provides, however, an explicit dependency on the number of games and players involved. For $n = 1$, we obtain a bound of $\frac{3\mathcal{K}+1}{2\mathcal{K}+2}$ for atomic network games with affine latency functions; this bound has previously been established by Cominetti, Correa and Stier-Moses [10]. For $\mathcal{K} \rightarrow \infty$ we can establish a bound of $\frac{6n}{n+3}$ that only depends on the number of games.

For latency functions in \mathcal{L}_d , we can show the following bounds.

Corollary 7. *If the latency functions of the ONLINEENG are in \mathcal{L}_d , the competitive ratio of the online algorithm ASEQNASH is at most $(1 + \frac{5}{4}d + \frac{1}{4}d^2)^{d+1}$.*

3.3 Lower Bounds

Based on an instance presented in Harks, Heinz, and Pfetsch [21], we can easily show that any deterministic online algorithm for the ONLINEENG has a competitive ratio greater than or equal to $\frac{4}{3}$ even for linear latencies. In the following, we present an increased lower bound for NSEQNASH. Note that all lower bounds for NSEQNASH also provide lower bounds for ASEQNASH since we can simulate a nonatomic player by infinitely many atomic players each controlling a negligible fraction of the demand.

Proposition 2. *In case of affine latency functions, the online algorithm NSEQNASH for the ONLINEENG has a competitive ratio greater than or equal to $\frac{3n-2}{n}$, where n is the number of games.*

Proof. We consider the network presented in Figure 2 with the latency functions: $\ell_{(s_i,s)}(z) = 0$, $\ell_{(t,t_i)}(z) = 0$, $\ell_{(s_i,t_i)}(z) = i$, $i = 1, \dots, k$, and $\ell_{(s,t)}(z) = z$. We consecutively release a sequence of games $(1, \dots, k)$, where in each game j , there is a single player type $j1$. The demand of player type $j1$ is 1 that has to be routed from s_j to t_j , for $i = 1, \dots, k$. Due to the choice of the affine terms i , NSEQNASH routes for every game the corresponding demand over the arc from s to t . Then

we release the $(k + 1)$ -th game with demand d from s to t . Thus, the total cost for the sequence $\sigma = (1, \dots, k + 1)$ for NSEQNASH with the new cost function is given by: $\text{NSEQNASH}(\sigma) = (k + d)^2$. The optimal offline algorithm OPT routes the demands of the first k games along the direct arcs from s_i to t_i incurring cost of: $\sum_{i=1}^k i = \frac{k(k+1)}{2}$. The last demand in game $k + 1$ is routed from s to t with cost d^2 . The total cost for OPT is given by: $\text{OPT}(\sigma) = \frac{k(k+1)}{2} + d^2$. Replacing $k = n - 1$ and setting $d = \frac{n}{2}$ yields

$$\frac{\text{NSEQNASH}(\sigma)}{\text{OPT}(\sigma)} = \frac{2(k + d)^2}{k(k + 1) + 2d^2} = \frac{3n - 2}{n}, \quad (9)$$

which proves the theorem.

Remark 1. For $n = 2$, the upper bound given in Corollary 1 is tight.

Based on the same instance, except using linear latency functions, we can prove a lower bound for purely linear latency functions.

Corollary 8. *For linear latency functions, the online algorithm NSEQNASH for ONLINENG has a competitive ratio greater than or equal to $\frac{33+5\sqrt{33}}{33+\sqrt{33}}$.*

Corollary 9. *For latency functions in \mathcal{L}_d , the online algorithm NSEQNASH for ONLINENG has a competitive ratio greater than or equal to $\frac{d+1}{d+2} 2^{d+1}$.*

The proof is again based on the instance in Figure 2 except that we use a monomial z^d for the (s, t) arc and the constant terms i become i^d . The rest of the proof then consists of technical calculations that are omitted.

The construction of the lower bounds show that the first player that routes its demand along the arc (s, t) experiences individual cost of $(k + x)$ after routing all commodities. In the common Nash equilibrium, where all players are adaptive and reroute their demand, this player would route its demand along the direct arc incurring cost of 1. Thus, the ratio of the individual cost of a nonadaptive player and that of an adaptive player is unbounded.

4 Parallel Arcs

For graphs that consist of two nodes and parallel arcs, we can show that NSEQNASH performs not worse than a Nash flow for the entire game sequence that is played in parallel. In other words, for a given sequence of games, we compare the cost of NSEQNASH to the cost of a Nash flow of a parallel game, where all players of the entire game sequence route their demands simultaneously.

For a given instance of the ONLINENG involving a sequence of games σ , we define the *parallel game* $\bar{\sigma}$ as a single game that contains all players of the sequence σ simultaneously.

Recall from the Wardrop condition [32] that a flow \mathbf{f} is at Nash equilibrium if and only if the following condition is satisfied:




Fig. 3. Bad Example 1 based on the Graph of the Braess Paradox

Lemma 6. A feasible flow \mathbf{f} for the game $\bar{\sigma}$ is a Nash equilibrium if and only if:

$$\ell_a(f_a) \leq \ell_{\hat{a}}(f_{\hat{a}}), \text{ for all arcs } a, \hat{a} \in A \text{ such that } f_a > 0. \quad (10)$$

Note that for nonatomic network games, Nash equilibria and Wardrop equilibria are the same. A similar condition holds for the flow that is produced by NSEQNASH.

Lemma 7. A feasible flow \mathbf{f} for the sequence of games σ is produced by NSEQNASH if and only if for all $k \in [n]$:

$$\ell_a\left(\sum_{i=1}^k f_a^i\right) \leq \ell_{\hat{a}}\left(\sum_{i=1}^k f_{\hat{a}}^i\right), \text{ for all edges } a, \hat{a} \in A, \text{ such that } f_a^k > 0. \quad (11)$$

Theorem 3. Let $D = (V, A)$ with $V = \{s, t\}$ and A a set of edges from s to t . We are given a sequence of games $\sigma = 1, \dots, n$. Let \mathbf{f} be a flow produced by NSEQNASH for the nonatomic ONLINEING with a single nonatomic player routing d_i from s to t in every game $i \in [n]$. Let \mathbf{f}^* be a flow at Nash equilibrium for the corresponding game $\bar{\sigma}$ with a single player routing $\sum_{i=1}^n d_i$ from s to t . Then, $C(\mathbf{f}) = C(\mathbf{f}^*)$.

Proof. We prove that the flow \mathbf{f} satisfies all conditions of Lemma 6 for the game $\bar{\sigma}$. By the uniqueness of the cost of a Nash equilibrium the claim is proven. The latency of the flow \mathbf{f} on edge a is equal $\ell_a(f_a)$. By contradiction assume that there exist edges $a, \hat{a} \in A$ with $\ell_a(f_a) > \ell_{\hat{a}}(f_{\hat{a}})$, with $f_a > 0$. Let $k \in [n]$ be the largest index with $f_a^k > 0$. The existence of such an index k is granted since $f_a = \sum_{i=1}^n f_a^i > 0$ is assumed. As in games $k+1, \dots, n$, the edge a is not used any more, we have that $\ell_a(f_a) = \ell_a(\sum_{i=1}^k f_a^i)$. Using the assumption that latency functions are nondecreasing it follows that $\ell_{\hat{a}}(f_{\hat{a}}) \geq \ell_{\hat{a}}(\sum_{i=1}^k f_{\hat{a}}^i)$. By Lemma 7 for game k , we have $\ell_{\hat{a}}(\sum_{i=1}^k f_{\hat{a}}^i) \geq \ell_a(\sum_{i=1}^k f_a^i)$, thus $\ell_{\hat{a}}(f_{\hat{a}}) \geq \ell_{\hat{a}}(\sum_{i=1}^k f_{\hat{a}}^i) \geq \ell_a(\sum_{i=1}^k f_a^i) = \ell_a(f_a)$, a contradiction.

The intuition of the above proof fails, however, for general networks with a single source and a single destination. To see this, we present an instance, where the cost of a flow \mathbf{f} produced by NSEQNASH is larger than that of the corresponding Nash flow \mathbf{f}^* for the game $\bar{\sigma}$.

Example 1. Consider the graph of Braess's paradox in Figure 3 and two games that are released consecutively. Each game has a single nonatomic player routing one unit $d_1 = 1$, $d_2 = 1$ from s to t . The path system \mathcal{P}_1 for the first player contains $P_1 = (s, a, t)$, $P_2 = (s, a, b, t)$, $P_3 = (s, b, t)$. A flow that is at Nash equilibrium for the first game routes 1 unit of flow on P_2 , having path latency $\ell_1(\mathbf{f}^1) = 2$. In the second game, we route $\frac{1}{2}$ unit on P_1 and $\frac{1}{2}$ on P_3 , both having path latency $\ell_2(\mathbf{f}) = 2.5$. Now $\ell_{P_2}^2(\mathbf{f}) = 3$. Thus, the total cost is $C(\mathbf{f}) = 1 \times 2.5 + 1 \times 3 = 5.5$. However, for the game $\bar{\sigma}$ we route 2 units of flow from s to t . Then, a flow \mathbf{f}^* at Nash equilibrium routes one unit along paths P_1 and P_3 . The path latencies are $\ell_{P_1}(\mathbf{f}) = \ell_{P_3}(\mathbf{f}) = 2$, thus the total cost is $C(\mathbf{f}^*) = 2 \times 2 = 4$.

In the following, we investigate the parallel arc setting for ASEQNASH. For the atomic case in a parallel arc network, Hayrapetyan, Tardos and, Wexler [22] have proved a similar result. Assume we have a single game with an arbitrary number of atomic players, and $x \ell_a(x)$ is a convex function for all edges a . Then, the equilibrium cost is at most as much as the cost of the nonatomic Nash equilibrium of the total demand. Now we generalize this result for the ONLINENG.

We will compare a sequence of games with an arbitrary number of atomic players with a single game, where nonatomic players routes the flow \mathbf{g} such that the total demand of the entire sequence is satisfied. Instead of comparing the usual costs, our result involves another cost function $C_2(\mathbf{f})$:

$$C_2(\mathbf{f}) = \sum_{i=1}^n \sum_{a \in A} \ell_a \left(\sum_{j=1}^i f_a^j \right) f_a^i. \quad (12)$$

$C_2(\mathbf{f})$ means that players in game i have to pay only after their current latency in game i , and no cost according to the games $i+1, \dots, n$. Note that the relation $C_2(\mathbf{f}) \leq C(\mathbf{f})$ holds because of the monotonicity of ℓ_a . In the following, we will show that $C_2(\mathbf{f}) \leq C(\mathbf{g})$. Then, by defining $\delta := \sup_{\mathbf{f}} (C(\mathbf{f})/C_2(\mathbf{f}))$, we are still able to bound the usual cost: $C(\mathbf{f}) \leq C_2(\mathbf{g}) \leq \delta C(\mathbf{g})$. For latency functions $\ell \in \mathcal{L}_d$, we have for instance $\delta \leq d + 1$. Now we are ready to state the following theorem:

Theorem 4. *Consider an instance of ONLINENG with a sequence of games $\sigma = 1, \dots, n$ and the underlying graph $D = (V, A)$ with $V = \{s, t\}$ and A consisting of parallel st -edges. Let \mathbf{f} be a flow produced by ASEQNASH for this game with n_i atomic players in game i . Let $d = \sum_{i=1}^n \sum_{j=1}^{n_i} d_{ij}$ denote the total demand over all games. Let \mathbf{g} be at Nash equilibrium for a single nonatomic game with a single player routing d units from s to t . Then $C_2(\mathbf{f}) \leq C(\mathbf{g})$.*

For $n = 1$, this gives Theorem 2.3 in [22]. To prove the theorem, we introduce $d_i = \sum_{j=1}^{n_i} d_{ij}$, the total demand in game i . The essence of the proof is the following lemma:

Lemma 8. *Assume we have a sequence of two games with $D = (V, A)$ satisfying the conditions of Theorem 4. In the first game, n_1 atomic players route each d_{1j}*

units of flow from s to t . In the second game, there is a nonatomic player routing z units. Let $\mathbf{h} = (h^1, h^2)$ be in equilibrium for this sequence, and let \mathbf{m} be in equilibrium for a single game with a single nonatomic player routing $d_1 + z$ from s to t . Then $C_2(\mathbf{h}) \leq C(\mathbf{m})$.

Before proving this lemma, we show how it implies Theorem 4. The proof is by induction. If $n = 1$, then we apply the lemma for $z = 0$. If $n > 1$, suppose we have proved the theorem for $n - 1$. Fix the flows of the first game, and modify the cost function ℓ_a to $q_a(x) = \ell_a(x + f_a^1)$. Consider the games 2, \dots , $n - 1$ with cost function q_a , and let C_2^q denote the modified cost function. By definition, (f^2, \dots, f^n) is in equilibrium for this sequence, and $C_2(\mathbf{f}) = \sum_{a \in A} \ell_a(f_a^1) f_a^1 + C_2^q(f^2, \dots, f^n)$.

Let \mathbf{g}^q be a routing in Nash equilibrium of $z = \sum_{i=2}^n d_i$ units for the cost function q_a ; let C^q denote its cost. By induction, $C_2^q(f^2, \dots, f^n) \leq C^q$. Consider now the game described in the lemma. $\mathbf{h} = (f^1, \mathbf{g}^q)$ is in equilibrium, and the game \mathbf{m} is identical to \mathbf{g} as $d_1 + z = d$. By the lemma, $C_2(\mathbf{h}) \leq C(\mathbf{m}) = C(\mathbf{g})$.

On the other hand,

$$C_2(\mathbf{h}) = \sum_{a \in A} \ell_a(f_a^1) f_a^1 + C^q \geq \sum_{a \in A} \ell_a^1(f_a^1) f_a^1 + C_2^q(f^2, \dots, f^n) = C_2(\mathbf{f}),$$

and this is what we wanted to prove.

Before proving Lemma 8, we prove two other lemmas, which are motivated by [22].

Lemma 9. *For the game as in Lemma 8, \mathbf{h} is in equilibrium if and only if for any $j \in [n_1]$, and any edges $a, \hat{a} \in A$ with $h_a^{1j} > 0$,*

$$\ell_a(h_a^1) + h_a^{1j} \ell'_a(h_a^1) \leq \ell_{\hat{a}}(h_{\hat{a}}^1) + h_{\hat{a}}^{1j} \ell'_{\hat{a}}(h_{\hat{a}}^1). \quad (13)$$

Furthermore, for any edges $a, \hat{a} \in A$, if $h_a^2 > 0$, then $\ell_a(h_a) \leq \ell_{\hat{a}}(h_{\hat{a}})$.

This easily follows as player 1j wants to minimize $\sum_{a \in A} \ell_a(h_a^1) h_a^{1j}$. Let T denote the minimum edge latency of game \mathbf{m} in Lemma 8. We call an edge a overloaded, if $\ell_a(h_a) > T$ and underloaded if $\ell_a(h_a) \leq T$. The idea is, following [22], that moving a small flow from an overloaded edge to an underloaded increases the cost.

Lemma 10. *Assume a is overloaded and \hat{a} is underloaded with $h_a > 0$. Then $h_a^1 > 0$, and modifying h^1 by moving a small amount of flow from a to \hat{a} does not decrease $C_2(\mathbf{h})$.*

Proof. Let T_2 denote the minimum edge latency in the Nash equilibrium of h^2 . We show that $T_2 \leq T$. Assume by contradiction that $T_2 > T$. Then by moving flows from edges with latency higher than T_2 to edges of lower latency, we can finally arrive in a Nash-equilibrium of edge latency at least T_2 , contradicting the fact that the edge latency is the same in any two different Nash-equilibria, see Roughgarden [28].

By $T_2 \leq T$, if a is overloaded, then $h_a^2 = 0$. This implies $h_a^1 > 0$, and $\ell_a(h_a) = \ell_a(h_a^1)$. Our aim is to prove, that decreasing h_a^1 a little bit and increasing $h_{\hat{a}}^1$ with the same amount, $C_2(\mathbf{h})$ does not decrease. The statement follows if we can prove that $\frac{\partial C_2}{\partial h_a^1} \leq \frac{\partial C_2}{\partial h_{\hat{a}}^1}$. This is equivalent with

$$\ell_a(h_a^1) + h_a^1 \ell'_a(h_a^1) + h_a^2 \ell'_a(h_a^1 + h_a^2) \leq \ell_{\hat{a}}(h_{\hat{a}}^1) + h_{\hat{a}}^1 \ell'_{\hat{a}}(h_{\hat{a}}^1) + h_{\hat{a}}^2 \ell'_{\hat{a}}(h_{\hat{a}}^1 + h_{\hat{a}}^2). \quad (14)$$

Let $B = \{j : h_a^{1j} > 0\}$. If we sum (13) for $j \in B$, we get

$$|B| \ell_a(h_a^1) + h_a^1 \ell'_a(h_a^1) \leq |B| \ell_{\hat{a}}(h_{\hat{a}}^1) + \sum_{j \in B} h_{\hat{a}}^{1j} \ell'_{\hat{a}}(h_{\hat{a}}^1)$$

Increasing the right hand side by $\sum_{j \notin B} h_{\hat{a}}^{1j} \ell'_{\hat{a}}(h_{\hat{a}}^1) + h_{\hat{a}}^2 \ell'_{\hat{a}}(h_{\hat{a}}^1 + h_{\hat{a}}^2)$, and adding $h_a^2 \ell'_a(h_a^1 + h_a^2) = 0$ to the left hand side, we get

$$|B| \ell_a(h_a^1) + h_a \ell'_a(h_a) + h_a^2 \ell'_a(h_a^1 + h_a^2) \leq |B| \ell_{\hat{a}}(h_{\hat{a}}^1) + h_{\hat{a}}^1 \ell'_{\hat{a}}(h_{\hat{a}}^1) + h_{\hat{a}}^2 \ell'_{\hat{a}}(h_{\hat{a}}^1 + h_{\hat{a}}^2)$$

As a is overloaded and \hat{a} is underloaded, we have $\ell_a(h_a^1) = \ell_a(h) > \ell_{\hat{a}}(h) \geq \ell_{\hat{a}}(h_{\hat{a}}^1)$. This in turn implies (14).

Now we are ready to prove Lemma 8. We modify \mathbf{h} by moving flow amounts from overloaded to underloaded links. We avoid creating new overloaded links: we increase the flow on the underloaded edge \hat{a} so that $\ell_{\hat{a}}(h_{\hat{a}})$ should not exceed T . Applying such a modification maintains $h_a^2 = 0$ on any overloaded edge a . Observe that (14) holds not only for h , but for any \bar{h} satisfying $\bar{h}_a^1 \leq h_a^1$ and $\bar{h}_{\hat{a}}^1 \geq h_{\hat{a}}^1$ and $\bar{h}^2 = h^2$. This is since the monotonicity and convexity of $\ell_a(x)$ implies that $x \ell_a(x + c)$ is as well convex for $c, x \geq 0$.

This ensures that we can always go on by moving flow from overloaded to underloaded links, as far as some modifications are applicable. Suppose no more modifications can be applied. In this case for each edge a , $\ell_a(h_a) \geq T$. If $\ell_a(h_a) = T$ for all edges, then Lemma 8 follows. Otherwise we have some edge a with $\ell_a(h_a) > T$, and for all edges \hat{a} with $\ell_{\hat{a}}(h_{\hat{a}}) = T$, increasing $h_{\hat{a}}$ by an arbitrary small positive amount would result in $\ell_{\hat{a}}(h) > T$. But this is a contradiction, as now the flows can be rerouted to obtain a Nash-equilibrium with edge latency strictly larger than T . Again we use that the edge latency is the same in any two different Nash-equilibria, and T was the latency of the equilibrium state \mathbf{m} .

References

1. Altman, E., Basar, T., Jimenez, T., Shimkin, N.: Competitive routing in networks with polynomial costs. *IEEE Trans. Automat. Control* 47(1), 92–96 (2002)
2. 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)
3. Awerbuch, B., Azar, Y., Grove, E.F., Kao, M.-Y., Krishnan, P., Vitter, J.S.: Load balancing in the L_p norm. In: *IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 383–391 (1995)

4. Awerbuch, B., Azar, Y., Plotkin, S.: Throughput-competitive on-line routing. In: Proc. 34th Annual IEEE Symposium on Foundations of Computer Science (FOCS), Palo Alto, pp. 32–40 (1993)
5. Borodin, A., El-Yaniv, R.: *Online Computation and Competitive Analysis*. Cambridge University Press, Cambridge (1998)
6. Braess, D.: Über ein Paradoxon der Verkehrsplanung. *Unternehmensforschung* 11, 258–268 (1968)
7. Caragiannis, I., Flammini, M., Kaklamanis, C., Kanellopoulos, P., Moscardelli, L.: Tight bounds for selfish and greedy load balancing. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4051, pp. 311–322. Springer, Heidelberg (2006)
8. Catoni, S., Pallotino, S.: Traffic equilibrium paradoxes. *Transportation Science* 25, 240–244 (1991)
9. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing (STOC), pp. 67–73 (2005)
10. Cominetti, R., Correa, J.R., Stier-Moses, N.: Network games with atomic players. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4051, Springer, Heidelberg (2006)
11. Correa, J.R., Schulz, A.S., Stier-Moses, N.E.: Selfish routing in capacitated networks. *Math. Oper. Res.* 29, 961–976 (2004)
12. Correa, J.R., Schulz, A.S., Stier-Moses, N.E.: On the inefficiency of equilibria in congestion games. In: Jünger, M., Kaibel, V. (eds.) *Integer Programming and Combinatorial Optimization*. LNCS, vol. 3509, pp. 167–181. Springer, Heidelberg (2005)
13. Czumaj, A., Vöcking, B.: Tight bounds for worst-case equilibria. In: Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms (SODA), pp. 413–420 (2002)
14. 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)
15. Dubey, P.: Inefficiency of Nash Equilibria. *Math. Oper. Res.* 11, 1–8 (1986)
16. Even-Dar, E., Mansour, Y.: Fast convergence of selfish rerouting. In: Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 772–781 (2005)
17. Fiat, A. (ed.): *Online Algorithms*. LNCS, vol. 1442. Springer, Heidelberg (1998)
18. Fischer, S., Vöcking, B.: Adaptive routing with stale information. In: Aguilera, M.K., Aspnes, J. (eds.) Proc. 24th Ann. ACM SIGACT-SIGOPS Symp. on Principles of Distributed Computing (PODC), Las Vegas, NV, USA, July 2005, pp. 276–283. ACM Press, New York (2005)
19. Friedman, E.J.: Genericity and congestion control in selfish routing. In: *Decision and Control, CDC. 43rd IEEE Conference on*, pp. 4667–4672 (2004)
20. Harks, T.: On the price of anarchy of network games with nonatomic and atomic players. Technical report, available at Optimization Online (January 2007)
21. Harks, T., Heinz, S., Pfetsch, M.E.: Online multicommodity routing problem. In: Erlebach, T., Kaklamanis, C. (eds.) WAOA 2006. LNCS, vol. 4368, Springer, Heidelberg (2007)
22. Hayrapetyan, A., Tardos, E., Wexler, T.: The effect of collusion in congestion games. In: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing (STOC), pp. 89–98 (2006)
23. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)

24. La, R.J., Walrand, J., Anantharam, V.: Issues in TCP Vegas. Electronics Research Laboratory, University of California, Berkeley, UCB/ERL Memorandum, No. M99/3 (January 1999)
25. Qiu, L., Yang, R.Y., Zhang, Y., Shenker, S.: On selfish routing in internet-like environments. *IEEE/ACM Trans. on Netw.* 14(4), 725–738 (2006)
26. Rosen, J.B.: Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica* 33, 520–534 (1965)
27. Roughgarden, T.: The price of anarchy is independent of the network topology. *Journal of Computer and System Science* 67, 341–364 (2002)
28. Roughgarden, T.: *Selfish Routing and the Price of Anarchy*. The MIT Press, Cambridge (2005)
29. Roughgarden, T.: Selfish routing with atomic players. In: *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 973–974 (2005)
30. Roughgarden, T., Tardos, E.: How bad is selfish routing? *Journal of the ACM* 49(2), 236–259 (2002)
31. Suri, S., Toth, C., Zhou, Y.: Selfish load balancing and atomic congestion games. *Algorithmica* 47(1), 79–96 (2007)
32. Wardrop, J.G.: Some theoretical aspects of road traffic research. In: *Proceedings of the Institute of Civil Engineers*, 1(Part II), pp. 325–378 (1952)