Computing Approximate Nash Equilibria in Network Congestion Games*

Andreas Emil Feldmann[1], Heiko Röglin[2], and Berthold Vöcking[2]

[1] Institute of Theoretical Computer Science
ETH Zürich, Switzerland
andreas.feldmann@inf.ethz.ch
[2] Department of Computer Science
RWTH Aachen, Germany
{roeglin,voecking}@cs.rwth-aachen.de

Abstract. We consider the problem of computing ε-approximate Nash equilibria in network congestion games. The general problem is known to be PLS-complete for every $\varepsilon > 0$, but the reductions are based on artificial and steep delay functions with the property that already two players using the same resource cause a delay that is significantly larger than the delay for a single player.

We consider network congestion games with delay functions such as polynomials, exponential functions, and functions from queuing theory. We analyse which approximation guarantees can be achieved for such congestion games by the method of randomised rounding. Our results show that the success of this method depends on different criteria depending on the class of functions considered. For example, queuing theoretical functions admit good approximations if the equilibrium load of every resource is bounded away appropriately from its capacity.

1 Introduction

In recent years, there has been an increased interest in understanding selfish routing in large networks like the Internet. Since the Internet is operated by different economic entities with varying interests, it is natural to model these entities as selfish agents who are only interested in maximising their own benefit. *Congestion games* are a classical model for resource allocation among selfish agents. We consider the special case of *network congestion games*, in which the resources are the edges of a graph and every player wants to allocate a path between her designated source and target node. The delay of an edge increases with the number of players allocating it, and every player is interested in allocating a routing path with minimum delay.

Rosenthal [10] shows that congestion games are potential games and hence they always possess *Nash equilibria*[3], i.e. allocations of resources from which no

* This work was supported by DFG grant VO 889/2 and by the Ultra High-Speed Mobile Information and Communication Research Cluster (UMIC) established under the excellence initiative of the German government.

[3] In this paper, the term *Nash equilibrium* always refers to a pure equilibrium.

player wants to deviate unilaterally. Fabrikant et al. [5] show that the problem of computing a pure Nash equilibrium can be phrased as a local search problem belonging to the complexity class PLS. They show that it is already PLS-complete for the special case of network congestion games if different players are allowed to have different source and target nodes. Ackermann et al. [1] show that this result can even be extended to network congestion games with linear delay functions. This implies that there is no efficient algorithm for computing pure Nash equilibria, unless $\mathsf{PLS} \subseteq \mathsf{P}$. On the other hand, for *symmetric network congestion games*, in which all players have the same source and the same target node, Nash equilibria can be computed efficiently by solving a min-cost flow problem [5].

In many applications players incur some costs when they change their strategy. Hence, it is reasonable to assume that a player is only interested in changing her strategy if this decreases her delay significantly. This assumption leads to the notion of an *ε-approximate Nash equilibrium*, which is a state in which no player can decrease her delay by more than a factor of $1+\varepsilon$ by unilaterally changing her strategy. For symmetric congestion games, in which the strategy spaces of the players coincide, Chien and Sinclair [3] show that ε-approximate equilibria can be computed by simulating the best response dynamics for a polynomial (in the number of players and ε^{-1}) number of steps. Unfortunately, the problem of computing an ε-approximate Nash equilibrium is still PLS-complete for every constant $\varepsilon > 0$ (and even every polynomial-time computable function ε) for general congestion games [12] and even for network congestion games[4]. The delay functions used in these reductions are quite artificial and steep with the property that already two players using the same resource cause a delay that is significantly larger than the delay for a single player. In this paper, we study natural classes of delay functions such as polynomials and functions from queuing theory, and we analyse which approximation guarantees can be achieved for these functions by the method of randomised rounding [9].

1.1 Models and Method

A network congestion game is described by a directed graph $G = (V, E)$ with m edges, a set $\mathcal{N}$ of n players, a pair $(s_i, t_i) \in V \times V$ of source and target node for each player $i \in \mathcal{N}$, and a non-decreasing delay function $d_e \colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ for each edge $e \in E$. For $i \in \mathcal{N}$ we denote by $\mathcal{P}_i$ the set of all paths from node s_i to node t_i. Every player i has to choose one path P_i from the set $\mathcal{P}_i$ and to allocate all edges on this path. For a *state* $S = (P_1, \ldots, P_n) \in \mathcal{P}_1 \times \ldots \times \mathcal{P}_n$ and an edge $e \in E$, we denote by $n_e(S)$ the number of players allocating edge e in state S, i.e. $n_e(S) = |\{i \in \mathcal{N} \mid e \in P_i\}|$. The *delay* $\delta_i(S)$ to a player $i \in \mathcal{N}$ in state S is defined as equal to the delay $d_{P_i}(S) := \sum_{e \in P_i} d_e(n_e(S))$ of the chosen path P_i in S and every player wants to allocate a path with minimum delay. We say that a state S is a *Nash equilibrium* if no player can decrease her delay by changing her strategy, i.e. if state S' is obtained from S by letting one player $i \in \mathcal{N}$ change her strategy, then the delay $\delta_i(S')$ is at least as large as the delay $\delta_i(S)$. A state

[4] Alexander Skopalik, personal communication.

S is said to be an *ε-approximate Nash equilibrium* if $\delta_i(S) \leq (1+\varepsilon) \cdot \delta_i(S')$ for every state S' that is obtained from S by letting one player $i \in \mathcal{N}$ change her strategy.

In order to compute approximate Nash equilibria, we use the method of randomised rounding. Therefore, we first relax the network congestion game by replacing each player by an infinite set of agents, each of which controlling an infinitesimal amount of flow. To be more precise, we transform the network congestion game into a multi-commodity flow problem and we introduce a flow demand of 1 that is to be routed from node s_i to node t_i for every player $i \in \mathcal{N}$. A flow vector $f \in \mathbb{R}_{\geq 0}^{|E|}$ induces delays on the edges. The delay of edge $e \in E$ is $d_e(f) = d_e(f_e)$, and the delay on a path P is the sum of the delays of its edges, i.e. $d_P(f) = \sum_{e \in P} d_e(f)$. A flow vector f is called a *Wardrop equilibrium* [13] if for all commodities $i \in \mathcal{N}$ and all paths $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$ it holds that $d_{P_1}(f) \leq d_{P_2}(f)$. It is well-known that Wardrop equilibria can be computed in polynomial time using convex programming [2].

After relaxing the network congestion game and computing a Wardrop equilibrium f, we compute a decomposition of the flow f into polynomially many paths. For a commodity $i \in \mathcal{N}$ let $\mathcal{D}_i \subseteq \mathcal{P}_i$ denote the set of paths used in this decomposition, and for $P \in \mathcal{D}_i$, let f_P^i denote the flow of commodity i that is sent along path P. For fixed i the flows f_P^i can be interpreted as a probability distribution on the set $\mathcal{D}_i$. Following the method of randomised rounding, we choose, according to these probability distributions, independently for each player i a routing path from $\mathcal{D}_i \subseteq \mathcal{P}_i$. In the following, we analyse for several classes of delay functions the approximation guarantee of this approach.

1.2 Our Results

After the randomised rounding the congestion on an edge is a sum of independent Bernoulli random variables whose expectation equals the flow on that edge in the Wardrop equilibrium. By applying Chernoff bounds, we can find for each edge a small interval such that it is unlikely that any congestion takes a value outside the corresponding interval. If the delay functions are not too steep in these intervals, then the delays on the edges after the rounding are neither much smaller nor much larger than the delays in the Wardrop equilibrium, implying that the resulting state is an ε-approximate Nash equilibrium for some ε depending on the steepness of the delay functions. Due to the multiplicative definition of approximate Nash equilibria, delay functions can be steeper in intervals in which they take larger values in order to achieve the same ε.

In the literature on selfish routing, it is a common assumption that the delay functions are polynomials with nonnegative coefficients [4, 11]. Hence, we start our investigations with the question which properties polynomial delay functions need to satisfy in order to guarantee that randomised rounding yields an ε-approximate Nash equilibrium with high probability. We have argued that the delay functions must not grow too fast relative to their values. For polynomials $d_e(x) = \sum_{j=0}^{g} a_j^e x^j$ with $a_j^e \geq 0$, this implies that the offset a_0^e must not be too small. If all delay functions are polynomials of some constant degree g and if, for each edge $e \in E$, the offset a_0^e satisfies

$$a_0^e \geq \frac{((1+\varepsilon)g^2 \cdot 6\ln(4m))^g}{((\sqrt{1+\varepsilon})-1)^{2g+1}} \sum_{j=1}^{g} a_j^e = \Theta\left(\frac{\ln^g m}{\varepsilon^{2g+1}}\right) \sum_{j=1}^{g} a_j^e \ ,$$

then an ε-approximate Nash equilibrium can be computed by randomised rounding with high probability in polynomial time. In the above asymptotic estimate ε tends to zero while m approaches infinity. If, for example, all delay functions are linear and one wants to obtain an ε-approximate equilibrium for some constant $\varepsilon > 0$, then all delay functions must have the form $d_e(x) = a_1^e x + a_0^e$ where a_0^e is sufficiently large in $\Omega(a_1^e \cdot \ln(m))$. A lower bound on a_0^e is not unrealistic since most network links have a non-negligible delay even if they are relatively uncrowded. For example, in communication networks the offset corresponds to the sum of *packet-propagation delay* and *packet-processing delay*, which should dominate the *packet-queuing delay* if an edge is not dramatically overloaded.

The second class that we study are exponential delay functions of the form $d_e(x) = \alpha_e \cdot \exp(x/\beta_e) + \gamma_e$. We show that for these functions an ε-approximation can be achieved if β_e is lower bounded by some function $f(\chi_e, m, \varepsilon)$ growing in the order of

$$f(\chi_e, m, \varepsilon) \in O\left(\frac{\ln(m) \cdot \sqrt{\chi_e}}{\varepsilon}\right) \ ,$$

where χ_e denotes the load of edge e in the Wardrop equilibrium and hence corresponds to the expected congestion on e after the randomised rounding. Such exponential functions grow very slowly as long as less than β_e players allocate the edge, but they start growing rapidly beyond this point. This reflects typical behaviour in practice, because one often observes that the delay on a network link grows rather slowly with its congestion until some overload point is reached after which the quality of the link deteriorates quickly. We show that it is even possible to replace the exponential function up until β_e by a polynomial. To be precise, we show that if an ε-approximate equilibrium can be computed for two functions by randomised rounding, then this is also the case for the function that takes for every input the maximum of these functions.

Finally, we consider functions that arise when using queuing theory for modelling the behaviour of network links. We consider the $M/M/1$ queuing model in which there is one queue and the inter-arrival and service times of the packets are exponentially distributed. A network of such queues constitutes a so-called Jackson network [6]. We take as delay function the limiting behaviour of the sojourn time of a packet on the network link. We interpret the congestion on an edge as the rate of the packet arrival process and we show that, in order to obtain an ε-approximate equilibrium, it suffices that the rate μ_e of the service process (corresponding to the capacity of the edge) is lower bounded by some function $g(\chi_e, m, \varepsilon)$ growing in the order of

$$g(\chi_e, m, \varepsilon) \in O\left(\chi_e + \frac{\ln(m) \cdot \sqrt{\chi_e}}{\varepsilon}\right) \ .$$

Our result implies that it is sufficient if the equilibrium load χ_e is bounded away from the capacity μ_e by some additive term of order only $O(\ln(m) \cdot \sqrt{\mu_e}/\varepsilon)$.

Outline. In the remainder of this paper we first state some preliminaries and illustrate our approach. After that we give a sufficient condition on the delay functions that guarantees that randomised rounding computes an ε-approximate equilibrium in polynomial time with high probability. Then we analyse which restrictions this condition imposes when applied to polynomial delay functions, exponential delay functions, and delay functions from queuing theory. Finally, we prove a theorem for combined delay functions.

2 Preliminaries

The number N_e of players that use an edge $e \in E$ after the rounding is a sum of independent Bernoulli random variables whose expectation χ_e equals the flow on e in the Wardrop equilibrium. We can use Chernoff bounds to identify, for each edge e, an interval $[l_e, u_e]$ such that it is unlikely that N_e takes a value outside this interval. We choose these intervals such that for $x \in [l_e, u_e]$ the delay $d_e(x)$ of edge e lies between $d_e(\chi_e)/\sqrt{1+\varepsilon}$ and $\sqrt{1+\varepsilon} \cdot d_e(\chi_e)$.

Lemma 1. *If for all $e \in E$ it holds that $N_e \in [l_e, u_e]$ and $d_e(\chi_e)/\sqrt{1+\varepsilon} \leq d_e(x) \leq \sqrt{1+\varepsilon} \cdot d_e(\chi_e)$ for any $x \in [l_e, u_e]$, then the resulting state is an ε-approximate Nash equilibrium.*

Proof. Let S denote the state computed by the randomised rounding. Assume that a path P_i is chosen for player i by the randomised rounding and that P_i' is a path with minimum delay after the randomised rounding. From the definition of a Wardrop equilibrium it follows that in the computed flow the delay L_i on P_i is at most as large as the delay on P_i' because flow is sent along P_i (otherwise the probability that path P_i is chosen would equal 0). Since the delay on P_i increases at most by a factor of $\sqrt{1+\varepsilon}$ and the delay on P_i' decreases at most by a factor of $\sqrt{1+\varepsilon}$ during the randomised rounding, we obtain

$$\frac{d_{P_i}(S)}{d_{P_i'}(S)} \leq \frac{\sqrt{1+\varepsilon} \cdot L_i}{L_i/\sqrt{1+\varepsilon}} = 1+\varepsilon \ ,$$

which proves the lemma. □

3 A Sufficient Condition on the Delay Functions

In this section, we present a sufficient condition on the delay functions that guarantees that an ε-approximate Nash equilibrium can be computed by randomised rounding in polynomial time. We will make use of the following Chernoff bounds.

Lemma 2 ([8]). *Let $X_1, \ldots, X_n$ be independent random variables with $\mathbf{Pr}[X_i = 1] = p_i$ and $\mathbf{Pr}[X_i = 0] = 1 - p_i$ for each $i \in \{1, \ldots, n\}$ and let the random variable X be defined as $\sum_{i=1}^n X_i$.*

– *If* $\mu \geq \mathbf{E}[X]$ *and* $0 \leq \delta \leq 1$*, then* $\mathbf{Pr}[X > (1+\delta)\cdot\mu] \leq \exp\left(-\frac{\delta^2\mu}{3}\right)$.
– *If* $\mu \leq \mathbf{E}[X]$ *and* $0 \leq \delta \leq 1$*, then* $\mathbf{Pr}[X < (1-\delta)\cdot\mu] \leq \exp\left(-\frac{\delta^2\mu}{3}\right)$.

We make two assumptions on the delay functions to avoid case distinctions and to keep the statement of the next theorem simple. We assume that each delay function is defined on $\mathbb{R}$ and w.l.o.g. we set $d_e(x) = d_e(0)$ for $x < 0$. Additionally, we assume that the delay function never equals zero, i.e. $d_e(x) > 0$ for all $x \in \mathbb{R}$. The latter condition is reasonable since in practice the delay of a network link never drops to zero.

For an edge e, let χ_e in the following denote the expected congestion, which equals the flow on edge e in the Wardrop equilibrium.

Theorem 3. *Using the method of randomised rounding, it is possible to compute an* ε*-approximate Nash equilibrium of a network congestion game with high probability in polynomial time if for each edge* $e \in E$ *and for all* $x \in [0, \max\{6\ln(4m), \chi_e + \sqrt{3\ln(4m)\cdot\chi_e}\}]$

$$\frac{d_e(x)}{d_e(x - \sqrt{6\ln(4m)\cdot x})} \leq \sqrt{1+\varepsilon}\ . \tag{1}$$

Proof. Following the arguments in Section 2, we define an interval $\mathfrak{I}_e := [l_e, u_e]$ for each edge e such that, after the randomised rounding, the congestion N_e lies in this interval with probability at least $1 - \frac{1}{2m}$ and such that $d_e(u_e)/d_e(\chi_e) \leq \sqrt{1+\varepsilon}$ and $d_e(\chi_e)/d_e(l_e) \leq \sqrt{1+\varepsilon}$. Given these properties, one can easily see that after the randomised rounding and with probability at least $1/2$, $N_e \in \mathfrak{I}_e$ for all edges e. If this event occurs, then the resulting state is an ε-approximate Nash equilibrium due to Lemma 1. Since the failure probability is at most $1/2$, repeating the randomised rounding, say, n times independently yields an exponentially small failure probability.

Since we assume $d_e(x) = d_e(0)$ for $x < 0$, we can assume that Inequality (1) holds for all $x \in \mathbb{R}_{\leq 0}$. For an edge e, we set

$$\mathfrak{I}_e = [l_e, u_e] = \left[\chi_e - \sqrt{3\ln(4m)\cdot\chi_e}, \max\{6\ln(4m), \chi_e + \sqrt{3\ln(4m)\cdot\chi_e}\}\right]\ .$$

Since $l_e \geq \chi_e - \sqrt{6\ln(4m)\cdot\chi_e}$, Inequality (1) and the monotonicity of d_e imply

$$\frac{d_e(\chi_e)}{d_e(l_e)} \leq \frac{d_e(\chi_e)}{d_e(\chi_e - \sqrt{6\ln(4m)\cdot\chi_e})} \leq \sqrt{1+\varepsilon}\ .$$

If $\chi_e \geq 3\ln(4m)$, then $u_e = \chi_e + \sqrt{3\ln(4m)\cdot\chi_e}$ and

$$\frac{d_e(u_e)}{d_e(\chi_e)} \leq \frac{d_e(\chi_e + \sqrt{3\ln(4m)\cdot\chi_e})}{d_e(\chi_e)} \leq \sqrt{1+\varepsilon}\ ,$$

where the last inequality follows from substituting x by $x + \sqrt{3\ln(4m)\cdot x}$ in Inequality (1) and by using the monotonicity of the delay function d_e. If $\chi_e \leq$

$3\ln(4m)$, then $u_e = 6\ln(4m)$ and

$$\frac{d_e(u_e)}{d_e(\chi_e)} \leq \frac{d_e(6\ln(4m))}{d_e(0)} \leq \sqrt{1+\varepsilon} \ ,$$

where the last inequality follows directly from (1) by setting $x = 6\ln(4m)$. Altogether, this implies that we have achieved the desired properties that $d_e(u_e)/d_e(\chi_e) \leq \sqrt{1+\varepsilon}$ and $d_e(\chi_e)/d_e(l_e) \leq \sqrt{1+\varepsilon}$.

It remains to analyse the probability with which the congestion N_e of an edge e takes on a value in the interval $\mathfrak{I}_e$ defined above. Since the congestion N_e is the sum of independent Bernoulli random variables, we can apply the Chernoff bound stated in Lemma 2, yielding

$$\begin{aligned}\mathbf{Pr}\left[N_e < l_e\right] =& \mathbf{Pr}\left[N_e < \left(1 - \sqrt{\frac{3\ln(4m)}{\chi_e}}\right)\chi_e\right] \\ \leq& \exp\left(-\frac{1}{3}\left(\sqrt{\frac{3\ln(4m)}{\chi_e}}\right)^2 \chi_e\right) = \frac{1}{4m} \ .\end{aligned}$$

If $\chi_e \geq 3\ln(4m)$, then $u_e = \chi_e + \sqrt{3\ln(4m)\cdot\chi_e}$, for which we obtain

$$\begin{aligned}\mathbf{Pr}\left[N_e > u_e\right] =& \mathbf{Pr}\left[N_e > \left(1 + \sqrt{\frac{3\ln(4m)}{\chi_e}}\right)\chi_e\right] \\ \leq& \exp\left(-\frac{1}{3}\left(\sqrt{\frac{3\ln(4m)}{\chi_e}}\right)^2 \chi_e\right) = \frac{1}{4m} \ .\end{aligned}$$

If $\chi_e \leq 3\ln(4m)$, then $u_e = 6\ln(4m)$, for which we obtain

$$\begin{aligned}\mathbf{Pr}\left[N_e > u_e\right] =& \mathbf{Pr}\left[N_e > (1+1)\cdot 3\ln(4m)\right] \\ \leq& \exp\left(-\frac{1}{3}\cdot 3\ln(4m)\right) = \frac{1}{4m} \ .\end{aligned}$$

Altogether, this implies that $\mathbf{Pr}\left[N_e \notin \mathfrak{I}_e\right] \leq \mathbf{Pr}\left[N_e < u_e\right] + \mathbf{Pr}\left[N_e > l_e\right] \leq 1/2m$, as desired. □

4 Analysis of Classes of Delay Functions

In this section we analyse which conditions Theorem 3 imposes when applied to polynomial delay functions, exponential delay functions, and delay functions from queuing theory.

4.1 Polynomial Delay Functions

We consider polynomial delay functions with nonnegative coefficients and constant degree g. That is, the delay function has the form $d(x) = \sum_{j=0}^{g} a_j x^j$, where $a_j \geq 0$ for $j \in \{0, \ldots, g-1\}$ and $a_g > 0$. Since the coefficients are nonnegative, the function d is non-decreasing. To fulfil the assumption that the delay function never equals zero, we also assume that $a_0 > 0$.

Theorem 4. *A polynomial delay function d with degree g and nonnegative coefficients satisfies Condition (1) in Theorem 3 for all $x \in \mathbb{R}_{\geq 0}$ if*

$$a_0 \geq \frac{((1+\varepsilon)g^2 \cdot 6\ln(4m))^g}{(\sqrt{1+\varepsilon}-1)^{2g+1}} \sum_{j=1}^{g} a_j = \Theta\left(\frac{\ln^g m}{\varepsilon^{2g+1}}\right) \sum_{j=1}^{g} a_j^e \ . \tag{2}$$

Proof. To establish the theorem we show that (2) implies Inequality (1) from Theorem 3 for any $x \geq 0$. In the following, we assume $g > 0$ because for constant functions Inequality (1) is trivially satisfied.

In order to show that Inequality (1) is satisfied we use two upper bounds on the function

$$f(x) = \frac{d(x)}{d(x - \sqrt{6\ln(4m) \cdot x})} \ ,$$

of which one is monotonically increasing and the other is monotonically decreasing. We show that the upper bounds are chosen such that their minimum is bounded from above by $\sqrt{1+\varepsilon}$ for every $x \geq 0$. Since d is non-decreasing and we assumed that $d(x)$ equals $d(0)$ for any $x \leq 0$, we obtain, for every $x \geq 0$,

$$f(x) = \frac{d(x)}{d(x - \sqrt{6\ln(4m) \cdot x})} \leq \frac{d(x)}{d(0)} = \frac{1}{a_0} d(x) \ . \tag{3}$$

The second upper bound on $f(x)$ is presented in the following lemma.

Lemma 5. *For all $x > g^2 \cdot 6\ln(4m)$,*

$$f(x) \leq \frac{1}{1 - g\sqrt{\frac{6\ln(4m)}{x}}} \ .$$

Proof. Since the second derivative of any polynomial with nonnegative coefficients is greater or equal to 0, the delay function is convex. The fact that the first order Taylor approximation of a convex function is always a global underestimator yields, for $x \geq 0$,

$$d(x - \sqrt{6\ln(4m) \cdot x}) \geq d(x) - \sqrt{6\ln(4m) \cdot x} \cdot d'(x) \ . \tag{4}$$

The lower bound in (4) is positive for $x > g^2 \cdot 6\ln(4m)$ because

$$\begin{aligned} d(x) - \sqrt{6\ln(4m) \cdot x} \cdot d'(x) &= \frac{d(x)}{\sqrt{x}}\left(\sqrt{x} - \sqrt{6\ln(4m)} \cdot \frac{xd'(x)}{d(x)}\right) \\ &\geq \frac{d(x)}{\sqrt{x}}\left(\sqrt{x} - \sqrt{g^2 \cdot 6\ln(4m)}\right) > 0 \ , \end{aligned}$$

where the second to the last inequality follows because $xd'(x)/d(x)$ is the so-called *elasticity* of d, which can readily be seen to be upper bounded by g for polynomials with degree g and nonnegative coefficients. Hence, for $x > g^2 \cdot 6\ln(4m)$, we obtain

$$\frac{d(x)}{d(x-\sqrt{6\ln(4m)\cdot x})} \leq \frac{d(x)}{d(x)-\sqrt{6\ln(4m)\cdot x}\cdot d'(x)} = \frac{1}{1-\frac{xd'(x)}{d(x)}\sqrt{\frac{6\ln(4m)}{x}}} \leq \frac{1}{1-g\sqrt{\frac{6\ln(4m)}{x}}} \ ,$$

which concludes the proof of the lemma. □

Let $x_\varepsilon = \frac{(1+\varepsilon)g^2\cdot 6\ln(4m)}{((\sqrt{1+\varepsilon})-1)^2}$. We show that, for $x \leq x_\varepsilon$, the upper bound in (3) yields $f(x) \leq \sqrt{1+\varepsilon}$ and that, for $x \geq x_\varepsilon$, Lemma 5 implies $f(x) \leq \sqrt{1+\varepsilon}$. Since the upper bound in (3) is non-decreasing, it suffices to observe that $d(x_\varepsilon)/a_0 \leq \sqrt{1+\varepsilon}$, which follows from

$$\frac{1}{a_0}d(x_\varepsilon) = \frac{1}{a_0}\sum_{j=0}^{g} a_j x_\varepsilon^j = 1 + \frac{1}{a_0}\sum_{j=1}^{g} a_j x_\varepsilon^j \leq 1 + \frac{x_\varepsilon^g}{a_0}\sum_{j=1}^{g} a_j \leq \sqrt{1+\varepsilon} \ ,$$

where the last inequality follows from (2) and we used the fact that $x_\varepsilon \geq 1$ if $g \geq 1$. Since the upper bound on $f(x)$ given in Lemma 5 is non-increasing and

$$\frac{1}{1-g\sqrt{\frac{6\ln(4m)}{x_\varepsilon}}} = \sqrt{1+\varepsilon} \ ,$$

the theorem follows. □

4.2 Exponential Delay Functions

The lower bound on a_0 in Theorem 4 is determined by the fact that we allow any input value from the domain $\mathbb{R}_{\geq 0}$, which is a natural assumption. However, this means that the bound is too restrictive in the case that the congestion is large, since then also the interval in which the congestion falls is located at some point far to the right of the abscissa. From the fact that the upper bound given in Lemma 5 is decreasing, we can see that in these intervals only smaller values than needed in order to guarantee the approximation factor ε are reached. Since this seems to be a typical characteristic of polynomials, this raises the question whether Theorem 3 can also be applied to delay functions that grow superpolynomially from a certain point on. The next theorem gives an affirmative answer.

Theorem 6. *A delay function d of the form*

$$d(x) = \alpha \cdot \exp\left(\frac{x}{\beta}\right) + \gamma$$

satisfies Condition (1) *in Theorem 3 in some interval* $[0, u]$ *if* $\alpha > 0$*,* $\gamma \geq 0$*, and* $\beta \geq 2\sqrt{u \cdot 6\ln(4m)} / \ln(1+\varepsilon)$*.*

Proof. We have to show that all functions of the suggested form comply with (1). This follows because, for $x \in [0, u]$,

$$\frac{d(x)}{d(x - \sqrt{6\ln(4m) \cdot x})} \leq \exp\left(\frac{\sqrt{6\ln(4m) \cdot x}}{\beta}\right) \leq \exp\left(\frac{\ln(1+\varepsilon)}{2}\right) = \sqrt{1+\varepsilon}\ ,$$

which concludes the proof. □

In Theorem 3 the upper bound u_e is set to $\max\{6\ln(4m), \chi_e + \sqrt{3\ln(4m) \cdot \chi_e}\}$. When substituting u accordingly in the lower bound on β we obtain a bound $f(\chi_e, m, \varepsilon)$ growing in the order of

$$f(\chi_e, m, \varepsilon) \in O\left(\frac{\ln(m) \cdot \sqrt{\chi_e}}{\varepsilon}\right) \quad .$$

4.3 Delay Functions from Queuing Theory

In Kendall's notation [7], we consider the $M/M/1$ queuing model. This means that the queue is processed in a first-come first-served manner, the inter-arrival times at the queue as well as the service times of the packets are exponentially distributed, and each network link can process only one packet at each point in time.

The following basic theorem from queuing theory describes the limiting behaviour of the sojourn time of a packet on a network link. This is the time that packet k, where k tends to infinity, spends on that link in total until it arrives at the next node, i.e. it includes the waiting time in the queue plus the service time of the packet. In the $M/M/1$ queuing model, the arrival of jobs is a Poisson process whose rate is denoted by λ, and the processing of jobs is a Poisson process whose rate is denoted by μ. A basic assumption that has to be fulfilled in order for the theorem to hold is that the occupation rate $\rho = \frac{\lambda}{\mu}$ is strictly smaller than 1. Otherwise there would be more arrivals than the link can handle, which would result in an unbounded growth of the queue.

Theorem 7 ([7]). *In an $M/M/1$ queuing system with arrival rate λ, service rate μ, and in which $\rho < 1$ the limiting behaviour of the expected sojourn time is*

$$E[S] = \frac{1}{\mu - \lambda}.$$

In the following theorem we interpret the congestion as the arrival rate λ and we assume that the considered link has a certain service rate μ.

Theorem 8. *A delay function d of the form*

$$d(x) = \begin{cases} \frac{1}{\mu - x} & \text{if } x < \mu \\ \infty & \text{if } x \geq \mu \end{cases}$$

satisfies Condition (1) *in Theorem 3 in some interval* $[0, u]$ *if*

$$\mu \geq u + \frac{\sqrt{6 \ln(4m) u}}{(\sqrt{1+\varepsilon}) - 1} \enspace .$$

Proof. Since $x \leq u$ and $\mu \geq u$ we can use the finite part of $d(x)$ to obtain

$$\frac{d(x)}{d(x - \sqrt{6\ln(4m) \cdot x})} = 1 + \frac{\sqrt{6\ln(4m) \cdot x}}{\mu - x} \leq 1 + \frac{\sqrt{6\ln(4m) \cdot u}}{\mu - u} \leq \sqrt{1+\varepsilon} \enspace .$$

The first inequality follows from the fact that the function is monotonically increasing in x and $x \leq u$, while the second inequality follows directly from the lower bound on μ. □

When setting u to u_e, analogous to the case of the exponential functions, we obtain a lower bound $g(\chi_e, m, \varepsilon)$ on μ growing in the order of

$$g(\chi_e, m, \varepsilon) \in O\left(\chi_e + \frac{\ln(m) \cdot \sqrt{\chi_e}}{\varepsilon}\right) \enspace .$$

5 Combined Delay Functions

In the previous section we applied Theorem 3 to several classes of functions. In this section we prove a general theorem showing that if two delay functions satisfy (1) in Theorem 3, then also their maximum satisfies this property. This allows us to combine different types of delay functions. One weak point of Theorem 6 concerning exponential functions is that it works only for exponential functions that grow slowly up until $2\sqrt{u \cdot 6\ln(4m)}/\ln(1+\varepsilon)$. The following theorem allows us to combine such an exponential function with a polynomial that satisfies Theorem 4. If we take the maximum over these two functions, we obtain a function that grows polynomially until some point and exponentially thereafter.

Theorem 9. *Let p and q denote two delay functions that satisfy Condition* (1) *in Theorem 3 in some interval $x \in [0, u]$. Then the function*

$$d(x) = \max\{p(x), q(x)\}$$

also satisfies this condition for $x \in [0, u]$.

Proof. Let $x \in [0, u]$, $x' = x - \sqrt{6\ln(4m) \cdot x}$, and without loss of generality assume $d(x) = p(x)$. Since p satisfies (1), we know that $p(x)/p(x') \leq \sqrt{1+\varepsilon}$. Hence, if $d(x') = p(x')$, then (1) follows immediately. If, however, $d(x') = q(x')$ then $q(x') \geq p(x')$ and by the definition of $d(x)$ we obtain

$$\frac{d(x)}{d(x')} = \frac{p(x)}{q(x')} \leq \frac{p(x)}{p(x')} \leq \sqrt{1+\varepsilon} \enspace .$$

The last inequality holds again because p satisfies (1). □

6 Conclusions

In this paper, we have considered network congestion games with delay functions from several different classes. We have identified properties that delay functions from these classes have to satisfy in order to guarantee that an approximate Nash equilibrium can be computed by randomised rounding in polynomial time. Additionally, we have presented a method of combining these delay functions.

It remains an interesting open question to explore the limits of approximability further and to close the gap between the PLS-completeness results and the positive results presented in this paper. This could be done by either proving PLS-completeness of computing approximate equilibria for more natural delay functions or by extending the positive results to larger classes of functions. We believe that other techniques than randomised rounding are required for the latter.

References

1. H. Ackermann, H. Röglin, and B. Vöcking. On the impact of combinatorial structure on congestion games. In *Proc. of the 47th Ann. IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 613–622, 2006.
2. M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
3. S. Chien and A. Sinclair. Convergence to approximate Nash equilibria in congestion games. In *Proc. of the 18th Ann. ACM–SIAM Symp. on Discrete Algorithms (SODA)*, pages 169–178, 2007.
4. G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proc. of the 37th Ann. ACM Symp. on Theory of Computing (STOC)*, pages 67–73, 2005.
5. A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *Proc. of the 36th Ann. ACM Symp. on Theory of Computing (STOC)*, pages 604–612, 2004.
6. J. Jackson. Jobshop-like queueing systems. *Management Science*, 10(1):131–142, 1963.
7. D. G. Kendall. Some problems in the theory of queues. *Royal Statistic Society*, 13(2):151–182, 1951.
8. M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005.
9. P. Raghavan and C. D. Thompson. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7:365–374, 1987.
10. R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *Int. Journal of Game Theory*, 2:65–67, 1973.
11. T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
12. A. Skopalik and B. Vöcking. Inapproximability of pure Nash equilibria. In *Proc. of the 40th Ann. ACM Symp. on Theory of Computing (STOC)*, 2008. to appear.
13. J. G. Wardrop. Some theoretical aspects of road traffic research. In *Proc. of the Institute of Civil Engineers, Pt. II*, volume 1, pages 325–378, 1952.