Improving Selfish Routing for Risk-Averse Players*

Dimitris Fotakis[1], Dimitris Kalimeris[2], and Thanasis Lianeas[3]

[1] School of Electrical and Computer Engineering,
National Technical University of Athens, Greece. fotakis@cs.ntua.gr
[2] Department of Informatics and Telecommunications,
National and Kapodistrian University of Athens, Greece. sdi1000049@di.uoa.gr
[3] Department of Electrical and Computer Engineering,
University of Texas at Austin, Austin, TX 78701. tlianeas@austin.utexas.edu

Abstract. We investigate how and to which extent one can exploit risk-aversion and modify the perceived cost of the players in selfish routing so that the Price of Anarchy (PoA) is improved. We introduce small random perturbations to the edge latencies so that the expected latency does not change, but the perceived cost of the players increases due to risk-aversion. We adopt the model of γ-modifiable routing games, a variant of routing games with restricted tolls. We prove that computing the best γ-enforceable flow is NP-hard for parallel-link networks with affine latencies and two classes of heterogeneous risk-averse players. On the positive side, we show that for parallel-link networks with heterogeneous players and for series-parallel networks with homogeneous players, there exists a nicely structured γ-enforceable flow whose PoA improves fast as γ increases. We show that the complexity of computing such a γ-enforceable flow is determined by the complexity of computing a Nash flow of the original game. Moreover, we prove that the PoA of this flow is best possible in the worst-case, in the sense that there are instances where (i) the best γ-enforceable flow has the same PoA, and (ii) considering more flexible modifications does not lead to any further improvement.

1 Introduction

Routing games provide an elegant and practically useful model of selfish resource allocation in transportation and communication networks and have been extensively studied (see e.g., [17]). The majority of previous work assumes that the players select their routes based on precise knowledge of edge delays. In practical applications however, the players cannot accurately predict the actual delays due to their limited knowledge about the traffic conditions and due to unpredictable

* This research was supported by the project Algorithmic Game Theory, co-financed by the European Union (European Social Fund) and Greek national funds, through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework - Research Funding Program: THALES, investing in knowledge society through the European Social Fund, and by grant NSF CCF 1216103.

events that affect the edge delays and introduce uncertainty (see e.g., [14,12,1,13] for examples). Hence, the players select their routes based only on delay estimations and are aware of the uncertainty and the potential inaccuracy of them. Therefore, to secure themselves from increased delays, whenever this may have a considerable influence, the players select their routes taking uncertainty into account (e.g., people take a safe route or plan for a longer-than-usual delay when they head to an important meeting or to catch a long-distance flight).

Recent work (see e.g., [12,15,1,13] and the references therein) considers routing games with *stochastic delays* and *risk-averse players*, where instead of the route that minimizes her expected delay, each player selects a route that guarantees a reasonably low actual delay with a reasonably high confidence. There have been different models of stochastic routing games, each modeling the individual cost of risk-averse players in a slightly different way. In all cases, the actual delay is modeled as a random variable and the perceived cost of the players is either a combination of the expectation and the standard deviation (or the variance) of their delay [12,13] or a player-specific quantile of the delay distribution [14,1] (see also [18,4] about the perceived cost of risk-averse players).

No matter the precise modeling, we should expect that stochastic delays and risk-aversion cannot improve the network performance at equilibrium. Interestingly, [15,13] indicate that in certain settings, stochastic delays and risk-aversion can actually improve the network performance at equilibrium. Motivated by these results, we consider routing games on parallel-link and series-parallel networks and investigate how one can exploit risk-aversion in order to modify the perceived cost of the (possibly heterogeneous) players so that the PoA is significantly improved.

Routing Games. To discuss our approach more precisely, we introduce the basic notation and terminology about routing games. A (non-atomic) *selfish routing game* (or instance) is a tuple $\mathcal{G} = (G(V,E), (\ell_e)_{e\in E}, r)$, where $G(V,E)$ is a directed network with a source s and a sink t, $\ell_e : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is a non-decreasing delay (or latency) function associated with edge e and $r > 0$ is the traffic rate. We let $\mathcal{P}$ denote the set of simple $s-t$ paths in G. We say that G is a parallel-link network if each $s-t$ path is a single edge (or link).

A (feasible) *flow* f is a non-negative vector on $\mathcal{P}$ such that $\sum_{p\in\mathcal{P}} f_p = r$. We let $f_e = \sum_{p:e\in p} f_p$ be flow routed by f on edge e. Given a flow f, the latency of each edge e is $\ell_e(f) = \ell_e(f_e)$, the latency of each path p is $\ell_p(f) = \sum_{e\in p} \ell_e(f)$ and the latency of f is $L(f) = \max_{p:f_p>0} \ell_p(f)$.

The traffic r is divided among infinitely many players, each trying to minimize her latency. A flow f is a *Wardrop-Nash flow* (or a *Nash flow*, for brevity), if all traffic is routed on minimum latency paths, i.e., for any $p \in \mathcal{P}$ with $f_p > 0$ and for all $p' \in \mathcal{P}$, $\ell_p(f) \leq \ell_{p'}(f)$. Therefore, in a Wardrop-Nash flow f, all players incur a minimum common latency $\min_p \ell_p(f) = L(f)$. Under weak assumptions on delay functions, a Nash flow exists and is essentially unique (see e.g., [17]).

The efficiency of a flow f is measured by the *total latency* $C(f)$ of the players, i.e., by $C(f) = \sum_{e\in E} f_e \ell_e(f)$. The *optimal flow*, denoted o, minimizes the total latency among all feasible flows. The *Price of Anarchy* (PoA) quantifies the performance degradation due to selfishness. The PoA($\mathcal{G}$) of a routing game $\mathcal{G}$ is the ratio $C(f)/C(o)$, where f is the Nash flow and o is the optimal flow o of $\mathcal{G}$. The PoA of a class of routing games is the maximum PoA over all games in the class. For routing games with latency functions in a class $\mathcal{D}$, the PoA is equal to $\mathrm{PoA}(\mathcal{D}) = (1 - \beta(\mathcal{D}))^{-1}$, where $\beta(\mathcal{D}) = \sup_{\ell \in \mathcal{D}, x \geq y \geq 0} \frac{y(l(x)-l(y))}{xl(x)}$ only depends on the class of latency functions $\mathcal{D}$ [17,3].

Using Risk-Aversion to Modify Edge Latencies. The starting point of our work is that in some practical applications, we may intentionally introduce variance to edge delays so that the expected delay does not change, but the risk-averse cost of the players increases. E.g., in a transportation network, we can randomly increase or decrease the proportion of time allocated to the green traffic light for short periods or we can open or close an auxiliary traffic lane. In a communication network, we might randomly increase or decrease the link capacity allocated to a particular type of traffic or change its priority. At the intuitive level, we expect that the effect of such random changes to risk-averse players is similar to that of refundable tolls (see e.g., [6,11]), albeit restricted in magnitude due to the bounded variance in edge delays that we can afford.

E.g., let e be an edge with latency $\ell_e(x)$ where we can increase the latency temporarily to $(1+\alpha_1)\ell_e(x)$ and decrease it temporarily to $(1-\alpha_2)\ell_e(x)$. If we implement the former change with probability p_1 and the latter with probability $p_2 < 1-p_1$, the latency function of e becomes a random variable with expectation $[p_1(1+\alpha_1)+p_2(1-\alpha_2)+(1-p_1-p_2)]\ell_e(x)$. Adjusting p_1 and p_2 (and possibly α_1 and α_2) so that $p_1\alpha_1 = p_2\alpha_2$, we achieve an expected latency of $\ell_e(x)$. However, if the players are (homogeneously) risk-averse and their perceived delay is given by an $(1-p_1+\varepsilon)$-quantile of the delay distribution (e.g., as in [14,1]), the perceived latency on e is $(1+\alpha_1)\ell_e(x)$. Similarly, if the individual cost of the risk-averse players are given by the expectation plus the standard deviation of the delay distribution (e.g., as in [12]), the perceived latency is $(1+\sqrt{p_1\alpha_1^2+p_2\alpha_2^2})\ell_e(x)$. In both cases, we can achieve a significant increase in the delay perceived by risk-averse players, while the expected delay remains unchanged.

In most practical situations, the feasible changes in the latency functions are bounded (and relatively small). Combined with the particular form of risk-averse individual cost, this determines an upper bound γ_e on the multiplicative increase of the delay on each edge e. Moreover, the players may evaluate risk differently and be *heterogeneous* wrt. their risk-aversion factors. So, in general, the traffic rate r is partitioned into k risk-averse classes, where each class i consists of the players with risk-aversion factor a^i and includes a traffic rate r^i. If we implement a multiplicative increase γ_e on the perceived latency of each edge e, the players in class i have perceived cost $(1+a^i\gamma_e)\ell_e(f)$ on each e and $\sum_{e\in p}(1+a^i\gamma_e)\ell_e(f)$ on each path[1] p. If the players are *homogeneous* wrt. their risk-aversion, there is a single class of players with traffic rate r and risk-aversion factor $a = 1$.

[1] To simplify the model and make it easily applicable to general networks, we assume that the perceived cost of the players under latency modifications is separable. This is a reasonable simplifying assumption on the structure of risk-averse costs (see also [15,13]) and only affects the extension of our results to series-parallel networks.

Contribution. In this work, we assume a given upper bound γ on the maximum increase in the latency functions and refer to the corresponding routing game as a γ-*modifiable game*. We consider both homogeneous and heterogeneous risk-averse players. We adopt this model as a simple and general abstraction of how one can exploit risk-aversion to improve the PoA of routing games. Technically, our model is a variant of restricted refundable tolls considered in [9,2] for homogeneous players and in [10] for heterogeneous players. However, on the conceptual side and to the best of our knowledge, this is the first time that risk-aversion is proposed as a means of implementing restricted tolls, and through this, as a potential remedy to the inefficiency of selfish routing.

A flow f is γ-*enforceable* if there is γ_e-modification on each edge e, with $0 \leq \gamma_e \leq \gamma$, so that f is a Nash flow of the modified game, i.e., for each player class i, for every path p used by class i, and for all paths p', $\sum_{e\in p}(1+a^i\gamma_e)\ell_e(f) \leq \sum_{e\in p'}(1+a^i\gamma_e)\ell_e(f)$. In this work, we are interested in computing either the best γ-enforceable flow, which minimizes total latency among all γ-enforceable flows, or a γ-enforceable flow with low PoA. We measure the PoA in terms of the total expected latency (instead of the total perceived delay of the players). In practical applications, the total expected latency is directly related to many crucial performance parameters (e.g., to the expected pollution in a transportation network or to the expected throughput in a communication network) and thus, it is the quantity that a central planner usually seeks to minimize.

In Section 3, we consider routing games on parallel links with homogeneous players and show that for every $\gamma > 0$, there is a nicely structured γ-enforceable flow whose PoA improves significantly as γ increases. More specifically, based on a careful rerouting procedure, we show that given an optimal flow o, we can find a γ-enforceable flow f (along with the corresponding γ-modification) that "mimics" o in the sense that if $f_e < o_e$, e gets a 0-modification, while if $f_e > o_e$, e gets a γ-modification (Lemma 1). The proof of Lemma 1 implies that given o, we can compute such a flow f and the corresponding γ-modification in time $O(|E|\,T_{\mathrm{NE}})$, where T_{NE} is the complexity of computing a Nash flow in the original instance. Generalizing the variational inequality approach of [3], similarly to [2, Section 4], we prove (Theorem 1) that the PoA of the γ-enforceable flow f constructed in Lemma 1 is at most $(1-\beta_\gamma(\mathcal{D}))^{-1}$, where $\mathcal{D}$ is the class of latency functions and $\beta_\gamma(\mathcal{D}) = \sup_{\ell\in\mathcal{D},x\geq y\geq 0} \frac{y(\ell(x)-\ell(y))-\gamma(x-y)\ell(x)}{x\ell(x)}$ is a natural generalization of the quantity $\beta(\mathcal{D})$ introduced in [3]. E.g., for affine latency functions, the PoA of the γ-enforceable f is at most $(1-(1-\gamma)^2/4)^{-1}$ (Corollary 1), which is significantly less that $4/3$ even for small values of γ. We also show that the PoA of such γ-enforceable flows is best possible in the worst-case for γ-modifiable games with latency functions in class $\mathcal{D}$ (Theorem 2).

In Section 4, we switch to parallel-link games with heterogenous players. We prove that computing the best γ-enforceable flow is NP-hard for parallel-link games with affine latencies and only two classes of heterogeneous risk-averse players (Theorem 3). The proof modifies the construction in [16, Section 6], which shows that the best Stackelberg modification of parallel-link instances is NP-hard. Our result significantly strengthens [10, Theorem 1], which establishes

NP-hardness of best restricted tolls in general $s-t$ networks with affine latencies. On the positive side, we apply [10, Algorithm 1] and show (Theorem 5) that the γ-enforceable flow f of Lemma 1 can be turned into a γ-enforceable flow for parallel-link instances with heterogeneous players. Since only the γ-modifications are adjusted for heterogeneous players, but the flow itself does not change, the PoA of f is bounded as above and remains best possible in the worst case.

In Section 5, we extend our approach of finding a γ-enforceable flow that "mimics" the optimal flow to series-parallel networks. Series-parallel networks have received considerable attention in the literature of refundable tolls, see e.g., [5,7], but to the best of our knowledge, they have not been explicitly considered in the setting of restricted tolls. Extending the rerouting procedure of Lemma 1, we show that for routing games in series-parallel networks with homogeneous players, there is a γ-enforceable flow with PoA at most $(1-\beta_\gamma(\mathcal{D}))^{-1}$ (Lemma 2 and Theorem 6). Such a γ-enforceable flow and the corresponding γ-modifications can be computed in time polynomially related to the time needed for computing Nash flows in series-parallel networks (Lemma 3).

In Section 6, we consider (p,γ)-modifiable games, where the p-norm of the edge modifications vector $(\gamma_e)_{e\in E}$ is at most γ. This generalization captures applications where the total variance introduced in the network should be bounded by γ and could potentially lead to an improved PoA. We prove that the worst-case PoA under (p,γ)-modifications is essentially identical to the worst-case PoA under $\gamma/\sqrt[p]{m}$-modifications (Theorem 8). Therefore, even for (p,γ)-modifiable games, the PoA of the $\gamma/\sqrt[p]{m}$-enforceable flow of Lemma 1 is essentially best possible. Due to space constraints, we only sketch the main ideas behind our results and defer the technical details to the full version of this work.

Previous Work. On the conceptual side, our work is closest to those considering the PoA of stochastic routing games with risk-averse players [12,1,15]. Nikolova and Stier-Moses [13] recently introduced the *price of risk-aversion* (PRA), which is the worst-case ratio of the total latency of the Nash flow for risk-averse players to the total latency of the Nash flow for risk-neutral players. Interestingly, PRA can be smaller than 1 and as low as $1-\beta(\mathcal{D})$ for stochastic routing games on parallel-links (i.e., risk-aversion can improve the PoA to 1 for certain instances).

On the technical side, our work is closest to those investigating the properties of restricted refundable tolls for routing games [9,2,10]. Bonifaci et al. [2] proved that for parallel-link networks with homogeneous players, computing the best γ-enforceable flow reduces to the solution of a convex program. Moreover, they presented a tight bound of $(1-\beta_\gamma(\mathcal{D}))^{-1}$ on the PoA of a γ-enforceable flow for routing games with latency functions in class $\mathcal{D}$. Jelinek et al. [10] considered restricted tolls for heterogeneous players and proved that computing the best γ-enforceable flow for $s-t$ networks with affine latencies is NP-hard. On the positive side, they proved that for parallel-link games with heterogeneous players, deciding whether a given flow is γ-enforceable (and finding the corresponding γ-modification) can be performed in polynomial time. Moreover, they showed how to compute the best γ-enforceable flow for parallel-link games with heterogeneous players if the maximum allowable modification on each edge is either 0 or infinite.

2 The Model and Preliminaries

The basic model of routing games is introduced in Section 1. Next, we introduce some more notation and the classes of γ-modifiable and (p,γ)-modifiable games.

γ-Modifiable Routing Games. A selfish routing game with heterogeneous players in k classes is a tuple $\mathcal{G} = (G(V,E), (\ell_e)_{e\in E}, (a^i)_{i\in[k]}, (r^i)_{i\in[k]})$, where G is a directed $s-t$ network with m edges, a^i is the aversion factor of the players in class i and r_i is the amount of traffic with aversion a^i. We assume that $a^1 = 1$ and $a^1 < a^2 < \ldots < a^k$. If the players are homogeneous, there is a single class with risk aversion $a^1 = 1$ and traffic rate r. Then, an instance is $\mathcal{G} = (G, \ell, r)$.

A flow f is a non-negative vector on $\mathcal{P} \times \{1, \ldots, k\}$. We let $f_p^{a^i}$ be the flow with aversion a^i on path p and $f_p = \sum_i f_p^{a^i}$ be the total flow on path p. Similarly, $f_e^{a^i} = \sum_{p:e\in p} f_p^{a^i}$ is the flow with aversion a^i on edge e and $f_e = \sum_i f_p^{a^i}$ is the total flow on edge e. We let $a_e^{\min}(f)$ (resp. $a_e^{\max}(f)$) be the smallest (resp. largest) aversion factor in e under f. If e is not used by f, we let $a_e^{\min}(f) = a_e^{\max}(f) = a^k$. We say that an edge e (resp. path p) is used by players of type a^i if $f_e^{a^i} > 0$ (resp. for all $e \in p$). To simplify notation, we may write ℓ_e, instead of $\ell_e(f)$.

We say that a routing game $\mathcal{G}$ is γ-*modifiable* if we can select a $\gamma_e \in [0, \gamma]$ for each edge e and change the edge latencies perceived by the players of type a^i from $\ell_e(x)$ to $(1 + a^i\gamma_e)\ell_e(x)$ using small random perturbations. Any vector $\boldsymbol{\Gamma} = (\gamma_e)_{e\in E}$, where $\gamma_e \in [0, \gamma]$ for each edge e, is a γ-*modification* of $\mathcal{G}$. Given a γ-modification $\boldsymbol{\Gamma}$, we let $\mathcal{G}^{\boldsymbol{\Gamma}}$ denote the γ-modified routing game where the perceived cost of the players is changed according to the modification $\boldsymbol{\Gamma}$.

A flow f is a *Nash flow* of $\mathcal{G}^{\boldsymbol{\Gamma}}$, if for any path p and any type a^i with $f_p^{a^i} > 0$ and for all paths p', $\sum_{e\in p}(1 + a^i\gamma_e)\ell_e(f) \le \sum_{e\in p'}(1 + a^i\gamma_{e'})\ell_{e'}(f)$. Given a routing game $\mathcal{G}$, we say that a flow f is γ-*enforceable*, or simply *enforceable*, if there exists a γ-modification $\boldsymbol{\Gamma}$ of $\mathcal{G}$ such that f is a Nash flow of $\mathcal{G}^{\boldsymbol{\Gamma}}$.

Our assumption is that γ-modifications do not change the expected latency. Therefore, the total latency of f in both $\mathcal{G}^{\boldsymbol{\Gamma}}$ and $\mathcal{G}$ is $C(f) = \sum_{e\in E} f_e\ell_e(f)$. Hence, the optimal flow o of $\mathcal{G}$ is also an optimal flow of $\mathcal{G}^{\boldsymbol{\Gamma}}$. A flow f is the *best* γ-*enforceable* flow of $\mathcal{G}$ if for any other γ-enforceable flow f' of $\mathcal{G}$, $C(f) \le C(f')$. The Price of Anarchy $\mathrm{PoA}(\mathcal{G}^{\boldsymbol{\Gamma}})$ of the modified game $\mathcal{G}^{\boldsymbol{\Gamma}}$ is equal to $C(f)/C(o)$, where f is the Nash flow of $\mathcal{G}^{\boldsymbol{\Gamma}}$. For a γ-modifiable game $\mathcal{G}$, the PoA of $\mathcal{G}$ under γ-modifications, denoted $\mathrm{PoA}_\gamma(\mathcal{G})$, is $C(f)/C(o)$, where f is the best γ-enforceable flow of $\mathcal{G}$. For routing games with latency functions in class $\mathcal{D}$, $\mathrm{PoA}_\gamma(\mathcal{D})$ denotes the maximum $\mathrm{PoA}_\gamma(\mathcal{G})$ over all γ-modifiable games $\mathcal{G}$ with latencies in $\mathcal{D}$.

(p,γ)-Modifiable Routing Games. Generalizing γ-modifiable games, we select a modification $\gamma_e \ge 0$ for each edge e so that $\|(\gamma_e)_{e\in E}\|_p = \sqrt[p]{\sum_{e\in E}\gamma_e^p} \le \gamma$, for some given integer $p \ge 1$, and change the perceived edge latencies as above. We refer to such games as (p,γ)-modifiable. All the notation above naturally generalizes to (p,γ)-modifiable games. The PoA of a game $\mathcal{G}$ under (p,γ)-modifications, denoted $\mathrm{PoA}_\gamma^p(\mathcal{G})$, is $C(f)/C(o)$, where f is the best (p,γ)-enforceable flow of $\mathcal{G}$. Similarly, $\mathrm{PoA}_\gamma^p(\mathcal{D})$ is the maximum PoA of all (p,γ)-modifiable games with latency functions in class $\mathcal{D}$.

Series-Parallel Networks. A directed $s - t$ network $G(V, E)$ is *series-parallel* if it either consists of a single edge (s, t) or can be obtained from two series-parallel networks with terminals (s_1, t_1) and (s_2, t_2) composed either in series or in parallel. In a *series composition*, t_1 is identified with s_2, s_1 becomes s, and t_2 becomes t. In a *parallel composition*, s_1 is identified with s_2 and becomes s, and t_1 is identified with t_2 and becomes t (see also [19] for computing the decomposition of a series-parallel network in linear time).

3 Modifying Routing Games in Parallel-Link Networks

In this section, we study γ-modifiable games on parallel-link networks with homogeneous risk-averse players. The following is a corollary of [2, Theorem 1] (see also the main result of [6,11]) and characterizes γ-enforceable optimal flows.

Proposition 1. *Let $\mathcal{G}$ be a γ-modifiable game on parallel links and let o be the optimal flow of $\mathcal{G}$. Then, o is γ-enforceable in $\mathcal{G}$ if and only if for any link e with $o_e > 0$ and all links $e' \in E$, $\ell_e(o) \le (1+\gamma)\ell_{e'}(o)$.*

Next, we show that for any instance $\mathcal{G}$, there exist a flow f mimicking o and a γ-modification enforcing f as the Nash flow of the modified instance.

Lemma 1. *Let $\mathcal{G} = (G, \ell, r)$ be a γ-modifiable instance on parallel-links with homogeneous risk-averse players and let o be the optimal flow of $\mathcal{G}$. There is a feasible flow f and a γ-modification $\boldsymbol{\Gamma}$ of $\mathcal{G}$ such that*

(i) f is a Nash flow of the modified instance $G^{\boldsymbol{\Gamma}}$.

(ii) for any link e, if $f_e < o_e$, then $\gamma_e = 0$, and if $f_e > o_e$, then $\gamma_e = \gamma$.

Moreover, given o, we can compute f and $\boldsymbol{\Gamma}$ in time $O(mT_{NE})$, where T_{NE} is the complexity of computing the Nash flow of any given γ-modification of $\mathcal{G}$.

Proof sketch. The interesting case is where o is not γ-enforceable. Then, we use induction on the number of links. The base case is obvious. For the inductive step, let m be a used link with maximum latency in o. Removing m and decreasing the total traffic rate by $o_m > 0$, we obtain an instance $\mathcal{G}_{-m} = (G_{-m}, \ell, r - o_m)$ with one link less than $\mathcal{G}$. By induction hypothesis, there are a flow f' and a γ-modification $\boldsymbol{\Gamma}' = (\gamma'_e)_{e \in E_{-m}}$ so that properties (i) and (ii) hold for $\mathcal{G}_{-m}$.

Now we restore link m and the traffic rate to r. The lemma follows directly from the hypothesis if there is a modification γ_m so that $(1+\gamma_m)\ell_m(o) = L(f')$.

Otherwise, we have that $\ell_m(o) > L(f')$. Then, we carefully reroute flow from link m to the remaining links while maintaining properties (i) and (ii) in $\mathcal{G}_{-m}$. We do so until the latency of m becomes equal to the cost of the equilibrium flow that we maintain (under rerouting) in $\mathcal{G}_{-m}$. In order to maintain property (ii), we pay attention to links e where the flow f'_e reaches o_e for the first time and to links e' where $\gamma'_{e'}$ reaches γ for the first time. For the former, we stop increasing flow and start increasing γ'_e, so that the equilibrium property is maintained. For the latter, we stop increasing $\gamma'_{e'}$ and start increasing the flow again.

More formally, we partition the links in E_{-m} in three classes, according to property (ii) and to the current equilibrium flow f' and modification $\boldsymbol{\Gamma}'$. We let

$E_1 = \{e \in E_{-m} : f'_e < o_e \text{ and } \gamma'_e = 0\}$, $E_2 = \{e \in E_{-m} : f'_e = o_e \text{ and } \gamma'_e < \gamma\}$ and $E_3 = \{e \in E_{-m} : f'_e \geq o_e \text{ and } \gamma'_e = \gamma\}$. We let $L = (1 + \gamma'_e)\ell_e(f')$, where e is any link with $f'_e > 0$, be the cost of the current equilibrium flow f' in $\mathcal{G}_{-m}$. Moreover, we let $L_1 = \min_{e \in E_1} \ell_e(o)$ be the minimum cost of an equilibrium flow in $\mathcal{G}_{-m}$ that causes some links of E_1 to move to E_2, let $L_2 = \min_{e \in E_2}(1+\gamma)\ell_e(o)$ be the minimum cost of an equilibrium flow in $\mathcal{G}_{-m}$ that causes some links of E_2 to move to E_3, and let $L' = \min\{L_1, L_2\} \geq L$.

We reroute flow from link m to the links in $E_1 \cup E_3$ and increase γ'_e's for the links in E_2 so that we obtain an equilibrium flow in E_{-m} with cost L'. To this end, we let x_e be such that $L' = (1 + \gamma'_e)\ell_e(f'_e + x_e)$, for all $e \in E_1 \cup E_3$. Namely, x_e is the amount of flow we need to reroute to a link $e \in E_1 \cup E_3$ so that its cost becomes L'. For each link $e \in E_2$, we let $x_e = 0$ and increase its modification factor so that $L' = (1 + \gamma'_e)\ell_e(o)$. So the total amount of flow that we need to reroute from E_{-m} is $x = \sum_{e \in E_{-m}} x_e$. Next, we distinguish between different cases depending on the flow and the latency in link m after rerouting.

If $x < o_m$ and $\ell_m(o_m - x) \geq L'$, we update the flow on link m to $o_m - x$, the flow on each link $e \in E_{-m}$ to $f'_e + x_e$, and the modification factors of all links in E_2 and apply the rerouting procedure again (in fact, if $\ell_m(o_m - x) = L'$, the procedure terminates). By the definition of L', every time we apply the rerouting procedure, either some links e move from E_1 to E_2 (because after the update $f'_e = o_e$) or some links e' move from E_2 to E_3 (because after the update $\gamma'_e = \gamma$). Since links in E_3 cannot move to a different class, this rerouting procedure can be applied at most $2m$ times (in total, for all induction steps).

If $x < o_m$ and $\ell_m(o_m - x) < L'$, by continuity (see also [8, Section 3]), there is some $L'' \in (L, L')$ such that updating the flow and the modification factors with target equilibrium cost L'' (instead of L') reroutes flow $x' \leq x < o_m$ from link m to the links in E_{-m} so that $\ell_m(o_m - x') = L''$ and L'' is the cost of any used link in E_{-m}. Hence, we obtain the desired γ-enforceable flow f and the corresponding modification Γ. Such a value L'' can be found by computing the (unique) equilibrium flow f for the links in $E_1 \cup E_3 \cup \{m\}$ with total traffic rate $o_m + \sum_{e \in E_1 \cup E_3} f'_e$ and modifications $\gamma_e = 0$ for all links $e \in E_1 \cup \{m\}$ and $\gamma_e = \gamma$ for all links $e \in E_3$. Moreover, for all links $e \in E_2$, we let $f_e = o_e$ and set γ_e so that $L'' = (1 + \gamma_e)\ell_e(o_e)$, where $\gamma_e \leq \gamma$, because $L'' \leq L'$.

If $x = o_m$ and $\ell_m(0) < L'$, the target equilibrium cost L'' lies between L and L' and we apply the same procedure as above. If $x = o_m$ and $\ell_m(0) \geq L'$, we let $\gamma_m = 0$ and $f_m = 0$. Then, we apply rerouting as above and set $f_e = f'_e$ and $\gamma_e = \gamma'_e$ for the remaining links $e \in E_{-m}$.

If $x > o_m$ and $\ell_m(0) \geq L'$, the target equilibrium cost L'' lies between L and L' and link m is not used at equilibrium. So, we let $\gamma_m = 0$ and $f_m = 0$, compute the equilibrium flow f for the links in $E_1 \cup E_3$ with traffic rate $r - \sum_{e \in E_2} o_e$ and modifications $\gamma_e = 0$ for all $e \in E_1$ and $\gamma_e = \gamma$ for all $e \in E_3$. If $L'' \in (L, L')$ is the cost of this equilibrium flow, for all links $e \in E_2$, we let $f_e = o_e$ and set γ_e so that $L'' = (1 + \gamma_e)\ell_e(o_e)$. If $x > o_m$ and $\ell_m(0) < L'$, the target equilibrium cost L'' again lies between L and L', but now link m may be used at equilibrium. Hence, we apply the same procedure but with link m now included in E_1. □

Price of Anarchy Analysis. We next prove an upper bound on the PoA of the γ-enforceable flow f of Lemma 1. This also serves as an upper bound on the PoA_γ of the best γ-enforceable flow. The approach is conceptually similar to that of [3] and exploits the properties (i) and (ii) of Lemma 1. The results are similar to the results in [2, Section 4], although our approach and the γ-modification that we consider here are different.

Theorem 1. *For γ-modifiable instances on parallel-links with latency functions in class $\mathcal{D}$, $\mathrm{PoA}_\gamma(\mathcal{D}) \leq (1-\beta_\gamma(\mathcal{D}))^{-1}$, where*

$$\beta_\gamma(\mathcal{D}) = \sup_{\ell\in\mathcal{D}, x\geq y\geq 0} \frac{y(\ell(x)-\ell(y)) - \gamma(x-y)\ell(x)}{x\ell(x)}$$

Proof sketch. Let $\mathcal{G} = (G, \ell, r)$ be an instance on parallel-links with latency functions in class $\mathcal{D}$ and let o be the optimal solution of $\mathcal{G}$. We consider the γ-enforceable flow f and the corresponding modification $\boldsymbol{\Gamma} = (\gamma_e)_{e\in E}$ of Lemma 1. By definition, $\mathrm{PoA}_\gamma(\mathcal{G}) \leq \mathrm{PoA}(\mathcal{G}^{\boldsymbol{\Gamma}})$. We next show an upper bound on $\mathrm{PoA}(\mathcal{G}^{\boldsymbol{\Gamma}})$.

Similarly to the proof of Lemma 1, we partition the links used by f into sets $E_1 = \{e \in E : 0 < f_e < o_e\}$, $E_2 = \{e \in E : f_e = o_e > 0\}$ and $E_3 = \{e \in E : f_e > o_e\}$. Using the fact that f is a Nash flow of $\mathcal{G}^{\boldsymbol{\Gamma}}$, we obtain that

$$\sum_{e\in E} f_e\ell_e(f) \leq \sum_{e\in E} o_e\ell_e(o) + \sum_{e\in E_3}\Big(o_e(\ell_e(f)-\ell_e(o)) - \gamma(f_e-o_e)\ell_e(f)\Big) \quad (1)$$

Using the definition of $\beta_\gamma(\mathcal{D})$, we obtain that:

$$\sum_{e\in E} f_e\ell_e(f) \leq \sum_{e\in E} o_e\ell_e(o) + \beta_\gamma(\mathcal{D})\sum_{e\in E_3} f_e\ell_e(f)$$

Therefore, $\mathrm{PoA}(\mathcal{G}^{\boldsymbol{\Gamma}}) \leq (1-\beta_\gamma(\mathcal{D}))^{-1}$. □

Next we give upper bounds on the $\mathrm{PoA}_\gamma(\mathcal{D})$ for γ-modifiable instances with polynomial latency functions. These bounds apply to the γ-enforceable flow f of Lemma 1 and to the best γ-enforceable flow.

Corollary 1. *For γ-modifiable instances on parallel links with polynomial latency functions of degree d, we have that $\mathrm{PoA}_\gamma(d) = 1$, for all $\gamma \geq d$, and*

$$\mathrm{PoA}_\gamma(d) \leq \left(1 - d\big(\tfrac{\gamma+1}{d+1}\big)^{\frac{d+1}{d}} + \gamma\right)^{-1}, \text{ for all } \gamma \in [0, d)\,.$$

For affine latency functions, in particular, $\mathrm{PoA}_\gamma(1) = 1$, for all $\gamma \geq 1$, and

$$\mathrm{PoA}_\gamma(1) \leq \big(1 - (1-\gamma)^2/4\big)^{-1}, \text{ for all } \gamma \in [0, 1)\,.$$

Furthermore, we can show that bounds on the PoA_γ of Theorem 1 and Corollary 1 are best possible in the worst-case.

Theorem 2. *For any class of latency functions $\mathcal{D}$ and for any $\epsilon > 0$, there is a γ-modifiable instance $\mathcal{G}$ on parallel links with homogeneous risk-averse players and latencies in class $\mathcal{D}$ so that $\mathrm{PoA}_\gamma(\mathcal{G}) \geq (1-\beta_\gamma(\mathcal{D}))^{-1} - \epsilon$.*

4 Parallel-Link Games with Heterogeneous Players

In contrast to the case of homogeneous players, computing the best γ-enforceable flow for heterogeneous risk-averse players is NP-hard, even for affine latencies.

Theorem 3. *Given an instance $\mathcal{G}$ on parallel links with affine latencies and two classes of risk-averse players, a $\gamma > 0$ and a target cost $C > 0$, it is* NP*-hard to determine whether there is a γ-enforceable flow with total latency at most C.*

Proof sketch. The proof is a modification of the construction in [16, Section 6], which shows that the best Stackelberg modification for parallel links with affine latencies is NP-hard. Intuitively, the players with low aversion factor a^1 (resp. high aversion factor a^2) correspond to selfish (resp. coordinated) players in [16].

Specifically, we reduce $(1/3, 2/3)$-Partition to the best γ-enforceable flow. An instance of $(1/3, 2/3)$-Partition consists of n positive integers $s_1, s_2, \ldots, s_n$, so that $S = \sum_{i=1}^{n} s_i$ is a multiple of 3. The goal is to determine whether there exists a subset X so that $\sum_{i \in X} s_i = 2S/3$.

Given an instance $\mathcal{I}$ of $(1/3, 2/3)$-Partition, we create a routing game $\mathcal{G}$ with $n+1$ parallel links and latencies $\ell_i(x) = (x/s_i) + 4$, $1 \leq i \leq n$, and $\ell_{n+1}(x) = 3x/S$. The traffic rate is $r = 2S$, partitioned into two classes with traffic $r^1 = 3S/2$ and $r^2 = S/2$. We set $\gamma = 2/17$. Working similarly to [16, Section 6], we show that if $a^1 = 0$ and $a^2 = 1$, $\mathcal{I}$ admits a $(1/3, 2/3)$-partition if and only if the routing game $\mathcal{G}$ admits a γ-enforceable flow f of total latency at most $35S/4$. We show that this holds if a^1 is small enough, e.g., if $a^1 = O(1/S^3)$. So, we can extend the NP-hardness proof to the case where $1 = a^1 < a^2$. $\square$

γ-Enforceable Flows with Good Price of Anarchy. Since the best enforceable flow is NP-hard, we next establish the existence of an enforceable flow that "mimics" the optimal flow o, as described by the properties (i) and (ii) in Lemma 1 and achieves a PoA as low as that in Theorem 1. In the following, we assume that the links are indexed in increasing order of $\ell_i(f)$, i.e. $i < j \Rightarrow \ell_i(f) \leq \ell_j(f)$, with ties broken in favor of links with $f_e > 0$. We start with a necessary and sufficient condition for a flow f to be γ-enforceable. [10, Algorithm 1] shows how to efficiently compute a γ-modification for any flow f that satisfies the following.

Theorem 4. *([10, Theorem 5]) Let $\mathcal{G}$ be a γ-modifiable instance on parallel links with heterogeneous players, let f be a feasible flow and let μ be the maximum index of a link used by f. Then, f is γ-enforceable if and only if (i) for any used link i, $\gamma \ell_i(f) \geq \sum_{l=i}^{\mu-1} \frac{\ell_{l+1}(f) - \ell_l(f)}{a_{l+1}^{\min}}$ and (ii) for all links i and j, if $\ell_i(f) < \ell_j(f)$, then $a_i^{\max}(f) \leq a_j^{\min}(f)$ (more risk-averse players use links of higher latency).*

To obtain a γ-enforceable flow f for an instance with heterogeneous players, we combine Lemma 1 with Theorem 4 and apply [10, Algorithm 1]. Specifically, we first ignore player heterogeneity and compute, using Lemma 1, a γ-enforceable flow f and the corresponding modification $\boldsymbol{\Gamma}$ so that f is a Nash flow of the modified game $\mathcal{G}^{\boldsymbol{\Gamma}}$ when all players have the minimum risk-aversion factor $a^1 =$

1. Assuming that the links are indexed in increasing order of their latencies in f, since f is γ-enforceable with risk-aversion factor $a^1 = 1$ for all players, Theorem 4 implies that for any used link i, $(1+\gamma)\ell_i(f) \geq \ell_\mu(f)$.

Next, we greedily allocate the heterogeneous risk-averse players to f, taking their risk-averse factors into account, so that each link i receives flow f_i and property (ii) in Theorem 4 is satisfied. Finally, we use [10, Algorithm 1] and compute a γ-modification that turns f into an equilibrium flow for the modified instance with heterogeneous players. This is possible because, by construction, f satisfies condition (i) of Theorem 4. Moreover, since f satisfies the properties of (i) and (ii) in Lemma 1, the PoA of f can be bounded as in Theorem 1 and (in Corollary 1, for polynomial and affine latencies). Hence, we obtain the following.

Theorem 5. *Let $\mathcal{G}$ be a γ-modifiable instance on parallel-links with heterogeneous risk-averse players. Given the optimal flow of $\mathcal{G}$, we can compute a feasible flow f and a γ-modification $\boldsymbol{\Gamma}$ of $\mathcal{G}$ in time $O(mT_{NE})$, where T_{NE} is the complexity of computing the Nash flow of any given γ-modification of $\mathcal{G}$ with homogeneous risk-averse players. Moreover, the* PoA_γ*, under γ-modifications, achieved by f is upper bounded as in Theorem 1 and Corollary 1.*

5 Modifying Routing Games in Series-Parallel Networks

In this section, we consider γ-modifiable instances on series-parallel networks with homogeneous players and generalize the results of Section 3. We start with a sufficient and necessary condition for the optimal flow o to be γ-enforceable. The following generalizes Proposition 1 and is a corollary of [2, Theorem 1].

Proposition 2. *Let $\mathcal{G}$ be a γ-modifiable instance on a series-parallel network and let o be the optimal flow of $\mathcal{G}$. Then, o is γ-enforceable if and only if for any pair of internally vertex-disjoint paths p and q with common endpoints (possibly different from s and t) and with $o_e > 0$ for all edges $e \in p$, $\ell_p(o) \leq (1+\gamma)\ell_q(o)$.*

We proceed to generalized Lemma 1 to series-parallel networks. The proof is based on an extension of the rerouting procedure in Lemma 1 combined with a continuity property of γ-enforceable flows in series-parallel networks.

Lemma 2. *Let $\mathcal{G} = (G, \ell, r)$ be a γ-modifiable instance with homogeneous risk-averse players on a series-parallel network G and let o be the optimal flow of $\mathcal{G}$. There is a feasible flow f and a γ-modification $\boldsymbol{\Gamma}$ of $\mathcal{G}$ such that*
(i) f is a Nash flow of the modified instance $G^{\boldsymbol{\Gamma}}$.
(ii) for any edge e, if $f_e < o_e$, then $\gamma_e = 0$, and if $f_e > o_e$, then $\gamma_e = \gamma$.

Proof sketch. The proof is by induction on the series-parallel structure of G. For the base case of a single edge e, the lemma holds without any modifications.

The induction step follows directly from the induction hypothesis if G is obtained as a series composition of two series-parallel networks. The interesting case is where G is the result of a parallel composition of series-parallel networks

G_1 and G_2. By induction hypothesis, for $i \in \{1,2\}$, we let f_i be a γ-enforceable flow of rate r_i, with $r_1 + r_2 = r$, and $\boldsymbol{\Gamma}_i$ be a γ-modification of $\mathcal{G}_i$ such that f_i is the Nash flow of $\mathcal{G}_i^{\boldsymbol{\Gamma}_i}$. In the following, we let $L_i = L(f_i)$ be the equilibrium cost of flow f_i through network G_i with latency functions modified according to $\boldsymbol{\Gamma}_i$.

If $L_1 = L_2$, the claim follows directly from the induction hypothesis. Otherwise, we assume wlog. that $L_1 > L_2$. In this case, we generalize the rerouting procedure of Lemma 1. Starting with f_1 and f_2, we reroute flow from used paths of $G_1^{\boldsymbol{\Gamma}_1}$ to $G_2^{\boldsymbol{\Gamma}_2}$, maintaining the equilibrium property on both $G_1^{\boldsymbol{\Gamma}_1}$ and $G_2^{\boldsymbol{\Gamma}_2}$ and trying to equalize their equilibrium cost. As in Lemma 1, we have also to maintain property (ii), by paying attention to edges e where f_e reaches o_e for the first time and to edges e' where $\gamma_{e'}$ reaches γ for the first time. For the former, we stop increasing the flow through any paths including e and start increasing γ_e, so that the equilibrium property is maintained. For the latter, we stop increasing $\gamma_{e'}$ and start increasing again the flow through paths that include e'.

The idea of the proof is similar to the induction step in Lemma 1. However, since G_1 and G_2 are general series-parallel networks connected in parallel, we need a continuity property, shown in [8, Section 3], about the changes in the equilibrium flow when the traffic rate slightly increases or decreases. □

Using the properties (i) and (ii), we show that the upper bound on the PoA in Theorem 1 extends to the γ-enforceable flow f of Lemma 2 and to the PoA_γ of the best γ-enforceable flow in series-parallel networks with homogeneous players.

Theorem 6. *For γ-modifiable instances on series-parallel networks with homogeneous players and latency functions in class $\mathcal{D}$,* $\mathrm{PoA}_\gamma(\mathcal{D}) \leq (1 - \beta_\gamma(\mathcal{D}))^{-1}$.

Given the optimal flow of an instance $\mathcal{G}$ on a series-parallel network, we show how to compute a γ-enforceable flow f and the corresponding modification so that we achieve a PoA at most $(1 - \beta_\gamma(\mathcal{D}))^{-1}$. Given o, the running time is determined by the time required to compute a Nash flow of the original instance.

We first determine whether the optimal flow o is γ-enforceable. To this end, we remove from G all edges unused by o and check the feasibility of the following:

$$\begin{array}{cl} 0 \leq \gamma_e \leq \gamma & \forall \text{ used edges } e \\ \sum_{e \in p}(1+\gamma_e)\ell_e(o) = \max_{p: o_p > 0} \ell_p(o) & \forall \text{ used paths } p \end{array} \qquad (O_\gamma)$$

If the linear system (O_γ) is not feasible, then o is not γ-enforceable, by Proposition 2. Otherwise, using the solution of (O_γ) as γ_e's for the edges of G used by o and setting $\gamma_e = 0$ for the unused edges e, we enforce o as a Nash flow of the modified game $\mathcal{G}^{\boldsymbol{\Gamma}}$.

If (O_γ) is not feasible and o is not γ-enforceable, we exploit the constructive nature of the proof of Lemma 2 and find a γ-enforceable flow in time dominated by the time required to compute a Nash flow in series-parallel networks.

Lemma 3. *Let $\mathcal{G}$ be a γ-modifiable instance on a series-parallel network with homogeneous players. Given the optimal flow of $\mathcal{G}$ and any $\epsilon > 0$, we can compute a feasible flow f and a γ-modification $\boldsymbol{\Gamma}$ of $\mathcal{G}$ with the properties (i) and (ii) of Lemma 2 in time $O(m^2 T_{NE} \log(r/\epsilon))$, where T_{NE} is the complexity of computing the Nash flow of any given γ-modification of $\mathcal{G}$ and ϵ is an accuracy parameter.*

6 Parallel-Link Games with Relaxed Restrictions

In this section, we consider (p, γ)-modifiable games on parallel links with heterogeneous risk-averse players. Observing that any $\gamma/\sqrt[p]{m}$-modification is a (p, γ)-modification for a (p, γ)-modifiable game, we next show an upper bound on the PoA under such modifications.

Theorem 7. *For any (p, γ)-modifiable instance $\mathcal{G}$ on m parallel links with heterogeneous risk-averse players and latency functions in class $\mathcal{D}$, we have that* $\mathrm{PoA}^p_\gamma(\mathcal{G}) \leq \mathrm{PoA}_{\gamma_0}(\mathcal{G}) \leq (1 - \beta_{\gamma_0}(\mathcal{D}))^{-1}$*, where* $\gamma_0 = \gamma/\sqrt[p]{m}$.

The above bound is tight under weak assumptions on the class $\mathcal{D}$ of latency functions. More specifically, we say that a class of latency functions $\mathcal{D}$ is of the form $\mathcal{D}_0$ if (a) ℓ is continuous and twice differentiable in $(0, +\infty)$, (b) $\ell'(x) > 0, \forall x \in (0, +\infty)$ or ℓ is constant, (c) ℓ is semi-convex, i.e. $x\ell(x)$ is convex in $[0, +\infty)$ and (d) if $\ell \in \mathcal{D}$, then $(\ell + c) \in \mathcal{D}$, for all constants $c \in \mathbb{R}$ such that for all $x \in \mathbb{R}_{\geq 0}$, $\ell(x) + c \geq 0$ [2]. Then we obtain the following.

Theorem 8. *For any class $\mathcal{D}$ of the form $\mathcal{D}_0$ and any $\epsilon > 0$, there is an instance $\mathcal{G}$ on m parallel links with homogeneous players and latency functions in class $\mathcal{D}$, so that* $\mathrm{PoA}^p_\gamma(\mathcal{G}) \geq (1 - \beta_{\gamma_0}(\mathcal{D}))^{-1} - \epsilon$*, where* $\gamma_0 = \gamma/\sqrt[p]{m}$.

Proof sketch. We consider an instance $\mathcal{I}_m$, with m parallel links, where the first $m - 1$ links have the same latency function $\ell \in \mathcal{D}$ (to be fixed later) and link m has constant latency $(1 + \gamma_1)\ell(\frac{r}{m-1})$, where $\gamma_1 = \gamma/\sqrt[p]{m-1}$. The instance has homogeneous risk-averse players with risk-aversion $a^1 = 1$. Also we let $\gamma_0 = \gamma/\sqrt[p]{m}$. The proof is an immediate consequence of the following three claims:

Claim 1. For every $m \geq 2$ and any latency function $\ell \in \mathcal{D}$ with $\ell(0) = 0$, $\mathrm{PoA}^p_\gamma(\mathcal{I}_m) = \mathrm{PoA}_{\gamma_1}(\mathcal{I}_m)$. I.e., Claim 1 states that the best (p, γ)-modification for the instance $\mathcal{I}_m$ is the modification that splits γ evenly among the first $m-1$ edges. The proof follows from an application of KKT optimality conditions.

Claim 2. For every $m \geq 2$ and any $\epsilon > 0$, there is a latency function $\ell_{\epsilon,m}$ with $\ell_{\epsilon,m}(0) = 0$ such that setting $\ell = \ell_{\epsilon,m}$ in the instance $\mathcal{I}_m$ results in $\mathrm{PoA}_{\gamma_1}(\mathcal{I}_m) \geq (1 - \beta_{\gamma_1}(\mathcal{D}))^{-1} - \epsilon/2$. The proof of Claim 2 is similar to the proof of Theorem 2.

Since $\ell_{\epsilon,m}(0) = 0$, we can combine claims 1 and 2 and obtain that for any $m \geq 2$ and any $\epsilon > 0$, $\mathrm{PoA}^p_\gamma(\mathcal{I}_m) \geq (1 - \beta_{\gamma_1}(\mathcal{D}))^{-1} - \epsilon/2$, if we use the latency function $\ell_{\epsilon,m}$.

Claim 3. For every class of latency functions $\mathcal{D}$, any $\epsilon > 0$ and any γ, there exists an $m_\epsilon \geq 2$ such that $(1 - \beta_{\gamma_1}(\mathcal{D}))^{-1} \geq (1 - \beta_{\gamma_0}(\mathcal{D}))^{-1} - \epsilon/2$.

The proof is based on the fact that γ_1 tends to γ_0 as the number of parallel links m grows. Therefore, for any $\epsilon > 0$, there are an m_ϵ and a latency function $\ell_{\epsilon,m_\epsilon}$ such that $\mathrm{PoA}^p_\gamma(\mathcal{I}_{m_\epsilon}) \geq (1 - \beta_{\gamma_0}(\mathcal{D}))^{-1} - \epsilon$. □

[2] Property (d) requires that $\mathcal{D}$ should be closed under addition of constants, as long as the resulting function remains nonnegative.

References

1. H. Angelidakis, D. Fotakis, and T. Lianeas. Stochastic congestion games with risk-averse players. In *Proc. of the 6th Symposium on Algorithmic Game Theory (SAGT '13)*, LNCS 8146, pp. 86-97, 2013.
2. V. Bonifaci, M. Salek, and G. Schäfer. Efficiency of restricted tolls in non-atomic network routing games. In *Proc. of the 4th Symposium on Algorithmic Game Theory (SAGT '10)*, LNCS 6982, pp. 302-313, 2011.
3. J.R. Correa, A.S. Schulz, and N.E. Stier Moses. Selfish Routing in Capacitated Networks. *Mathematics of Operations Research*, 29(4):961-976, 2004.
4. A. Fiat and C.H. Papadimitriou. When the players are not expectation maximizers. In *Proc. of the 3th Symposium on Algorithmic Game Theory (SAGT '10)*, LNCS 6386, pp. 1-14, 2010.
5. L. Fleischer. Linear Tolls Suffice: New Bounds and Algorithms for Tolls in Single Source Networks. *Theoretical Computer Science*, 348:217-225, 2005.
6. L. Fleischer, K. Jain, and M. Mahdian. Tolls for Heterogeneous Selfish Users in Multicommodity Networks and Generalized Congestion Games. In *Proc. of the 45th IEEE Symp. on Foundations of Computer Science (FOCS '04)*, pp. 277-285, 2004.
7. D. Fotakis and P. Spirakis. Cost-Balancing Tolls for Atomic Network Congestion Games. In *Proc. of the 3rd Workshop on Internet and Network Economics (WINE '07)*, pp. 179-190, 2007.
8. M.A. Hall. Properties of the Equilibrium State in Transportation Networks. *Transportation Science*, 12(3):208-216, 1978.
9. M. Hoefer, L. Olbrich, and A. Skopalik. Taxing subnetworks. In *Proc. of the 4th Workshop on Internet and Network Economics (WINE '08)*, LNCS 5385, pp. 286-294, 2008.
10. T. Jelinek, M. Klaas, and G. Schäfer. Computing optimal tolls with arc restrictions and heterogeneous players. In *Proc. of the 31st Symposium on Theoretical Aspects of Computer Science (STACS '14)*, LIPIcs 25, pp. 433-444, 2014.
11. G. Karakostas and S. Kolliopoulos. Edge Pricing of Multicommodity Networks for Heterogeneous Users. In *Proc. of the 45th IEEE Symp. on Foundations of Computer Science (FOCS '04)*, pp. 268-276, 2004.
12. E. Nikolova and N. Stier Moses. Stochastic selfish routing. In *Proc. of the 4th Symposium on Algorithmic Game Theory (SAGT '11)*, LNCS 6982, pp. 314-325, 2011.
13. E. Nikolova and N. Stier-Moses. The burden of risk aversion in mean-risk selfish routing. In *Proc. of the 16th ACM Conference on Electronic Commerce (EC '15)*, pp. 489-506, 2015.
14. F. Ordóñez and N. Stier Moses. Wardrop equilibria with risk-averse users. *Transportation Science*, 44(1):63-86, 2010.
15. G. Piliouras, E. Nikolova, and J.S. Shamma. Risk Sensitivity of Price of Anarchy under Uncertainty. In *Proc. of the 14th ACM Conference on Electronic Commerce (EC '13)*, pp. 715-732, 2013.
16. T. Roughgarden. Stackelberg Scheduling Strategies. *SIAM Journal on Computing*, 33(2):332-350, 2004.
17. Tim Roughgarden. *Selfish routing and the Price of Anarchy.* MIT press, 2005.
18. A. Tversky and D. Kahneman. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263-291, 1979.
19. J. Valdez, R.E. Tarjan, and E.L. Lawler. The Recognition of Series-Parallel Digraphs. *SIAM Journal on Computing*, 11(2):298-313, 1982.