Stackelberg Strategies for Atomic Congestion Games

Dimitris Fotakis

Dept. of Information and Communication Systems Engineering,
University of the Aegean, 83200 Samos, Greece. fotakis@aegean.gr

Abstract. We investigate the effectiveness of Stackelberg strategies for atomic congestion games with unsplittable demands. In our setting, only a fraction of the players are selfish, while the rest are willing to follow a predetermined strategy. A *Stackelberg strategy* assigns the coordinated players to appropriately selected strategies trying to minimize the performance degradation due to the selfish players. We consider two orthogonal cases, namely linear congestion games with arbitrary strategies and congestion games on parallel links with arbitrary non-negative and non-decreasing latency functions. We restrict our attention to pure Nash equilibria and derive strong upper and lower bounds on the Price of Anarchy under different Stackelberg strategies.

1 Introduction

Congestion games provide a natural model for non-cooperative resource allocation in large-scale communication networks and have been the subject of intensive research in algorithmic game theory. In a *congestion game* [15], a finite set of non-cooperative players, each controlling an unsplittable unit of load, compete over a finite set of resources. All players using a particular resource experience a cost (or latency) given by a non-negative and non-decreasing function of the resource's load (or congestion). Each player selects her strategy selfishly trying to minimize her *individual cost*, that is the sum of the costs for the resources in her strategy. A natural solution concept is that of a *pure Nash equilibrium*, a configuration where no player can decrease her cost by unilaterally changing her strategy.

At the other end, the network manager cares about the public benefit and aims to minimize the *total cost* incurred by all players. Since a Nash equilibrium does not need to optimize the total cost, one seeks to quantify the inefficiency due to selfish behaviour. The *Price of Anarchy* was introduced in [12] and has become a widely accepted measure of the performance degradation due to the players' selfish behaviour. The (pure) Price of Anarchy is the worst-case ratio of the total cost of a (pure) Nash equilibrium to the optimal total cost. Many recent contributions have provided strong upper and lower bounds on the Price of Anarchy (PoA) for several classes of congestion games, mostly linear congestion games and congestion games on parallel links (see e.g. [14, 7, 2, 3, 1]).

In many cases however, only a fraction of the players are selfish, while the rest are willing to follow a strategy suggested by the network manager (see e.g. [11, 17] for motivation and examples). Korilis *et al.* [11] introduced the notion of *Stackelberg routing* as a theoretical framework for this setting. In Stackelberg routing, a central authority coordinates a fixed fraction of the players and assigns them to appropriately selected strategies trying to minimize the performance degradation due to the selfish behaviour of the remaining players. A *Stackelberg strategy* is an algorithm that determines the strategies of the coordinated players. Given the strategies of the coordinated players, the selfish players lead the system to a configuration where they are at a Nash equilibrium. The goal is to find a Stackelberg strategy of minimum PoA, that is the worst-case ratio of the total cost of all (coordinated and selfish) players in such a configuration to the optimal total cost. Now the PoA is a non-increasing function of the fraction of coordinated players and ideally is given by a continuous curve decreasing from the PoA when all players are selfish to 1 if all players are coordinated (aka a *normal* curve [10]).

In this work, we investigate the effectiveness of Stackelberg routing for atomic congestion games. We consider two essentially orthogonal settings, namely linear congestion games with arbitrary strategies and congestion games on parallel links with arbitrary non-negative and non-decreasing latency functions. We restrict our attention to pure Nash equilibria and obtain strong upper and lower bounds on the pure PoA under different Stackelberg strategies. To the best of our knowledge, this is the first work on Stackelberg routing for atomic congestion games with unsplittable demands.

Related Work. Lücking *et al.* [14] were the first to consider the PoA of atomic congestion games for the objective of total cost[1]. For the special case of uniformly related parallel links, they proved that the PoA is $4/3$. For parallel links with polynomial latency functions of degree d, Gairing *et al.* [7] proved that the PoA is at most $d+1$. Awerbuch *et al.* [2] and Christodoulou and Koutsoupias [3] proved independently that the PoA of congestion games is $5/2$ for linear latencies and $d^{\Theta(d)}$ for polynomial latencies of degree d. Subsequently, Aland *et al.* [1] presented exact bounds on the PoA of congestion games with polynomial latencies. For non-atomic congestion games, where the number of players is infinite and each player controls an infinitesimal amount of load, Roughgarden [16] proved that the PoA is independent of the strategy space and equal to $\rho(\mathcal{D})$, where ρ depends only on the class of latency functions $\mathcal{D}$. Subsequently, Correa *et al.* [4] gave a simple proof of the same bound by introducing $\beta(\mathcal{D}) = 1 - \frac{1}{\rho(\mathcal{D})}$.

To the best of our knowledge, Stackelberg routing has been investigated only in the context of non-atomic games. Focusing on parallel links, Roughgarden [17] proved that it is NP-hard to compute an optimal Stackelberg configuration and investigated the performance of two natural strategies, Scale and Largest Latency First (LLF). Scale uses the optimal configuration scaled by the fraction of coordinated players, denoted α. LLF assigns the coordinated players to the largest cost strategies in the optimal configuration. Roughgarden proved that the PoA of LLF is $1/\alpha$ for arbitrary latencies and $4/(3+\alpha)$ for linear latencies. Kumar and Marathe [13] presented an approximation scheme for the best Stackelberg configuration on parallel links with polynomial latencies.

Some recent papers [19, 5, 10] extended the results of Roughgarden [17] in several directions. Swamy [19] and independently Correa and Stier-Moses [5] proved that the PoA of LLF is at most $1 + 1/\alpha$ for series-parallel networks with arbitrary latency functions. In addition, Swamy proved that the PoA of LLF is at most $\alpha + (1-\alpha)\rho(\mathcal{D})$ for parallel links with latency functions in class $\mathcal{D}$ and obtained upper bounds on the

[1] We cite only the most relevant results on the pure PoA of linear congestion games and congestion games on parallel links for the objective of total cost. For a survey on the PoA of congestion games for total and max cost, see e.g. [6, 8].

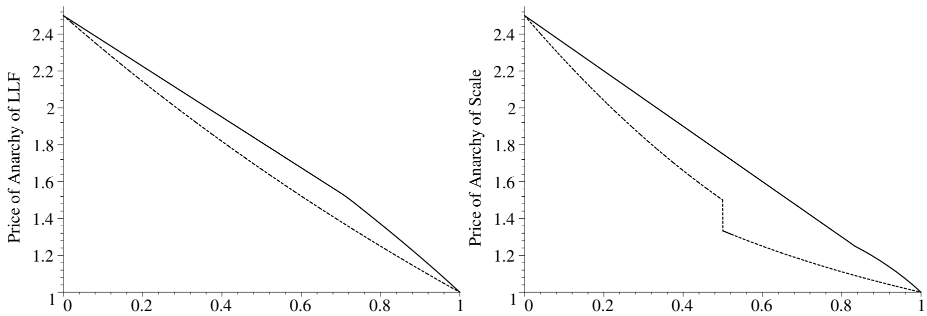

Fig. 1. The upper (solid curves) and lower (dotted curnes) bounds on the PoA of LLF and Scale for linear congestion games as functions of the fraction of coordinated players α. The bounds for LLF are on the left and the bounds for Scale are on the right. The lower bound on the right holds for any optimal-restricted strategy.

PoA of LLF and Scale for general networks. Karakostas and Kolliopoulos [10] considered non-atomic linear congestion games with arbitrary strategies and presented the best known upper and lower bounds on the PoA of LLF and Scale.

Other recent work on Stackelberg routing for non-atomic games includes [9, 18]. In particular, Kaporis and Spirakis [9] showed how to compute efficiently the smallest fraction of coordinated players required to induce an optimal configuration. For the related question of determining the smallest fraction of coordinated players required to improve the cost of a Nash equilibrium, Sharma and Williamson [18] derived a closed expression for parallel links with linear latencies.

Contribution. Motivated by the recent interest in bounding the PoA of LLF and Scale in the non-atomic setting, we investigate the effectiveness of Stackelberg routing in the context of atomic congestion games with unsplittable demands.

For linear congestion games, we derive strong upper and lower bounds on the PoA of LLF and Scale expressed as decreasing functions of the fraction of coordinated players, denoted α (see the plots in Fig. 1). For LLF, we obtain an upper bound of $\min\{(20-11\alpha)/8, (3-2\alpha+\sqrt{5-4\alpha})/2\}$ and a lower bound of $5(2-\alpha)/(4+\alpha)$, whose ratio is less than 1.1131. We use a randomized version of Scale, because scaling the optimal configuration by α may be infeasible in the atomic setting. We prove that the expected total cost of the worst configuration induced by Scale is at most $\max\{(5-3\alpha)/2, (5-4\alpha)/(3-2\alpha)\}$ times the optimal total cost. On the negative side, we present a lower bound that holds not only for Scale, but also for any randomized Stackelberg strategy that assigns the coordinated players to their optimal strategies.

An interesting case arises when the number of players is large and the number of coordinated players is considerably larger than the number of resources, even if α is small. To take advantage of this possibility, we introduce a simple Stackelberg strategy called *Cover*. Assuming that the ratio of the number of coordinated players to the number of resources is no less than a positive integer λ, Cover assigns to every resource either at least λ or as many coordinated players as the resource has in the optimal configuration. We prove that the PoA of Cover tends to the PoA of the corresponding non-atomic

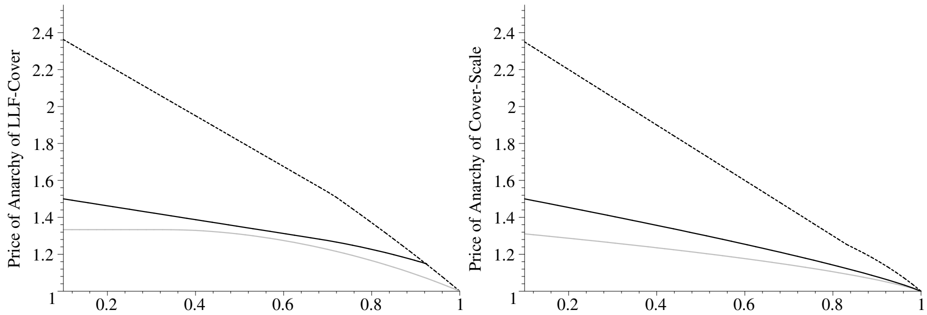

Fig. 2. The PoA of combined LLF-Cover and Cover-Scale vs. the PoA of LLF and Scale for atomic and non-atomic games. Let n be the total number of players, n_s the number of coordinated players, and m the number of resources. For these plots, we let $n = 10m$, assume that $n_s \geq m$, and use 1-Cover. On the x-axis, we have the fraction of coordinated players n_s/n (since $n_s \geq m$, the x-axis starts at 0.1). The solid black curves are the upper bounds on the PoA of atomic linear games under LLF-Cover and Cover-Scale. The dotted black curves are the upper bounds on the PoA of atomic games under LLF and Scale. The solid grey curves are the upper bounds on the PoA of non-atomic linear games under LLF ([10, Theorem 2]) and Scale ([10, Theorem 1]).

linear congestion game as λ grows[2]. More precisely, for linear latencies without constant term, the PoA of Cover is at most $1 + \frac{1}{2\lambda}$ for arbitrary strategies and at most $1 + \frac{1}{4(\lambda+1)^2-1}$ for parallel links. For arbitrary linear latencies, the PoA of Cover is at most $(4\lambda-1)/(3\lambda-1)$. Furthermore, if the ratio of the total number of players n to the number of resources m is large enough (e.g. $n/m \geq 10$), combining Cover with either LLF or Scale gives considerably stronger bounds on the PoA than when using LLF or Scale alone. These bounds are quite close to the best known bounds for non-atomic linear games [10] (see the plots in Fig. 2).

For parallel links, we prove that the PoA of LLF matches that for non-atomic games. In particular, we show that the PoA of LLF is at most $1/\alpha$ for arbitrary latencies and at most $\alpha + (1-\alpha)\rho(\mathcal{D})$ for latency functions in class $\mathcal{D}$.

2 Model, Definitions, and Notation

Congestion Games. A *congestion game* is a tuple $\Gamma(N, E, (\Sigma_i)_{i\in N}, (d_e)_{e\in E})$, where N denotes the set of players, E denotes the set of resources, $\Sigma_i \subseteq 2^E$ denotes the strategy space of each player $i \in N$, and $d_e : \mathbb{N} \mapsto \mathbb{N}$ is a non-negative and non-decreasing latency function associated with each resource $e \in E$. A congestion game is *symmetric* if all players share the same strategy space. A vector $\sigma = (\sigma_1, \dots, \sigma_n)$ consisting of a strategy $\sigma_i \in \Sigma_i$ for each player $i \in N$ is a *configuration*. For each resource e, let $\sigma_e \equiv |\{i \in N : e \in \sigma_i\}|$ denote the congestion (or load) induced on e by σ. The individual cost of each player i in the configuration σ is $c_i(\sigma) = \sum_{e\in\sigma_i} d_e(\sigma_e)$.

[2] If all players are selfish however, the PoA of a non-symmetric linear congestion game can be 2.5 even if the ratio of the number of players to the number of resources is arbitrarily large.

A configuration σ is a pure *Nash equilibrium* if no player can improve her individual cost by unilaterally changing her strategy. Formally, σ is a pure Nash equilibrium if for every player i and all strategies $s_i \in \Sigma_i, c_i(\sigma) \leq c_i(\sigma_{-i}, s_i)$. Rosenthal [15] proved that the pure Nash equilibria of congestion games correspond to the local optima of a natural potential function. Hence every congestion game admits a pure Nash equilibrium.

We mostly consider *linear* congestion games, where every resource e is associated with a linear latency function $d_e(x) = a_e x + b_e$, $a_e, b_e \geq 0$. In the special case of linear functions without constant term (i.e. if $b_e = 0$ for all $e \in E$), we say that the resources are *uniformly related*. We also pay special attention to congestion games on *parallel links*, where the game is symmetric and the common strategy space consists of m singleton strategies, one for each resource.

Social Cost. We evaluate configurations using the objective of *total cost*. The total cost $C(\sigma)$ of a configuration σ is the sum of players' costs in σ. Formally, $C(\sigma) = \sum_{i=1}^{n} c_i(\sigma) = \sum_{e \in E} \sigma_e d_e(\sigma_e)$. An optimal configuration, usually denoted o, minimizes the total cost $C(o)$ among all configurations in $\times_{i \in N} \Sigma_i$. Even though this work is not concerned with the complexity of computing an optimal configuration, we remark that an optimal configuration can be computed in polynomial time for symmetric network congestion games if $x d_e(x)$'s are convex.

Price of Anarchy. The pure *Price of Anarchy* (PoA) is the maximum ratio $C(\sigma)/C(o)$ over all pure Nash equilibria σ. In other words, the PoA is equal to $C(\sigma)/C(o)$, where σ is the pure Nash equilibrium of maximum total cost.

Stackelberg Strategies. We consider a scenario where a *leader* coordinates n_s players, and only $n - n_s$ players are selfish. The leader assigns the *coordinated* players to appropriately selected strategies trying to minimize the performance degradation due to the selfish behaviour of the remaining players. A *Stackelberg configuration* consists of the strategies to which the coordinated players are assigned. A *Stackelberg strategy* is an algorithm computing a Stackelberg configuration[3].

We restrict our attention to *optimal-restricted* Stackelberg strategies that assign the coordinated players to their strategies in the optimal configuration. Given an optimal configuration $o = (o_1, \ldots, o_n)$, an optimal-restricted Stackelberg strategy selects a (possibly random) set $L \subseteq N$, $|L| = n_s$, and assigns the coordinated players to the corresponding strategies in o. For non-symmetric games, L also determines the identities of the coordinated players[4]. The Stackelberg configuration corresponding to L, denoted $s(L)$, is $s(L) = (o_i)_{i \in L}$. By definition, for every (optimal-restricted) Stackelberg strategy and every L, $s_e(L) \leq o_e$ for all $e \in E$. In the following, we let $\alpha = n_s/n$ denote the fraction of players coordinated by the Stackelberg strategy, and let $k = n - n_s$ denote the number of selfish players.

Let $\Gamma(N, E, (\Sigma_i)_{i \in N}, (d_e)_{e \in E})$ be the original congestion game. The congestion game induced by L is $\tilde{\Gamma}_L(N \setminus L, E, (\Sigma_i)_{i \in N \setminus L}, (\tilde{d}_e)_{e \in E})$, where $\tilde{d}_e(x) = d_e(x + s_e(L))$

[3] We highlight the distinction between a Stackelberg strategy, that is an algorithm, and a strategy of some player i, that is an element of Σ_i.

[4] In the terminology of [10], we employ *strong* Stackelberg strategies for non-symmetric games. A strong Stackelberg strategy is free to choose the identities of the coordinated players in addition to their strategies.

for each $e \in E$. The players in $N \setminus L$ are selfish and reach a pure Nash equilibrium of $\tilde{\Gamma}_L$. Since there may be many pure Nash equilibria, we assume that the selfish players reach the *worst* pure Nash equilibrium of $\tilde{\Gamma}_L$, namely the equilibrium $\sigma(L)$ that maximizes the total cost of $\sigma(L) + s(L)$[5]. We let $\sigma(L)$ denote the worst pure Nash equilibrium of $\tilde{\Gamma}_L$ that maximizes $C(\sigma(L) + s(L)) = \sum_{e \in E}(\sigma_e(L) + s_e(L))d_e(\sigma_e(L) + s_e(L))$. We usually refer to $\sigma(L)$ as the worst Nash equilibrium induced by L (or by $s(L)$). In addition, we let $f(L) = \sigma(L) + s(L)$ denote the worst configuration induced by L (or by $s(L)$)[6]. With this notation in place, the PoA is equal to $C(f(L))/C(o)$.

Largest Latency First (LLF). LLF assigns the coordinated players to the largest cost strategies in o. More precisely, if we index the players in non-decreasing order of their cost in the optimal configuration o, i.e. $c_1(o) \leq \cdots \leq c_n(o)$, the set of coordinated players selected by LLF is $L = \{k + 1, \ldots, n\}$.

Scale. We use a randomized version of Scale that selects each different set $L \subseteq N$, $|L| = n_s$, with probability $1/\binom{n}{n_s}$ and adopts the configuration $s(L) = (o_i)_{i \in L}$. Every player $i \in N$ is selected in L with probability $\alpha = n_s/n$ and the expected number of coordinated players on every resource e is αo_e.

Cover. Cover can be applied only if the number of coordinated players is so large that every resource e with $o_e \geq 1$ can be "covered" by at least one coordinated player. Even though this may be achieved with less coordinated players than m, we assume that $n_s \geq m$ for simplicity, and let $\lambda = \lfloor n_s/m \rfloor \geq 1$. Cover selects a set $L \subseteq N$ such that $|L| \leq n_s$ and either $s_e(L) \geq \lambda$ or $s_e(L) = o_e$ for every $e \in E$. Hence $o_e \geq s_e(L) \geq \min\{\lambda, o_e\}$ for all $e \in E$. Given an optimal configuration o, such a set L can be computed efficiently by the greedy λ-covering algorithm. Despite its simplicity, there are instances where Cover outperforms Scale and LLF.

When L is clear from the context, we omit the dependence on L and use s instead of L. In the following, s denotes the Stackelberg configuration, and σ (resp. f) denotes the worst pure Nash equilibrium (resp. configuration) induced by s. We sometimes write α-LLF, (resp. α-Scale, λ-Cover) to denote LLF (resp. Scale, Cover) coordinating at least αn (resp. αn, λm) players.

3 Stackelberg Strategies for Linear Congestion Games

In this section, we establish upper and lower bounds on the PoA of linear congestion games under different Stackelberg strategies. The upper bounds are based on the following lemma.

Lemma 1. *Let o be an optimal configuration, let s be any optimal-restricted Stackelberg configuration, and let f be the worst configuration induced by s. For all $\nu \in (0, 1)$,*

[5] If $s \in \times_{i \in L}\Sigma_i$ is a configuration for $L \subseteq N$ and $\sigma \in \times_{i \in N \setminus L}\Sigma_i$ is a configuration for $N \setminus L$, $f = \sigma + s$ denotes a configuration for N with $f_i = s_i$ if $i \in L$, and $f_i = \sigma_i$ if $i \in N \setminus L$. The notation is motivated by the fact that $f_e = \sigma_e + s_e$ for all $e \in E$.

[6] $\tilde{\Gamma}_L$ may have many different worst pure Nash equilibria. Since we are interested in bounding the PoA, we can assume wlog that for every L, there is a unique worst pure Nash equilibrium $\sigma(L)$ and a unique worst configuration $f(L) = \sigma(L) + s(L)$ induced by L.

(a) $C(f) \leq \sum_{e \in E}[a_e f_e o_e + a_e(o_e - s_e) + b_e o_e]$.
(b) $(1-\nu)C(f) \leq \frac{1}{4\nu} \sum_{e \in E} a_e o_e^2 + \sum_{e \in E} a_e(o_e - s_e) + \sum_{e \in E} b_e o_e - \nu \sum_{e \in E} b_e f_e$.
(c) $(1-\nu)C(f) \leq \frac{1}{4\nu} C(o) + \sum_{e \in E} a_e(o_e - s_e) + (1 - \frac{1}{4\nu}) \sum_{e \in E} b_e(o_e - s_e)$.

Proof sketch. To obtain (a), we apply the approach of [3, Theorem 1] and [2, Theorem 3.2] for the selfish players and use the fact that s is optimal-restricted. Then (b) is obtained from (a) by applying the inequality $xy \leq \nu x^2 + \frac{1}{4\nu} y^2$, which holds for all $x, y \in \mathbb{R}$ and all $\nu \in (0,1)$, to the terms $a_e f_e o_e$. Finally, (c) is obtained from (b) using $s_e \leq f_e$ and the inequality $\nu + \frac{1}{4\nu} \geq 1$, which holds for all $\nu \in (0,1)$. $\square$

Largest Latency First. We first prove an upper and a lower bound on the PoA of LLF.

Theorem 1. *The PoA of α-LLF is at most* $\min\{\frac{20-11\alpha}{8}, \frac{3-2\alpha+\sqrt{5-4\alpha}}{2}\}$.

Proof. Starting from Lemma 1.a, observing that for all non-negative integers x, y, z with $z \leq y$, $xy + y - z \leq \frac{1}{3}x^2 + \frac{5}{3}y^2 - \frac{11}{12}yz$, and using that $s_e \leq f_e$ and $s_e \leq o_e$ for all $e \in E$, we obtain that

$$C(f) \leq \frac{1}{3}C(f) + \frac{5}{3}C(o) - \frac{11}{12}\sum_{e \in E} s_e(a_e o_e + b_e) \leq \frac{1}{3}C(f) + \frac{20-11\alpha}{12}C(o)$$

For the last inequality, we use that $\sum_{e \in E} s_e(a_e o_e + b_e) = \sum_{i \in L} c_i(o) \geq \alpha C(o)$, because LLF assigns the coordinated players to the largest cost strategies in o. Therefore, the PoA is at most $\frac{20-11\alpha}{8}$.

For the second bound, we start from Lemma 1.c and observe that for all $e \in E$, $o_e - s_e \leq o_e(o_e - s_e)$, because o_e and $o_e - s_e$ are non-negative integers. Thus we obtain that for all $\nu \in (0,1)$,

$$(1-\nu)C(f) \leq \frac{1}{4\nu}C(o) + \sum_{e \in E}(o_e - s_e)(a_e o_e + b_e) \leq (\frac{1}{4\nu} + 1 - \alpha)C(o)$$

For the last inequality, we use that $\sum_{e \in E}(o_e - s_e)(a_e o_e + b_e) = C(o) - \sum_{i \in L} c_i(o) \leq (1-\alpha)C(o)$. Therefore, we obtain an upper bound of $\min_{\nu \in (0,1)}(\frac{1}{4\nu} + 1 - \alpha)/(1-\nu)$. Using $\nu = \frac{\sqrt{5-4\alpha}-1}{4(1-\alpha)}$, we conclude that the PoA is at most $\frac{3-2\alpha+\sqrt{5-4\alpha}}{2}$. $\square$

The following theorem gives a lower bound on the PoA of LLF.

Theorem 2. *For every $\alpha \in [0,1)$ and $\varepsilon > 0$, there is a symmetric linear congestion game for which the PoA under α-LLF is at least* $\frac{5(2-\alpha)}{4+\alpha} - \varepsilon$.

Proof. For any fixed $\alpha \in [0,1)$, let n be a positive integer chosen sufficiently large. Wlog we assume that $n_s = \alpha n$ is an integer. Let $k = n - n_s$, and let $L = \{k+1, \ldots, n\}$. We construct a symmetric game with n_s coordinated and k selfish players.

The construction consists of two parts, one for the selfish players and one for the coordinated players. For the selfish players, we employ the instance of [3, Theorem 4]. For the coordinated players, we use n_s singleton parallel strategies. Formally, there are k^2 resources $g_{i,j}$, $i, j \in [k]$, $k^2(k-1)/2$ resources $h_{i,(j,q)}$, $i \in [k]$, $1 \leq j < q \leq k$, and n_s resources r_i, $i \in L$. The latency function of each g-resource is $d_g(x) = x$, the latency function of each h-resource is $d_h(x) = \frac{2}{k+2}x$, and the latency function of each r-resource is $d_r(x) = \frac{(5k-2)k}{2(k+2)}x$.

The g-resources are partitioned in k rows $G_i = \{g_{i,j} : j \in [k]\}$, $i \in [k]$, and in k columns $G^i = \{g_{j,i} : j \in [k]\}$, $i \in [k]$. The h-resources are partitioned in k rows $H_i = \{h_{i,(j,q)} : 1 \leq j < q \leq k\}$, $i \in [k]$. For h-resources, we have k sets of columns $H^i = \{h_{j,(i,q)} : j \in [k], i < q \leq k\} \cup \{h_{j,(q,i)} : j \in [k], 1 \leq q < i\}$, $i \in [k]$. Every H^i contains $k(k-1)$ resources and every resource $h_{i,(j,q)}$ is included in H^j and H^q. The r-resources are partitioned in n_s singleton sets $R_i = \{r_i\}$, $i \in L$. The common strategy space of all players is $\Sigma = \{G_i \cup H_i : i \in [k]\} \cup \{G^i \cup H^i : i \in [k]\} \cup \{R_i : i \in L\}$, i.e. a player can choose either a row strategy $G_i \cup H_i$, or a column strategy $G^i \cup H^i$, or a parallel strategy R_i.

In the optimal configuration o, every player $i \in [k]$ uses the row strategy $G_i \cup H_i$ and every player $i \in L$ uses the parallel strategy R_i. The cost of the optimal configuration is $C(o) = \frac{(5k-2)kn - k^2(k-4)}{2(k+2)}$. Assuming that n is so large that $k > 4$, the cost of the parallel strategies in o is greater than the cost of the row strategies. Hence the configuration of LLF is $s = (R_i)_{i \in L}$. In the worst Nash equilibrium σ induced by s, every selfish player $i \in [k]$ uses the column strategy $G^i \cup H^i$ and the total cost is $C(\sigma + s) = \frac{(5k-2)k(n+k)}{2(k+2)}$. Therefore, the PoA is at least $\frac{(5k-2)(n+k)}{(5k-2)n - k(k-4)}$. Using $k = (1-\alpha)n$, we obtain the lower bound of $5(2-\alpha)/(4+\alpha) - \varepsilon$, where $\varepsilon \leq 3/n$. □

Scale. We proceed to obtain an upper bound on the PoA of Scale. The configuration of Scale s, the worst Nash equilibrium σ induced by s, and the worst configuration f induced by s are random variables uniquely determined by Scale's random choice L. Also for every resource e, the congestion s_e (resp. σ_e, f_e) induced on e by s (resp. σ, f) is a non-negative integral random variable uniquely determined by L.

Let $\mathrm{I\!Pr}[L] = 1/\binom{n}{n_s}$ be the probability that each $L \subseteq N$, $|L| = n_s$, occurs as the choice of Scale. The expected number of coordinated players on every resource e is:

$$\mathrm{I\!E}[s_e] = \sum_{L \subseteq N, |L| = n_s} \mathrm{I\!Pr}[L] s_e(L) = \alpha o_e \,, \tag{1}$$

because every player $i \in N$ is selected in L with probability $\alpha = n_s/n$. The following theorem gives an upper bound on $\mathrm{I\!E}[C(f)] = \sum_{L \subseteq N, |L| = n_s} \mathrm{I\!Pr}[L] C(f(L))$, namely the expected cost of the worst configuration f induced by Scale.

Theorem 3. *Let o be an optimal configuration. The expected cost of the worst configuration f induced by α-Scale is $\mathrm{I\!E}[C(f)] \leq \max\{\frac{5-3\alpha}{2}, \frac{5-4\alpha}{3-2\alpha}\}\, C(o)$.*

Proof sketch. Applying Lemma 1.a for any fixed choice of Scale L, multiplying by $\mathrm{I\!Pr}[L]$, and using linearity of expectation and (1), we obtain that

$$\mathrm{I\!E}[C(f)] \leq \textstyle\sum_{e \in E} [a_e(\mathrm{I\!E}[f_e] + 1 - \alpha) o_e + b_e o_e] \,, \tag{2}$$

where $\mathrm{I\!E}[f_e] = \sum_{L \subseteq N, |L| = n_s} \mathrm{I\!Pr}[L] f_e(L)$ is the expected number of players on resource e in the worst configuration induced by Scale. To complete the proof, we apply the following proposition to the right-hand side of (2).

Proposition 1. *Let y be a non-negative integer, let $\alpha \in [0,1]$, and let X be a non-negative integral random variable. Then,*

$$(\mathrm{I\!E}[X] + 1 - \alpha) y \leq \begin{cases} \frac{1}{3}\, \mathrm{I\!E}[X^2] + (\frac{5}{3} - \alpha) y^2 & \text{for all } \alpha \in [0, \frac{5}{6}] \\ (\alpha - \frac{1}{2}) \mathrm{I\!E}[X^2] + (\frac{5}{2} - 2\alpha) y^2 & \text{for all } \alpha \in (\frac{5}{6}, 1] \end{cases}$$

We observe that $(5-3\alpha)/2$ is greater (resp. less) than or equal to $(5-4\alpha)/(3-2\alpha)$ for all $\alpha \in [0, 5/6]$ (resp. $\alpha \in [5/6, 1]$). First we show that $\mathbb{E}[C(f)] \leq \frac{5-3\alpha}{2} C(o)$, for all $\alpha \in [0, 5/6]$. Then we show that $\mathbb{E}[C(f)] \leq \frac{5-4\alpha}{3-2\alpha} C(o)$, for all $\alpha \in (5/6, 1]$. $\square$

A Lower Bound for Optimal-Restricted Strategies. The following theorems give a lower bound on the PoA of Scale and any (even randomized) *optimal-restricted* strategy.

Theorem 4. *For every $\alpha \in [0, 1)$ and $\varepsilon > 0$, there is a symmetric linear congestion game for which the PoA under any randomized optimal-restricted Stackelberg strategy coordinating a fraction α of the players is at least $\frac{2}{1+\alpha} - \varepsilon$.*

Proof sketch. For any fixed $\alpha \in [0, 1)$, we construct a symmetric game with $n_s = \alpha n$ coordinated players and $k = n - n_s$ selfish players. There are $n^2(n-1)/2$ resources $h_{i,(j,q)}$, $i \in [n]$, $1 \leq j < q \leq n$, each with latency function $d_h(x) = x$, and a resource r of constant latency $d_r(x) = (n-1)(\frac{3}{2}n - k) + k - 1$. The common strategy space for all players consists of n row strategies $\{r\} \cup H_i$ and n column strategies H^i, where H_i's and H^i's are defined as in the proof of Theorem 2.

The optimal configuration o assigns every player i to the corresponding row strategy $\{r\} \cup H_i$. The total cost is $C(o) = n(n-1)(2n-k) + O(n^2)$. By symmetry, we can assume that any optimal-restricted Stackelberg strategy selects $L = \{k+1, \ldots, n\}$. In the worst Nash equilibrium σ induced by $s = (\{r\} \cup H_i)_{i \in L}$, every selfish player $i \in [k]$ uses the column strategy H^i and the total cost is $C(\sigma + s) \geq 2n^2(n-1)$. $\square$

Theorem 5. *For every $\alpha \in [0, \frac{1}{2})$ and $\varepsilon > 0$, there is a symmetric linear congestion game for which the PoA under any randomized optimal-restricted Stackelberg strategy coordinating a fraction α of the players is at least $\frac{5-5\alpha+2\alpha^2}{2} - \varepsilon$.*

Proof sketch. Instead of the resource r in the proof of Theorem 4, we use a grid of g-resources as in the proof of Theorem 2, which yields the same strategy space as in [3, Theorem 4]. The latency function of each h-resource is $d_h(x) = x$ and the latency function of each g-resource is $d_g(x) = \gamma x$, where $\gamma = \frac{(n-1)(3n/2-k)+k-1}{2k-n-1}$. The rest of the proof is similar to the proof of Theorem 4. $\square$

Cover. The PoA of Cover tends to the PoA of non-atomic linear congestion games as the ratio of the number of coordinated players to the number of resources grows.

Theorem 6. *If $n_s \geq m$, the PoA of λ-Cover is at most $\frac{4\lambda-1}{3\lambda-1}$, where $\lambda = \lfloor n_s/m \rfloor$. For uniformly related resources, the PoA of λ-Cover is at most $1 + \frac{1}{2\lambda}$.*

Proof. Let s be the configuration of λ-Cover, and let f be the worst configuration induced by s. Since either $s_e = o_e$ or $s_e \geq \lambda$, it holds that $o_e - s_e \leq \frac{1}{4\lambda} o_e^2$ for every resource e. Applying this inequality to Lemma 1.c, we obtain that for all $\nu \in (0, 1)$,

$$(1-\nu)C(f) \leq \tfrac{1}{4\nu} C(o) + \max\{\tfrac{1}{4\lambda}, 1 - \tfrac{1}{4\nu}\} C(o) = \max\{\tfrac{1}{4\lambda} + \tfrac{1}{4\nu}, 1\} C(o) \quad (3)$$

Using $\nu = \frac{\lambda}{4\lambda-1}$, we conclude that the PoA of λ-Cover is at most $\frac{4\lambda-1}{3\lambda-1}$.

For uniformly related resources, (3) becomes $(1-\nu)C(f) \leq (\frac{1}{4\nu} + \frac{1}{4\lambda})C(o)$. Using $\nu = \sqrt{\lambda^2 + \lambda} - \lambda$, we conclude that the PoA is at most $\frac{2\lambda + 2\sqrt{\lambda^2+\lambda}+1}{4\lambda} \leq 1 + \frac{1}{2\lambda}$. $\square$

A careful analysis gives a stronger upper bound for congestion games on uniformly related parallel links. The proof of the following lemma is omitted due to lack of space.

Lemma 2. *If $n_s \geq m$, the PoA of λ-Cover for congestion games on uniformly related parallel links is at most $1 + \frac{1}{4(\lambda+1)^2-1}$, where $\lambda = \lfloor n_s/m \rfloor$.*

Combining Cover with LLF and Scale. If the ratio of the number of players n to the number of resources m is large enough (e.g. $n/m \geq 10$), combining Cover with either LLF or Scale gives considerably stronger upper bounds on the PoA than when using LLF or Scale alone. Throughout this section, we assume that $n_s \geq m$ and let λ be any positive integer not exceeding $\lfloor n_s/m \rfloor$. Also the definition of α is different and does not take the players coordinated by Cover into account.

Combining Cover with LLF. First LLF assigns $n_s - \lambda m$ coordinated players to the largest cost strategies in the optimal configuration o. Let $L^L \subseteq N$, $|L^L| = n_s - \lambda m$, be the set of players assigned by LLF. Then Cover assigns the remaining λm coordinated players wrt $(o_i)_{i \in N \setminus L^L}$. Let $L^C \subseteq N \setminus L^L$, $|L^C| \leq \lambda m$, be the set of players assigned by Cover such that $\min\{\lambda, o_e - s_e(L^L)\} \leq s_e(L^C) \leq o_e - s_e(L^L)$ for all $e \in E$. The joint Stackelberg configuration is $s = (o_i)_{i \in L^L \cup L^C}$.

Theorem 7. *If $n_s \geq m$, let λ be any positive integer not exceeding $\lfloor n_s/m \rfloor$, and let $\alpha = \frac{n_s - \lambda m}{n}$. The PoA of combined α-LLF and λ-Cover is at most*

$$\begin{cases} \frac{4\lambda-1}{3\lambda-1}(1 - \frac{\alpha}{4\lambda}) & \textit{for all } \alpha \in [0, \frac{4\lambda^2}{12\lambda^2-6\lambda+1}] \\ \frac{2-\alpha+\sqrt{4\alpha-3\alpha^2}}{2} & \textit{for all } \alpha \in [\frac{4\lambda^2}{12\lambda^2-6\lambda+1}, 1] \end{cases}$$

Proof sketch. Combining the proofs of Theorem 6 and Theorem 1, we obtain that the PoA is at most $\min_{\nu \in (0,1)} [\frac{1}{4\nu} + \max\{\frac{1}{4\lambda}, 1 - \frac{1}{4\nu}\}(1-\alpha)]/(1-\nu)$. The theorem follows by choosing ν appropriately. $\square$

Remark 1. In Theorem 7, $\alpha \leq 1 - \lambda\frac{m}{n}$ and the upper bound remains greater than 1 even if $n_s = n$. To obtain a normal curve, we replace combined LLF-Cover with LLF as soon as n_s/n is so large that the bound of Theorem 1 becomes stronger than the bound of Theorem 7 (see also the left plot of Fig. 2).

Combining Cover with Scale. First Cover assigns (at most) λm coordinated players. Let $L^C \subseteq N$, $|L^C| \leq \lambda m$, be the set of players assigned by Cover such that for all $e \in E$, $\min\{\lambda, o_e\} \leq s_e(L^C) \leq o_e$. Then Scale selects a random set $L^S \subseteq N \setminus L^C$, $|L^S| = n_s - \lambda m$, and assigns the remaining $n_s - \lambda m$ coordinated players to the corresponding optimal strategies. The joint Stackelberg configuration is $s(L^C \cup L^S) = (o_i)_{i \in L^C \cup L^S}$. Hence the joint Stackelberg configuration s and the worst configuration f induced by s are random variables uniquely determined by $L^C \cup L^S$.

For the analysis, we fix an arbitrary choice L^C of Cover. Let $\mathrm{I\!Pr}[L^S] = 1/\binom{n-\lambda m}{n_s - \lambda m}$ be the probability that each set $L^S \subseteq N \setminus L^C$, $|L^S| = n_s - \lambda m$, occurs as the choice of Scale. The following theorem establishes an upper bound on the expected cost $\mathrm{I\!E}[C(f)]$ of the worst configuration f induced by the combined Cover-Scale strategy s, where $\mathrm{I\!E}[C(f)]$ is given by

$$\mathrm{I\!E}[C(f)] = \sum_{L^S \subseteq N \setminus L^C, |L^S| = n_s - \lambda m} \mathrm{I\!Pr}[L^S] C(f(L^C \cup L^S))$$

Theorem 8. *If $n_s \geq m$, let λ be any positive integer not exceeding $\lfloor n_s/m \rfloor$, let $\alpha = \frac{n_s - \lambda m}{n - \lambda m}$, and let o be an optimal configuration. The expected cost of the worst configuration f induced by combining λ-Cover and α-Scale is bounded as follows:*

$$\mathbb{E}[C(f)] \leq \frac{2\lambda(3\alpha-2)+1-\alpha^2-(1-\alpha)\sqrt{\alpha^2-2\alpha(8\lambda^2-4\lambda+1)+16\lambda^2-8\lambda+1}}{2\lambda(4\alpha-3)+2(1-\alpha)}\, C(o)$$

Proof sketch. Combining the proofs of Theorem 3 and Theorem 6, we show that

$$\mathbb{E}[C(f)] \leq \min_{\nu\in(0,1)}[\max\{\tfrac{1}{4\nu} + \tfrac{1-\alpha}{4\lambda}, 1-\alpha\nu\}/(1-\nu)]\, C(o)$$

The theorem follows by choosing ν appropriately. $\square$

4 Largest Latency First for Congestion Games on Parallel Links

LLF becomes particularly simple when restricted to parallel links (see also [17, Section 3.2]). LLF indexes the links in non-decreasing order of their latencies in the optimal configuration o, i.e. $d_1(o_1) \leq \cdots \leq d_m(o_m)$, and finds the largest index q with $\sum_{\ell=q}^{m} o_\ell \geq n_s$. In the configuration of LLF, $s_\ell = 0$ for all $\ell < q$, $s_\ell = o_\ell$ for all $\ell > q$, and $s_q = n_s - \sum_{\ell=q+1}^{m}$. Hence q is the first link to which some coordinated players are assigned and $\Lambda = d_q(o_q)$ is a lower bound on the cost of the coordinated players in o.

Theorem 9. *The PoA of LLF for atomic congestion games on parallel links is at most $1/\alpha$, where α is the fraction of players coordinated by LLF.*

Proof sketch. The proof is similar to that of [19, Theorem 3.4]. Let o be the optimal configuration, let s be the configuration of LLF, let σ be the worst Nash equilibrium induced by s, and let $f = \sigma + s$ be the worst configuration induced by s.

Every coordinated player has cost at least Λ in o and $C(o) \geq n_s\Lambda$. The crucial observation is that for every link ℓ with $\sigma_\ell > 0$, $d_\ell(f_\ell) \leq \Lambda$. Hence the total cost of selfish players in f is at most $(n - n_s)\Lambda$, and for every link ℓ with $s_\ell > 0$, $d_\ell(f_\ell) \leq d_\ell(o_\ell)$. The latter holds because either $\sigma_\ell = 0$ and $f_\ell = s_\ell \leq o_\ell$, or both $\sigma_\ell > 0$ and $s_\ell > 0$, in which case $d_\ell(f_\ell) \leq \Lambda \leq d_\ell(o_\ell)$. Therefore, the total cost of coordinated players in f is at most $C(o)$ and $C(f) \leq C(o) + (n - n_s)\Lambda \leq (n/n_s)C(o)$. $\square$

Let $\mathcal{D}$ be a non-empty class of non-negative and non-decreasing latency functions. In [16, 4], it is shown that the PoA of non-atomic congestion games with latency functions in class $\mathcal{D}$ is at most $\rho(\mathcal{D}) = \sup_{d\in\mathcal{D}} \sup_{x\geq y\geq 0} \frac{x\,d(x)}{y\,d(y)+(x-y)d(x)}$. The following theorem establishes the same upper bound on the PoA of atomic congestion games on parallel links. The proof is omitted due to lack of space.

Theorem 10. *The PoA of atomic congestion games on parallel links with latency functions in class $\mathcal{D}$ is at most $\rho(\mathcal{D})$.*

The next theorem follows easily from Theorem 10 and Theorem 9. The proof is similar to that of [19, Theorem 3.5].

Theorem 11. *For atomic congestion games on parallel links, the PoA of LLF is at most $\alpha + (1-\alpha)\rho(\mathcal{D})$, where α is the fraction of players coordinated by LLF and $\mathcal{D}$ is a class of latency functions containing $\{d_\ell(x + s_\ell)\}_{\ell\in[m]}$.*

Remark 2. For linear latencies, Theorem 11 gives an upper bound of $(4 - \alpha)/3$ on the PoA of LLF for atomic congestion games on parallel links, which is quite close to the tight bound of $4/(3 + \alpha)$ for the corresponding class of non-atomic games [17].

References

1. S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact Price of Anarchy for Polynomial Congestion Games. *STACS '06*, LNCS 3884, pp. 218–229, 2006.
2. B. Awerbuch, Y. Azar, and A. Epstein. The Price of Routing Unsplittable Flow. *STOC '05*, pp. 57–66, 2005.
3. G. Christodoulou and E. Koutsoupias. The Price of Anarchy of Finite Congestion Games. *STOC '05*, pp. 67–73, 2005.
4. 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.
5. J.R. Correa and N.E. Stier-Moses. Stackelberg Routing in Atomic Network Games. *Technical Report DRO-2007-03*, Columbia Bussiness School, 2007.
6. A. Czumaj. Selfish Routing on the Internet. In *Handbook of Scheduling: Algorithms, Models, and Performance Analysis*, J. Leung (ed.), CRC Press, 2004.
7. M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. Nash Equilibria in Discrete Routing Games with Convex Latency Functions. *ICALP '04*, LNCS 3142, pp. 645–657, 2004.
8. M. Gairing, T. Lücking, B. Monien, and K. Tiemann. Nash Equilibria, the Price of Anarchy and the Fully Mixed Nash Equilibrium Conjecture. *ICALP '05*, LNCS 3580, pp. 51–65, 2005.
9. A.C. Kaporis and P.G. Spirakis. The Price of Optimum in Stackelberg Games on Arbitrary Single Commodity Networks and Latency Functions. *SPAA '06*, pp. 19–28, 2006.
10. G. Karakostas and S. Kolliopoulos. Stackelberg Strategies for Selfish Routing in General Multicommodity Networks. *Technical Report AdvOL 2006/08*, McMaster University, 2006.
11. Y.A. Korilis, A.A. Lazar, and A. Orda. Achieving Network Optima Using Stackelberg Routing Strategies. *IEEE/ACM Transactions on Networking* 5(1):161–173, 1997.
12. E. Koutsoupias and C. Papadimitriou. Worst-case Equilibria. *STACS '99*, LNCS 1563, pp. 404–413, 1999.
13. V.S.A. Kumar and M.V. Marathe. Improved Results for Stackelberg Strategies. *ICALP '02*, LNCS 2380, pp. 776–787, 2002.
14. T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A New Model for Selfish Routing. *STACS '04*, LNCS 2996, pp. 547–558, 2004.
15. R.W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory* 2:65–67, 1973.
16. T. Roughgarden. The Price of Anarchy is Independent of the Network Topology. *Journal of Computer and System Sciences* 67(2):341–364, 2003.
17. T. Roughgarden. Stackelberg Scheduling Strategies. *SIAM Journal on Computing* 33(2):332–350, 2004.
18. Y. Sharma and D.P. Williamson. Stackelberg Thresholds in Network Routing Games. *EC '07*, to appear, 2007.
19. C. Swamy. The Effectiveness of Stackelberg Strategies and Tolls for Network Congestion Games. *SODA '07*, 2007.