Weighted Congestion Games: The Price of Anarchy, Universal Worst-Case Examples, and Tightness

KSHIPRA BHAWALKAR, Stanford University, Stanford, CA, USA
MARTIN GAIRING, University of Liverpool, Liverpool, U.K.
TIM ROUGHGARDEN, Stanford University, Stanford, CA, USA.

We characterize the price of anarchy (POA) in weighted congestion games, as a function of the allowable resource cost functions. Our results provide as thorough an understanding of this quantity as is already known for nonatomic and unweighted congestion games, and take the form of universal (cost function-independent) worst-case examples. One noteworthy byproduct of our proofs is the fact that weighted congestion games are "tight," which implies that the worst-case price of anarchy with respect to pure Nash equilibria, mixed Nash equilibria, correlated equilibria, and coarse correlated equilibria are always equal (under mild conditions on the allowable cost functions). Another is the fact that, like nonatomic but unlike atomic (unweighted) congestion games, weighted congestion games with trivial structure already realize the worst-case POA, at least for polynomial cost functions.

We also prove a new result about unweighted congestion games: the worst-case price of anarchy in symmetric games is as large as in their more general asymmetric counterparts.

Categories and Subject Descriptors: F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Price of Anarchy, weighted congestion games, Nash equilibria

ACM Reference Format:
Kshipra Bhawalkar, Martin Gairing, and Tim Roughgarden, 2014. Weighted Congestion Games: The Price of Anarchy, Universal Worst-Case Examples and Tightness. *ACM Trans. Econ. Comp.* X, X, Article X (June 2014), 23 pages.
DOI: http://dx.doi.org/10.1145/0000000.0000000

1. INTRODUCTION

The class of *congestion games* is expressive enough to capture a number of otherwise unrelated applications — including routing, network design, and the migration of species (see references in Roughgarden [2006]) — yet structured enough to permit a useful theory. Such a game has a ground set of resources, and each strategy of a player is to select a subset of them (e.g., a path in a network). Each resource has a univariate cost function that depends on the load induced by the players that use it, and each

---

A previous version of this work appeared in the proceedings of the European Symposium on Algorithms(ESA), 2010.
The first author was supported by a Stanford Graduate Fellowship and NSF grant CCF-1016885. The second author was supported by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD). The third author was supported in part by NSF grants CCF-0448664 and CCF-1016885, an ONR Young Investigator Award, an ONR PECASE Award, an AFOSR MURI grant, and an Alfred P. Sloan Fellowship.

DOI: http://dx.doi.org/10.1145/0000000.0000000

player aspires to minimize the sum of the resources' costs in its chosen strategy (given the strategies chosen by the other players).

Congestion games have played a starring role in recent research on quantifying the inefficiency of game-theoretic equilibria. They are rich enough to encode the Prisoner's Dilemma, and more generally can have Nash equilibria in which the sum of the players' costs is arbitrarily larger than that in a minimum-cost outcome. Thus the research goal is to understand how the parameters of a congestion game govern the inefficiency of its equilibria, and in particular to establish useful sufficient conditions that guarantee near-optimal equilibria.

A simple observation is that the inefficiency of equilibria in a congestion game depends fundamentally on the "degree of nonlinearity" of the cost functions. Because of this, we identify a "thorough understanding" of the inefficiency of equilibria in congestion games with a simultaneous solution to every possible special case of cost functions. In more detail, an ideal theory would include the following ingredients.

(1) For every set $\mathcal{C}$ of allowable resource cost functions, a relatively simple recipe for computing the largest-possible *price of anarchy (POA)* — the ratio between the sum of players' costs in an equilibrium and in a minimum-cost outcome — in congestion games with cost functions in $\mathcal{C}$.
(2) For analytically simple classes $\mathcal{C}$ like bounded-degree polynomials, an exact formula for the worst-case POA in congestion games with cost functions in $\mathcal{C}$.
(3) An understanding of the "game complexity" required for the worst-case POA to be realized. Ideally, such a result should refer only to the strategy sets and be independent of the allowable cost functions $\mathcal{C}$.
(4) An understanding of the equilibrium concepts — roughly equivalently, the rationality assumptions needed — to which the POA guarantees apply.

The earliest example of such a theory is for *nonatomic* congestion games, where there is a continuum of players, each of which is infinitesimally small. (The load on a resource is defined as the mass or measure of players that choose a strategy that includes it.) For every class $\mathcal{C}$ that satisfies mild conditions, the worst-case POA of nonatomic congestion games with cost functions in $\mathcal{C}$ is already realized in the simplest of examples — symmetric two-resource games (i.e., two-node two-link networks) — where one resource has a constant cost function and the other resource has, intuitively, the "steepest" cost function in $\mathcal{C}$ [Roughgarden 2003; Roughgarden and Tardos 2004]. This result can be interpreted as the third ingredient above — a cost function-independent characterization of the (essentially trivial) necessary and sufficient conditions on the strategy sets needed to realize the worst-case POA over *all* congestion games with cost functions from the permitted class. This characterization immediately implies, in principle, a recipe for computing the worst-case POA with respect to a set $\mathcal{C}$ of allowable cost functions — just try each cost function from $\mathcal{C}$ in a trivial two-resource example, and remember the worst example found. It also easily leads to precise numerical bounds on the worst-case POA for simple classes of cost functions, such as a bound of $\approx d/\ln d$ when $\mathcal{C}$ is the set of polynomials with degree at most d and nonnegative coefficients [Correa et al. 2004; Roughgarden 2003]. Blum et al. [2010] later supplied the final component of this research agenda by showing that all of these POA bounds apply under relatively weak behavioral assumptions — specifically, that all users achieve vanishing average regret over repeated plays of the game. More recently, Aland et al. [2011], Awerbuch et al. [2005], Christodoulou and Koutsoupias [2005b], and Roughgarden [2009] provided the four ingredients above for *atomic* congestion games, where there is a finite number of players and the load on a resource is defined as the number of players that use it.

This paper is about the fundamental model of *weighted* congestion games [Milchtaich 1996; Rosenthal 1973], where each player i has a weight w_i and the load on a resource is defined as the sum of the weights of the players that use it. Such weights model non-uniform resource consumption among the players and can, of course, be relevant for many reasons: for modeling different amounts of traffic (e.g., by Internet Service Providers from different "tiers"); for different durations of resource usage [Shapley 1953]; and even for collusion among several identical users, who can be thought of as a single "virtual" player with weight equal to the number of colluding players [Fotakis et al. 2008; Hayrapetyan et al. 2006].

Our main contribution is a thorough understanding of the worst-case POA in weighted congestion games, in the form of the four ingredients listed above. Interesting byproducts of our proofs include the fact that weighted congestion games are *tight* in the sense of Roughgarden [2009], and thus the worst-case POA with respect to pure Nash equilibria, mixed Nash equilibria, correlated equilibria, and coarse correlated equilibria are always equal (under mild conditions on $\mathcal{C}$); and the fact that, like nonatomic but unlike atomic (unweighted) congestion games, weighted congestion games with trivial structure already realize the worst-case POA, at least for polynomial cost functions.

1.1. Overview of Results

Result 1: Exact POA of general weighted congestion games with general cost functions. We provide the first characterization of the exact POA of general weighted congestion games with general cost functions. For a given set of cost functions $\mathcal{C}$, the properties of the functions in this class determine certain "feasible values" of two parameters λ, μ, which lead to upper bounds of the form $\lambda/(1-\mu)$. The best upper bound that can be obtained using this two-parameter approach is denoted by $\zeta(\mathcal{C})$. The hard work then lies in proving that there always exists a weighted congestion game that realizes this upper bound. The abstract approach is to make use of the inequalities used in the upper bound proof — in the spirit of complementary slackness arguments in linear programming — with the cost functions and loads on the resources that make these inequalities tight employed in the worst-case example. Ultimately, we can exhibit examples with POA arbitrarily close to our upper bound of $\zeta(\mathcal{C})$.

This approach is similar to the one taken by Roughgarden [2009] for unweighted congestion games, although non-uniform player weights create additional technical issues that necessitate completely different constructions.

A side effect of our upper bound proof is that $\zeta(\mathcal{C})$ is actually the "Robust POA" defined by Roughgarden [2009]. Every bound on the robust POA of a game automatically has numerous consequences [Roughgarden 2009], and in particular upper bounds on the POA of mixed Nash equilibria, correlated equilibria, and coarse correlated equilibria. (See Section 2 for definitions.) The bound of $\zeta(\mathcal{C})$ is valid under much weaker rationality assumptions on players than are needed to justify convergence to pure or mixed Nash equilibria. Since we establish a matching lower bound for the pure POA — which obviously applies also to the three more general equilibrium concepts — our bound of $\zeta(\mathcal{C})$ is the exact POA with respect to all of these equilibrium concepts.

Exact robust POA bounds were previously established for nonatomic and atomic unweighted congestion games [Roughgarden 2009], which are potential games and hence possess pure Nash equilibria. Weighted congestion games do not have these nice properties. They are not potential games and a pure Nash equilibrium does not always exist (see Harks et al. [2011] and Harks and Klimm [2012] for characterizations), yet our results show that the robust POA is exact.

Table I. Values of exact POA bounds for polynomial cost functions. $\alpha(d)$ is the non-atomic POA [Roughgarden and Tardos 2004], $\gamma(d)$ is the unweighted atomic POA [Christodoulou and Koutsoupias 2005b; Aland et al. 2011], and $\zeta(d)$ is the weighted atomic POA [Awerbuch et al. 2005; Aland et al. 2011].

d	$\alpha(d)$	$\gamma(d)$	$\zeta(d)$
1	4/3	2.5	2.618
2	1.626	9.583	9.909
3	1.896	41.54	47.82
4	2.151	267.6	277.0
d	$O(d/\log d)$	$O((d/\log d)^d)$	$O((d/\log d)^d)$

Result 2: Exact POA of weighted congestion games on parallel links with polynomial cost functions. We prove that, for polynomial cost functions with nonnegative coefficients and maximum degree d, the worst-case POA is realized on a network of parallel links (for each d). We note that even for affine cost functions, only partial results were previously known about the worst-case POA of weighted congestion games in networks of parallel links. Thus our work implies, for the first time for weighted congestion games, that the worst-case POA is essentially independent of the allowable network topologies in the sense of Roughgarden [2003], at least for polynomial cost functions. This result stands in contrast to unweighted congestion games, where the worst-case POA in networks of parallel links (as well as some slight generalizations) is the same as that in nonatomic congestion games, denoted $\alpha(\mathcal{C})$, which in turn is provably smaller than the worst-case POA in general (atomic) congestion games [Anshelevich et al. 2008; Fotakis 2010; Holzman and Law-Yone 2003] [1]. See Table I for some concrete bounds for polynomial cost functions.

Result 3: POA of symmetric unweighted congestion games is as large as asymmetric ones. Our final result contributes to understanding how the worst-case POA of *unweighted* congestion games depends on the game complexity. We show that the POA of symmetric unweighted congestion games with general cost functions is the same value $\gamma(\mathcal{C})$ as that for asymmetric unweighted congestion games. A result of this form was previously known only for affine cost functions [Christodoulou and Koutsoupias 2005b]. With this result, we can conclude that the known gap between the worst-case POA $\alpha(\mathcal{C})$ of networks of parallel links and the worst-case POA $\gamma(\mathcal{C})$ of general unweighted atomic congestion games is located inside the class of symmetric games, rather than between symmetric and asymmetric games.

1.2. Further Related Work

Koutsoupias and Papadimitriou [1999] initiated the study of the POA of mixed Nash equilibria in weighted congestion games on parallel links, but with a different objective function: the expected maximum of the players' costs. Lücking et al. [2008] then developed results for the pure POA of unweighted congestion games on parallel links (with the standard objective function) They showed that the exact POA of unweighted congestion games on parallel links with linear cost functions is 4/3, and when all the links have the same linear cost function it drops to 9/8. The first results for general networks were obtained independently by Christodoulou and Koutsoupias [2005b] and Awerbuch et al. [2005]. Christodoulou and Koutsoupias [2005b,a] established an exact POA bound of 5/2 for unweighted congestion games with affine cost functions. They obtained the same bound for the POA of pure Nash equilibria, mixed Nash equilibria,

[1]The price of anarchy bound for unweighted congestion games on parallel links follows from results in Anshelevich et al. [2008] and Holzman and Law-Yone [2003]. Anshelevich et al. [2008] establish a Price of Stability bound for unweighted congestion games which is the same as the nonatomic POA bound. Holzman and Law-Yone [2003] show that for a slight generalization of parallel link networks the pure Nash equilibrium is unique. The bound was also independently established by Fotakis [2010].

and correlated equilibria [Christodoulou and Koutsoupias 2005a]. They also provided an asymptotic POA upper bound of $d^{\Theta(d)}$ for cost functions that are polynomials with nonnegative coefficients and degree at most d.

In their concurrent paper, Awerbuch et al. [2005] provided the exact POA of weighted congestion games with affine cost function (namely, $1+\phi \approx 2.618$ where ϕ is golden ratio). They also provided an upper bound of $d^{\Theta(d)}$ for the POA of weighted congestion games with cost functions that are polynomials with nonnegative coefficients and degree at most d. Later Aland et al. [2011] obtained exact POA bounds for both weighted and unweighted congestion games with cost functions that are polynomials with nonnegative coefficients.

Caragiannis et al. [2011] analyzed asymmetric singleton congestion games with affine cost functions. They gave lower bounds which established a worst-case POA of $5/2$ for unweighted congestion games and a POA of $1+\phi$ for the weighted case. Gairing and Schoppmann [2007] provided a detailed analysis of the POA of singleton congestion games. They generalized the results in Caragiannis et al. [2011] to polynomial cost functions, showing that the worst-case POA in asymmetric singleton games is as large as in general games. For symmetric singleton congestion games (i.e., networks of parallel links) and polynomial cost functions of maximum degree d, they show a lower bound of the $(d+1)$th Bell number B_{d+1}. This last result is subsumed by our second contribution.

Finally using the results presented in the conference version of this paper as a building block, Roughgarden [2012] established an upper bound on the Bayes-Nash POA in an incomplete information congestion game where the players' weights or their sets of permitted strategies are drawn from a product distribution. The lower bound construction presented in this paper proves that the Bayes-Nash POA upper bound established there is tight.

2. PRELIMINARIES

Congestion Games. A *general weighted congestion game* Γ is composed of a set of N players $\mathcal{N}$, a set of resources E and a set $\mathcal{C}$ of non-negative, non-decreasing cost functions from $\mathbb{R}^+ \to \mathbb{R}^+$. For each player i in $\mathcal{N}$ a weight w_i and a strategy set $S_i \subseteq 2^E$ are specified. The congestion on a resource is the total weight of all players using that resource and the associated congestion cost is specified by a function $c_e \in \mathcal{C}$ of the congestion on that edge.

An outcome is a choice of strategies $\mathsf{s} = (s_1, s_2, \ldots, s_N)$ by players with $s_i \in S_i$. Then congestion on a resource e in the outcome s is given by $x_e = \sum_{i \in \mathcal{N}: e \in s_i} w_i$. The cost of a player i is $C_i(\mathsf{s}) = w_i \sum_{e \in s_i} c_e(x_e)$. The *social cost* of an outcome is the sum of player's costs i.e. $C(\mathsf{s}) = \sum_{i=1}^{N} C_i(\mathsf{s})$. The total cost can also be written as $C(\mathsf{s}) = \sum_{e \in E} x_e c_e(x_e)$.

A special type of congestion games is *unweighted congestion games* where all players have unit weight. A congestion game is *symmetric* when all players have the same set of strategies, so $S_i = S \subseteq 2^E$ for all i. A game is called *singleton* if every strategy for every player has a single resource in it. A *symmetric singleton congestion game* is where all players have access to the same set of singleton strategies. This is commonly known as a *congestion game on parallel links*. A *network congestion game* is a congestion game in which the resources correspond to edges in an underlying network and strategies for a player i are given by paths from a vertex s_i to another vertex t_i. Figure 1 illustrates the relationships between these different classes.

Pure Nash Equilibrium. A *pure Nash equilibrium* is an outcome where each player chooses a single strategy to play and no player has an incentive to deviate from its current strategy. In the context of weighted congestion games an outcome s is a pure Nash equilibrium if for each player i and $s_i^* \in S_i$, an alternative strategy for

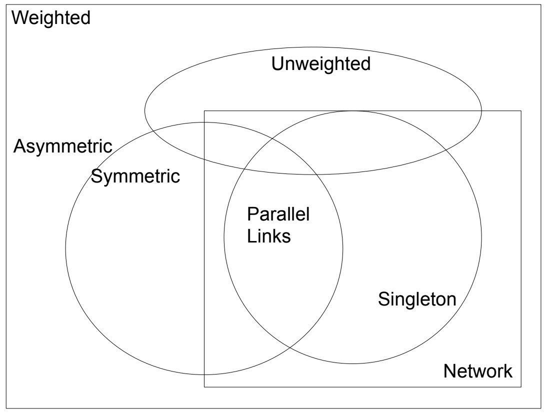

Fig. 1. Classes of congestion games based on structure

player i, the following holds:

$$C_i(\mathsf{s}) \leq C_i(\mathsf{s}_{-i}, s_i^*). \quad (1)$$

Here (s_{-i}, s_i^*) denotes the outcome that results when player i changes its strategy in s from s_i to s_i^*

A pure Nash equilibrium need not have the minimum-possible social cost. The *POA* [Koutsoupias and Papadimitriou 1999] captures how much worse Nash equilibria are compared to the cost of the best social outcome. For a congestion game Γ, if s denotes a Nash equilibrium with the worst social cost and s^* an outcome with the best social cost, then the POA is defined as $C(s)/C(s^*)$. The POA of a class of games is the worst POA among all games in the class.

Other Equilibrium Concepts. As pure Nash equilibria do not always exist in weighted congestion games, more general equilibrium concepts have been introduced in the literature to remedy such situations, and also to weaken the rationality assumptions required to justify convergence to equilibrium. We next review the concepts of mixed Nash, correlated, and coarse correlated equilibria.

A set $(\sigma_1, \cdots, \sigma_N)$ of independent probability distributions over players' strategy sets is a *mixed Nash equilibrium* if

$$E_{s\sim\sigma}[C_i(\mathsf{s})] \leq E_{s_{-i}\sim\sigma_{-i}}[C_i(s_{-i}, s_i')]$$

holds for every i and $s_i' \in S_i$. Here σ_{-i} is a product distribution of all σ_j's other than σ_i and s_{-i} denotes a strategy drawn from this distribution. A *correlated equilibrium* is a joint probability distribution σ over the outcomes of the game satisfying

$$E_{s\sim\sigma}[C_i(\mathsf{s})|s_i] \leq E_{s\sim\sigma}[C_i(s_{-i}, s_i')|s_i]$$

for every i and $s_i, s_i' \in S_i$. A *coarse correlated equilibrium* is given by a joint distribution σ over the outcomes of the game satisfying

$$E_{s\sim\sigma}[C_i(\mathsf{s})] \leq E_{s\sim\sigma}[C_i(s_{-i}, s_i')]$$

for all i and $s_i' \in S_i$.

For each of these equilibrium concepts the corresponding POAs (mixed POA, correlated POA, coarse correlated POA) are defined as the ratio of the cost of the worst equilibrium outcome to the optimal social outcome. All pure Nash equilibria can be represented as mixed Nash equilibria, all mixed Nash equilibria are correlated equilibria, and all correlated equilibria are coarse correlated equilibria. Thus the corresponding POAs can only be nondecreasing.

Roadmap. The rest of the paper is organized as follows. Section 3 describes the upper bound for weighted congestion games. The lower bound constructions for asymmetric weighted congestion games and weighted congestion games on parallel links are included in Section 4. Section 5 describes the lower bound construction for symmetric unweighted congestion games. We conclude in Section 6.

3. POA OF WEIGHTED CONGESTION GAMES

In this section we provide an upper bound on the POA of weighted congestion games with general cost functions.

For a class of functions $\mathcal{C}$ the upper bound is parameterized by two parameters λ and μ. Consider

$$\mathcal{A}(\mathcal{C}) = \{(\lambda, \mu) : \mu < 1; x^* c(x + x^*) \leq \lambda x^* c(x^*) + \mu x c(x)\}. \tag{2}$$

Here the constraints range over all functions $c \in \mathcal{C}$ and reals $x \geq 0$ and $x^* > 0$. Each pair (λ, μ) in $\mathcal{A}(\mathcal{C})$ yields an upper bound on the pure POA of weighted congestion games. We establish the following,

PROPOSITION 3.1 (POA UPPER BOUND). *For a class of functions $\mathcal{C}$, if $(\lambda, \mu) \in \mathcal{A}(\mathcal{C})$ then every weighted congestion game with cost functions in $\mathcal{C}$ has pure POA at most $\lambda/(1-\mu)$.*

PROOF. For a congestion game, let s denote a Nash equilibrium outcome and s^* denote an outcome that minimizes the social cost. Let x_e and x_e^* denote the loads on edge e in outcomes s and s^* respectively. The Nash condition (1) implies that

$$\forall i, \;\; C_i(\mathsf{s}) \leq C_i(\mathsf{s}_{-i}, \mathsf{s}_i^*).$$

Summing over all players we get

$$C(\mathsf{s}) = \sum_{i=1}^{N} C_i(\mathsf{s}) \leq \sum_{i=1}^{N} C_i(\mathsf{s}_{-i}, s_i^*). \tag{3}$$

Since λ, μ satisfy equation (2),

$$\begin{aligned}
\sum_{i=1}^{k} C_i(\mathsf{s}_{-i}, s_i^*) &\leq \sum_{e \in E} \sum_{i: e \in s_i^*} w_i c_e(x_e + w_i) \\
&\leq \sum_{e \in E} x_e^* c_e(x_e + x_e^*) \\
&\leq \sum_{e \in E} \lambda x_e^* c_e(x_e^*) + \mu x_e c_e(x_e) \\
&= \lambda C(\mathsf{s}^*) + \mu C(\mathsf{s}).
\end{aligned} \tag{4}$$

Here the second inequality follows from the fact that for every edge $\sum_{i: e \in s_i^*} w_i = x_e^*$ and hence $w_i \leq x_e^*$. Combining equations (3) and (4) we get

$$C(\mathsf{s}) \leq \lambda C(\mathsf{s}^*) + \mu C(\mathsf{s}). \tag{5}$$

Rearranging,

$$\mathsf{POA} = \frac{C(\mathsf{s})}{C(\mathsf{s}^*)} \leq \frac{\lambda}{1-\mu}.$$

□

Every upper bound proved using Proposition 3.1 is a “smoothness argument” in the sense of Roughgarden [2009], and thus automatically applies to (among other things) all of the equilibrium concepts defined in Section 2.

We denote the best upper bound implied by Proposition 3.1 by $\zeta(\mathcal{C})$.

Definition 3.2. For a class of functions $\mathcal{C}$, define

$$\zeta(\mathcal{C}) := \inf\left\{\frac{\lambda}{1-\mu} : (\lambda,\mu) \in \mathcal{A}(\mathcal{C})\right\}. \tag{6}$$

Define $\zeta(\mathcal{C}) := +\infty$ if $\mathcal{A}(\mathcal{C})$ is empty.

4. LOWER BOUNDS FOR WEIGHTED CONGESTION GAMES

In this section we describe two different lower bounds that match the upper bound $\zeta(\mathcal{C})$ given in Definition 3.2. The first lower bound applies to every class $\mathcal{C}$ of allowable cost functions satisfying a mild technical condition, and makes use of asymmetric congestion games. The second lower bound applies only to polynomial cost functions with nonnegative coefficients, but uses only networks of parallel links.

For each lower bound, we are given a class of cost functions $\mathcal{C}$, and we will describe a series of games with POA approaching $\zeta(\mathcal{C})$. For each game we specify player weights, player strategies, and congestion cost functions on the resources. Additionally, we describe two outcomes s and s^*. Justifying the example as a lower bound requires checking that s is a (pure) Nash equilibrium and that the ratio of costs of the outcomes s and s^* is close to $\zeta(\mathcal{C})$.

Our lower bound constructions are guided by the aspiration to simultaneously satisfy all of the inequalities in the proof of Proposition 3.1 exactly. This goal translates to the following conditions.

(a) In the outcome s, each player is indifferent between its strategy s_i and the deviation s_i^*.
(b) For each player i, the strategies s_i and s_i^* are disjoint.
(c) Cost functions and congestion on resources in the outcomes s and s^* correspond to tuples (c, x, x^*) that correspond to binding constraints in the infimum in (6).
(d) In the outcome s^*, each resource is used by a single player.

We believe that satisfying all of these conditions simultaneously is impossible (and can prove it for congestion games on parallel links). Nevertheless, we are able to “mostly” satisfy these conditions, which permits an asymptotic lower bound of $\zeta(\mathcal{C})$ as the number of players and resources tend to infinity.

4.1. Weighted Congestion Games with General Cost Functions

Now we present the lower bound examples that obtain POA arbitrarily close to $\zeta(\mathcal{C})$ for most classes of cost functions $\mathcal{C}$. We assume that the class $\mathcal{C}$ is *closed under scaling and dilation*, meaning that if $c(x) \in \mathcal{C}$ and $r \in \mathbb{R}^+$, then $rc(x)$ and $c(rx)$ are also in $\mathcal{C}$. Standard scaling and replication tricks (see [Roughgarden 2003]) imply that the first assumption is without loss of generality. The second assumption is not without loss but is satisfied by most natural classes of cost functions.

Because our assumptions on $\mathcal{C}$ are so weak, we have to proceed at a quite abstract level. The first step is to study deeply the optimization problem implicit in Definition 3.2, where constraints on the feasible values of λ, μ prohibit the value of $\zeta(\mathcal{C})$ from getting arbitrarily low. To generically construct lower bound examples, we then examine the set of constraints to find ones that are binding, and use them to get close to the values of (λ, μ) that yield $\zeta(\mathcal{C})$.

First, we observe that scaling and dilation does not change the set of constraints in the definition of the set $\mathcal{A}(\mathcal{C})$. The set of cost functions $\mathcal{C}$ can then be seen as composed of a number of disjoint equivalence classes (where the relation is differing by scaling and dilation). Henceforth whenever we speak of a cost function, we think of it as a representative of its equivalence class.

The next lemma identifies one or two of the defining constraints of $\mathcal{A}(\mathcal{C})$ that captures the quantity $\zeta(\mathcal{C})$ to arbitrary precision.

LEMMA 4.1 (CHARACTERIZATION OF BINDING HALF-PLANES). *If $\zeta(\mathcal{C}) > 1$, then for every ϵ such that $0 < \epsilon < \zeta(\mathcal{C}) - 1$, one of the following two holds true.*

(1) *There exists a cost function c and values $x \geq 0, x^* > 0$ such that $x^* c(x + x^*) \geq xc(x)$ and*

$$\frac{xc(x)}{x^* c(x^*)} \geq \zeta(\mathcal{C}) - \epsilon.$$

(2) *There exist cost functions c_1, c_2 and values $x_1, x_2 \geq 0$ and $x_1^*, x_2^* > 0$ such that, if (λ, μ) satisfy*

$$x_1^* c_1(x_1 + x_1^*) = \lambda x_1^* c_1(x_1^*) + \mu x_1 c_1(x_1)$$
$$x_2^* c_2(x_2 + x_2^*) = \lambda x_2^* c_2(x_2^*) + \mu x_2 c_2(x_2),$$

then $\lambda/(1-\mu) > \zeta(\mathcal{C}) - \epsilon$. Moreover, $x_1^* c_1(x_1 + x_1^*) \leq x_1 c_1(x_1)$ and $x_2^* c_2(x_2 + x_2^*) \geq x_2 c_2(x_2)$.

PROOF. For a cost function $c \in \mathcal{C}$, $x \geq 0$, and $x^* > 0$, let $\mathcal{H}_{c,x,x^*}$ denote the half-plane

$$x^* \cdot c(x + x^*) \leq \lambda \cdot x^* \cdot c(x^*) + \mu \cdot x \cdot c(x)$$

and $\partial\mathcal{H}_{c,x,x^*}$ the boundary of this half-plane. Recall from (2) that these are the half-planes that define the set $\mathcal{A}(\mathcal{C})$ of feasible pairs (λ, μ) for the set $\mathcal{C}$ of cost functions. Also, define

$$\beta_{c,x,x^*} = \frac{x \cdot c(x)}{x^* \cdot c(x + x^*)} \text{ and } \zeta_{c,x,x^*} = \frac{x \cdot c(x)}{x^* \cdot c(x^*)}.$$

Fix a positive $\epsilon < \zeta(\mathcal{C}) - 1$ and let $\zeta' = \zeta(\mathcal{C}) - \epsilon/2$. If $\zeta(\mathcal{C})$ is not finite, set $\zeta' = 1/\epsilon$. We write $\mathcal{L}_{\zeta'}$ for the line $\lambda + \zeta' \cdot \mu = \zeta'$ in the λ, μ plane.

If we think of a boundary line $\partial\mathcal{H}_{c,x,x^*}$ as specifying μ as a function of λ, then this line has slope $-1/\zeta_{c,x,x^*}$ and μ-intercept $1/\beta_{c,x,x^*}$. The half-space $\mathcal{H}_{c,x,x^*}$ consists of everything "northeast" of its boundary.

Consider the half-planes with $\beta_{c,x,x^*} \leq 1$. In the lucky event that there is such a half-plane with $\zeta_{c,x,x^*} \geq \zeta'$, we are done: this choice of c, x, x^* satisfies the conditions of the first case of the lemma. For the rest of the proof, we assume that $\zeta_{c,x,x^*} < \zeta'$ for every half-plane with $\beta_{c,x,x^*} \leq 1$.

We consider two cases. To define them, pick an arbitrary cost function c_1 with $c_1(1) > 0$ — since $\mathcal{C}$ is closed under dilation, such a function exists — and a sufficiently large value of x_1 so that $\zeta_{c_1,x_1,1} > \zeta'$. Our standing assumption implies that $\beta_{c_1,x_1,1} > 1$. Define $(\hat{\lambda}, \hat{\mu})$ as the unique point of intersection of $\partial H_{c_1,x_1,1}$ and $\mathcal{L}_{\zeta'}$. Since the former line has a larger slope ($-1/\zeta_{c_1,x_1,1}$ vs.$-1/\zeta'$) and a smaller μ-intercept ($1/\beta_{c_1,x_1,1}$ vs. 1) than the latter, $\hat{\lambda} > 0$ and hence $\hat{\mu} < 1$.

For the first case, we assume that there exists a half-plane $\mathcal{H}_{c_2,x_2,x_2^*}$ with $\beta_{c_2,x_2,x_2^*} < 1$ whose boundary intersects the line $\mathcal{L}_{\zeta'}$ at a point (λ_2, μ_2) with $\mu_2 < \hat{\mu}$. Equivalently, the line $\partial\mathcal{H}_{c_2,x_2,x_2^*}$ intersects $\mathcal{L}_{\zeta'}$ to the right of where $\partial\mathcal{H}_{c_1,x_1,1}$ intersects $\mathcal{L}_{\zeta'}$. Since the μ-intercepts of $\partial\mathcal{H}_{c_2,x_2,x_2^*}$ and $\partial\mathcal{H}_{c_1,x_1,1}$ (namely, $1/\beta_{c_2,x_2,x_2^*} > 1$ and $1/\beta_{c_1,x_1,1} < 1$) are on either side of that of $L_{\zeta'}$ (namely, 1) and $\hat{\lambda} > 0$, this implies that the intersection (λ, μ) of $\partial\mathcal{H}_{c_1,x_1,1}$ and $\partial\mathcal{H}_{c_2,x_2,x_2^*}$ lies on the "northeast side" of $L_{\zeta'}$. It follows that $\lambda + \zeta'\mu \geq \zeta'$. Thus, $c_1, c_2, x_1, x_2, 1, x_2^*, \lambda, \mu$ satisfy the conditions in the second case of the lemma.

Finally, assume that all half-planes $\mathcal{H}_{c,x,x^*}$ with $\beta_{c,x,x^*} < 1$ have boundaries that intersect the line $\mathcal{L}_{\zeta'}$ at points (λ, μ) with $\mu \geq \hat{\mu}$. Let μ^* denote the infimum of all μ-coordinates of such intersections. Under our standing assumption, every such boundary $\partial\mathcal{H}_{c,x,x^*}$ has a smaller slope ($-1/\zeta_{c,x,x^*}$ vs. $-1/\zeta'$) and a larger μ-intercept ($1/\beta_{c_1,x_1,1}$ vs. 1) than $\mathcal{L}_{\zeta'}$, and hence intersects $\mathcal{L}_{\zeta'}$ at a point (λ, μ) with $1 > \mu \geq \hat{\mu}$. Thus, $1 > \mu^* \geq \hat{\mu}$.

We now find appropriate (c_1, x_1, x_1^*) and (c_2, x_2, x_2^*) with $\beta_{c_1,x_1,x_1^*} \geq 1$ and $\beta_{c_2,x_2,x_2^*} < 1$, such that the corresponding half-plane boundaries intersect $\mathcal{L}_{\zeta'}$ at points (λ_1, μ_1) and (λ_2, μ_2) with μ_1, μ_2 very close to μ^*. Let $\delta = \frac{\epsilon\cdot(1-\mu^*)}{4\cdot\zeta'-\epsilon} > 0$. Consider the point $(\zeta' \cdot (1 - \mu^* + \delta), \mu^* - \delta)$ of $\mathcal{L}_{\zeta'}$. This point is feasible for all constraints corresponding to (c, x, x^*) with $\beta_{c,x,x^*} < 1$. Since $\zeta' < \zeta(\mathcal{C})$, this point cannot belong to the feasible set $\mathcal{A}(\mathcal{C})$ and hence there exists (c_1, x_1, x_1^*) with $\beta_{c_1,x_1,x_1^*} \geq 1$ such that the point $(\zeta' \cdot (1 - \mu^* + \delta), \mu^* - \delta)$ violates the corresponding constraint. Note that the point $(0, 1)$ of $\mathcal{L}_{\zeta'}$ lies in $\mathcal{H}_{c_1,x_1,x_1^*}$. This implies that $\partial\mathcal{H}_{c_1,x_1,x_1^*}$ intersects $\mathcal{L}_{\zeta'}$ at a point (λ_1, μ_1) with $\mu_1 \geq \mu^* - \delta$. Moreover, $\lambda_1 + \zeta' \cdot \mu_1 = \zeta'$.

If $\mu_1 > \mu^*$, then we can find (c_2, x_2, x_2^*) with $\beta_{c_2,x_2,x_2^*} < 1$ that intersects $\mathcal{L}_{\zeta'}$ at (λ_2, μ_2) with $\mu^* \leq \mu_2 \leq \mu_1$. Then, similarly to the previous case, $\partial\mathcal{H}_{c_1,x_1,x_1^*}$ and $\partial\mathcal{H}_{c_2,x_2,x_2^*}$ intersect at a point (λ, μ) such that $\lambda/(1-\mu) \geq \zeta'$, completing the proof.

We can now assume that $\mu^* - \delta \leq \mu_1 \leq \mu^*$. By the definition of μ^*, there exist (c_2, x_2, x_2^*) such that $\partial\mathcal{H}_{c_2,x_2,x_2^*}$ intersects $\mathcal{L}_{\zeta'}$ at (λ_2, μ_2), with $\mu^* \leq \mu_2 \leq \mu^* + \delta$. Note that $\mu_2 \geq \mu_1$ and $\lambda_2 + \zeta' \cdot \mu_2 = \zeta'$.

Let (λ, μ) be the point where $\partial\mathcal{H}_{c_1,x_1,x_1^*}$ and $\partial\mathcal{H}_{c_2,x_2,x_2^*}$ intersect. Both these boundaries have negative slopes, which means (λ, μ) lies in the triangle formed by the points (λ_1, μ_1), (λ_2, μ_2), and (λ_2, μ_1). Then $\lambda/(1-\mu) \geq \lambda_2/(1-\mu_1)$. Since $\lambda_1 - \lambda_2 = \zeta' \cdot (\mu_2 - \mu_1) \leq 2 \cdot \zeta' \cdot \delta$, we have

$$\begin{aligned} \frac{\lambda_2}{1-\mu_1} &= \frac{\lambda_1}{1-\mu_1} - \frac{\lambda_1 - \lambda_2}{1-\mu_1} \\ &\geq \zeta' - \frac{2\cdot\zeta'\cdot\delta}{1-\mu^*+\delta} \\ &\geq \zeta' - \frac{\epsilon}{2}. \end{aligned}$$

This proves that the conditions of the second case in the statement of the lemma hold. □

As long as $\zeta(\mathcal{C}) > 1$, the above Lemma identifies one or two binding constraints that closely approximate the upper bound. We use these in our lower bound construction.

Fix an $\epsilon > 0$. First consider the case when there are two binding constraints. The lemma guarantees triples $(\bar{c}_1, x_1, x_1^*)$, $(\bar{c}_2, x_2, x_2^*)$ such that the corresponding half-planes intersect at (λ, μ) with $\lambda/(1-\mu) = \zeta_\epsilon > \zeta(\mathcal{C}) - \epsilon$.

Let $z_1 = x_1/x_1^*$ and $z_2 = x_2/x_2^*$. For any $w > 0$, we can identify functions c_1, c_2 that are dilated versions of $\bar{c}_1, \bar{c}_2$ such that (λ, μ) satisfy:

$$\begin{aligned} c_1(w\cdot(z_1+1)) &= \lambda\cdot c_1(w) + \mu\cdot z_1\cdot c_1(w\cdot z_1) \\ c_2(w\cdot(z_2+1)) &= \lambda\cdot c_2(w) + \mu\cdot z_2\cdot c_2(w\cdot z_2). \end{aligned} \quad (7)$$

Finally, since $c_1(w \cdot (z_1 + 1)) \le z_1 c_1(z_1 w)$ and $c_2(w \cdot (z_2 + 1)) \ge z_2 c_2(z_2 w)$, for every w, there exists a constant $\eta \in [0, 1]$ such that

$$\begin{aligned} &\eta \cdot c_1(w \cdot (z_1 + 1)) + (1 - \eta) \cdot c_2(w \cdot (z_2 + 1)) \\ &= \eta \cdot c_1(w \cdot z_1) \cdot z_1 + (1 - \eta) \cdot c_2(w \cdot z_2) \cdot z_2 \end{aligned} \tag{8}$$

The following example serves as a lower bound.

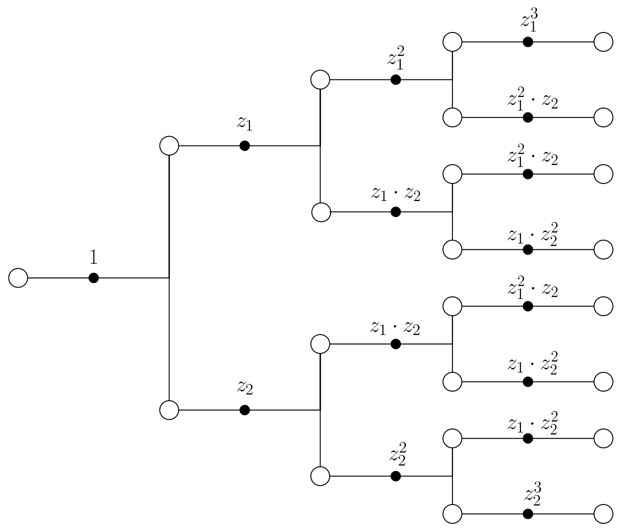

Fig. 2. Structure of the asymmetric lower bound for $k = 4$. Hollow nodes represent resources and solid nodes represent players.

LOWER BOUND 1 (ASYMMETRIC WEIGHTED CONGESTION GAMES). *For a parameter $k \in \mathbb{N}$ (chosen later) we construct a weighted congestion game with player set $\mathcal{N}$ and resource set E as follows (see Figure 2).*

Player strategies:. Organize the resources in a tree of depth k, which is a complete binary tree of depth $k - 1$ with each leaf extended by a path of length 1. For each non-leaf node i in the tree there is a player i with 2 strategies: either choose node i or all children of i.
Player weights:. If i is the root then $w_i = 1$. Otherwise, if node i is the left (right) child of some node j, then $w_i = w_j \cdot z_1$ ($w_i = w_j \cdot z_2$), where z_1, z_2 are chosen for (λ, μ) as noted above. Let $\mathcal{N}_L \subset \mathcal{N}$ be the set of players connected to a leaf.
Cost functions: . The cost functions of the resources are defined recursively as follows:
- *For the root we can choose any cost function $c \in \mathcal{C}$ with $c(1) = 1$. By a scaling argument such a cost function exists.*
- *Every leaf resource gets the same cost function as its parent.*
- *Consider an arbitrary resource e which is not a leaf nor its children are leaves. Let c_e be its cost function and let w_e be the weight of the corresponding player e. Let l, r be the left and right child of e respectively. By construction the corresponding players have weights $w_l = z_1 \cdot w_e$ and $w_r = z_2 \cdot w_e$. Among all pairs of cost functions c_1, c_2 that satisfy (7) for $w = w_e$, choose those that also satisfy*

$$c_1(w_e \cdot z_1) \cdot z_1 = c_2(w_e \cdot z_2) \cdot z_2 = c_e(w_e). \tag{9}$$

By a scaling argument such a pair always exists. Let η_e be the corresponding value for η in (8) *and define*

$$c_l = \eta_e \cdot c_1 \quad \textit{and } c_r = (1 - \eta_e) \cdot c_2. \tag{10}$$

Nash strategy:. The Nash outcome in this example is the outcome s *where each player chooses the resource closer to the root.*
Optimal strategy:. the optimal outcome s^* *is the outcome where each player chooses its strategy further from the root.*

We now prove that, in the construction above, the outcome s is indeed a Nash equilibrium and that the ratio $C(s)/C(s^*)$ is at least $\lambda/(1-\mu)$. We check those conditions in the proof of the following lemma.

LEMMA 4.2. *The congestion game in Lower Bound 1 has POA $\lambda/(1-\mu)$.*

PROOF. We claim that s is a pure Nash equilibrium. Observe that by construction no player in $\mathcal{N}_L$ can improve by choosing its leaf strategy since the leaf resource has the same cost function as the resource in its current strategy. Now, fix an arbitrary player $e \in \mathcal{N} \setminus \mathcal{N}_L$ and let l, r be her left and right child, respectively. Then,

$$\begin{aligned} C_e(\mathsf{s}) = c_e(w_e) &\overset{(9)}{=} \eta_e \cdot c_1(w_e \cdot z_1) \cdot z_1 + (1-\eta_e) \cdot c_2(w_e \cdot z_2) \cdot z_2 \\ &\overset{(8)}{=} \eta_e \cdot c_1(w_e \cdot (z_1 + 1)) + (1-\eta_e) \cdot c_2(w_e \cdot (z_2+1)) \\ &\overset{(10)}{=} c_l(w_e \cdot (z_1+1)) + c_r(w_e \cdot (z_2+1)) \\ &= C_e(\mathsf{s}_{-e}, s_e^*). \end{aligned} \tag{11}$$

Thus, players in $\mathcal{N} \setminus \mathcal{N}_L$ can also not improve. So s is a pure Nash equilibrium.
By (11) and (7) we also get that

$$\begin{aligned} C_e(\mathsf{s}) &= \eta_e \cdot (\lambda \cdot c_1(w_e) + \mu \cdot c_1(w_e \cdot z_1) \cdot z_1) \\ &\quad + (1-\eta_e) \cdot (\lambda \cdot c_2(w_e) + \mu \cdot c_2(w_e \cdot z_1) \cdot z_1) \\ &\overset{(9)}{=} \eta_e \cdot (\lambda \cdot c_1(w_e) + \mu \cdot c_e(w_e)) + (1-\eta_e) \cdot (\lambda \cdot c_2(w_e) + \mu \cdot c_e(w_e)) \\ &= \lambda \cdot (\eta_e \cdot c_1(w_e) + (1-\eta_e) \cdot c_2(w_e)) + \mu \cdot c_e(w_e) \\ &= \lambda \cdot C_e(\mathsf{s}^*) + \mu \cdot C_e(\mathsf{s}), \end{aligned}$$

or equivalently

$$C_e(\mathsf{s}) = \frac{\lambda}{1-\mu} \cdot C_e(\mathsf{s}^*) = \zeta_\epsilon \cdot C_e(\mathsf{s}^*). \tag{12}$$

In the following, we show that $C(\mathsf{s}) = k$ and $C(\mathsf{s}^*) = 1 + (k-1)/\zeta_\epsilon$. To see that $C(\mathsf{s}) = k$ we show that the contribution of each non-leaf player $e \in \mathcal{N} \setminus \mathcal{N}_L$ to the total latency is the same as the combined contribution of both of her left child and right children l, r. The contribution of e to $C(\mathsf{s})$ is $c_e(w_e) \cdot w_e$ while the combined contribution of l and r is

$$\begin{aligned} &c_l(w_e \cdot z_1) \cdot w_e \cdot z_1 + c_r(w_e \cdot z_2) \cdot w_e \cdot z_2 \\ &\quad \overset{(10)}{=} \eta_e \cdot c_1(w_e \cdot z_1) \cdot w_e \cdot z_1 + (1-\eta_e) \cdot c_2(w_e \cdot z_2) \cdot w_e \cdot z_2 \\ &\quad \overset{(9)}{=} \eta_e \cdot c_e(w_e) \cdot w_e + (1-\eta_e) \cdot c_e(w_e) \cdot w_e \\ &\quad = c_e(w_e) \cdot w_e. \end{aligned}$$

It follows that the combined contribution of each layer of players in our tree is the same. Since the root player contributes 1 to the total latency and our tree has depth k, we conclude that $C(\mathsf{s}) = k$.

On the other hand

$$\begin{aligned} C(\mathsf{s}^*) &= \sum_{i \in \mathcal{N}} w_i \cdot C_i(\mathsf{s}^*) \\ &= \sum_{i \in \mathcal{N}_L} w_i \cdot C_i(\mathsf{s}^*) + \sum_{i \in \mathcal{N} \setminus \mathcal{N}_L} w_i \cdot C_i(\mathsf{s}^*) \\ &\overset{(12)}{=} \sum_{i \in \mathcal{N}_L} w_i \cdot C_i(\mathsf{s}) + \sum_{i \in \mathcal{N} \setminus \mathcal{N}_L} w_i \cdot \frac{C_i(\mathsf{s})}{\zeta_\epsilon} \\ &= 1 + \frac{k-1}{\zeta_\epsilon}. \end{aligned}$$

The claim of the theorem now follows since

$$\lim_{k \to \infty} \frac{C(\mathsf{s})}{C(\mathsf{s}^*)} = \lim_{k \to \infty} \frac{k}{1 + \frac{k-1}{\zeta_\epsilon}} = \zeta_\epsilon.$$

□

Next we handle the case where there is exactly one binding half-plane and it corresponds to the cost function c and values $x, x^* > 0$. Let $z = x/x^*$. Using scaling and dilation, it is possible to obtain for each $w > 0$, a cost function c_w that is a dilated version of c such that $c_w(w(z+1)) > zc_w(zw)$ and $zc_w(zw)/c_w(w) = \zeta_\epsilon$. We construct a slightly modified lower bound example for this case.

Represent the resources as the vertices of simple line graph with a single player on each edge. Each player has two strategies - use exactly one of the two adjacent resources. One end of the line is identified as the root. The weight of the player that is closest to the root is set to 1. For each subsequent player, we set its weight to equal z times the weight of the previous player. For the root resource, we choose a cost function $c_0 \in \mathcal{C}$ with $c_0(1) = 1$. For each other resource, we define its cost function based on the cost function of the adjacent resource that is closer to the root. More concretely, consider two resources e and e' that are adjacent to each other with resource e being closer to the root. Let w denote the weight of the player on the adjoining edge and $c_e, c_{e'}$ denote the cost functions corresponding e and e'. We define the cost function $c_{e'}$ as a scaled or dilated version of c such that $zc_{e'}(zw) = c_e(w)$ and $zc_{e'}(zw)/c_{e'}(w) = \zeta_\epsilon$. When e' is the last resource, we set $c_{e'}$ to be the same as c_e. Let s denote the outcome where all players play the resource that is closer to the root, and let s^* denote the outcome where all players play the resource that is further from the root.

We will first claim that the outcome s is a Nash equilibrium. Consider a player (that is not the last player) with weight w. Let e denote that resource with cost c_e that this player plays on in this outcome. Note that it is the only player that is playing on this resource. Its cost is thus $c_e(w)$. The player's other option is to play on a resource e' with cost function $c_{e'}$ chosen such that $zc_{e'}(zw) = c_e(w)$. In the outcome s, another player with weight zw is already playing on that resource e'. Thus if the player was to deviate to play on the resource e', its cost will be $c_{e'}(zw + w)$. By the choice of $c_{e'}$ as a scaled/dilated version of c, this is at least $zc_{e'}(zw)$ which in turn equals $c_e(w)$. Thus the player's cost on deviation is more than its cost in the current outcome. For the last player, its two strategies have the same cost function and no other player is playing on either of the resources thus it will have no incentive to deviate from playing on the resource closer to the root.

We will next compute the ratio of the costs of the two outcomes s and s^*. First we will show that the weighted cost in the outcome s for all players is the same. Consider a player P_1 with weight w that is not the last player. This player plays on a resource e with cost function c_e in the outcome s. Its weighted cost in the outcome s is thus $wc_e(w)$. The player P_2 adjacent to this player that is further from the root has weight zw. In the outcome s, this player plays on a resource e' with cost function $c_{e'}$. Moreover, $c_{e'}$ is chosen such that $zc_{e'}(zw) = c_e(w)$. Player P_2's weighted cost is $zwc_{e'}(zw) = wc_e(w)$. Thus for all players, their weighted costs in the outcome s are the same.

Next we show that, for all players other than the last, the ratio of their weighted costs in the outcomes s and s^* is ζ_ϵ. Consider a player with weight w that has access to resources e and e' with resource e being closer to the root. The player's cost in the outcome s is $c_e(w)$. Unless this is the last player, this cost equals $zc_{e'}(zw)$. On the other hand the player's cost in the outcome s^* is $c_{e'}(w)$. We thus see that the ratio of the two costs for this player is ζ_ϵ.

For the last player, cost functions for the two resources that it has access to are the same. In either outcome, it is also the only player playing on these resources. Thus the ratio of its costs in the outcomes s, s^* is 1. We conclude that as the number of resources in the path grows, the ratio $C(\mathsf{s})/C(\mathsf{s}^*)$, which is a lower bound on the POA, will approach ζ_ϵ.

Combining all of the analysis above we obtain the following result.

THEOREM 4.3 (ASYMMETRIC WEIGHTED POA LOWER BOUND). *For every class $\mathcal{C}$ of cost functions that is closed under dilation, the worst-case POA of weighted congestion games with cost functions in $\mathcal{C}$ is precisely $\zeta(\mathcal{C})$.*

4.2. A Lower Bound for Parallel-Link Networks

We now focus on weighted congestion games with cost functions that are polynomials with nonnegative coefficients and maximum degree d. For such games we show that the POA does not change if we restrict to (symmetric) weighted congestion games on parallel links. Recall that for general case the POA was shown to be ϕ_d^{d+1} [Aland et al. 2011] where ϕ_d satisfies $(\phi_d + 1)^d = \phi_d^{d+1}$. We establish the following theorem.

THEOREM 4.4 (POA LOWER BOUND FOR PARALLEL LINK NETWORKS). *For weighted congestion games on parallel-link networks with cost functions that are polynomials with nonnegative coefficients and maximum degree d, the worst-case POA is precisely ϕ_d^{d+1}.*

We reiterate that the lower bound in Theorem 4.4 is new even for affine cost functions. We describe the example below.

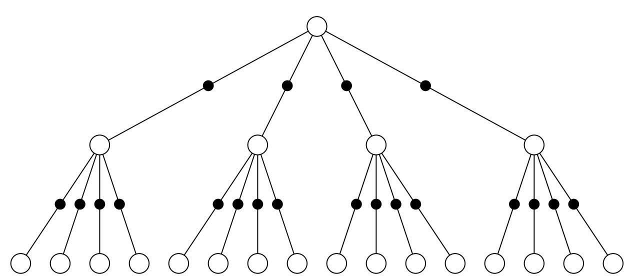

Fig. 3. Structure of the symmetric lower bound example for $\alpha = 4$ and $k = 3$. α is the branching factor, k is number of levels. Hollow circles are resources and solid circles are players.

LOWER BOUND 2 (WEIGHTED CONGESTION GAMES ON PARALLEL LINKS). *Let k be an integer. We construct the following congestion game on parallel links. (See Figure 3 for reference.) Let ϕ_d satisfy $(\phi_d+1)^d = \phi_d^{d+1}$. Let α be an integer satisfying $\alpha^d \geq \phi_d^{2d+2} \geq \phi_d^{d+1}$.*

Player strategies:. Player strategies are single resources and all players have access to all resources.
Cost functions:. Group the resources in groups $A_0, A_1 \cdots A_k$. For each $i = 0, 1, \ldots, k-1$, group A_i contains α^i resources with cost function $c_i(x) = \left(\alpha^d/\phi_d^{d+1}\right)^i x^d$. The last group A_k contains α^k resources with cost functions $c_k(x) = \alpha^d \left(\alpha^d/\phi_d^{d+1}\right)^{k-1} x^d$. These resources are arranged in a tree with resources from group A_i at level i of the tree.
Player weights:. Group players into groups $P_1, P_2, \ldots, P_k$. For $i = 1, 2, \ldots, k$, group P_i contains one player for each resource in A_i with player weight $w_i = (\alpha/\phi_d)^{k-i}$.
Optimal strategy:. Players in group P_i play on resources in group A_i in the optimal strategy. Denote this outcome by s^*.
Nash strategy:. Players in group P_i play on resources in group A_{i-1} in the Nash strategy with α players on each resource. Denote this outcome by s.

The example described above is symmetric which means that all players have access to all of the strategies. Thus we have to check that no player would gain from switching to any other strategy in the Nash outcome s. We verify these conditions in the following lemma. Theorem 4.4 then follows immediately from this Lemma.

LEMMA 4.5. *For the games in Lower Bound 2, outcome s is a pure Nash equilibrium and $\lim_{k\to\infty} \frac{C(\mathsf{s})}{C(\mathsf{s}^*)} = \phi_d^{d+1}$.*

PROOF. Recall from Lower Bound 2 that we have two outcomes s, s^*. We will establish that s is a Nash equilibrium and then show that $C(\mathsf{s})/C(\mathsf{s}^*)$ approaches ϕ_d^{d+1} as k tends to infinity. Note that since this is a symmetric congestion game — that is, all players have access to all strategies — we have to show that no player wants to switch to any other strategy.

Throughout these calculations $C_i(s)$ denotes the total weighted cost of resources in A_i in outcome s, while $c_j(s)$ denotes the cost of a single player j in outcome s. s_{-j} denotes the outcome when all players other than j play their strategy in outcome s. We first verify that the outcome s is indeed a Nash equilibrium.

Case (i) Players from groups $i = 1 \cdots k-1$.
In the outcome s, α players from group P_i play on resources in group A_{i-1}. The load on a resource in A_{i-1} is $\alpha \cdot (\alpha/\phi_d)^{k-i}$ and the cost of a player j in that group is

$$c_j(\mathsf{s}) = \left(\frac{\alpha^d}{\phi_d^{d+1}}\right)^{i-1} \left(\alpha \cdot \left(\frac{\alpha}{\phi_d}\right)^{k-i}\right)^d = \frac{\alpha^{kd}}{\phi^{d(k-1)+i-1}}. \tag{13}$$

Notice that this cost is decreasing in i, so players closer to the root pay more in the Nash strategy than players lower in the tree. Since all players on a resource pay the same cost no player would wish to switch to a strategy that already costs more than his current cost. Hence it remains to consider players wishing to switch to resources with a bigger index i.

Fix an index i. We verify that for any player j from group P_i, using a resource in A_{i-1} in s, switching to a resource $(A_i)_l$ in A_i is at least as bad. Recall this is the group in which this player's strategy in the optimal outcome s^* lies.

$$\begin{aligned} c_j(\mathsf{s}_{-j},(A_i)_l) &= \left(\frac{\alpha^d}{\phi_d^{d+1}}\right)^i \left(\alpha\left(\frac{\alpha}{\phi_d}\right)^{k-i-1} + \left(\frac{\alpha}{\phi_d}\right)^{k-i}\right)^d \\ &= \left(\frac{\alpha^d}{\phi_d^{d+1}}\right)^i \left(\frac{\alpha}{\phi_d}\right)^{(k-i)d} (\phi_d+1)^d \\ &= \left(\frac{\alpha^d}{\phi_d^{d+1}}\right)^i \left(\frac{\alpha}{\phi_d}\right)^{(k-i)d} \phi_d^{d+1} = c_j(\mathsf{s}). \end{aligned}$$

Next we verify that for any player switching to a strategy $(A_{i+t})_l$ in group A_{i+t} for $t > 0$ can only increase a player's cost. Even if the resource is empty in the Nash equilibrium, the player does not want to switch to it:

$$\begin{aligned} c_j(\mathsf{s}_{-j},(A_{i+t})_l) &= \left(\frac{\alpha^d}{\phi_d^{d+1}}\right)^{i+t} \left(\alpha * \left(\frac{\alpha}{\phi_d}\right)^{k-i-t-1} + \left(\frac{\alpha}{\phi_d}\right)^{k-i}\right)^d \\ &> \left(\frac{\alpha^d}{\phi_d^{d+1}}\right)^{i+t} \left(\left(\frac{\alpha}{\phi_d}\right)^{k-i}\right)^d \geq \left(\frac{\alpha^d}{\phi_d^{d+1}}\right)^{i+1} \left(\left(\frac{\alpha}{\phi_d}\right)^{k-i}\right)^d \\ &\geq c_j(\mathsf{s}) \cdot \left(\frac{\alpha^d}{\phi_d^{2d+2}}\right) \geq c_j(\mathsf{s}). \end{aligned}$$

Here the last inequality follows from the fact that α is chosen such that $\alpha^d \geq \phi_d^{2d+2}$. Now if we consider the player switching to a resource $(A_k)_l$ in the kth group,

$$\begin{aligned} c_j(\mathsf{s}_{-j},(A_k)_l) &= \alpha^d \left(\frac{\alpha^d}{\phi_d^{d+1}}\right)^{k-1} \left(\frac{\alpha}{\phi_d}\right)^{(k-i)d} \\ &= \alpha^d \frac{\alpha^{d(k-i-1)}}{\phi_d^{(d+1)(k-i)}} c_j(\mathsf{s}) \\ &\geq c_j(\mathsf{s}). \end{aligned}$$

Here the last inequality follows from the choice of α. This establishes that the players in groups $1, 2, \ldots, k-1$ will not wish to change their strategy in the outcome s.

Case(ii) Players in P_k *.*
Players in group P_k have unit weight and play on resources in A_{k-1}, α at a time in s. The cost incurred by such a player is

$$c_j(\mathsf{s}) = \left(\frac{\alpha^d}{\phi_d^{d+1}}\right)^{k-1} \alpha^d = \frac{\alpha^{kd}}{\phi^{(d+1)(k-1)}}.$$

Comparing with equation (13), this is less than the cost in s of any other group of resources. Now consider moving a player j from his strategy in A_{k-1} to an empty resource $(A_k)_l$ in A_k:

$$c_j(\mathsf{s}_{-j}, A_{kj}) = \alpha^d \cdot \left(\frac{\alpha^d}{\phi_d^{d+1}}\right)^{k-1} = c_j(\mathsf{s}).$$

These cases thus establish that the outcome s is a Nash equilibrium.

Now we calculate the POA. For each group of resources A_i we calculate its cost in the outcomes s and s^*. All the resources in a group have the same cost function and the same load in each of s and s^*. If the cost function is of the form $a_i x^d$, then the cost of the group is given by ((# of resources) $\cdot a_i \cdot$ (load)$^{d+1}$).

Case(i) Cost of group A_i for $i = 1 \cdots k-1$.
The cost in the outcome s is

$$C_i(\mathsf{s}) = \alpha^i \cdot \left(\frac{\alpha^d}{\phi_d^{d+1}}\right)^i \cdot \left(\alpha \cdot \left(\frac{\alpha}{\phi_d}\right)^{k-i-1}\right)^{d+1} = \frac{\alpha^{k(d+1)}}{\phi_d^{(d+1)(k-1)}}.$$

In the outcome s^*, this cost is

$$C_i(\mathsf{s}^*) = \alpha^i \cdot \left(\frac{\alpha^d}{\phi_d^{d+1}}\right)^i \cdot \left(\frac{\alpha}{\phi_d}\right)^{(k-i)(d+1)} = \frac{\alpha^{k(d+1)}}{\phi_d^{(d+1)k}}.$$

Case(ii) Cost of group A_0 .
Resources in group A_0 are not used in the outcome s^*, so $C_0(\mathsf{s}^*) = 0$. $C_0(\mathsf{s})$ is

$$C_0(\mathsf{s}) = 1 \cdot 1 \cdot \left(\alpha \cdot \left(\frac{\alpha}{\phi_d}\right)^{(k-1)}\right)^{(d+1)} = \frac{\alpha^{k(d+1)}}{\phi_d^{(k-1)(d+1)}}.$$

Case(iii) Cost of group A_k .
Resources in group A_k are not used in the outcome s, so $C_k(\mathsf{s}) = 0$. $C_k(\mathsf{s}^*)$ is

$$C_k(\mathsf{s}^*) = (\alpha^k) \cdot \alpha^d \cdot \left(\frac{\alpha^d}{\phi_d^{d+1}}\right)^{k-1} = \frac{\alpha^{k(d+1)}}{\phi_d^{(d+1)(k-1)}}.$$

Combining all these expressions,

$$C(\mathsf{s}) = \sum_{i=0}^{k} C_i(\mathsf{s}) = k \cdot \frac{\alpha^{k(d+1)}}{\phi_d^{(k-1)(d+1)}}, \text{ and}$$

$$\begin{aligned} C(\mathsf{s}^*) &= \sum_{i=0}^{k} C_i(\mathsf{s}^*) = (k-1)\frac{\alpha^{k(d+1)}}{\phi_d^{(d+1)k}} + \frac{\alpha^{k(d+1)}}{\phi_d^{(d+1)(k-1)}} \\ &= \frac{\alpha^{k(d+1)}}{\phi_d^{(k-1)(d+1)}} \cdot \left((k-1)(1/\phi_d^{d+1}) + 1\right). \end{aligned}$$

Hence,

$$\frac{C(s)}{C(s^*)} = \frac{k}{(k-1)\frac{1}{\phi_d^{d+1}} + 1}.$$

The claim follows since $\lim_{k \leftarrow \infty} \frac{k}{(k-1)\frac{1}{\phi_d^{d+1}}+1} = \phi_d^{d+1}$.

□

5. UNWEIGHTED CONGESTION GAMES

We show that for symmetric unweighted congestion games the POA is the same as that for general unweighted congestion games. A result in Roughgarden [2009] gives an upper bound on the POA of unweighted congestion games. For a class of cost functions $\mathcal{C}$, one uses the set of parameters $\mathcal{A}(\mathcal{C})$ defined as

$$\mathcal{A}(\mathcal{C}) = \{(\lambda, \mu) : \mu < 1; x^* c(x + 1) \leq \lambda x^* c(x^*) + \mu x c(x)\}, \tag{14}$$

where the constraints are for cost functions $c \in \mathcal{C}$ and integers $x \geq 0, x^* > 0$. Then for each $(\lambda, \mu) \in \mathcal{A}(\mathcal{C})$, a simple derivation shows that the pure POA is at most $\lambda/(1 - \mu)$. Let $\gamma(\mathcal{C})$ denote the best such upper bound:

$$\gamma(\mathcal{C}) = \inf \left\{ \frac{\lambda}{1 - \mu} : (\lambda, \mu) \in \mathcal{A}(\mathcal{C}) \right\}.$$

Roughgarden [2009] showed how to construct an *asymmetric* unweighted congestion game that matches the upper bound $\gamma(\mathcal{C})$ (for each $\mathcal{C}$). Here we show that even symmetric games can be used to achieve this upper bound, in the limit as the number of players and resources tend to infinity.

THEOREM 5.1 (SYMMETRIC UNWEIGHTED POA LOWER BOUND). *For every set of cost functions $\mathcal{C}$ there exist symmetric congestion games with cost functions in $\mathcal{C}$ and POA arbitrarily close to $\gamma(\mathcal{C})$.*

We define $\gamma(\mathcal{C}, n)$ as the minimal value of $\lambda/(1 - \mu)$ that can be obtained when the load on each edge is restricted to be at most n. Then $\gamma(\mathcal{C}, n)$ approaches $\gamma(\mathcal{C})$ as n approaches ∞. For any finite n the feasible region for (λ, μ) is then the intersection of a finite number of half planes, one for each value of x and x^*. We also maintain one additional constraint on the feasible region that is $\mu < 1$. We now state the following lemma which was presented in Roughgarden [2009].

LEMMA 5.2 (CHARACTERIZATION OF BINDING HALF-PLANES). *Fix finite n and a set of functions $\mathcal{C}$ and suppose there exists $(\hat{\lambda}, \hat{\mu})$ such that, $\frac{\hat{\lambda}}{1-\hat{\mu}} = \gamma(\mathcal{C}, n)$. Then there exist $c_1, c_2 \in \mathcal{C}$, $x_1, x_2 \in \{0, 1, \ldots, n\}$, $x_1^*, x_2^* \in \{1, 2, \ldots, n\}$ such that,*

$$c_j(x_j + 1)x_j^* = \hat{\lambda} c_j(x_j^*)x_j^* + \hat{\mu} c_j(x_j)x_j \tag{15}$$

for $j \in \{1, 2\}$. Additionally, $\beta_{c_1,x_1,x_1^*} < 1$ and $\beta_{c_2,x_2,x_2^*} \geq 1$ for $\beta_{c,x,x^*} = \frac{xc(x)}{x^* c(x+x^*)}$.

Note that the lemma as stated here differs from the one in Roughgarden [2009] in the additional condition on $\beta_{c_1,x_1,x_1^*}, \beta_{c_2,x_2,x_2^*}$. However the modified version can be easily obtained by noting that we are always guaranteed a half plane with $\beta < 1$ and as long as the value of $\gamma(\mathcal{C}, n)$ is attained there is another half plane with $\beta \geq 1$.

We now describe the lower bound.

LOWER BOUND 3 (SYMMETRIC UNWEIGHTED CONGESTION GAMES). *Let N be an integer which will denote the number of players. Let c_1, x_1, x_1^* and c_2, x_2, x_2^* be defined as in the above lemma.*

Resources and Cost functions. There are two groups of resources A_1 and A_2. For $j = 1, 2$, group A_j contains $\binom{N}{x_j}\binom{N}{x_j^*}$ resources with cost function $\alpha_j c_j(x)$, α_1 and α_2 will be chosen later. Arrange resources in A_j in a $\binom{N}{x_j} \times \binom{N}{x_j^*}$ grid.

Strategies. There are two main types of strategies: column strategies and row strategies. There are N column strategies . For $i \in \{1, 2, \ldots, N\}$, the i'th column strategy $O(i)$ is composed of sets $O(i, 1)$ and $O(i, 2)$ from resource sets A_1, A_2 respectively.

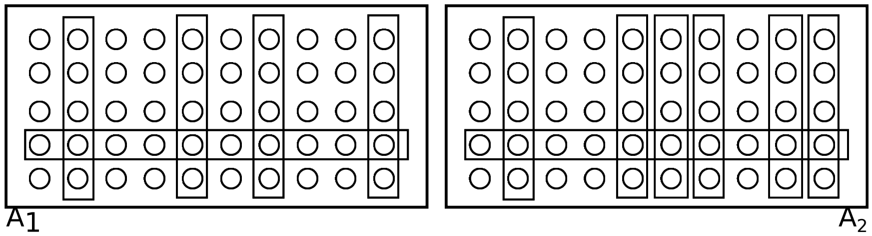

Fig. 4. Example of column and row strategies with N = 5, $x_1^* = x_2^* = 1$, $x_1 = 2$ and $x_2 = 3$. Hollow circles denote resources. One column strategy includes all columns marked by vertical blocks. An example row strategy includes all resources marked by horizontal blocks.

Set $O(i,j)$ contains $\binom{N-1}{x_j^*-1}$ of the columns from set A_j, resulting in $\Delta_j = \binom{N}{x_j}\binom{N-1}{x_j^*-1}$ resources in all. The columns are chosen such that each resource is contained in exactly x_j^* sets.[2]
Row strategies denoted by $E(i)$ for $i \in \{1,2,\ldots,N\}$ are composed of sets $E(i,1)$ and $E(i,2)$ from A_1 and A_2, respectively. Set $E(i,j)$ is composed of $\binom{N-1}{x_j-1}$ rows from A_j, chosen such that each resource is contained in x_j different row strategies. The number of resources in E(i,j) is $\frac{x_j}{x_j^*}\Delta_j$. The sets $E(i,j)$ and $O(i',j)$ share $\frac{x_j}{N}\Delta_j$ resources.
Players. There are N players each with unit weight. Players have access to all of the strategies.
Optimal strategy. The optimal strategy s_i^* for player i is the strategy $O(i)$.
Equilibrium strategy. The equilibrium strategy s_i for player i is the strategy $E(i)$.

We choose $\alpha_1, \alpha_2 \geq 0$ such that each player is indifferent between his strategy in the equilibrium and optimal outcomes while all other players play their equilibrium strategy. The following lemma establishes that such α_1, α_2 exist.

LEMMA 5.3 (CHOICE OF α_1, α_2). *If tuples c_1, x_1, x_1^* and c_2, x_2, x_2^* with $x_1, x_2 \geq 0, x_1^*, x_2^* > 0$ are such that $\beta_{c_1,x_1,x_1^*} < 1$ and $\beta_{c_2,x_2,x_2^*} \geq 1$, then for the game instance described in Lower Bound 3 there exist $\alpha_1, \alpha_2 \geq 0$ s.t. for each player i, $C_i(s_i, s_{-i}) = C_i(s_i^*, s_{-i})$.*

PROOF. We start with the required condition, $C_i(\mathsf{s}_{-i}, \mathsf{s}_i^*) = C_i(\mathsf{s})$, to derive the values of α_1 and α_2. Recall that player i's equilibrium strategy is composed of sets of resources $E(i,1)$ and $E(i,2)$ that in s cost $\alpha_1 c_1(x_1)$ and $\alpha_2 c_2(x_2)$, respectively. Player i's optimal strategy is composed of sets of resources $O(i,1), O(i,2)$. Hence we have

$$\begin{aligned}
&C_i(\mathsf{s}) = C_i(\mathsf{s}_{-i}, \mathsf{s}_i^*) \iff \\
&|E(i,1)|\alpha_1 c_1(x_1) + |E(i,2)|\alpha_2 c_2(x_2) \\
=&|E(i,1) \cap O(i,1)|\alpha_1 c_1(x_1) + (|O(i,1)| - |E(i,1) \cap O(i,1)|)\,\alpha_1 c_1(x_1+1) \\
&+ |E(i,2) \cap O(i,2)|\alpha_2 c_2(x_2) + (|O(i,2)| - |E(i,2) \cap O(i,2)|)\,\alpha_2 c_2(x_2+1).
\end{aligned}$$

Plugging in the sizes of the strategy sets, we get

$$\begin{aligned}
\frac{x_1}{x_1^*}\Delta_1\alpha_1 c_1(x_1) + \frac{x_2}{x_2^*}\Delta_2\alpha_2 c_2(x_2) =& \frac{x_1}{N}\Delta_1\alpha_1 c_1(x_1) + (1 - \frac{x_1}{N})\Delta_1\alpha_1 c_1(x_1+1) \\
&+ \frac{x_2}{N}\Delta_2\alpha_2 c_2(x_2) + (1 - \frac{x_2}{N})\Delta_2\alpha_2 c_2(x_2+1).
\end{aligned}$$

[2]A simple way of defining these strategies is to denote each column by a N bit number containing exactly x_j^* ones. The ith strategy then contains columns whose i'th bit is 1.

Regrouping this equation,

$$\Delta_1\alpha_1\left[\left(1-\frac{x_1}{N}\right)c_1(x_1+1)-\left(\frac{x_1}{x_1^*}-\frac{x_1}{N}\right)c_1(x_1)\right]$$
$$=\Delta_2\alpha_2\left[\left(\frac{x_2}{x_2^*}-\frac{x_2}{N}\right)c_2(x_2)-\left(1-\frac{x_2}{N}\right)c_2(x_2+1)\right].$$

Further regrouping and some simplification yields

$$\Delta_1\alpha_1[N\,(c_1(x_1+1)-x_1c_1(x_1)/x_1^*)+(x_1c_1(x_1)-x_1c_1(x_1+1))]$$
$$=\Delta_2\alpha_2[N\,(x_2c_2(x_2)/x_2^*-c_2(x_2+1))+(x_2c_2(x_2+1)-x_2c_2(x_2))].$$

Choosing α_1, α_2 such that

$$a_2=\frac{x_2c_2(x_2)}{x_2^*}-c_2(x_2+1), \qquad b_2=x_2c_2(x_2+1)-x_2c_2(x_2),$$
$$a_1=c_1(x_1+1)-\frac{x_1c_1(x_1)}{x_1^*}, \qquad b_1=x_1c_1(x_1)-x_1c_1(x_1+1),$$
$$\alpha_1=(a_2N+b_2)/\Delta_1,\ \text{and} \qquad \alpha_2=(a_1N+b_1)/\Delta_2 \tag{16}$$

ensures that the relation holds. Note that since the cost functions are increasing and $\beta_{c_2,x_2,x_2^*}\geq 1$, both $a_2, b_2 \geq 0$ and α_1 is non-negative. We have $b_1 < 0$ but since $\beta_{c_1,x_1,x_1^*} < 1$ and $a_1 > 0$, for sufficiently large N, α_2 is also non-negative. Due to the symmetry between the players we conclude that the above choice of α_1, α_2 renders each player indifferent between their Nash and optimal strategies in the equilibrium. □

It now remains to prove that the Lower Bound 3 has the desired POA. We first note a few easy identities.

LEMMA 5.4. *The following two identities hold:*

$$\frac{a_2x_1c_1(x_1)}{x_1^*}+\frac{a_1x_2c_2(x_2)}{x_2^*}=\hat{\lambda}\left[\frac{x_2c_1(x_1^*)c_2(x_2)}{x_2^*}-\frac{x_1c_1(x_1)c_2(x_2^*)}{x_1^*};\right]$$
$$a_2c_1(x_1)+a_1c_2(x_2)=(1-\hat{\mu})\left[\frac{x_2c_2(x_2)c_1(x_1^*)}{x_2^*}-\frac{x_1c_1(x_1)c_2(x_2^*)}{x_1^*}.\right]$$

PROOF. From Lemma 5.2 we know that, for $j\in\{1,2\}$,

$$x_j^*c_j(x_j+1)=\hat{\lambda}x_j^*c_j(x_j^*)+\hat{\mu}x_jc_j(x_j).$$

Hence,

$$a_2=\frac{x_2c_2(x_2)}{x_2^*}-c_2(x_2+1)=\frac{x_2c_2(x_2)}{x_2^*}-\left(\hat{\lambda}c_2(x_2^*)+\hat{\mu}\frac{x_2c_2(x_2)}{x_2^*}\right)$$
$$=(1-\hat{\mu})\frac{x_2c_2(x_2)}{x_2^*}-\hat{\lambda}c_2(x_2^*).$$

Similarly,

$$a_1=\hat{\lambda}c_1(x_1^*)-(1-\hat{\mu})\frac{x_1c_1(x_1)}{x_1^*}.$$

Then,

$$
\begin{aligned}
&\frac{a_2 x_1 c_1(x_1)}{x_1^*} + \frac{a_1 x_2 c_2(x_2)}{x_2^*} \\
=&\left((1-\hat{\mu})\frac{x_2 c_2(x_2)}{x_2^*} - \hat{\lambda} c_2(x_2^*)\right)\frac{x_1 c_1(x_1)}{x_1^*} + \left(\hat{\lambda} c_1(x_1^*) - (1-\hat{\mu})\frac{x_1 c_1(x_1)}{x_1^*}\right)\frac{x_2 c_2(x_2)}{x_2^*} \\
=&\hat{\lambda}\left[\frac{x_2 c_1(x_1^*) c_2(x_2)}{x_2^*} - \frac{x_1 c_1(x_1) c_2(x_2^*)}{x_1^*}\right].
\end{aligned}
$$

A similar proof verifies the second identity. □

Using these identities, we prove the following lemma. Recall that $\hat{\lambda}$ and $\hat{\mu}$ were fixed in the statement of Lemma 5.2.

LEMMA 5.5. *The* POA *of the game in Lower Bound 3 approaches* $\hat{\lambda}/(1-\hat{\mu})$ *as* $N \to \infty$.

PROOF. We first establish that the outcome s is a Nash equilibrium. The equilibrium cost for each player is the same, hence switching to another player's equilibrium strategy will only increase the player's cost. Next we check that switching to another player j's optimal strategy is also more expensive. From a player i's perspective all the optimal strategies are similar, hence by the choice of α_1, α_2 we know that $C_i(s_i, s_{-i}) = C_i(s_j^*, s_{-i})$. Thus, a player will not gain by switching to another player's optimal strategy.

Next we calculate the POA. Recall that α_1, α_2 are defined as in equation (16). We have

$$
\begin{aligned}
C_i(\mathsf{s}) =& \frac{x_1}{x_1^*}\Delta_1 \alpha_1 c_1(x_1) + \frac{x_2}{x_2^*}\Delta_2 \alpha_2 c_2(x_2) \\
=& \frac{x_1}{x_1^*}(a_2 N + b_2)\alpha_1 c_1(x_1) + \frac{x_2}{x_2^*}(a_1 N + b_1)\alpha_2 c_2(x_2) \\
=& \left(\frac{a_2 x_1 c_1(x_1)}{x_1^*} + \frac{a_1 x_2 c_2(x_2)}{x_2^*}\right) N + \left(\frac{b_2 x_1 c_1(x_1)}{x_1^*} + \frac{b_1 c_2(x_2)}{x_2^*}\right) \\
=& \hat{\lambda} N \left[\frac{x_2 c_1(x_1^*) c_2(x_2)}{x_2^*} - \frac{x_1 c_1(x_1) c_2(x_2^*)}{x_1^*}\right] + \left(\frac{b_2 x_1 c_1(x_1)}{x_1^*} + \frac{b_1 c_2(x_2)}{x_2^*}.\right)
\end{aligned}
$$

Here the final step follows from the identities in Lemma 5.4. Similar calculations show that

$$
\begin{aligned}
C_i(\mathsf{s}^*) =& \Delta_1 \alpha_1 c_1(x_1^*) + \Delta_2 \alpha_2 c_2(x_2^*) \\
=& (1-\hat{\mu}) N \left[\frac{x_2 c_1(x_1^*) c_2(x_2)}{x_2^*} - \frac{x_1 c_1(x_1) c_2(x_2^*)}{x_1^*}\right] + (b_2 c_1(x_1^*) + b_1 c_2(x_2^*).)
\end{aligned}
$$

Since the coefficients of $\hat{\lambda}$ and $1-\hat{\mu}$ are strictly positive, we have

$$
\mathsf{POA} \geq \frac{N \cdot C_i(\mathsf{s})}{N \cdot C_i(\mathsf{s}^*)} \to \frac{\hat{\lambda}}{1-\hat{\mu}} \text{ as } N \to \infty.
$$

□

When $\gamma(\mathcal{C}, n)$ is not attained by any pair (λ, μ), a similar construction yields instances with POA arbitrarily close to $\gamma(\mathcal{C}, n)$. We establish the following lemma.

LEMMA 5.6. *If $\mathcal{C}$ is a class of functions and n an integer such that the value $\gamma(\mathcal{C}, n)$ is not attained by any feasible pair (λ, μ), then there exist symmetric unweighted congestion games with POA arbitrarily close to $\gamma(\mathcal{C}, n)$.*

PROOF. The hypotheses imply that the optimal value of $\lambda/(1-\mu)$ is approached on the last face of the feasible region as $\lambda \to \infty$ and $\mu \to -\infty$. The exact value of the POA is then $nc(n)/c(1)$ and we also have that $c(n+1) > nc(n)$.

We construct an example similar to one in Lower Bound 3 for this special case. We choose a k such that

$$nc(n)(1 + \frac{1 - 1/n}{k-1}) \leq c(n+1). \tag{17}$$

Such a k exists as $c(n+1) > n \cdot c(n)$. The instance is then composed of kn players. There are $\binom{kn}{n}kn$ resources, each with cost function $c(x)$. In the equilibrium, each player plays $\binom{kn-1}{n-1}kn$ resources while sharing each resource with n other players. For every player we specify an optimal strategy composed of $\binom{kn-1}{n-1}$ resources from each player's equilibrium strategy. No resources are shared in the optimal outcome. The optimal strategy has $\binom{kn}{n}$ resources.

We first verify that the equilibrium conditions are met. Since all the players' equilibrium strategies are identical, no player has an incentive to deviate to another player's equilibrium strategy. Now suppose a player i deviates to another player j's optimal strategy. The two strategies share $\binom{kn-1}{n-1}$ resources. We have

$$\begin{aligned} C_i(\mathsf{s}_{-i}, \mathsf{s}_j^*) &= \binom{kn-1}{n-1} \cdot c(n) + \binom{kn-1}{n} c(n+1) \\ &= \binom{kn-1}{n-1} \left(c(n) + \frac{kn-n}{n} \cdot c(n+1) \right) \\ &\geq \binom{kn-1}{n-1} \left(c(n) + (kn-n)(1 + \frac{n-1}{n(k-1)}) \cdot c(n) \right) \\ &= \binom{kn-1}{n-1} \cdot kn \cdot c(n) = C_i(\mathsf{s}), \end{aligned}$$

where the inequality follows from the choice of k.

The social cost of the equilibrium is $kn \cdot kn \cdot \binom{kn-1}{n-1} \cdot c(n)$ while the cost of the optimal outcome is $kn \cdot \binom{kn}{n} \cdot c(1)$. The POA is then $nc(n)/c(1)$. □

Theorem 5.1 follows from Lemmas 5.5 and 5.6.

6. CONCLUSION AND OPEN QUESTIONS

This paper determines the exact POA in several classes of congestion games. There remain some interesting open questions.

— **Symmetric Unweighted Congestion Games on Networks -** It is not clear if our lower bound for the POA of symmetric unweighted congestion games (Theorem 5.1) extends to *network* congestion games, where each player's strategies correspond to the paths between a source and a sink vertex.

— **Symmetric Weighted Congestion Games -** The lower bound for parallel link networks in Section 4.2 applies only to polynomial cost functions. It would be interesting to determine the exact POA for such games with general cost functions.

REFERENCES

S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. 2011. Exact Price of Anarchy for Polynomial Congestion Games. *SIAM J. Comput.* 40, 5 (2011), 1211–1233.

E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. 2008. The Price of Stability for Network Design with Fair Cost Allocation. *SIAM J. Comput.* 38, 4 (2008), 1602–1623.

B. Awerbuch, Y. Azar, and L. Epstein. 2005. The Price of Routing Unsplittable Flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*. 57–66.

A. Blum, E. Even-Dar, and K. Ligett. 2010. Routing without Regret: On Convergence to Nash Equilibria of Regret-Minimizing Algorithms in Routing Games. *Theory of Computing* 6 (2010), 179–199.

I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli. 2011. Tight Bounds for Selfish and Greedy Load Balancing. *Algorithmica* 61, 3 (2011), 606–637.

G. Christodoulou and E. Koutsoupias. 2005a. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Proceedings of the 13th Annual European Symposium on Algorithms (ESA)*. 59–70.

G. Christodoulou and E. Koutsoupias. 2005b. The Price of Anarchy of Finite Congestion Games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*. 67–73.

J. R. Correa, A. S. Schulz, and N. E. Stier Moses. 2004. Selfish Routing in Capacitated Networks. *Mathematics of Operations Research* 29, 4 (2004), 961–976.

D. Fotakis. 2010. Congestion Games with Linearly Independent Paths: Convergence Time and Price of Anarchy. *Theory of Computing Systems* 47, 1 (2010), 113–136.

D. Fotakis, S. C. Kontogiannis, and P. G. Spirakis. 2008. Atomic congestion games among coalitions. *ACM Transactions on Algorithms* 4, 4 (2008).

M. Gairing and F. Schoppmann. 2007. Total Latency in Singleton Congestion Games. In *Proceedings of the 3rd International Workshop on Internet and Network Economics (WINE)*. 381–387.

T. Harks and M. Klimm. 2012. On the Existence of Pure Nash Equilibria in Weighted Congestion Games. *Mathematics of Operations Research* 37, 3 (2012), 419–436.

T. Harks, M. Klimm, and R. H. Möhring. 2011. Characterizing the Existence of Potential Functions in Weighted Congestion Games. *Theory of Computing Systems* 49, 1 (2011), 46–70.

A. Hayrapetyan, É. Tardos, and T. Wexler. 2006. The Effect of Collusion in Congestion Games. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*. 89–98.

R. Holzman and N. Law-Yone. 2003. Network Structure and Strong Equilibrium in Route Selection Games. *Mathematical Social Sciences* 46 (2003), 193–205.

E. Koutsoupias and C. H. Papadimitriou. 1999. Worst-case Equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*. 404–413.

T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. 2008. A New Model for Selfish Routing. *Theoretical Computer Science* 406, 3 (2008), 187–206.

I. Milchtaich. 1996. Congestion Games with Player-Specific Payoff Functions. *Games and Economic Behavior* 13, 1 (1996), 111–124.

R. W. Rosenthal. 1973. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory* 2, 1 (1973), 65–67.

T. Roughgarden. 2003. The Price of Anarchy Is Independent of the Network Topology. *J. Comput. System Sci.* 67, 2 (2003), 341–364.

T. Roughgarden. 2006. Potential Functions and the Inefficiency of Equilibria. In *Proceedings of the ICM*, Vol. III. 1071–1094.

T. Roughgarden. 2009. Intrinsic Robustness of the Price of Anarchy. In *Proceedings of the 41st ACM Symposium on Theory of Computing (STOC)*. 513–522.

T. Roughgarden. 2012. The Price of Anarchy in Games of Incomplete Information. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*. 862–879.

T. Roughgarden and É. Tardos. 2004. Bounding the Inefficiency of Equilibria in Nonatomic Congestion Games. *Games and Economic Behavior* 49, 2 (2004), 389–403.

L. S. Shapley. 1953. *Additive and Non-Additive Set Functions*. Ph.D. Dissertation. Department of Mathematics, Princeton University.