Optimal Cost-Sharing in Weighted Congestion Games

Vasilis Gkatzelis, Konstantinos Kollias, and Tim Roughgarden

Stanford University, Stanford CA 94305, USA

Abstract. We identify how to share costs locally in weighted congestion games with polynomial cost functions in order to minimize the worst-case price of anarchy (PoA). First, we prove that among all cost-sharing methods that guarantee the existence of pure Nash equilibria, the Shapley value minimizes the worst-case PoA. Second, if the guaranteed existence condition is dropped, then the proportional cost-sharing method minimizes the worst-case PoA over all cost-sharing methods. As a byproduct of our results, we obtain the first PoA analysis of the simple marginal contribution cost-sharing rule, and prove that marginal cost taxes are ineffective for improving equilibria in (atomic) congestion games.

Keywords: cost-sharing, selfish routing, congestion games.

1 Introduction

Sharing Costs to Optimize Equilibria. Weighted congestion games [35] are a simple class of competitive games that are flexible enough to model diverse settings (e.g., routing [38], network design [3], and scheduling [26]). These games consist of a ground set of resources E and a set of players N who are competing for the use of these resources. Every player $i \in N$ is associated with a weight w_i, and has a set of strategies $\mathcal{P}_i \subseteq 2^E$, each of which corresponds to a subset of the resources. Given a strategy profile $P = (P_i)_{i \in N}$, where $P_i \in \mathcal{P}_i$. The set of players $S_j = \{i : j \in P_i\}$ using some resource $j \in E$ generates a social cost $C_j(f_j)$ on this resource (e.g., $C_j(f_j) = \alpha \cdot f_j^d$), which is a function of their total weight $f_j = \sum_{i \in S_j} w_i$; this joint cost could represent monetary cost, or a physical cost such as aggregate queueing delay [40]. The social cost of every resource can be distributed among the players using it. As a result, each player i suffers some cost $c_{ij}(P)$ from each of the resources $j \in P_i$, and its goal is to choose a strategy that minimizes the cost $\sum_{j \in P_i} c_{ij}(P)$ that it suffers across all the resources that it uses.

How can a system designer minimize the sum of the players' costs

$$\sum_{i \in N} \sum_{j \in P_i} c_{ij}(P) \;=\; \sum_{j \in E} C_j(f_j) \qquad (1)$$

over all possible outcomes P? The answer depends on what the designer can do. We are interested in settings where centralized optimization is infeasible, and

T.-Y. Liu et al. (Eds.): WINE 2014, LNCS 8877, pp. 72–88, 2014.

the designer can only influence players' behavior through local design decisions. Precisely, we allow a designer to choose a *cost-sharing rule* that defines what share $\xi_j(i, S_j)$ of the joint cost $C_j(f_j)$ each player $i \in S_j$ is responsible for (so $c_{ij}(P) = \xi_j(i, S_j)$). Such cost-sharing rules are "local" in the sense that they depend only on the set of players using the resource, and not on the users of other resources. Given a choice of a cost-sharing rule, we can quantify the inefficiency in the resulting game via the price of anarchy (PoA) — the ratio of the total cost at the worst equilibrium and the optimal cost. The set of equilibria and hence the PoA are a complex function of the chosen cost-sharing rule.

The goal of this paper is to answer the following question.

> *Which cost-sharing rule minimizes the worst-case PoA in polynomial weighted congestion games?*

In other words, when a system designer can only indirectly influence the game outcome through local design decisions, what should he or she do?

The present work is the first to study this question. Previous work on the PoA in weighted congestion games, reviewed next, has focused exclusively on evaluating the worst-case PoA with respect to a single cost-sharing rule. Previous work on how to best locally influence game outcomes in other models is discussed in Section 1.2.

Example: Proportional Cost Sharing. Almost all previous work on weighted congestion games has studied the proportional cost-sharing rule [4, 6, 18, 20, 21, 29, 30, 36]. According to this rule, at each resource j, each player $i \in S_j$ is responsible for a w_i/f_j fraction of the joint cost, i.e., a fraction proportional to its weight. The worst-case PoA is well understood in games with proportional cost sharing. For polynomial cost functions $C_j(f_j)$ of maximum degree d, the worst-case PoA is ϕ_d^d, where $\phi_d = \Theta(d/\ln d)$ is the positive root of $f_d(x) = x^d - (x+1)^{d-1}$ [2]. One of the main disadvantages of this cost-sharing rule is that it does not guarantee the existence of a pure Nash equilibrium (PNE), but the upper bounds are still meaningful since they apply to much more general equilibrium concepts, like coarse correlated equilibria, which do exist [37].

Example: Shapley Cost Sharing. The only other cost-sharing rule for which the worst-case PoA of weighted congestion games is known is the Shapley cost sharing rule. The cost shares defined by the Shapley value can be derived in the following manner: given an ordering over the users of a resource, these users are introduced to the resource in that order and each user is responsible for the marginal increase of the cost caused by its arrival. The Shapley cost share of each user is then defined as its average marginal increase over all orderings. Unlike proportional cost sharing, this cost-sharing rule guarantees the existence of a PNE in weighted congestion games [25]. The worst-case PoA of this rule for polynomials of maximum degree d is χ_d^d, where $\chi_d \approx 0.9 \cdot d$ is the root of $g_d(x) = 3 \cdot x^d - (x+1)^d - 1$ [25].

Other Cost-Sharing Rules. Without the condition of guaranteed PNE existence, there is obviously a wide range of possible cost-sharing rules. The space of rules that guarantee the existence of PNE is much more limited but still quite rich. [19] showed that the space of such rules correspond exactly to the weighted Shapley values. This class of cost-sharing rules generalizes the Shapley cost sharing rule by assigning a sampling weight λ_i to each player i. The cost shares of the players are then defined to be an appropriately *weighted* average of their marginal increases over different orderings (see Section 2.1); hence the "design space" is $(k-1)$-dimensional, where k is the number of players.[1]

1.1 Our Results

Our two main results resolve the question of how to optimally share cost in weighted congestion games to minimize the worst-case PoA.

Main Result 1 (Informal): Among all cost-sharing rules that guarantee the existence of PNE, the worst-case PoA of weighted congestion games is minimized by the Shapley cost sharing rule.

For example, the plot of Figure 1 shows how the worst-case PoA varies for a well-motivated subclass of the weighted Shapley rules parameterized by a variable γ (details are in Section 2.1). The PoA of these cost-sharing rules varies and it exhibits discontinuous behavior, but in all cases it is at least as large as the PoA when the parameter value is $\gamma = 0$, which corresponds to the (unweighted) Shapley cost-sharing rule.

Main Result 2 (Informal): Among all cost-sharing rules, the worst-case PoA of weighted congestion games is minimized by the proportional cost sharing rule.

In the second result, we generously measure the PoA of pure Nash equilibria only in instances where such equilbria exist. That is, our lower bounds construct games that have a PNE that is far from optimal. For the optimal rule (proportional cost-sharing), however, the PoA upper bound applies more generally to equilibrium concepts that are guaranteed to exist, including coarse correlated equilibria.

As a byproduct of our results, we also obtain tight bounds for the worst-case PoA of the marginal contribution cost-sharing rule (see Section 2.2). The marginal contribution rule defines individual cost shares that may in general add up to more than the total joint costs, but we show that, even if any additional costs are disregarded, which reduces this policy to marginal cost taxes, the worst-case PoA remains high.

[1] The sampling weights λ_i can be chosen to be related to the players' weights w_i, or not. The joint cost is a function of players' weights w_i; the sampling weights only affect how this joint cost is shared amongst them.

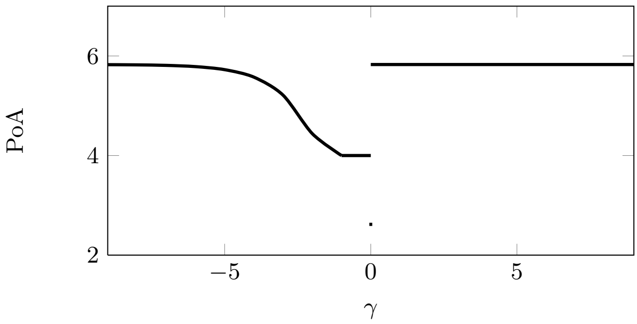

Fig. 1. PoA of parameterized weighted Shapley values for quadratic resource costs

1.2 Further Related Work

This paper contributes to the literature on how to design and modify games to minimize the inefficiency of equilibria. Several previous works consider how the choice of a cost-sharing rule affects this inefficiency in other models: [33] in participation games; [11, 15] in the network cost-sharing games of [3]; [32, 31, 22] in queueing games; and [28] in distributed welfare games. Closely related in spirit is previous work on coordination mechanisms, beginning with [12] and more recently in [23, 5, 24, 8, 13, 1, 14]. Most work on coordination mechanisms concerns scheduling games, and how the price of anarchy varies with the choice of local machine policies (i.e., the order in which to process jobs assigned to the same machine). Some of the recent work comparing the price of anarchy of different auction formats, such as [27, 7, 41], also has a similar flavor.

1.3 Organization of the Paper

In Section 2, we restrict ourselves to cost-sharing rules that guarantee the existence of PNE and we prove worst-case PoA bounds for such policies. In Section 3, we remove this restriction and we provide lower bounds for the worst-case PoA of arbitrary cost-sharing rules.

2 Cost-Sharing Rules That Guarantee PNE Existence

In this section we restrict our attention to cost-sharing rules that guarantee the existence of a PNE. This class of cost-sharing rules corresponds to weighted Shapley values, parameterized by a set of sampling weights λ_i, one for each player i [19]. We focus on congestion games with resource cost functions $C_j(f_j)$ that are polynomials with positive coefficients and maximum degree $d \geq 1$, and we provide worst-case PoA bounds parameterized by d. Note that every game with such cost functions has an equivalent game such that all cost functions have the form $\alpha \cdot x^k$ for $\alpha \geq 0$ and $k \in [1, d]$ (each resource is decomposed into many resources of this form that can only be used as a group). Hence, we will be assuming that all games studied in what follows have cost functions of the form $C_j(x) = \alpha_j \cdot x^{k_j}$. We conclude this section with a byproduct of our results: tight worst-case PoA bounds for the marginal contribution cost-sharing rule.

2.1 Weighted Shapley Values

We first describe how a weighted Shapley value defines payments for the players S that share a cost given by function $C(\cdot)$. For a given ordering π of the players in S, the marginal cost increase caused by player i is $C(f_i^\pi + w_i) - C(f_i^\pi)$, where f_i^π denotes the total weight of the players preceding i in π. Given a distribution over orderings, the cost share of player i is given by

$$\mathrm{E}_\pi \left[C(f_i^\pi + w_i) - C(f_i^\pi) \right]. \tag{2}$$

A weighted Shapley value then defines a distribution over orderings by assigning to each player i a sampling parameter $\lambda_i > 0$: the last player of the ordering is picked with probability proportional to its λ_i; given this choice, the penultimate player is chosen randomly from the remaining ones, again with probability proportional to its λ_i, and so on. Below we present an example of how the values of the λ_i's lead to the distribution over orderings of the players.

Example 1. Consider players a, b, c, with sampling parameters $1, 2, 3$, respectively. The probability that a is the last in the ordering is $1/(1+2+3)$. Similarly we get $2/(1+2+3)$ for player b and $3/(1+2+3)$ for player c. Also, suppose c was chosen to be the last, then the probability that b is the second is $2/(1+2)$, while the probability that a is the second is $1/(1+2)$. This yields the following distribution over orderings: $1/3$ probability for ordering (a, b, c), $1/4$ for (a, c, b), $1/6$ for (b, a, c), $1/10$ for (b, c, a), $1/12$ for (c, a, b), and $1/15$ for (c, b, a).

Defining a weighted Shapley value reduces to choosing a λ_i value for each player i. In weighted congestion games the weight w_i of each player fully defines the impact that this player has on the social cost of the resources it uses. Hence, this is the only pertinent attribute of the player and it would be natural to assume that λ_i depends only on the value of w_i. Nevertheless, our results hold even if we allow the value of λ_i to also depend on the ID of player i, which enables treating players with the same weight w_i differently in an arbitrary fashion.

We first focus on an interesting subclass of weighted Shapley values for which λ_i is a function of w_i parameterized by a real number γ. In particular, we let the sampling parameter λ_i of each player i be $\lambda_i = \lambda(w_i) = w_i^\gamma$. Within this class, which contains all the previously studied weighted Shapley value variants, we prove that the unweighted Shapley value is optimal with respect to the PoA. We then extend this result, proving that the optimality of the unweighted Shapley value remains true even if we let λ_i be an arbitrary continuous function.

A parameterized class of weighted Shapley values. For the weighted Shapley values induced by a function of the form $\lambda(w_i) = w_i^\gamma$, we show that their PoA lies between that of the (unweighted) Shapley value, i.e., approximately $(0.9 \cdot d)^d$, and $(1.4 \cdot d)^d$. This class of λ_i values is interesting because it contains all the well-known weighted Shapley values: If $\gamma = 0$, then the induced cost-sharing rule is equivalent to the (unweighted) Shapley value. With $\gamma = -1$, we recover the most common weighted Shapley value [39]. For $\gamma \to +\infty$ (resp. $\gamma \to -\infty$) we get the order-based cost-sharing rule that introduces players to the resource from smallest to largest (resp. largest to smallest) and charges them the increase they cause to the joint cost when they are introduced. Also, as we show with Lemma 1, this class of λ_i values is natural because these are the only ones that induce scale-independent cost-sharing rules when the cost functions of the resources are homogeneous (e.g., in our setting which has $C_j(x) = \alpha_j \cdot x_j^k$).

Lemma 1. *For any homogeneous cost function, the weighted Shapley value cost-sharing rule is scale-independent if and only if* $\lambda(w) = w^\gamma$ *for some* $\gamma \in \mathbb{R}$.

The parameter γ fully determines our cost-sharing rule. Higher values of λ_i for some player i imply higher cost shares so, if $\gamma > 0$, this benefits lower weight players, and if $\gamma < 0$, this benefits higher weight players. We therefore use γ in order to parameterize the PoA that the cost-sharing rules of this class yield. The following theorem, which follows directly from Lemma 2 and Lemma 3, shows that, for any value of γ other than 0, the PoA of the induced cost-sharing rule is strictly worse. Figure 1 plots the PoA for $d = 2$.

Theorem 1. *The optimal PoA among weighted Shapley values of the form* $\lambda(w_i) = w_i^\gamma$ *is achieved for* $\gamma = 0$*, which recovers the (unweighted) Shapley value. Hence, the optimal PoA is approximately* $(0.9 \cdot d)^d$.

Before presenting our lower bounds in Lemma 2 and Lemma 3, we begin with an upper bound on the PoA of any weighted Shapley value, for polynomials with maximum degree d.

Theorem 2. *The PoA of any weighted Shapley value is at most*

$$\left(2^{\frac{1}{d}} - 1\right)^{-d} \approx (1.4 \cdot d)^d .$$

Proof. Let P be a PNE and P^* the optimal profile. We get

$$\sum_{j \in E} C_j(f_j) = \sum_{j \in E} \sum_{i \in N} \xi_j(i, S_j) = \sum_{i \in N} \sum_{j \in P_i} \xi_j(i, S_j) \leq \sum_{i \in N} \sum_{j \in P_i^*} \xi_j(i, S_j \cup \{i\}). \quad (3)$$

The inequality follows from the equilibrium condition on P. Note that, when the cost-sharing method is a weighted Shapley value and the resource costs are convex, the cost share of any player on any resource is upper bounded by the increase that would be caused to the joint resource cost if that player was be the last in the ordering. This means that for every $j \in P_i^*$ we get

$$\xi_j(i, S_j \cup \{i\}) \leq C_j(f_j + w_i) - C_j(f_j). \quad (4)$$

Combining (3) with (4), we get

$$\sum_{j \in E} C_j(f_j) \le \sum_{i \in N} \sum_{j \in P_i^*} C_j(f_j + w_i) - C_j(f_j) \tag{5}$$

$$= \sum_{j \in E} \sum_{i : j \in P_i^*} C_j(f_j + w_i) - C_j(f_j) \tag{6}$$

$$\le \sum_{j \in E} C_j(f_j + f_j^*) - C_j(f_j), \tag{7}$$

where f_j^* is the total weight on j in P^*. The last inequality follows by convexity of the expression as a function of w_i. We now claim that the following is true, for any $x, y > 0$, and $d \ge 1$:

$$(x + y)^d - x^d \le \hat{\lambda} \cdot y^d + \hat{\mu} \cdot x^d, \tag{8}$$

with

$$\hat{\lambda} = 2^{(d-1)/d} \cdot \left(2^{1/d} - 1\right)^{-(d-1)} \text{ and } \hat{\mu} = 2^{(d-1)/d} - 1. \tag{9}$$

We can verify this as follows. Note that, without loss of generality, we can set $y = 1$ (equivalent to dividing both sides of (8) with y^d and renaming x/y to q). We can then see that the value of q that maximizes $(q+1)^d - (\hat{\mu}+1) \cdot q^d$, and, hence, is the worst case for (8), is $q = 1/(2^{1/d} - 1)$, for which inequality (8) is tight. Also, note that the expressions for $\hat{\lambda}$ and $\hat{\mu}$ are increasing as functions of d, which implies that the given values for degree d, satisfy (8) for smaller degrees as well. This means we can combine (7), (8), and (9), to get

$$\sum_{j \in E} C_j(f_j) \le \sum_{j \in E} \hat{\lambda} \cdot C_j(f_j^*) + \hat{\mu} \cdot C_j(f_j). \tag{10}$$

Rearranging, we get $\sum_{j \in E} C_j(f_j) / \sum_{j \in E} C_j(f_j^*) \le \hat{\lambda}/(\hat{\mu}+1) = \left(2^{1/d} - 1\right)^{-d}$, which completes the proof. □

We now proceed with our lower bounds for $\gamma \neq 0$.

Lemma 2. *The PoA of any weighted Shapley value of the form* $\lambda(w_i) = w_i^\gamma$, *for* $\gamma > 0$ *is at least*

$$\left(2^{\frac{1}{d}} - 1\right)^{-d} \approx (1.4 \cdot d)^d .$$

Proof. Define $\rho = (2^{1/d} - 1)^{-1}$ and let T be a set of ρ/ϵ players with weight ϵ each, where $\epsilon > 0$ is an arbitrarily small parameter. Consider a player i with weight $w_i = 1$ and suppose she uses a resource j with cost function $C_j(x) = x^d$ with the players in T, for our weighted Shapley value with $\gamma > 0$. We now argue that, as we let $\epsilon \to 0$, the cost share of i in j becomes $(\rho+1)^d - \rho^d$. Consider the probability p that i is not among the last $\delta \cdot |T|$ players of the random ordering generated by our sampling weights (i.e., i is not among the first $\delta \cdot |T|$ players sampled), for some $\delta < 1$. This probability is upper bounded by the probability that i is not drawn, using our sampling weights, among everyone in $T \cup \{i\}$, $\delta \cdot |T| = \delta \cdot \rho/\epsilon$ times. Note that the sampling weight of i is 1 and the total sampling weight of the players in T is $\rho \cdot \epsilon^{\gamma-1}$. Hence, if $\gamma \geq 1$, we get:

$$p \leq \left(1 - \frac{1}{1+\rho \cdot \epsilon^{\gamma-1}}\right)^{\delta \cdot \rho/\epsilon} \leq \left(1 - \frac{1}{1+\rho}\right)^{\delta \cdot \rho/\epsilon}, \tag{11}$$

which goes to 0 as $\epsilon \to 0$. Similarly, if $\gamma < 1$, we get:

$$p \leq \left(1 - \frac{1}{1+\rho \cdot \epsilon^{\gamma-1}}\right)^{\delta \cdot \rho/\epsilon} \leq \exp\left(-\delta \cdot \frac{\rho}{\epsilon} \cdot \frac{1}{1+\rho \cdot \epsilon^{\gamma-1}}\right), \tag{12}$$

which always goes to 0 as $\epsilon \to 0$, for any arbitrarily small $\delta > 0$. Then, by letting $\delta \to 0$, our claim that the cost share of i is $(\rho+1)^d - \rho^d$ follows by the definition of the weighted Shapley value. Similarly, it follows that if a player with weight w shares a resource with cost function $a \cdot x^d$ with ρ/ϵ players with weight $w \cdot \epsilon$ each, her cost share will be $a \cdot w^d \cdot ((\rho+1)^d - \rho^d)$ (since scaling the cost function and the player weights does not change the fractions of the cost that are assigned to the players), which, for our choice of ρ is equal to $a \cdot w^d \cdot \rho^d$.

Using facts from the previous paragraph as building blocks, we construct a game such that the total cost in the worst equilibrium is ρ^d times the optimal. Suppose our resources are organized in a tree graph $G = (E, A)$, where each vertex corresponds to a resource. There is a one-to-one mapping between the set of edges of the tree, A, and the set of players of the game, N. The player i, that corresponds to edge (j, j'), has strategy set $\{\{j\}, \{j'\}\}$, i.e., she must choose one of the two endpoints of her designated edge. Tree G has branching factor ρ/ϵ and l levels, with the root positioned at level 1.

Player weight. The weight of every player (edge) between resources (vertices) at levels j and $j+1$ of the tree is ϵ^{j-1}.

Cost functions. The cost function of any resource (vertex) at level $j = 1, 2, \ldots, l-1$, is:

$$C^j(x) = \left(\frac{1}{\rho \cdot \epsilon^{d-1}}\right)^{j-1} \cdot x^d. \tag{13}$$

The cost functions of any resource (vertex) at level l is equal to:

$$C^l(x) = \frac{\rho^{d-l+1}}{\epsilon^{(d-1)\cdot(l-1)}} \cdot x^d. \tag{14}$$

Pure Nash equilibrium. Let P be the outcome that has all players play the resource closer to the root. We claim that this outcome is a PNE. The cost of every player, using a resource at level $j < l$, in P, is $(\rho/\epsilon)^{d-j}$. If one of the players that are adjacent to the leaves were to switch to her other strategy (play the leaf resource), she would incur a cost equal to $(\rho/\epsilon)^{d-l+1}$, which is the same as the one she has in P. Consider any other player and her potential deviation from the resource at level j, to the resource at level $j+1$. By the analysis in the first paragraph of this proof (she would be a player with weight ϵ^{j-1} sharing a resource with ρ/ϵ players with weight ϵ^j), her cost would be $(\rho \cdot \epsilon^{d-1})^j \cdot \epsilon^{d\cdot(j-1)} \cdot \rho^d = (\rho/\epsilon)^{d-j}$, which is her current cost in P. This proves that the equilibrium condition holds for all players in P.

Price of anarchy. As we have shown, every player using a resource at level j has cost $(\rho/\epsilon)^{d-j}$ in P. There are $(\rho/\epsilon)^j$ such players, which implies the total cost of P is $(l-1)\cdot(\rho/\epsilon)^d$, since there are $l-1$ levels of nonempty resources, and every level has the same total cost, $(\rho/\epsilon)^d$. Now, let P^* be the outcome that has all players play the resource further from the root. In this outcome, every player using a resource at level $j = 2, \ldots, l-1$, has cost $\rho^{-j+1}/\epsilon^{d-j+1}$. There are $(\rho/\epsilon)^{j-1}$ such players, hence, the total cost at level j is $(1/\epsilon)^d$. Similarly, we get that the total cost at level l is $(\rho/\epsilon)^d$. In total, the cost of P^* is $(l-2)\cdot(1/\epsilon)^d + (\rho/\epsilon)^d$. We can then see that, as $l \to +\infty$, the ratio between the cost of P and the cost of P^* becomes ρ^d. □

Lemma 3. *The PoA of any weighted Shapley value of the form* $\lambda(w_i) = w_i^\gamma$, *for* $\gamma < 0$ *is at least* d^d.

Proof. We construct a game such that the total cost in the worst equilibrium is d^d times the optimal. Suppose our resources are organized in a tree graph $G = (E, A)$, where each vertex corresponds to a resource. There is a one-to-one mapping between the set of edges of the tree, A, and the set of players of the game, N. The player i, that corresponds to edge (j, j'), has strategy set $\{\{j\}, \{j'\}\}$, i.e., she must choose one of the two endpoints of her designated edge. Tree G has branching factor $1/(d\cdot\epsilon)$, with $\epsilon > 0$ an arbitrarily small parameter, and l levels.

Player weights. The weight of every player (edge) between resources (vertices) at levels j and $j+1$ of the tree is ϵ^{j-1}.

Cost functions. The cost function of any resource (vertex) at level $j = 2, 3, \ldots, l$, is:

$$C^j(x) = \left(\frac{d}{\epsilon^{d-1}}\right)^{j-2} \cdot x^d. \tag{15}$$

The cost functions of the root is:

$$C^1(x) = x^d. \tag{16}$$

PNE. Let P be the outcome that has all players play the resource further from the root. We prove that this outcome is a PNE. The cost of every player that has played a resource at level j is $(d\cdot\epsilon)^{j-2}$. If one of the players that are adjacent to the root were to switch to her other strategy (play the root), she would incur a cost equal to 1, which is the same as the one she has in P. Consider any other player and her potential deviation from the resource at level j, to the resource at level $j-1$. Since our construction considers ϵ arbitrarily close to 0, the deviating player will go last with probability 1 in the Shapley ordering (since $\gamma < 0$) and her cost will be equal to $(d/\epsilon^{d-1})^{j-3} \cdot \left((\epsilon^{j-1} + \epsilon^{j-2})^d - \epsilon^{(j-1)\cdot d}\right) = (d \cdot \epsilon)^{j-2}$, which is equal to her current cost in P. Hence, the equilibrium condition holds for all players.

PoA. As we have shown, every player playing a resource at level j has cost $(d \cdot \epsilon)^{j-2}$ in P. There are $1/(d \cdot \epsilon)^{j-1}$ such players, hence, the total cost at level j is $1/(d \cdot \epsilon)$. Then, it follows that the total cost of P is $(l-1)/(d \cdot \epsilon)$. Now let P^* be the outcome that has all players play the resource closer to the root. Then the joint cost at the root is $1/(d \cdot \epsilon)^d$. The joint cost of every other resource at level j is $(d \cdot \epsilon)^{j-2}/d^d$, and the number of resources at level j is $1/(d \cdot \epsilon)^{j-1}$. Hence, we get in total, that the cost of P^* is $(l-2)/(d^{d+1} \cdot \epsilon) + 1/(d \cdot \epsilon)^d$. We can then see that, as $l \to +\infty$, the ratio of the cost of P to the cost of P^* becomes d^d. □

We defer details on deriving the plot in Figure 1 to our full version.

Main result. We now state the main result of the section. The proof applies arguments similar to the ones we used for Lemma 2 and Lemma 3 but on a more technical level. We defer the details to the full version.

Theorem 3. *The optimal PoA among weighted Shapley values induced by a collection of continuous functions $\lambda_i(\cdot)$ is achieved by the (unweighted) Shapley value.*

2.2 Marginal Contribution

We now focus on the marginal contribution cost-sharing rule which dictates that every player i is responsible for the marginal increase it causes to the joint resource cost. Namely, the cost share of player $i \in S_j$ on resource j is equal to $C_j(f_j) - C_j(f_j - w_i)$. Although the marginal contribution rule does induce games that always possess PNE, the total cost suffered by the players will, in general, be greater than the total cost that they generate, something that places the marginal contribution rule outside the scope of our model. To see this, note that in (1) the left hand side can be larger than the right hand side when the marginal contribution rule is used. In fact, for polynomial cost functions of maximum degree d the total cost suffered can be up to d times the generated cost. The following theorem, which is a byproduct of our previous results, provides the exact worst-case PoA of this cost-sharing rule.

Theorem 4. *The PoA of the marginal contribution rule is*

$$\left(2^{\frac{1}{d}} - 1\right)^{-d} \approx (1.4 \cdot d)^d .$$

Theorem 4, shows that the worst-case PoA of marginal contribution is equal to that of the worst weighted Shapley value (see Theorem 2). In fact, this holds even if the equilibrium cost on each resource j is measured as $\sum_{i \in S_j} \xi_j(i, S_j)$ instead of $C_j(f_j)$, i.e., even if the additional costs that the marginal contribution rule enforces are disregarded when evaluating the quality of the outcome. Measuring the PoA with respect to the cost generated by the players (instead of the cost that they actually suffer) can be motivated by thinking of the costs suffered by the players as tolls that the system uses in order to affect the incentives of the players. There has been a sequence of results focusing on designing tolls of this form in order to optimize this PoA measure for atomic congestion games [10, 17, 9, 16], and the marginal contribution rule is known as *marginal cost pricing* tolls in this literature. In this context, Theorem 4 leads to the following corollary.

Corollary 1. *There exists a weighted congestion game with marginal cost pricing tolls that has a PNE with joint cost* $(2^{1/d} - 1)^{-d}$ *times the optimal joint cost.*

3 Unrestricted Cost-Sharing Rules

In this section we consider any possible cost-sharing rule and show that the price of anarchy is always at least $\Theta(d/\ln d)^d$. In fact, our lower bound is approximately $(1.3 \cdot d/\ln d)^d$, which is also the approximate value of the PoA of proportional sharing. Before presenting this main result, we provide, as a warm-up, the proof of a weaker, but still exponential in d, lower bound for all cost-sharing rules; this simpler proof carries some of the ideas used in the more elaborate proof of Theorem 5. A key idea in the proof is to define a tournament amongst the players, with the winner of a match of players i, j corresponding to the player with the larger cost share when i and j are together. This tournament admits a Hamiltonian path [34], which we use to construct a bad example.

Proposition 1. *The PoA of any cost-sharing rule is at least* 2^{d-1}.

Proof. The structure of this proof resembles the proof structure of the main theorem of this section: we begin by partly defining the elements of the problem instance (the number of players and resources, as well as the cost functions of the resources), and then, using any given cost-sharing rule as input, we come up with a set of strategies for each player. This way, even though the cost-sharing rule may not be anonymous, we can still ensure that we "place" each players in a role such that some inefficient strategy profile is a PNE for the given cost-sharing rule.

Instance initialization: Our resources are organized as a line graph $G = (E, A)$. The edge between resources i and $i+1$ is a player that has to pick one of the two resources. All players have unit weight. Call vertex (resource) 1, which has cost $C_1(x) = x^d$, the root. As we move along the path, resources are getting better by a factor of 2^{d-1}, which means $C_i(x) = x^d/2^{(i-1)\cdot(d-1)}$. Resource n is the only exception and has the same cost function as its neighbor. For the instance to be complete, we now also need to define what the strategy set of each player is.

We are given identities of $n-1$ players and we must decide how to place them on the edges of G, i.e., determine who will be player i, the one that must choose between resources i and $i+1$.

Player placement: To determine the strategy sets of the players, we will first define a permutation π and then we will let player $\pi(i)$ be the $\pi(i)$-th edge of G. In choosing this permutation, we seek that the following property is satisfied for all $i = 1, 2, \ldots, n-1$: when players $\pi(i)$ and $\pi(i+1)$ share one of the resources in E, the cost-sharing rule[2] distributes at least half of the induced cost to player $\pi(i)$. We now show that a permutation satisfying this property always exists.

Consider a directed graph whose vertices correspond to the players of our instance and a directed edge from i to i' exists if and only if player i suffers at least half of the cost induced when it shares a resource with player i'. This is a tournament graph and hence it has at least one Hamiltonian path; starting from its first vertex and following this Hamiltonian path implies a desired permutation.

Now that we have established existence of such a permutation π, we let each player $\pi(i)$ pick between resources $\pi(i)$ and $\pi(i+1)$. Given the property that is guaranteed by the permutation, it is not hard to verify that the strategy profile P, which has total cost C, and according to which every player $\pi(i)$ chooses resource $\pi(i)$, is a PNE, while the strategy profile P^* according to which every player $\pi(i)$ chooses resource $\pi(i+1)$ is the one that achieves the optimal total cost. The corresponding total costs are

$$\sum_{j=1}^{n} 2^{-(d-1)\cdot(j-1)} \quad \text{and} \quad 2^{-(d-1)\cdot(j-1)} + \sum_{j=1}^{n-1} 2^{-(d-1)\cdot j}.$$

The ratio of the two approaches 2^{d-1} as $n \to +\infty$. □

The main result of this section strengthens this, more intuitive, lower bound further. Let

$$\Upsilon(d) = \frac{(\lfloor\phi_d\rfloor + 1)^{2\cdot(d-1)+1} - \lfloor\phi_d\rfloor^d \cdot (\lfloor\phi_d\rfloor + 2)^{d-1}}{(\lfloor\phi_d\rfloor + 1)^d - (\lfloor\phi_d\rfloor + 2)^{d-1} + (\lfloor\phi_d\rfloor + 1)^{d-1} - \lfloor\phi_d\rfloor^d},$$

where ϕ_d corresponds to the solution of $x^d = (x+1)^{d-1}$. The following theorem shows that the PoA of any cost-sharing rule that may depend on the weights and IDs of the players in an arbitrary fashion has a PoA of at least $\Upsilon(d)$, which is approximately $(1.3 \cdot d/\ln d)^d$ and, in fact, is at least $\lfloor\phi_d\rfloor^d$.

Theorem 5. *The PoA of any cost-sharing rule is at least* $\Upsilon(d)$*.*

---

[2] Even though it is not necessary for our results we assume for simplicity that if $C_{j'}(x) = \alpha \cdot C_j(x)$, then for every $i \in T$ we have $\xi_{j'}(i, T) = \alpha \cdot \xi_j(i, T)$. To see why this is not necessary note that in any game we can substitute a resource with cost function $\alpha \cdot C(x)$ with $\alpha \cdot M$ resources with cost functions $C(x)$ which can only be used as a group. Here M is a very large number and the incentive structure remains unaltered.

Proof. Our resources are organized in a graph $G = (E, A)$, which we describe below. A vertex of the tree corresponds to a resource, and an edge in A corresponds to a specific player, who must select one of the two endpoints of the edge as her strategy. All but one of the vertices of G are part of a tree as follows. At the root there is a complete d-ary tree with $k+1$ levels. The leaves of this tree are roots for complete $(d-1)$-ary trees with $k+1$ levels, and so on until the final stage with unary trees with $k+1$ levels. The final vertex of G is isolated and will have only self-loops (i.e., players that will only have this resource as their possible strategy). The main idea is that in the optimal profile P^*, all tree players move away from the root and are alone, while in the worst PNE P all tree players move towards the root and are congested. The purpose of the isolated resource is to cancel out any benefit the cost-sharing rule could extract by introducing major asymmetries on the players (using their IDs) as we will see in what follows. The construction of our lower bound is a three-stage process. During the first stage, we initialize our instance and set temporary cost functions for the resources that yield the required PoA. During the second stage we start with a very large number of players $|N|$, place a subset of them on the tree, and fix the rest on the isolated resource. This placement of players will happen in a way that will allow us to turn P into a PNE in the next stage. During the final stage we tweak the cost functions on the tree in order to maintain the fact that the profile with everyone moving towards the root is a PNE while, at the same time, we manage to keep the total costs of the PNE and the optimal profile intact.

Instance initialization: The cost function of resource j at level $(d-i) \cdot k + j$ of the tree is initialized as $(\prod_{l=i+1}^{d} l/(l+1))^{(d-1)\cdot(k-1)} \cdot (i/(i+1))^{(d-1)\cdot j} \cdot x^d$. The constant multipliers of x^d on the resources are not finalized yet and will be altered during Stage 3 of our construction. The cost function of the extra resource is constant $\delta \cdot x^d$, with δ arbitrarily small. With P and P^* as above we get a ratio of $\Upsilon(d)$ between the two total costs (see [18] for the detailed calculation).

Player placement: Suppose the game has a very large number of players $|N|$. Some of these players will be used to fill all slots of the tree and the rest will be fixed on the 0-cost resource. The players on the 0-cost resource will clearly have no impact on the PNE and optimal costs and their only purpose is to cancel out any benefit the cost-sharing rule could extract by using the player IDs in order to introduce asymmetry. Focus on a single resource of the tree. The structure of the tree clearly dictates how many players should have that resource as the top endpoint of the corresponding edge (we will call them the children players of the resource) and that one player should have that resource as the bottom endpoint of its corresponding edge (we will call this the parent player of the resource). Our claim is that given a large enough number of players, we can always find a subset of them to place on the tree such that for every vertex, the parent player covers for at least its proportional share when all children and the parent are on the resource. We will call this the *parent property*. Once we prove this fact, we will be ready to finalize the cost functions and prove our result.

Here we describe how the placement of players on the tree is performed. Let S be the set of available players which is initialized to all players N and is updated at each iteration. We will fill the edges of the tree in a bottom-up fashion. At each iteration we select a resource j that has not been assigned children players and has the maximum depth among such resources. We will select its children players from S simultaneously. Suppose j must have t children players. We must find a group T of t players, use them on these edges, and delete them from S. Then we look at each player i left in S. If i pays at least its proportional share when using j with the players in T, then i remains in S, otherwise i is assigned a unique possible strategy, which is the 0-cost resource and is removed from S. This way we ensure that no matter what our future choices are, the parent property will hold on resource j. We will pick the set T that maximizes the size of S after its placement.

The key question is how much smaller does S become at each iteration? The candidate sets of size t are $\binom{|S|}{t}$. The number of times a $t+1$-th player pays its fair share and can be a parent of a t-sized group is at least the number of groups of size $t+1$, i.e., $\binom{|S|}{t+1}$. This means there is a set T for which there exist at least $\binom{|S|}{t+1}/\binom{|S|}{t} = (|S|-t)/(t+1)$ players that can be used as its corresponding parent. So the set S is getting smaller by a factor that is a function of d at each iteration and the number of iterations is a function of d and k. This means we can take the initial $|N|$ large enough so that we can complete the process.

Cost function update: Recall P is the strategy profile such that every player on the tree picks the resource closer to the root. We want this to be a PNE. We know that, because of Stage 2, the parent property holds on every resource, hence every deviation has a player pay at least its proportional share on its new resource. We can also see that the initial cost functions are such that if a player pays its proportional share in P, then the deviation is at least tight (possibly the deviation costs more depending on what the cost-sharing rule does, but it is at least tight due to the parent property). To ensure that P is a PNE in our construction we make the following update on the cost functions. We start from the root and move towards the leaves examining every player i. If player i pays γ_i times its proportional share on its resource in P, then the cost function of the resource j it can deviate to, and the cost functions of the resources of the whole subtree rooted at j, are multiplied by γ_i. Given the facts described above with respect to the parent property and the cost functions, it follows that P is a PNE in our game.

Recall P^* is the profile such that every player picks the resource that is further from the root. At this point, all that is left to show is that the total cost of P and P^* do not change after we update the cost functions. Since, both in P and in P^*, at every level the total weight on every resource is the same, the initial multipliers of every resource cost function were identical, and the sum of the scaling factors we applied is equal to 1, it follows that the total costs remain unchanged. □

References

[1] Abed, F., Huang, C.-C.: Preemptive coordination mechanisms for unrelated machines. In: Epstein, L., Ferragina, P. (eds.) ESA 2012. LNCS, vol. 7501, pp. 12–23. Springer, Heidelberg (2012)
[2] Aland, S., Dumrauf, D., Gairing, M., Monien, B., Schoppmann, F.: Exact price of anarchy for polynomial congestion games. SIAM J. Comput. 40(5), 1211–1233 (2011)
[3] Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, É., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. SIAM Journal on Computing 38(4), 1602–1623 (2008)
[4] Awerbuch, B., Azar, Y., Epstein, L.: The price of routing unsplittable flow. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC), pp. 57–66 (2005)
[5] Azar, Y., Jain, K., Mirrokni, V.: (almost) optimal coordination mechanisms for unrelated machine scheduling. In: Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, pp. 323–332. Society for Industrial and Applied Mathematics, Philadelphia (2008)
[6] Bhawalkar, K., Gairing, M., Roughgarden, T.: Weighted congestion games: Price of anarchy, universal worst-case examples, and tightness. In: de Berg, M., Meyer, U. (eds.) ESA 2010, Part II. LNCS, vol. 6347, pp. 17–28. Springer, Heidelberg (2010)
[7] Bhawalkar, K., Roughgarden, T.: Simultaneous single-item auctions. In: Goldberg, P.W. (ed.) WINE 2012. LNCS, vol. 7695, pp. 337–349. Springer, Heidelberg (2012)
[8] Caragiannis, I.: Efficient coordination mechanisms for unrelated machine scheduling. In: Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2009, pp. 815–824. Society for Industrial and Applied Mathematics, Philadelphia (2009)
[9] Caragiannis, I., Kaklamanis, C., Kanellopoulos, P.: Improving the efficiency of load balancing games through taxes. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 374–385. Springer, Heidelberg (2008)
[10] Caragiannis, I., Kaklamanis, C., Kanellopoulos, P.: Taxes for linear atomic congestion games. ACM Transactions on Algorithms 7(1), 13 (2010)
[11] Chen, H., Roughgarden, T., Valiant, G.: Designing network protocols for good equilibria. SIAM Journal on Computing 39(5), 1799–1832 (2010)
[12] Christodoulou, G., Koutsoupias, E., Nanavati, A.: Coordination mechanisms. Theor. Comput. Sci. 410(36), 3327–3336 (2009)
[13] Christodoulou, G., Mehlhorn, K., Pyrga, E.: Improving the price of anarchy for selfish routing via coordination mechanisms. In: Demetrescu, C., Halldórsson, M.M. (eds.) ESA 2011. LNCS, vol. 6942, pp. 119–130. Springer, Heidelberg (2011)
[14] Cole, R., Correa, J.R., Gkatzelis, V., Mirrokni, V., Olver, N.: Decentralized utilitarian mechanisms for scheduling games. Games and Economic Behavior (2013)
[15] von Falkenhausen, P., Harks, T.: Optimal cost sharing for resource selection games. Math. Oper. Res. 38(1), 184–208 (2013)
[16] Fotakis, D., Karakostas, G., Kolliopoulos, S.G.: On the existence of optimal taxes for network congestion games with heterogeneous users. In: Kontogiannis, S., Koutsoupias, E., Spirakis, P.G. (eds.) SAGT 2010. LNCS, vol. 6386, pp. 162–173. Springer, Heidelberg (2010)

[17] Fotakis, D., Spirakis, P.G.: Cost-balancing tolls for atomic network congestion games. Internet Mathematics 5(4), 343–363 (2008)
[18] Gairing, M., Schoppmann, F.: Total latency in singleton congestion games. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 381–387. Springer, Heidelberg (2007)
[19] Gopalakrishnan, R., Marden, J.R., Wierman, A.: Potential games are necessary to ensure pure nash equilibria in cost sharing games. In: Proceedings of the Fourteenth ACM Conference on Electronic Commerce, EC 2013, pp. 563–564. ACM, New York (2013)
[20] Harks, T., Klimm, M.: On the existence of pure Nash equilibria in weighted congestion games. Mathematics of Operations Research 37(3), 419–436 (2012)
[21] Harks, T., Klimm, M., Möhring, R.H.: Characterizing the existence of potential functions in weighted congestion games. Theory of Computing Systems 49(1), 46–70 (2011)
[22] Harks, T., Miller, K.: The worst-case efficiency of cost sharing methods in resource allocation games. Operations Research 59(6), 1491–1503 (2011)
[23] Immorlica, N., Li, L.E., Mirrokni, V.S., Schulz, A.S.: Coordination mechanisms for selfish scheduling. Theor. Comput. Sci. 410(17), 1589–1598 (2009)
[24] Kollias, K.: Nonpreemptive coordination mechanisms for identical machines. Theory of Computing Systems 53(3), 424–440 (2013)
[25] Kollias, K., Roughgarden, T.: Restoring pure equilibria to weighted congestion games. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) ICALP 2011, Part II. LNCS, vol. 6756, pp. 539–551. Springer, Heidelberg (2011)
[26] Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. Computer Science Review 3(2), 65–69 (2009)
[27] Lucier, B., Borodin, A.: Price of anarchy for greedy auctions. In: SODA, pp. 537–553 (2010)
[28] Marden, J.R., Wierman, A.: Distributed welfare games. Operations Research 61(1), 155–168 (2013)
[29] Milchtaich, I.: Congestion games with player-specific payoff functions. Games and Economic Behavior 13(1), 111–124 (1996)
[30] Monderer, D., Shapley, L.S.: Potential games. Games and Economic Behavior 14(1), 124–143 (1996)
[31] Mosk-Aoyama, D., Roughgarden, T.: Worst-case efficiency analysis of queueing disciplines. In: Albers, S., Marchetti-Spaccamela, A., Matias, Y., Nikoletseas, S., Thomas, W. (eds.) ICALP 2009, Part II. LNCS, vol. 5556, pp. 546–557. Springer, Heidelberg (2009)
[32] Moulin, H.: The price of anarchy of serial, average and incremental cost sharing. Economic Theory 36(3), 379–405 (2008)
[33] Moulin, H., Shenker, S.: Strategyproof sharing of submodular costs: Budget balance versus efficiency. Economic Theory 18(3), 511–533 (2001)
[34] Redei, L.: Ein kombinatorischer satz. Acta Litteraria Szeged 7, 39–43 (1934)
[35] Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. International Journal of Game Theory 2(1), 65–67 (1973)
[36] Rosenthal, R.W.: The network equilibrium problem in integers. Networks 3(1), 53–59 (1973)

[37] Roughgarden, T.: Intrinsic robustness of the price of anarchy. In: 41st ACM Symposium on Theory of Computing (STOC), pp. 513–522 (2009)
[38] Roughgarden, T., Tardos, É.: How bad is selfish routing? Journal of the ACM 49(2), 236–259 (2002)
[39] Shapley, L.S.: Additive and Non-Additive Set Functions. Ph.D. thesis, Department of Mathematics, Princeton University (1953)
[40] Shenker, S.J.: Making greed work in networks: A game-theoretic analysis of switch service disciplines. IEEE/ACM Transactions on Networking 3(6), 819–831 (1995)
[41] Syrgkanis, V., Tardos, É.: Composable and efficient mechanisms. In: STOC, pp. 211–220 (2013)