Approximate Pure Nash Equilibria in Weighted Congestion Games*

Christoph Hansknecht[1], Max Klimm[1], and Alexander Skopalik[2]

1   Department of Mathematics, Technische Universität Berlin
    Straße des 17. Juni 136, 10623 Berlin, Germany
    {hansknecht,klimm}@math.tu-berlin.de
2   Department of Computer Science, University of Paderborn
    Fürstenallee 11, 33102 Paderborn, Germany
    skopalik@mail.uni-paderborn.de

── Abstract ──

We study the existence of approximate pure Nash equilibria in weighted congestion games and develop techniques to obtain approximate potential functions that prove the existence of α-approximate pure Nash equilibria and the convergence of α-improvement steps. Specifically, we show how to obtain upper bounds for approximation factor α for a given class of cost functions. For example for concave cost functions the factor is at most $3/2$, for quadratic cost functions it is at most $4/3$, and for polynomial cost functions of maximal degree ℓ it is at at most $\ell + 1$. For games with two players we obtain tight bounds which are as small as for example 1.054 in the case of quadratic cost functions.

1   Introduction

In many applications the state of a system depends on the behavior of individual participants that act selfishly in order to minimize their own private cost measured by individual objective functions. The framework of *non-cooperative games* has enveloped as the primary tool for the theoretical analysis of such systems. The central concept of game theory is that of a Nash equilibrium – a state in which no participant has an incentive to deviate to another strategy. While mixed Nash equilibria, i.e., Nash equilibria in randomized strategies, are guaranteed to exist under mild assumptions on the players' strategy spaces and the private cost functions (cf. Nash [17], Glicksberg [12]), they are often hard to interpret. As a consequence, attention is often restricted to pure Nash equilibria, i.e., Nash equilibria in deterministic strategies.

Rosenthal [19] introduced a rich class of games, called *congestion games* that models a wealth of strategic interactions and is guaranteed to have pure Nash equilibria. In a congestion game, we are given a finite set of players and a finite set of resources. A strategy

of each player is to choose a subset of the resources out of a set of subsets of resources allowable to her. In each strategy profile, each player pays for all used resources where the cost of a resource is a function of the number of players using it. In most applications, the set of resources corresponds to the set of edges of a directed or undirected graph and cost functions are used to model latencies or travel times that typically increase with congestion. In that way, congestion games can be used to model traffic in road networks and Internet routing applications, where streams of packets choose a chain of servers from their origin to their destination.

Unfortunately, several mild generalizations of congestion games may lack pure Nash equilibria. In a *weighted* congestion game, each player is associated with a positive *demand* and the cost of each resource depends on the aggregated demand rather than the mere cardinality of the set of its users. It is well known that weighted congestion games may fail to have a pure Nash equilibrium; examples of such games have been given by Libman and Orda [15], Goemans et al. [13], and Fotakis et al. [11]. The games of Fotakis et al. and Goemans et al. feature two players with demands one and two, respectively. While Fotakis et al. specify for each resource explicitly the cost for all possible aggregated demands, Goemans et al. use only polynomial cost functions with non-negative coefficients and maximal degree two. Full enumeration of all strategy profiles shows that no pure Nash equilibrium exists. Quite strikingly, however, both examples admit deterministic states that are *almost* stable, i.e., there are strategy profiles from which both players may only improve by a small factor. In the game of Fotakis et al. (with no structure on the cost functions), this factor is as little as 12/11, in the game of Goemans et al., this factor is 61/60. Put differently, if there is some friction in the system that prevents players from making deviations that improve their private costs only very little, then stable states exist.

Such an approximate stability is formally captured by the concept of an α-approximate pure Nash equilibrium, a state from which no player can improve her private cost by a factor of $\alpha \geq 1$. Besides mere existence, approximate equilibria are an appealing alternative solution concept from a computational point of view. While the computation of exact pure Nash equilibria in congestion games is PLS-complete, there has been some recent progress towards polynomial time algorithms to compute approximate equilibria in congestion games. Specifically, Caragiannis et al. [4] show how to compute a $2 + \epsilon$-approximate pure Nash equilibrium in congestion games with affine latencies. Subsequent work generalizes this approach to a polynomial algorithm for approximate pure Nash equilibria with constant approximation factor for weighted congestion games with polynomial latency functions [5]. They also show that weighted congestion games with polynomial cost functions with maximal degree ℓ have a $\ell!$-approximate pure Nash equilibrium.

Still, approximate pure Nash equilibria are only a reasonable concept if the approximation factor is sufficiently close to one. This motivates the main question of this paper: Given a set of cost functions, what is the minimal approximation factor α that one can allow in order to guarantee the existence of an α-approximate pure Nash equilibrium in all weighted congestion games?

1.1 Our Contribution

The main tool to answer this question are *approximate potential functions*. An α-approximate potential is a map from the space of all strategy profiles to the real numbers that has the property that it decreases if a player decreases his cost by a factor greater than α. Note that, unlike an exact potential function, for improvement steps of smaller relative size, an approximate potential function may actually increase. The existence of α-approximate

■ **Table 1** Approximation factors of approximate pure Nash equilibria in weighted congestion games for different sets of cost functions; results for two-player games are tight. Comparison with previous results [5] shows improvements by a factor exponential in ℓ.

| | ——— Our results ——— | | Previous work [5] |
Functions	≥ 3 players	2 players	≥ 3 players
concave	$\leq 3/2$		
polynomials of degree 2	$\leq 4/3$	≈ 1.054	≤ 2
— " — 3	≤ 1.785	≈ 1.074	≤ 6
— " — 4	≤ 2.326	≈ 1.153	≤ 24
— " — ℓ	$\leq \ell + 1$		$\leq \ell!$

potential functions immediately implies the existence of α-approximate equilibria and the convergence of α-improvement steps. We present two methods to obtain an α-approximate potential function and identify upper bounds for the value of α for a given class of cost functions.

Our technique yields small approximation factors for specific classes of cost functions summarized in Table 1. For concave cost functions we establish the existence of $\frac{3}{2}$-approximate equilibria. For quadratic cost functions the factor is at most $\frac{4}{3}$ in games with an arbitrary number of players. More surprisingly, in games with two players, we obtain a tight bound of about 1.054 using numerical methods. This shows that the factor of $61/60 \approx 1.017$ achieved by the two-player game of Goemans et al. is not so far from the worst-case bound for arbitrary two-player games with quadratic costs.

For polynomial cost functions of maximal degree ℓ, we obtain an upper bound of $\ell + 1$ which is a drastic improvement of the previously known bound of $\ell!$.

Our improved bounds on the minimal approximation factors for pure Nash equilibria may be used to design routing protocols with convergent behavior. While it is known that routing with distance vector computation causes flapping, our results suggest that for routers with quadratic latencies, routes should not be updated as long as the new route does not improve latency by a factor of at least 4/3.

1.2   Further Related Work

Rosenthal [19] proved that every congestion game has a pure Nash equilibrium using an elegant *potential function* argument. A potential function assigns a real value to each strategy profile such that for two profiles which differ only in the strategy choice of one player the cost difference for that player equals the difference of the two potential function values. This property implies that any sequence of improvement steps by single players converges to a pure Nash equilibrium. However, such a sequence might take exponentially many steps and computing a pure Nash equilibrium is a computationally hard task as it is PLS-hard [1, 3, 10].

In contrast to the original class of unweighted congestion games studied by Rosenthal, many natural generalizations do neither have a potential function nor a pure Nash equilibrium, in general. Milchtaich [16] introduced weighted congestion games and congestion games with player-specific cost functions. He restricts himself to the singleton case, where each strategy of each player contains a single resource only and showed that games with player-specific cost always have a pure Nash equilibrium if players are unweighted.

Fotakis et al. [11] study congestion games with weighted players and arbitrary strategy spaces. They show that pure Nash equilibria need not exist in general. This has been observed independently by Goemans et al. [13] and Libman and Orda [15]. The computation problem to decide whether a pure Nash equilibrium exists is NP-hard [9]. For the special case of only affine [11] or exponential [18] cost functions the existence of a pure Nash equilibrium is guaranteed. These results are complemented by the characterization of Harks and Klimm [14] who prove that these are the only cost functions that guarantee the existence. These results are entirely independent of the underlying structure of the games. In contrast, Ackermann et al. [2] consider weighted congestion games with arbitrary cost functions but restrictions on the combinatorial structure of strategy spaces. They prove existence of pure Nash equilibria for weighted congestion games that have the *matroid property*, i.e. the property that the set of possible strategies for different players forms a matroid.

In light of the negative results regarding existence and complexity of pure Nash equilibria, attention turned towards approximate equilibria. For symmetric and unweighted congestion games, Chien and Sinclair [7] showed fast convergence to approximate equilibria under some mild assumption on the cost functions. However, Skopalik and Vöcking [20] proved that in the asymmetric case it is PLS-hard to compute an α-approximate equilibrium for any polynomial time computable α. As it turns out, this seemingly devastating result relies on the use of cost functions that allow negative coefficients. Indeed Caragiannis et al. [4] presented a polynomial time algorithm to compute approximate equilibria in unweighted congestion games with linear and polynomial cost functions without negative coefficients. In subsequent work [5], the same authors study the existence and complexity of approximate equilibria in weighted congestion games. In particular they show that a game with polynomial delay functions of degree at most d, a $d!$-approximate equilibrium always exists. They introduce a new class of games called Ψ-games. A weighted congestion game is approximated by a corresponding Ψ-game. That is, the cost a player in a weighted congestion game is approximates by her cost in the corresponding Ψ-game up factor of $d!$. These Ψ-games are potential games which immediately proves the existence of $d!$-approximate equilibria. Using a similar algorithm as in [4] one can compute $d^{d+o(1)}$-approximate equilibria in games with polynomial delay functions and $\frac{3+\sqrt{5}}{2}$-approximate equilibria in the case of linear weighted congestion games.

Chen and Roughgarden [6] studied approximate equilibria in network design games with weighted players and showed existence of approximate equilibria using approximate potential functions. They were also used by Christodoulou et al. [8] in order to derive tight bounds on the price of anarchy and price of stability of approximate pure Nash equilibria in unweighted congestion games.

2   Preliminaries

We consider finite strategic games $G = (N, \mathbf{S}, \pi)$, where $N = \{1, \ldots, n\}$ is the non-empty and finite set of *players*, S_i is the set of *strategies* available to player i, $\mathbf{S} = S_1 \times \ldots \times S_n$ is the finite and non-empty set of *strategy profiles*, $\pi_i : \mathbf{S} \to \mathbb{R}^n$ is the *private cost function* player i strives to minimize, and $\pi : \mathbf{S} \to \mathbb{R}^n$, $\mathbf{s} \mapsto \pi_1(\mathbf{s}) \times \cdots \times \pi_n(\mathbf{s})$ is the *combined private cost function*.

Vectors of sets and vectors of real numbers are denoted with bold face. We use standard game theory notation, i.e., for a player i and a strategy profile $\mathbf{s}$, we write $\mathbf{s} = (s_i, \mathbf{s}_{-i})$ meaning that $s_i \in S_i$ and $\mathbf{s}_{-i} \in \mathbf{S}_{-i} = S_1 \times \cdots \times S_{i-1} \times S_{i+1} \times \cdots \times S_n$. For $\alpha \geq 1$, a strategy profile $\mathbf{s}$ is an α-approximate Nash equilibrium, if $\pi_i(\mathbf{s}) \leq \alpha \cdot \pi_i(t_i, \mathbf{s}_{-i})$ for all $i \in N$ and $t_i \in S_i$. For $\alpha = 1$, we call $\mathbf{s}$ a Nash equilibrium rather than a 1-approximate Nash equilibrium.

In a *weighted congestion game*, we are given a demand vector $\mathbf{d} = (d_1, \ldots, d_n) \in \mathbb{R}^n_{>0}$ specifying a positive demand d_i for each player i. The set of strategies of player i is a non-empty set $S_i \subseteq 2^R$ of subsets of a given set of resources R. Given a strategy profile $\mathbf{s} \in \mathbf{S}$, we denote by $N_r(\mathbf{s}) = \{i \in N : r \in s_i\}$ the set of players that use r in $\mathbf{s}$. The aggregated demand for resource r is denoted by $|\mathbf{d}_{N_r(\mathbf{s})}|$. Each resource is endowed with a cost function $c_r : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ that maps the aggregated demand $|\mathbf{d}_{N_r(\mathbf{s})}|$ to a cost value $c_r(|\mathbf{d}_{N_r(\mathbf{s})}|)$. The private cost of player i is then defined as $\pi_i(\mathbf{s}) = \sum_{r \in s_i} c_r(|\mathbf{d}_{N_r(\mathbf{s})}|)$.

A weighted congestion game in which all players have unit demand $d_i = 1$ is called an *unweighted congestion game*. Rosenthal proved that every unweighted congestion game has a pure Nash equilibrium by giving a potential function. For $\alpha \geq 1$, an α-approximate *potential function* is a map $P : \mathbf{S} \mapsto \mathbb{R}$ with the property that $P(t_i, \mathbf{s}_{-i}) < P(\mathbf{s})$ whenever $\alpha \pi_i(t_i, \mathbf{s}_{-i}) < \pi_i(\mathbf{s})$ for some $i \in N$, $t_i \in S_i$, and $\mathbf{s} \in \mathbf{S}$. In case $\alpha = 1$, we call P a potential function. Rosenthal [19] showed that $P : \mathbf{S} \to \mathbb{R}$ defined as $P(\mathbf{s}) = \sum_{r \in R} \sum_{x=0}^{|\mathbf{d}_{N_r(s)}|} c_r(x)$ is a potential function for unweighted congestion games.

3   Existence of an Approximate Potential Function

In this section, we show that weighted congestion games admit approximate potential functions with low approximation factor. Roughly speaking, we obtain an approximate potential as follows. For each resource, we choose an appropriate ordering of the players. Then, for each resource separately, we compute a *discrete integral*: We sum up the resource cost after introducing the first player multiplied with the first players' demand, the resource cost after introducing the first two players multiplied with the second players' demand, and so on. When the demands of all players go to zero while keeping the total demand of all players constant this value approaches the integral from zero to the total demand of the cost function, hence the name.

To abstract from the underlying set of players, we define the discrete integral of a function with respect to an arbitrary vector $\mathbf{v} \in \mathbb{R}^n_{>0}$.

▶ **Definition 1** (Discrete integral). Let $\mathbf{v} \in \mathbb{R}^n_{>0}$ and let $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. The $\mathbf{v}$-discrete integral is defined as $I_f(\mathbf{v}) = \sum_{i=1}^{n} v_i f(\sum_{j=1}^{i} v_j)$.

To obtain provably low approximation factors, we are interested in permutations of the vector $\mathbf{v}$ that minimize or maximize the value of the discrete integral among all permutations of the vector. For a vector $\mathbf{v} \in \mathbb{R}^n_{>0}$, and a function f we let denote $\hat{\sigma}_f(\mathbf{v})$ the vector in $\mathbb{R}^n_{>0}$, that contains the entries of $\mathbf{v} = (v_1, \ldots, v_n)$ in an order that minimizes the resulting discrete integral, i.e., $I_f(\hat{\sigma}_f(\mathbf{v})) \leq I_f((v_{p(1)}, \ldots, v_{p(n)}))$ for all permutations $p \in \Pi$. With a slight abuse of terminology, we call $\hat{\sigma}_f(\mathbf{v})$ a *permutation* of $\mathbf{v}$. If several permutations of $\mathbf{v}$ achieve the same value for the discrete integral, we assume that ties are broken according to an arbitrary, but fixed tie-breaking rule. Similarly, we denote by $\check{\sigma}_f(\mathbf{v})$ the permutation of $\mathbf{v}$ that maximizes the resulting value of the discrete integral, i.e., $I_f(\check{\sigma}_f(\mathbf{v})) \geq I_f((v_{p(1)}, \ldots, v_{p(n)}))$ for all permutations $p \in \Pi$. We sometimes call $\hat{\sigma}_f(\mathbf{v})$ and $\check{\sigma}_f(\mathbf{v})$ simply *minimizing* and *maximizing* permutations of $\mathbf{v}$ when the underlying function f is clear from the context.

Given a minimizing or maximizing permutation $\hat{\sigma}_f(\mathbf{v})$ or $\check{\sigma}_f(\mathbf{v})$ of a vector $\mathbf{v} = (v_1, \ldots, v_n)$, we are interested in the (relative) error in the minimization or maximization of the discrete integral when moving a given entry v_i to the end of the permutation. Formally, let $\mathbf{v} \in \mathbb{R}^n_{>0}$ and let $p \in \Pi$ be such that $\hat{\sigma}_f(\mathbf{v}) = (v_{p(1)}, \ldots, v_{p(n)})$. Further, let $i, j \in \{1, \ldots, n\}$ be such that $v_i = v_{p(j)}$. Then, we denote by $\hat{\sigma}^i_f(\mathbf{v})$ the vector that arises from $\hat{\sigma}_f(\mathbf{v})$ by moving entry $v_{p(j)}$ to the end of the vector, i.e., $\hat{\sigma}^i_f(\mathbf{v}) = (v_{p(1)}, \ldots, v_{p(j-1)}, v_{p(j+1)}, \ldots, v_{p(n)}, v_{p(j)})$.

Analogously, $\check{\sigma}_f^i(\mathbf{v})$ is the vector that arises from the maximizing permutation $\check{\sigma}_f(\mathbf{v})$ by moving the entry corresponding to v_i to the end of the vector.

The approximation factors for approximate equilibria that we are going to obtain depend on the following relative errors in the minimization respectively maximization of the discrete integral after moving entry v_i to the end of the vector.

▶ **Definition 2.** For a function $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, define

$$\hat{\mu}(f) = \sup_{\mathbf{v} \in \mathbb{R}_{>0}^n} \max_{i \in \{1,\dots,n\}} \frac{I_f(\hat{\sigma}_f^i(\mathbf{v})) - I_f(\hat{\sigma}_f(\mathbf{v}))}{v_i f(|\mathbf{v}|)},$$

$$\check{\mu}(f) = \sup_{\mathbf{v} \in \mathbb{R}_{>0}^n} \max_{i \in \{1,\dots,n\}} \frac{I_f(\check{\sigma}_f(\mathbf{v})) - I_f(\check{\sigma}_f^i(\mathbf{v}))}{v_i f(|\mathbf{v}|)}.$$

For a set $\mathcal{C}$ of functions, we set $\check{\mu}(\mathcal{C}) = \sup_{f \in \mathcal{C}} \check{\mu}(f)$ and $\hat{\mu}(\mathcal{C}) = \sup_{f \in \mathcal{C}} \hat{\mu}(f)$, respectively.

The following theorem relates the existence of an α-approximate pure Nash equilibrium in a weighted congestion game with cost functions in $\mathcal{C}$ with the value $\check{\mu}(\mathcal{C})$.

▶ **Theorem 3.** *Every weighted congestion game with cost functions in $\mathcal{C}$ has an α-approximate Nash equilibrium with $\alpha = 1 + \check{\mu}(\mathcal{C})$.*

Proof. To prove the result, we show that the function $P : \mathbf{S} \to \mathbb{R}$ defined as $P(\mathbf{s}) = \sum_{r \in R} I_{c_r}(\check{\sigma}_{c_r}(\mathbf{d}_{N_r(\mathbf{s})}))$ is an α-approximate potential function.

Consider an arbitrary strategy profile $\mathbf{s} \in S$ and let $i \in N$ and $\mathbf{t} \in S$ be such that $\mathbf{t} = (t_i, \mathbf{s}_{-i})$ for some $t_i \in S_i$ with $\alpha \cdot \pi_i(t_i, \mathbf{s}_{-i}) < \pi_i(\mathbf{s})$. We calculate

$$P(\mathbf{t}) - P(\mathbf{s}) = \sum_{r \in R} \Big(I_{c_r}(\check{\sigma}_{c_r}(\mathbf{d}_{N_r(\mathbf{t})})) - I_{c_r}(\check{\sigma}_{c_r}(\mathbf{d}_{N_r(\mathbf{s})})) \Big)$$

$$= \sum_{r \in R} \Big(I_{c_r}(\check{\sigma}_{c_r}^i(\mathbf{d}_{N_r(\mathbf{t})})) - I_{c_r}(\check{\sigma}_{c_r}^i(\mathbf{d}_{N_r(\mathbf{s})})) \Big) \tag{1a}$$

$$+ \sum_{r \in R} \Big(I_{c_r}(\check{\sigma}_{c_r}(\mathbf{d}_{N_r(\mathbf{t})})) - I_{c_r}(\check{\sigma}_{c_r}^i(\mathbf{d}_{N_r(\mathbf{t})})) \Big) \tag{1b}$$

$$+ \sum_{r \in R} \Big(I_{c_r}(\check{\sigma}_{c_r}^i(\mathbf{d}_{N_r(\mathbf{s})}))) - I_{c_r}(\check{\sigma}_{c_r}(\mathbf{d}_{N_r(\mathbf{s})})) \Big). \tag{1c}$$

We proceed to bound the expressions (1a), (1b), and (1c) separately, starting with (1a). Clearly, for all resources $r \in R \setminus (s_i \triangle t_i)$, the discrete integrals in (1a) cancel out, so we only have to consider the resources in $s_i \setminus t_i$ and $t_i \setminus s_i$. By definition, for all resources $r \in s_i \setminus t_i$, the demand d_i appears last in $\check{\sigma}_{c_r}^i(\mathbf{d}_{N_r(\mathbf{s})})$ but not at all in $\check{\sigma}_{c_r}^i(\mathbf{d}_{N_r(\mathbf{t})})$ and, thus, $I_{c_r}(\check{\sigma}_{c_r}^i(\mathbf{d}_{N_r(\mathbf{t})})) - I_{c_r}(\check{\sigma}_{c_r}^i(\mathbf{d}_{N_r(\mathbf{s})})) = -d_i c_r(|\mathbf{d}_{N_r(\mathbf{s})}|)$. Analogously, for a resource $r \in t_i \setminus s_i$, the demand d_i appears last in $\check{\sigma}_{c_r}^i(\mathbf{d}_{N_r(\mathbf{t})})$, but not at all in $\check{\sigma}_{c_r}^i(\mathbf{d}_{N_r(\mathbf{s})})$ and we obtain $I_{c_r}(\check{\sigma}_{c_r}^i(\mathbf{d}_{N_r(\mathbf{t})})) - I_{c_r}(\check{\sigma}_{c_r}^i(\mathbf{d}_{N_r(\mathbf{s})})) = d_i c_r(|\mathbf{d}_{N_r(\mathbf{t})}|)$. Thus, we may rewrite (1a) as

$$\sum_{r \in R} \Big(I_{c_r}(\check{\sigma}_{c_r}^i(\mathbf{d}_{N_r(\mathbf{t})})) - I_{c_r}(\check{\sigma}_{c_r}^i(\mathbf{d}_{N_r(\mathbf{s})})) \Big) = d_i \big(\pi_i(\mathbf{t}) - \pi_i(\mathbf{s}) \big).$$

Next, consider the expression (1b). Using the definition of $\check{\mu}(f)$, we bound (1b) from above by

$$\sum_{r \in R} \Big(I_{c_r}(\check{\sigma}_{c_r}(\mathbf{d}_{N_r(\mathbf{t})}))) - I_{c_r}(\check{\sigma}_{c_r}^i(\mathbf{d}_{N_r(\mathbf{t})})) \Big) \leq \sum_{r \in t_i} \check{\mu}_{c_r} \cdot d_i c_r(|\mathbf{d}_{N_r(\mathbf{t})}|) \leq \check{\mu}(\mathcal{C}) \cdot d_i \, \pi_i(\mathbf{t}).$$

Finally, for expression (1c), by definition of $\check{\sigma}_{c_r}$ we have $I_{c_r}\big(\check{\sigma}^i_{c_r}(\mathbf{d}_{N_r(\mathbf{s})})\big)) < I_{c_r}\big(\check{\sigma}_{c_r}(\mathbf{d}_{N_r(\mathbf{s})})\big)$ for all $r \in R$ and, thus,

$$\sum_{r \in R}\Big(I_{c_r}(\check{\sigma}^i_{c_r}(\mathbf{d}_{N_r(\mathbf{s})}))) - I_{c_r}(\check{\sigma}_{c_r}(\mathbf{d}_{N_r(\mathbf{s})}))\Big) < 0. \tag{2a}$$

Plugging everything together, we obtain

$$P(\mathbf{t}) - P(\mathbf{s}) \le d_i\big(\pi_i(\mathbf{t}) - \pi_i(\mathbf{s})\big) + \check{\mu}(\mathcal{C}) \cdot d_i\,\pi_i(\mathbf{t}) = d_i\big(\alpha\pi_i(\mathbf{t}) - \pi_i(\mathbf{s})\big) < 0.$$

We conclude that P is an α-approximate potential. As the set of strategy profiles is finite, P attains its minimum on $\mathbf{S}$ which is an α-approximate Nash equilibrium. ◀

We obtain a similar bound by choosing for each resource an order of the players' demands that *minimizes* the resulting discrete integral.

▶ **Theorem 4.** *Every weighted congestion game with cost functions in $\mathcal{C}$ and $\hat{\mu}(\mathcal{C}) < 1$ has an α-approximate Nash equilibrium with $\alpha = \big(1 - \hat{\mu}(\mathcal{C})\big)^{-1}$.*

Proof. Let $\alpha = \frac{1}{1-\hat{\mu}(\mathcal{C})}$. We consider the function $P : \mathbf{S} \to \mathbb{R}$ defined as $P(\mathbf{s}) = \sum_{r \in R} I_{c_r}(\hat{\sigma}_{c_r}(\mathbf{d}_{N_r(\mathbf{s})}))$ and show that it is an α-approximate potential function. To this end, let us again consider an arbitrary strategy profile $\mathbf{s} \in S$ and let $i \in N$ and $\mathbf{t} \in S$ be such that $\mathbf{t} = (t_i, \mathbf{s}_{-i})$ for some $t_i \in S_i$ with $\alpha \cdot \pi_i(t_i, \mathbf{s}_{-i}) \le \pi_i(\mathbf{s})$. We calculate

$$P(\mathbf{s}) - P(\mathbf{t}) = \sum_{r \in R}\Big(I_{c_r}(\hat{\sigma}_{c_r}(\mathbf{d}_{N_r(\mathbf{t})})) - I_{c_r}(\hat{\sigma}_{c_r}(\mathbf{d}_{N_r(\mathbf{s})}))\Big)$$

$$= \sum_{r \in R}\Big(I_{c_r}(\hat{\sigma}^i_{c_r}(\mathbf{d}_{N_r(\mathbf{t})})) - I_{c_r}(\hat{\sigma}^i_{c_r}(\mathbf{d}_{N_r(\mathbf{s})}))\Big) \tag{3a}$$

$$+ \sum_{r \in R}\Big(I_{c_r}(\hat{\sigma}_{c_r}(\mathbf{d}_{N_r(\mathbf{t})})) - I_{c_r}(\hat{\sigma}^i_{c_r}(\mathbf{d}_{N_r(\mathbf{t})}))\Big) \tag{3b}$$

$$+ \sum_{r \in R}\Big(I_{c_r}(\hat{\sigma}^i_{c_r}(\mathbf{d}_{N_r(\mathbf{s})}))) - I_{c_r}(\hat{\sigma}_{c_r}(\mathbf{d}_{N_r(\mathbf{s})}))\Big). \tag{3c}$$

We proceed along the same lines as in the proof of Theorem 3. Note that this time the expression (2a) does not exceed 0 as we chose an order that minimizes the discrete integral. Using the definition of $\hat{\mu}(\mathcal{C})$, we bound the expression in(3c) from above by $\hat{\mu}(\mathcal{C}) \cdot d_i\,\pi_i(\mathbf{s})$. Plugging everything together, we obtain

$$P(\mathbf{t}) - P(\mathbf{s}) \le d_i\big(\pi_i(\mathbf{t}) - \pi_i(\mathbf{s})\big) + \hat{\mu}(\mathcal{C}) \cdot d_i\,\pi_i(\mathbf{s})$$
$$= d_i\big(\pi_i(\mathbf{t}) - (1 - \hat{\mu}(\mathcal{C}))\pi_i(\mathbf{s})\big)$$
$$= \frac{d_i}{\alpha}\big(\alpha\pi_i(\mathbf{t}) - \pi_i(\mathbf{s})\big) < 0.$$

We conclude that P is an α-approximate potential. ◀

4   Bounding the Approximation Factor

To obtain explicit numerical bounds on the approximation factor of the approximate Nash equilibria, we want to compute $\check{\mu}(\mathcal{C})$ and $\hat{\mu}(\mathcal{C})$ for interesting sets $\mathcal{C}$ of cost functions. As a first result in this direction, we show that for a convex function f and a vector $\mathbf{v} \in \mathbb{R}^n_{>0}$, the discrete integral is maximized when $\mathbf{v}$ is sorted in non-decreasing order.

Formally, for a vector $\mathbf{v} \in \mathbb{R}^n_{>0}$, let us denote by $\overleftarrow{\sigma}(\mathbf{v}) = (v_{p(1)}, \ldots, v_{p(n)})$, $p \in \Pi$ the permutation of $\mathbf{v} = (v_1, \ldots, v_n)$ that contains the entries of $\mathbf{v}$ in non-increasing order, i.e, $v_{p(1)} \geq v_{p(2)} \geq \cdots \geq v_{p(n)}$. Equivalently, we denote by $\vec{\mathbf{v}}$ the permutation of $\mathbf{v}$ that contains the entries of $\mathbf{v}$ in non-decreasing order.

▶ **Lemma 5** (Sorting Lemma). *For all $\mathbf{v} \in \mathbb{R}^n_{>0}$ and $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ the following statements hold:*
1. *If f is convex, then $I_f(\hat{\sigma}_f(\mathbf{v})) = I_f(\overleftarrow{\sigma}(\mathbf{v}))$ and $I_f(\check{\sigma}_f(\mathbf{v})) = I_f(\vec{\sigma}(\mathbf{v}))$.*
2. *If f is concave, then $I_f(\hat{\sigma}_f(\mathbf{v})) = I_f(\vec{\sigma}(\mathbf{v}))$ and $I_f(\check{\sigma}(\mathbf{v})) = I_f(\overleftarrow{\sigma}(\mathbf{v}))$.*

Proof. We start proving *1*. Let $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be a convex function and let $\mathbf{v} \in \mathbb{R}^n_{>0}$ be arbitrary. It suffices to prove $I_f(\overleftarrow{\sigma}(\mathbf{v})) \leq I_f(\hat{\sigma}_f(\mathbf{v}))$. Let $\mathbf{w} = (w_1, \ldots, w_n) = \hat{\sigma}_f(\mathbf{v})$. If $\mathbf{w} = (w_1, w_2, \ldots, w_n)$ is already non-increasing, there is nothing left to show. Otherwise, there is $i \in \{1, \ldots, n-1\}$ with $w_i < w_{i+1}$. Consider the vector $\mathbf{w}'$ that is obtained from $\mathbf{w}$ by swapping entries w_i and w_{i+1}, i.e. $\mathbf{w}' = (w'_1, w'_2, \ldots, w'_n)$ where

$$w'_j = \begin{cases} w_j, & \text{if } j \notin \{i, i+1\}, \\ w_{i+1}, & \text{if } j = i, \\ w_i, & \text{if } j = i+1. \end{cases}$$

We claim that $I_f(\mathbf{w}') \leq I_f(\mathbf{w})$. To see this claim, let us denote by $x = \sum_{j=1}^{i-1} w_j$ the sum of the entries of $\mathbf{w}$ with index smaller than i. As f is convex, we have $f(\lambda a + (1-\lambda)b) \leq \lambda f(a) + (1-\lambda)f(b)$ for all $\lambda \in (0,1)$, $a, b \in \mathbb{R}_{\geq 0}$. For $\lambda = w_i/w_{i+1}$, $a = x + w_i$, and $b = x + w_i + w_{i+1}$, we obtain in particular

$$f(x + w_{i+1}) \leq \frac{w_i}{w_{i+1}} f(x + w_i) + \left(1 - \frac{w_i}{w_{i+1}}\right) f(x + w_i + w_{i+1}) \tag{4a}$$

which implies

$$w_{i+1} f(x + w_{i+1}) + v_i f(x + w_i + w_{i+1}) \leq w_i f(x + w_i) + w_{i+1} f(x + w_i + w_{i+1}). \tag{4b}$$

We derive that $I_f(\mathbf{w}') \leq I_f(\mathbf{w})$. By iteratively switching entries w_j, and w_{j+1} with $w_j < w_{j+1}$ we transform $\mathbf{w}$ into $\overleftarrow{\sigma}(\mathbf{v})$ without increasing the value of the discrete integral. This proves $I_f(\overleftarrow{\sigma}(\mathbf{v})) \leq I_f(\mathbf{w})$.

For concave functions, all inequality signs in (4a) and (4b) reverse and *2.* is obtained along the same lines. ◀

We proceed to provide a series of useful lemmas that help to identify the structure of the worst case vectors $\mathbf{v}$ that get arbitrarily close to the suprema

$$\hat{\mu}(f) = \sup_{\mathbf{v} \in \mathbb{R}^n_{>0}} \max_{i \in \{1, \ldots, n\}} \left(I_f(\hat{\sigma}^i_f(\mathbf{v})) - I_f(\hat{\sigma}_f(\mathbf{v}))\right) / \left(v_i f(|\mathbf{v}|)\right), \qquad \text{and}$$

$$\check{\mu}(f) = \sup_{\mathbf{v} \in \mathbb{R}^n_{>0}} \max_{i \in \{1, \ldots, n\}} \left(I_f(\check{\sigma}_f(\mathbf{v})) - I_f(\check{\sigma}^i_f(\mathbf{v}))\right) / \left(v_i f(|\mathbf{v}|)\right),$$

respectively.

Roughly speaking, in the next lemma, we show that for a convex function f and a vector $\mathbf{v} = (v_1, \ldots, v_i, \ldots, v_j, v_{j+1}, \ldots, v_n)$ the value $I_f(\check{\sigma}_f(\mathbf{v})) - I_f(\check{\sigma}^i_f(\mathbf{v}))$ cannot decrease when merging the two entries v_j and v_{j+1} to a single entry $v_j + v_{j+1}$. For a concave function, the same statement holds for $I_f(\hat{\sigma}^i_f(\mathbf{v})) - I_f(\hat{\sigma}_f(\mathbf{v}))$. To handle the convex and the concave case simultaneously, the statement of the lemma is slightly more general and holds for arbitrary vectors $\mathbf{v}$ rather than only those that are minimizing or maximizing.

▶ **Lemma 6** (Union Lemma). *Let $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, and let $\mathbf{v}, \mathbf{v}^i \in \mathbb{R}^n_{>0}$, $\tilde{\mathbf{v}}, \tilde{\mathbf{v}}^i \in \mathbb{R}^{n-1}_{>0}$ be such that*

$$\mathbf{v} = (v_1, \ldots, v_i, \ldots, v_j, v_{j+1} \ldots, v_n), \qquad \mathbf{v}^i = (v_1, \ldots, v_j, v_{j+1}, \ldots, v_n, v_i),$$
$$\tilde{\mathbf{v}} = (v_1, \ldots, v_i, \ldots, v_j + v_{j+1}, \ldots, v_n), \qquad \tilde{\mathbf{v}}^i = (v_1, \ldots, v_j + v_{j+1}, \ldots, v_n, v_i)$$

i.e., $\tilde{\mathbf{v}}$ is obtained from $\mathbf{v}$ by joining entry v_j and v_{j+1}, $j > i$ into one entry $v_j + v_{j+1}$ and $\mathbf{v}^i$ and $\tilde{\mathbf{v}}^i$ are obtained from $\mathbf{v}$ and $\tilde{\mathbf{v}}$, respectively, by moving entry v_i to the end of the vector. Then, $I_f(\mathbf{v}^i) - I_f(\mathbf{v}) - \big(I_f(\tilde{\mathbf{v}}^i) - I_f(\tilde{\mathbf{v}})\big)$ is non-negative, if f is convex and non-positive, if f is concave.

Proof. Let $x = \sum_{k=1}^{j-1} v_j$ denote the sum of the entries that appear before v_j in $\mathbf{v}$ and $\tilde{\mathbf{v}}$. We obtain

$$\begin{aligned}
&I_f(\mathbf{v}^i) - I_f(\tilde{\mathbf{v}}^i) - \big(I_f(\mathbf{v}) - I_f(\tilde{\mathbf{v}})\big) \\
&= v_j f(x - v_i + v_j) + v_{j+1}(x - v_i + v_j + v_{j+1}) \\
&\quad - (v_j + v_{j+1}) f(x - v_i + v_j + v_{j+1}) \\
&\quad - \big(v_j f(x + v_j) + v_{j+1} f(x + v_j + v_{j+1}) - (v_j + v_{j+1}) f(x + v_j + v_{j+1})\big) \\
&= v_j \Big(f(x - v_i + v_j) - f(x - v_i + v_j + v_{j+1}) - \big(f(x + v_j) - f(x + v_j + v_{j+1})\big)\Big),
\end{aligned}$$

which is clearly non-positive for convex f and non-negative for concave f. ◀

By reversing the roles of $\mathbf{v}$ and $\tilde{\mathbf{v}}$ in the Union Lemma (Lemma 6) we obtain the following Split Lemma as a direct corollary of the Union Lemma.

▶ **Lemma 7** (Split Lemma). *Let $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, and let $\mathbf{v}, \mathbf{v}^i \in \mathbb{R}^n_{>0}$, $\tilde{\mathbf{v}}, \tilde{\mathbf{v}}^i \in \mathbb{R}^{n+1}_{>0}$ be such that*

$$\mathbf{v} = (v_1, \ldots, v_i, \ldots, v_j, \ldots, v_n), \qquad \mathbf{v}^i = (v_1, \ldots, v_j, \ldots, v_n, v_i),$$
$$\tilde{\mathbf{v}} = (v_1, \ldots, v_i, \ldots, v_j/2, v_j/2, \ldots, v_n), \qquad \tilde{\mathbf{v}}^i = (v_1, \ldots, v_j/2, v_j/2 \ldots, v_n, v_i)$$

i.e., $\tilde{\mathbf{v}}$ is obtained from $\mathbf{v}$ by splitting entry v_j, $j > i$ into two entries of size $v_j/2$ and $\mathbf{v}^i$ and $\tilde{\mathbf{v}}^i$ are obtained from $\mathbf{v}$ and $\tilde{\mathbf{v}}$, respectively, by moving entry d_i to the end of the vector. Then, $I_f(\mathbf{v}^i) - I_f(\mathbf{v}) - \big(I_f(\tilde{\mathbf{v}}^i) - I_f(\tilde{\mathbf{v}})\big)$ is non-positive, if f is convex and non-negative, if f is concave.

We proceed to use the Union Lemma (Lemma 6) to devise a closed form expression for $\check{\mu}(f)$ for convex f respectively $\hat{\mu}(f)$ for concave f.

▶ **Corollary 8.** *For a function $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, let*

$$B_f(x, y, z) = \frac{x f(x + z) + y f(x + y + z) - y f(y + z) - x f(x + y + z)}{x f(x + y + z)}.$$

If f is convex, then $\check{\mu}(f) = \sup_{x,y,z \in \mathbb{R}_{>0}, y \geq x} B_f(x, y, z)$, and if f is concave, then $\hat{\mu}(f) = \sup_{x,y,z \in \mathbb{R}_{>0}, y \geq x} -B_f(x, y, z)$.

Proof. We start to prove $\check{\mu}(f) =_{x,y,z \in \mathbb{R}_{>0}, y \geq x} B(f)$, if f is convex. First, recall that $\check{\mu}(f)$ is defined as

$$\check{\mu}(f) = \sup_{\mathbf{v} \in \mathbb{R}^n_{>0}} \max_{i \in \{1, \ldots, n\}} \frac{I_f(\check{\sigma}_f(\mathbf{v})) - I_f(\check{\sigma}_f^i(\mathbf{v}))}{v_i f(|\mathbf{v}|)}.$$

By Lemma 5, it is without loss of generality to assume that $\check{\sigma}_f(\mathbf{v})$ contains the entries of $\mathbf{v}$ in non-decreasing order. Using this observation, we derive that there are sequences $(n_k)_{n \in \mathbb{N}}$, $n_k \in \mathbb{N}$ for all $k \in \mathbb{N}$, $(i_k)_{k \in \mathbb{N}}$, $i_k \in \{1, \ldots, n_k\}$ for all $k \in \mathbb{N}$ and a sequence of non-decreasing vectors $(\mathbf{v}_k)_{k \in \mathbb{N}}$, $\mathbf{v}_k \in \mathbb{R}_{>0}^{n_k}$ such that

$$\lim_{k \to \infty} \frac{I_f(\mathbf{v}_k) - I_f(\mathbf{v}_k^{i_k})}{v_i f(|\mathbf{v}_k|)},$$

where we denote by $\mathbf{v}_k^{i_k}$ the vector that arises from $\mathbf{v}_k$ by moving the i_kth entry to the end of the vector. For ease of exposition, let us set $n = n_k$, $i = i_k$, $\mathbf{v} = (v_1, \ldots, v_n) = \mathbf{v}_k$, $\mathbf{v}^i = (v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n, v_i) = v_k^{i_k}$ for some fixed $k \in \mathbb{N}$. Let $z = \sum_{j=1}^{i-1} v_j$ denote the sum of the entries in $\mathbf{v}$ before entry v_i, let $x = v_i$, and let $y = \sum_{j=i+1}^{n}$. Using the Union Lemma (Lemma 6), we derive that $I_f(\mathbf{v}) - I_f(\mathbf{v}^i)$ is maximal if $i + 1 = n$, i.e., there is only one entry $v_{i+1} = y$ that appears after v_i in $\mathbf{v}$. We then obtain

$$\frac{I_f(\mathbf{v}) - I_f(\mathbf{v}^i)}{v_i f(|\mathbf{v}|)} = \frac{xf(x+z) + yf(x+y+z) - yf(y+z) - xf(x+y+z)}{xf(x+y+z)},$$

which proves the result.

For concave f, recall that $\hat{\mu}(f)$ is defined as

$$\check{\mu}(f) = \sup_{\mathbf{v} \in \mathbb{R}_{>0}^n} \max_{i \in \{1, \ldots, n\}} \frac{I_f(\hat{\sigma}_f^i(\mathbf{v})) - I_f(\hat{\sigma}_f(\mathbf{v}))}{v_i f(|\mathbf{v}|)},$$

and that, by Lemma 5, it is without loss of generality to assume that $\hat{\sigma}_f(\mathbf{v})$ contains the entries of $\mathbf{v}$ in non-decreasing order. Analogously to the convex case, this implies the existence of sequences $(n_k)_{n \in \mathbb{N}}$, $n_k \in \mathbb{N}$ for all $k \in \mathbb{N}$, $(i_k)_{k \in \mathbb{N}}$, $i_k \in \{1, \ldots, n_k\}$ for all $k \in \mathbb{N}$ and a sequence of non-decreasing vectors $(\mathbf{v}_k)_{k \in \mathbb{N}}$, $\mathbf{v}_k \in \mathbb{R}_{>0}^{n_k}$ such that

$$\lim_{k \to \infty} \frac{I_f(\mathbf{v}_k^{i_k}) - I_f(\mathbf{v}_k)}{v_i f(|\mathbf{v}_k|)},$$

where we again denote by $\mathbf{v}_k^{i_k}$ the vector that arises from $\mathbf{v}_k$ by moving the i_kth entry to the end of the vector. We may again apply the union Lemma (Lemma 6) to derive that this ratio is maximal when $i + 1 = n$, i.e., there is only one entry after v_i. We then obtain

$$\frac{I_f(\mathbf{v}_k^{i_k}) - I_f(\mathbf{v}_k)}{v_i f(|\mathbf{v}_k|)} = \frac{yf(y+z) + xf(x+z) - xf(x+z) - yf(x+y+z)}{xf(x+y+z)},$$

which then establishes the result. ◄

Symmetrically to the previous corollary, we use the Split Lemma (Lemma 7) to derive closed form expressions for the two missing cases of $\hat{\mu}(f)$ for convex f and $\check{\mu}(f)$ for concave f. It is interesting to note that by repeatedly applying the Split Lemma the discrete integral approaches the integral, in the supremum.

▶ **Corollary 9.** *For a function $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, let*

$$A_f(x, y, z) = \frac{\int_z^y f(s)\, \mathrm{d}s + xf(x+y+z) - \int_{x+z}^{x+y+z} f(s)\, \mathrm{d}s - xf(x+z)}{xf(x+y+z)}.$$

If f is convex, then $\hat{\mu}(f) = \sup_{x,y \in \mathbb{R}_{>0}, z \in \{0\} \cup [x, \infty)} A_f(x, y, z)$, and if f is concave then $\check{\mu}(f) = \sup_{x,y \in \mathbb{R}_{>0}, z \in \{0\} \cup [x, \infty)} -A_f(x, y, z)$.

Proof. We start to prove $A(f) = \hat{\mu}(f)$, if f is convex. First, recall that $\hat{\mu}(f)$ is defined as

$$\hat{\mu}(f) = \sup_{\mathbf{v} \in \mathbb{R}^n_{>0}} \max_{i \in \{1,\dots,n\}} \frac{I_f(\hat{\sigma}^i_f(\mathbf{v})) - I_f(\hat{\sigma}_f(\mathbf{v}))}{v_i f(|\mathbf{v}|)}.$$

By Lemma 5, it is without loss of generality to assume that $\hat{\sigma}_f(\mathbf{v})$ sorts the entries of $\mathbf{v}$ in non-increasing order. This implies that there are sequences $(n_k)_{k \in \mathbb{N}}$, $n_k \in \mathbb{N}$ for all $k \in \mathbb{N}$, $(i_k)_{k \in \mathbb{N}}$, $i_k \in \{1, \dots, n_k\}$ for all $k \in \mathbb{N}$ and a sequence of non-increasing vectors $(\mathbf{v}_k)_{k \in \mathbb{N}}$, $\mathbf{v}_k \in \mathbb{R}^{n_k}_{>0}$ such that

$$\lim_{k \to \infty} \frac{I_f(\mathbf{v}^{i_k}_k) - I_f(\mathbf{v}_k)}{v_{i_k} f(|\mathbf{v}_k|)} = \hat{\mu}(f).$$

where we denote by $\mathbf{v}^{i_k}_k$ the vector that arises from $\mathbf{v}_k$ by moving the i_kth entry to the end of the vector. For ease of exposition, let us set $n = n_k$, $i = i_k$, and $\mathbf{v} = (v_1, \dots, v_n) = \mathbf{v}_k$, for some fixed $k \in \mathbb{N}$. Let $z = \sum_{j=1}^{i-1} v_j$ denote the sum of the entries in $\mathbf{v}$ before entry v_i, let $x = v_i$, and let $y = \sum_{j=i+1}^{n}$ denote the sum of the entries in $\mathbf{v}$ after entry v_i. As $\mathbf{v}$ is non-increasing, we have $z \in \{0\} \cup [x, \infty)$. As shown Lemma 7, splitting all entries v_j, $j > i$ into two entries $v_j/2$ does only increase the difference of the discrete integrals. Applying the procedure repeatedly on all entries v_j, $j > i$, their contribution to the discrete integral approaches the integral. We obtain

$$\frac{I_f(\mathbf{v}^i) - I_f(\mathbf{v})}{v_i f(|\mathbf{v}|)} \le \frac{\int_z^y f(s)\,ds + x f(x+y+z) - \int_{x+z}^{x+y+z} f(s)\,ds - x f(x+z)}{x f(x+y+z)}.$$

Applying this reasoning for all vectors $\mathbf{v}_k$, $k \in \mathbb{N}$, we obtain

$$\sup_{\mathbf{v} \in \mathbb{R}^n_{>0}} \max_{i \in \{1,\dots,n\}} \frac{I_f(\mathbf{v}^i) - I_f(\mathbf{v})}{v_i f(|\mathbf{v}|)}$$
$$= \sup_{\substack{x,y \in \mathbb{R}_{>0} \\ z \in \{0\} \cup [x,\infty)}} \frac{\int_z^y f(s)\,ds + x f(x+y+z) - \int_{x+z}^{x+y+z} f(s)\,ds - x f(x+z)}{x f(x+y+z)},$$

as claimed.

For a concave function f, we obtain along the same lines

$$\lim_{k \to \infty} \frac{I_f(\mathbf{v}_k) - I_f(\mathbf{v}^{i_k}_k)}{v_{i_k} f(|\mathbf{v}_k|)} = \check{\mu}(f)$$

for some sequences $(n_k)_{k \in \mathbb{N}}$, $n_k \in \mathbb{N}$ for all $k \in \mathbb{N}$, $(i_k)_{k \in \mathbb{N}}$, $i_k \in \{1, \dots, n_k\}$ for all $k \in \mathbb{N}$ and a sequence of non-increasing vectors $(\mathbf{v}_k)_{k \in \mathbb{N}}$, $\mathbf{v}_k \in \mathbb{R}^{n_k}_{>0}$. We here again denote by $\mathbf{v}^{i_k}_k$ the vector that arises from $\mathbf{v}_k$ by moving the i_kth entry to the end of the vector. Repeated application of the Split Lemma (Lemma 7) gives

$$\check{\mu}(f) = \frac{\int_{x+z}^{x+y+z} f(s)\,ds + x f(x+z) - \int_z^y f(s)\,ds - x f(x+y+z)}{x f(x+y+z)},$$

analogous to the convex case. ◀

We proceed by further simplifying the computation of $\hat{\mu}(f)$ and $\check{\mu}(f)$.

▶ **Lemma 10.** *Let $f : \mathbb{R}_{\geq 0} \to \mathbb{R}$ be a function. Then,*

$$\sup_{\substack{x,y \in \mathbb{R}_{>0}, \\ z \in \{0\} \cup [x,\infty)}} A_f(x,y,z) = \sup_{x,y \in \mathbb{R}_{>0}} A_f(x,y,0), \qquad \text{if } f \text{ is convex,}$$

$$\sup_{\substack{x,y \in \mathbb{R}_{>0}, \\ z \in \{0\} \cup [x,\infty)}} -A_f(x,y,z) = \sup_{x,y \in \mathbb{R}_{>0}} -A_f(x,y,0), \qquad \text{if } f \text{ is concave.}$$

Proof. We first prove the result for convex f. Let $(x_k)_{k \in \mathbb{N}}$, $(y_k)_{k \in \mathbb{N}}$, and $(z_k)_{k \in \mathbb{N}}$ sequences for which the supremum of $A_f(x,y,z)$ is attained. For fixed $k \in N$, let us set $x = x_k$, $y = y_k$, and $z = z_k$ and let $x' = x + z$, $y' = y$. We claim that

$$A_f(x,y,z) = \frac{\int_z^{y+z} f(s)\,\mathrm{d}s + xf(x+y+z) - \int_{x+z}^{x+y+z} f(s)\,\mathrm{d}s - xf(x+z)}{xf(x+y+z)}$$

$$\leq \frac{\int_0^{y'} f(s)\,\mathrm{d}s + x'f(x'+y') - \int_{x'}^{x'+y'} f(s)\,\mathrm{d}s - x'f(x')}{x'f(x'+y')} = A_f(x'_k, y'_k, 0). \qquad (5)$$

Using $f(x' + y') = f(x + y + z)$ and substituting x' and y', (5) is equivalent to

$$\frac{1}{x}\int_z^{y+z} f(s)\,\mathrm{d}s + f(x+y+z) - \frac{1}{x}\int_{x+z}^{x+y+z} f(s)\,\mathrm{d}s - f(x+z)$$

$$\leq \frac{1}{x+z}\int_0^{y} f(s)\,\mathrm{d}s + f(x+y+z) - \frac{1}{x+z}\int_{x+z}^{x+y+z} f(s)\,\mathrm{d}s - f(x+z). \qquad (6)$$

Rearranging terms (6) is equivalent to

$$\frac{1}{x}\int_z^{y+z} f(s)\,\mathrm{d}s - \frac{1}{x}\int_{x+z}^{x+y+z} f(s)\,\mathrm{d}s \leq \frac{1}{x+z}\int_0^{y} f(s)\,\mathrm{d}s - \frac{1}{x+z}\int_{x+z}^{x+y+z} f(s)\,\mathrm{d}s. \qquad (7)$$

Put differently, we need to show that the function $h : \mathbb{R}_{\geq 0} \to \mathbb{R}$ defined as

$$h(y) = \frac{1}{x}\int_z^{y+z} f(s)\,\mathrm{d}s - \left(\frac{1}{x} - \frac{1}{x+z}\right)\int_{x+z}^{x+y+z} f(s)\,\mathrm{d}s - \frac{1}{x+z}\int_0^{y} f(s)\,\mathrm{d}s$$

is non-positive on $\mathbb{R}_{\geq 0}$. It is easy to verify that this is true for $y = 0$. Thus it suffices to show that $h'(y) \leq 0$ for all $y \geq 0$. To this end, we derive

$$h'(y) = \frac{1}{x}f(y+z) - \left(\frac{1}{x} - \frac{1}{x+z}\right)f(x+y+z) - \frac{1}{x+z}f(y)$$

$$= \frac{1}{x}\left(f(y+z) - \left(1 - \frac{x}{x+z}\right)f(x+y+z) - \frac{x}{x+z}f(y)\right),$$

which is non-positive due to the convexity of f. We conclude that for each x_k, y_k, and z_k in the sequence that goes to the supremum of $A_f(x_k, y_k, z)$, there are x'_k and y'_k with $A_f(x_k, y_k, z_k) \leq A_f(x'_k, y'_k, 0)$. This implies the result for convex f.

To see the statement for concave f, note that since we seek to compute the supremum of $-A_f(x,y,z)$ rather than $A_f(x,y,z)$ all inequality signs in (5), (6), and (7) reverse and we shall show that the function h is non-negative on $\mathbb{R}_{\geq 0}$. It is easy to check that $h(y) = 0$ and $h'(y) \geq 0$ due to the concavity of f. This implies the result. ◀

▶ **Lemma 11.** *Let $f : \mathbb{R}_{\geq 0} \to \mathbb{R}$ be a function. Then,*

$$\sup_{\substack{x,y,z \in \mathbb{R}_{>0}, \\ y \geq x}} B_f(x,y,z) = \sup_{x,y \in \mathbb{R}_{>0}} B_f(x,y,0), \qquad \text{if } f \text{ is convex,}$$

$$\sup_{\substack{x,y,z \in \mathbb{R}_{>0}, \\ y \geq x}} -B_f(x,y,z) = \sup_{x,y \in \mathbb{R}_{>0}} -B_f(x,y,0), \qquad \text{if } f \text{ is concave.}$$

Proof. We first prove the result for convex f. Let $(x_k)_{k\in\mathbb{N}}$, $(y_k)_{k\in\mathbb{N}}$, and $(z_k)_{k\in\mathbb{N}}$ be sequences for which the supremum of B_f is attained in the limit. For fixed $k \in \mathbb{N}$, let us set $x = x_k$, $y = y_k$, and $z = z_k$ and let $x' = x$ and $y' = y + z$. We claim that

$$
\begin{aligned}
B_f(x,y,z) &= \frac{xf(x+z) + yf(x+y+z) - yf(y+z) - xf(x+y+z)}{xf(x+y+z)} \\
&\leq \frac{x'f(x') + y'f(x'+y') - y'f(y') - x'f(x'+y')}{x'f(x'+y')} = B_f(x',y',0)
\end{aligned} \tag{8}
$$

Substituting x' and y', (8) is equivalent to

$$
xf(x+z) \leq xf(x) + zf(x+y+z) - zf(y+z) \tag{9a}
$$

$$
\Leftrightarrow \qquad \frac{f(x+z) - f(x)}{z} \leq \frac{f(x+y+z) - f(y+z)}{x}, \tag{9b}
$$

which is satisfied since f is convex and $y \geq x$. For concave f, all inequality signs reverse. ◄

Combining all lemmas proven in this section, we obtain the following expressions for the computation of $\hat{\mu}(f)$ and $\check{\mu}(f)$, respectively.

▶ **Theorem 12.** *For a function $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. Then,*

$$
\hat{\mu}(f) = \begin{cases}
\displaystyle\sup_{x,y\in\mathbb{R}_{>0}} \frac{\int_0^y f(s)\,\mathrm{d}s + xf(x+y) - \int_x^{x+y} f(s)\,\mathrm{d}s - xf(x)}{xf(x+y)}, & \text{if } f \text{ is convex} \\[3ex]
\displaystyle\sup_{\substack{x,y\in\mathbb{R}_{>0} \\ y\geq x}} \frac{yf(y) + xf(x+y) - xf(x) - yf(x+y)}{xf(x+y)}, & \text{if } f \text{ is concave,}
\end{cases}
$$

$$
\check{\mu}(f) = \begin{cases}
\displaystyle\sup_{\substack{x,y\in\mathbb{R}_{>0} \\ y\geq x}} \frac{xf(x) + yf(x+y) - yf(y) - xf(x+y)}{xf(x+y)}, & \text{if } f \text{ is convex} \\[3ex]
\displaystyle\sup_{x,y\in\mathbb{R}_{>0}} \frac{\int_x^{x+y} f(s)\,\mathrm{d}s + xf(x) - \int_0^y f(s)\,\mathrm{d}s - xf(x+y)}{xf(x+y)}, & \text{if } f \text{ is concave.}
\end{cases}
$$

5 Polynomial Cost Functions

In this section, we consider the special case of polynomial functions. Formally, for $\ell \in \mathbb{N}$, let us denote by

$$
\mathcal{C}^\ell = \left\{ f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0} \mid \exists n \in \mathbb{N} \text{ with } f(x) = \sum_{j=1}^{n} a_j\, x^{b_j}, a_j \geq 0, b_j \in [1,\ell] \; \forall j \in \{1, \ldots, n\} \right\}
$$

the set of convex polynomial functions with non-negative coefficients and maximal degree ℓ. Combining Theorem 3 and Theorem 12, we show that weighted congestion with degree ℓ polynomials have an $(\ell+1)$-approximate pure Nash equilibrium.

▶ **Theorem 13.** *Every weighted congestion game with cost functions in $\mathcal{C}^\ell$ has an α-approximate pure Nash equilibrium with $\alpha \leq \ell + 1$.*

Proof. We first decompose the game into an isomorphic weighted congestion game G in which each resource r has a cost function of type $c_r(x) = a_r x^{b_r}$ with $a_r > 0$ and $b_r \in [1,\ell]$. Let us denote the set of these functions by $\bar{\mathcal{C}}^\ell$.

Theorem 3 shows that G has an $(1 + \check{\mu})$-approximate pure Nash equilibrium and Theorem 12 provides an explicit formula for computing $\check{\mu}$. We can thus compute the maximum approximation factor α as

$$\alpha = \sup_{f \in \bar{\mathcal{C}}^{\ell}} \sup_{\substack{x,y \in \mathbb{R}_{>0} \\ y \geq x}} \frac{xf(x) + yf(x+y) - yf(y) - xf(x+y)}{xf(x+y)} + 1$$

$$= \sup_{b \in [1,\ell]} \sup_{\substack{x,y \in \mathbb{R}_{>0} \\ y \geq x}} \frac{x^{b+1} + y(x+y)^b - y^{b+1} - x(x+y)^b}{x(x+y)^b} + 1$$

$$= \sup_{b \in [1,\ell]} \sup_{\substack{x,y \in \mathbb{R}_{>0} \\ y \geq x}} \frac{y\big((x+y)^b - y^b\big) - x\big((x+y)^b - x^b\big)}{x(x+y)^b} + 1$$

We use that for all $b \in [1,\ell]$ the function $s \mapsto s^b$ is convex with derivative $s \mapsto bs^{b-1}$ and obtain

$$\alpha \leq \sup_{b \in [1,\ell]} \sup_{\substack{x,y \in \mathbb{R}_{>0} \\ y \geq x}} \frac{xy\big(b \cdot (x+y)^{b-1}\big) - xy\big(bx^{b-1}\big)}{x(x+y)^b} + 1.$$

We simplify and obtain

$$\alpha \leq \sup_{b \in [1,\ell]} \sup_{\substack{x,y \in \mathbb{R}_{>0} \\ y \geq x}} b \cdot \left(\frac{y}{x+y} - \frac{yx^{b-1}}{(x+y)^b}\right) + 1 \leq \sup_{b \in [1,\ell]} b + 1 = \ell + 1,$$

which completes the proof. ◄

For low degrees, we use the minimizing permutation instead of the maximizing permutation and obtain better bounds.

▶ **Theorem 14.** *Every weighted congestion game with cost functions in $\mathcal{C}^{\ell}$ has an α-approximate pure Nash equilibrium, where*

$$\alpha = \sup_{b \in [1,\ell]} \sup_{x,y \in \mathbb{R}_{>0}} \frac{\lambda(1+\lambda)^b}{\frac{1}{b+1}(1+\lambda)^{b+1} + (1 - \frac{1}{b+1})\lambda^{b+1} - \frac{1}{b+1}}. \tag{10}$$

Proof. Using the same decomposition argument as in the proof of Theorem 13, we can assume without loss of generality that all cost functions are monomials. We again denote the set of monomials with maximum degree ℓ by $\bar{\mathcal{C}}^{\ell}$. Using Theorem 4 and Theorem 12, we obtain

$$\alpha \leq \sup_{f \in \bar{\mathcal{C}}^{\ell}} \sup_{x,y \in \mathbb{R}_{>0}} \frac{xf(x+y)}{\int_x^{x+y} f(s)\,\mathrm{d}s + xf(x) - \int_0^y f(s)\,\mathrm{d}s}$$

$$\leq \sup_{b \in [1,\ell]} \sup_{x,y \in \mathbb{R}_{>0}} \frac{x(x+y)^b}{\frac{1}{b+1}(x+y)^{b+1} + (1 - \frac{1}{b+1})x^{b+1} - \frac{1}{b+1}y^{b+1}}$$

Substituting $x = \lambda y$ for some $\lambda \in \mathbb{R}_{>0}$, we obtain

$$\alpha \leq \sup_{b \in [1,\ell]} \sup_{\lambda \in \mathbb{R}_{>0}} \frac{\lambda(1+\lambda)^b}{\frac{1}{b+1}(1+\lambda)^{b+1} + (1 - \frac{1}{b+1})\lambda^{b+1} - \frac{1}{b+1}},$$

as claimed. ◄

For polynomials with small maximal degrees ℓ, we solve (10) and obtain approximation guarantees strictly below the factor of $\ell + 1$ guaranteed by Theorem 13. Specifically, we obtain a factor of $4/3$ for (at most) quadratic cost functions, 1.785 for (at most) cubic cost functions, and a factor of 2.326 for polynomials with maximum degree 4.

6   Concave Cost Functions

As the main result of this section, we shown that weighted congestion games with concave cost functions have a 3/2-approximate pure Nash equilibrium.

▶ **Theorem 15.** *Every weighted congestion game with concave cost functions has 3/2-approximate pure Nash equilibrium.*

To prove this theorem we show that $\breve{\mu}(f) \leq 1/2$ for each concave function f. Due to space constraints the detailed proof is deferred to the full version of this paper.

7   Two-player Games

In the following we analyze two-player weighted congestion games and provide a tight upper bound for approximate equilibria. For this section, it is convenient to assume that the set of cost functions $\mathcal{C}$ is closed under addition, i.e, $f \in \mathcal{C}$ and $g \in \mathcal{C}$ implies $(f + g) \in \mathcal{C}$. However, it is straightforward to extend our results to sets of cost functions that do not have this property.

As a lower bound, consider the following game with two players $N = \{1, 2\}$ with demands $x, y > 0$, respectively, and resources $R = \big\{(0,0), (0,1), (1,0), (1,1)\big\}$ with symmetric cost functions $c_{(0,0)} = c_{(1,1)} = c_1(x)$ and $c_{(1,0)} = c_{(0,1)} = c_2(x)$ for some $c_1, c_2 \in \mathcal{C}$. Player 1 has the choice between the strategies $\{(0,0), (0,1)\}$ and $\{(1,0), (1,1)\}$ whereas the strategies for player 2 are given by $\{(0,0), (1,0)\}$ and $\{(0,1), (1,1)\}$. For this game not to have an α-approximate NE it is necessary for the players to be able to take turns in improving their private cost by factors of at least α. We can therefore bound the value of α from above:

$$\alpha \leq \min\left\{\frac{c_2(x+y) + c_1(x)}{c_1(x+y) + c_2(x)}, \frac{c_1(x+y) + c_2(y)}{c_2(x+y) + c_1(y)}\right\}$$

Clearly, for a given set of cost functions $\mathcal{C}$ and for every $\epsilon > 0$, by optimizing over $c_1, c_2 \in \mathcal{C}$ and $x, y \in \mathbb{R}_{>0}$, we can construct a game with no $(\alpha - \epsilon)$-approximate pure Nash equilibrium, where

$$\alpha = \sup_{x,y>0, c_1, c_2 \in \mathcal{C}} \min\left\{\frac{c_2(x+y) + c_1(x)}{c_1(x+y) + c_2(x)}, \frac{c_1(x+y) + c_2(y)}{c_2(x+y) + c_1(y)}\right\}.$$

As the main result of this section, we show that this bound tight. Due to space constraints the detailed proof is deferred to the full version of this paper.

▶ **Theorem 16.** *Every two-player weighted congestion game with cost functions in $\mathcal{C}$ has an α-approximate pure Nash equilibrium with*

$$\alpha \leq \sup_{x,y>0, c_1, c_2 \in \mathcal{C}} \min\left\{\frac{c_2(x+y) + c_1(x)}{c_1(x+y) + c_2(x)}, \frac{c_1(x+y) + c_2(y)}{c_2(x+y) + c_1(y)}\right\}.$$

For specific classes of cost functions, we solve for α and obtain the concrete numerical approximation factors shown in Table 1.

───── **References** ─────

1   Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. On the impact of combinatorial structure on congestion games. *Journal of the ACM*, 55(6), 2008.

2   Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. Pure Nash equilibria in player-specific and weighted congestion games. *Theoretical Computer Science*, 410(17):1552–1563, 2009.

3   Heiner Ackermann and Alexander Skopalik. Complexity of pure Nash equilibria in player-specific network congestion games. *Internet Mathematics*, 5(4):323–342, 2008.

4   Ioannis Caragiannis, Angelo Fanelli, Nick Gravin, and Alexander Skopalik. Efficient computation of approximate pure Nash equilibria in congestion games. In *Proceedings of the 52nd Annual Symposium on Foundations of Computer Science, (FOCS)*, pages 532–541, 2011.

5   Ioannis Caragiannis, Angelo Fanelli, Nick Gravin, and Alexander Skopalik. Approximate pure Nash equilibria in weighted congestion games: existence, efficient computation, and structure. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, pages 284–301, 2012.

6   Ho-Lin Chen and Tim Roughgarden. Network design with weighted players. In *Proceedings of the 18th Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 29–38, 2006.

7   Steve Chien and Alistair Sinclair. Convergence to approximate Nash equilibria in congestion games. In *Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 169–178, 2007.

8   George Christodoulou, Elias Koutsoupias, and Paul G. Spirakis. On the performance of approximate equilibria in congestion games. In *Proceedings of the 17th Annual European Symposium on Algorithms (ESA)*, pages 251–262. Springer, 2009.

9   Juliane Dunkel and Andreas S. Schulz. On the complexity of pure-strategy Nash equilibria in congestion and local-effect games. In *Proceedings of the 2nd International Workshop Internet & Network Economics (WINE)*, pages 62–73, 2006.

10   Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure Nash equilibria. In *Proceedings of the 36th ACM Symposium on Theory of Computing (STOC)*, pages 604–612, 2004.

11   Dimitris Fotakis, Spyros Kontogiannis, and Paul G. Spirakis. Selfish unsplittable flows. *Theoretical Computer Science*, 348(2-3):226–239, 2005.

12   Irving L. Glicksberg. A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points. *Proceedings of the AMS*, 3:170–174, 1952.

13   Michel X. Goemans, Vahab S. Mirrokni, and Adrian Vetta. Sink equilibria and convergence. In *Proceedings of the 46th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 142–154, 2005.

14   Tobias Harks and Max Klimm. On the existence of pure Nash equilibria in weighted congestion games. *Mathematics of Operations Research*, 37(3):419–436, 2012.

15   Lavy Libman and Ariel Orda. Atomic resource sharing in noncooperative networks. *Telecommunication Systems*, 17(4):385–409, 2001.

16   Igal Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124, 1996.

17   John F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences, USA*, 36:48–49, 1950.

18   Panagiota N. Panagopoulou and Paul G. Spirakis. Algorithms for pure Nash equilibria in weighted congestion games. *Journal of Experimental Algorithmics*, 11, 2007.

19   Robert Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

20   Alexander Skopalik and Berthold Vöcking. Inapproximability of pure Nash equilibria. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 355–364, 2008.