Existence and Complexity of Approximate Equilibria in Weighted Congestion Games*

George Christodoulou† Martin Gairing† Yiannis Giannakopoulos‡
Diogo Poças‡ Clara Waldmann‡

February 14, 2020

Abstract

We study the existence of approximate pure Nash equilibria (α-PNE) in weighted atomic congestion games with polynomial cost functions of maximum degree d. Previously it was known that d-approximate equilibria always exist, while nonexistence was established only for small constants, namely for 1.153-PNE. We improve significantly upon this gap, proving that such games in general do not have $\tilde{\Theta}(\sqrt{d})$-approximate PNE, which provides the first super-constant lower bound.

Furthermore, we provide a black-box gap-introducing method of combining such nonexistence results with a specific circuit gadget, in order to derive NP-completeness of the decision version of the problem. In particular, deploying this technique we are able to show that deciding whether a weighted congestion game has an $\tilde{O}(\sqrt{d})$-PNE is NP-complete. Previous hardness results were known only for the special case of *exact* equilibria and arbitrary cost functions.

The circuit gadget is of independent interest and it allows us to also prove hardness for a variety of problems related to the complexity of PNE in congestion games. For example, we demonstrate that the question of existence of α-PNE in which a certain set of players plays a specific strategy profile is NP-hard for any $\alpha < 3^{d/2}$, even for *unweighted* congestion games.

Finally, we study the existence of approximate equilibria in weighted congestion games with general (nondecreasing) costs, as a function of the number of players n. We show that n-PNE always exist, matched by an almost tight nonexistence bound of $\tilde{\Theta}(n)$ which we can again transform into an NP-completeness proof for the decision problem.

*Supported by the Alexander von Humboldt Foundation with funds from the German Federal Ministry of Education and Research (BMBF).

The third author is an associated researcher with the Research Training Group GRK 2201 "Advanced Optimization in a Networked Economy", funded by the German Research Foundation (DFG).

†Department of Computer Science, University of Liverpool. Emails: {g.christodoulou, gairing}@liverpool.ac.uk

‡Chair of Operations Research, TU Munich. Emails: {yiannis.giannakopoulos, diogo.pocas, clara.waldmann}@tum.de

1 Introduction

Congestion games constitute the standard framework to study settings where selfish players compete over common resources. They are one of the most well-studied classes of games within the field of *algorithmic game theory* [25, 31], covering a wide range of applications, including, e.g., traffic routing and load balancing. In their most general form, each player has her own weight and the latency on each resource is a nondecreasing function of the total weight of players that occupy it. The cost of a player on a given outcome is just the total latency that she is experiencing, summed over all the resources she is using.

The canonical approach to analysing such systems and predicting the behaviour of the participants is the ubiquitous game-theoretic tool of equilibrium analysis. More specifically, we are interested in the *pure Nash equilibria (PNE)* of those games; these are stable configurations from which no player would benefit from unilaterally deviating. However, it is a well-known fact that such desirable outcomes might not always exist, even in very simple weighted congestion games. A natural response, especially from a computer science perspective, is to relax the solution notion itself by considering *approximate* pure Nash equilibria (α-PNE); these are states from which, even if a player could improve her cost by deviating, this improvement could not be by more than a (multiplicative) factor of $\alpha \geq 1$. Allowing the parameter α to grow sufficiently large, existence of α-PNE is restored. But how large does α really *need* to be? And, perhaps more importantly from a computational perspective, how hard is it to check whether a specific game has indeed an α-PNE?

1.1 Related Work

The origins of the systematic study of (atomic) congestion games can be traced back to the influential work of Rosenthal [29, 30]. Although Rosenthal showed the existence of congestion games without PNE, he also proved that *unweighted* congestion games always possess such equilibria. His proof is based on a simple but ingenious *potential function* argument, which up to this day is essentially still the only general tool for establishing existence of pure equilibria.

In follow-up work [14, 17, 24], the nonexistence of PNE was demonstrated even for special simple classes of (weighted) games, including network congestion games with quadratic cost functions and games where the player weights are either 1 or 2. On the other hand, we know that equilibria do exist for affine or exponential latencies [14, 19, 27], as well as for the class of singleton[1] games [15, 20]. Dunkel and Schulz [11] were able to extend the nonexistence instance of Fotakis et al. [14] to a gadget in order to show that deciding whether a congestion game with step cost functions has a PNE is a (strongly) NP-hard problem, via a reduction from 3-Partition.

Regarding approximate equilibria, Hansknecht et al. [18] gave instances of very simple, two-player polynomial congestion games that do not have α-PNE, for $\alpha \approx 1.153$. This lower bound is achieved by numerically solving an optimization program, using polynomial latencies of maximum degree $d = 4$. On the positive side, Caragiannis et al. [4] proved that $d!$-PNE always exist; this upper bound on the existence of α-PNE was later improved to $\alpha = d + 1$ [9, 18] and $\alpha = d$ [3].

1.2 Our Results and Techniques

After formalizing our model in Section 2, in Section 3 we show the nonexistence of $\tilde{\Theta}(\sqrt{d})$-approximate equilibria for polynomial congestion games of degree d. This is the first super-constant lower bound on the nonexistence of α-PNE, significantly improving upon the previous constant of $\alpha \approx 1.153$ and reducing the gap with the currently best upper bound of d. More specifically (Theorem 1), for any integer d we construct congestion games with polynomial cost

[1] These are congestion games where the players can only occupy single resources.

functions of maximum degree d (and nonnegative coefficients) that do not have α-PNE, for any $\alpha < \alpha(d)$ where $\alpha(d)$ is a function that grows as $\alpha(d) = \Omega\left(\frac{\sqrt{d}}{\ln d}\right)$. To derive this bound, we had to use a novel construction with a number of players growing unboundedly as a function of d.

Next, in Section 4 we turn our attention to computational hardness constructions. Starting from a Boolean circuit, we create a gadget that transfers hard instances of the classic Circuit Satisfiability problem to (even unweighted) polynomial congestion games. Using this gadget we can immediately establish computational hardness for various computational questions of interest involving congestion games (Theorem 2). For example, we show that deciding whether a d-degree polynomial congestion game has an α-PNE in which a specific set of players play a specific strategy profile is NP-hard, even up to exponentially-approximate equilibria; more specifically, the hardness holds for *any* $\alpha < 3^{d/2}$. It is of interest to note here that our hardness gadget is *gap-introducing*, in the sense that the α-PNE and exact PNE of the game coincide.

In Section 5 we demonstrate how one can combine the hardness gadget of Section 4, in a black-box way, with any nonexistence instance for α-PNE, in order to derive hardness for the decision version of the existence of α-PNE (Lemma 2, Theorem 3). As a consequence, using the previous $\Omega\left(\frac{\sqrt{d}}{\ln d}\right)$ lower bound construction of Section 3, we can show that deciding whether a (weighted) polynomial congestion has an α-PNE is NP-hard, for any $\alpha < \alpha(d)$, where $\alpha(d) = \Omega\left(\frac{\sqrt{d}}{\ln d}\right)$ (Corollary 1). Since our hardness is established via a rather transparent, "master" reduction from Circuit Satisfiability, which in particular is parsimonious, one can derive hardness for a family of related computation problems; for example, we show that computing the number of α-approximate equilibria of a weighted polynomial congestion game is #P-hard (Corollary 2).

In Section 6 we drop the assumption on polynomial cost functions, and study the existence of approximate equilibria under arbitrary (nondecreasing) latencies as a function of the number of players n. We prove that n-player congestion games always have n-approximate PNE (Theorem 4). As a consequence, one cannot hope to derive super-constant nonexistence lower bounds by using just simple instances with a fixed number of players. In particular, this shows that the super-constant number of players in our construction in Theorem 1 is necessary. Furthermore, we pair this positive result with an almost matching lower bound (Theorem 5): we give examples of n-player congestion games (where latencies are simple step functions with a single breakpoint) that do not have α-PNE for all $\alpha < \alpha(n)$, where $\alpha(n)$ grows according to $\alpha(n) = \Omega\left(\frac{n}{\ln n}\right)$. Finally, inspired by our hardness construction for the polynomial case, we also give a new reduction that establishes NP-hardness for deciding whether an α-PNE exists, for any $\alpha < \alpha(n) = \Omega\left(\frac{n}{\ln n}\right)$. Notice that now the number of players n is part of the description of the game (i.e., part of the input) as opposed to the maximum degree d for the polynomial case (which was assumed to be fixed). On the other hand though, we have more flexibility on designing our gadget latencies, since they can be arbitrary functions.

Concluding, we would like to elaborate on a couple of points. First, the reader would have already noticed that in all our hardness results the (in)approximability parameter α ranges freely within an entire interval of the form $[1, \tilde{\alpha})$, where $\tilde{\alpha}$ is a function of the degree d (for polynomial congestion games) or of the number of players n; and that α, $\tilde{\alpha}$ are *not* part of the problem's input. It is easy to see that these features only make our results stronger, with respect to computational hardness, but also more robust. Secondly, although in this introductory section all our hardness results were presented in terms of NP-*hardness*, they immediately translate to NP-*completeness* under standard assumptions on the parameter α; e.g., if α is rational (for a more detailed discussion of this, see also the end of Section 2).

2 Model and Notation

A (weighted, atomic) *congestion game* is defined by: a finite (nonempty) set of *resources* E, each $e \in E$ having a nondecreasing *cost (or latency) function* $c_e : \mathbb{R}_{>0} \longrightarrow \mathbb{R}_{\geq 0}$; and a finite (nonempty) set of *players* N, $|N| = n$, each $i \in N$ having a *weight* $w_i > 0$ and a set of *strategies* $S_i \subseteq 2^E$. If all players have the same weight, $w_i = 1$ for all $i \in N$, the game is called *unweighted*. A *polynomial congestion game* of degree d, for d a nonnegative integer, is a congestion game such that all its cost functions are polynomials of degree at most d with nonnegative coefficients. A *strategy profile* (or *outcome*) $\mathbf{s} = (s_1, s_2, \dots, s_n)$ is a collection of strategies, one for each player, i.e. $\mathbf{s} \in \mathbf{S} = S_1 \times S_2 \times \cdots \times S_n$. Each strategy profile $\mathbf{s}$ induces a *cost* of $C_i(\mathbf{s}) = \sum_{e \in s_i} c_e(x_e(\mathbf{s}))$ to every player $i \in N$, where $x_e(\mathbf{s}) = \sum_{i:e \in s_i} w_i$ is the induced *load* on resource e. An outcome $\mathbf{s}$ will be called *α-approximate (pure Nash) equilibrium (α-PNE)*, where $\alpha \geq 1$, if no player can unilaterally improve her cost by more than a factor of α. Formally:

$$C_i(\mathbf{s}) \leq \alpha \cdot C_i(s_i', \mathbf{s}_{-i}) \qquad \text{for all } i \in N \text{ and all } s_i' \in S_i. \tag{1}$$

Here we have used the standard game-theoretic notation of $\mathbf{s}_{-i}$ to denote the vector of strategies resulting from $\mathbf{s}$ if we remove its i-th coordinate; in that way, one can write $\mathbf{s} = (s_i, \mathbf{s}_{-i})$. Notice that for the special case of $\alpha = 1$, (1) is equivalent to the classical definition of pure Nash equilibria; for emphasis, we will sometimes refer to such 1-PNE as *exact* equilibria.

If (1) does not hold, it means that player i could improve her cost by more than α by moving from s_i to some other strategy s_i'. We call such a move *α-improving*. Finally, strategy s_i is said to be *α-dominating* for player i (with respect to a fixed profile $\mathbf{s}_{-i}$) if

$$C_i(s_i', \mathbf{s}_{-i}) > \alpha \cdot C_i(\mathbf{s}) \qquad \text{for all } s_i' \neq s_i. \tag{2}$$

In other words, if a strategy s_i is α-dominating, every move from some other strategy s_i' to s_i is α-improving. Notice that each player i can have at most one α-dominating strategy (for $\mathbf{s}_{-i}$ fixed). In our proofs, we will employ a *gap-introducing* technique by constructing games with the property that, for any player i and any strategy profile $\mathbf{s}_{-i}$, there is always a (unique) α-dominating strategy for player i. As a consequence, the sets of α-PNE and exact PNE coincide.

Finally, for a positive integer n, we will use Φ_n to denote the unique positive solution of equation $(x+1)^n = x^{n+1}$. Then, Φ_n is strictly increasing with respect to n, with $\Phi_1 = \phi \approx 1.618$ (golden ratio) and asymptotically $\Phi_n \sim \frac{n}{\ln n}$ (see [9, Lemma A.3]).

Computational Complexity Most of the results in this paper involve complexity questions, regarding the existence of (approximate) equilibria. Whenever we deal with such statements, we will implicitly assume that the congestion game instances given as inputs to our problems can be succinctly represented in the following way:

- all player have *rational* weights;
- the resource cost functions are "efficiently computable"; for polynomial latencies in particular, we will assume that the coefficients are *rationals*; and for step functions we assume that their values and breakpoints are *rationals*;
- the strategy sets are given *explicitly*.[2]

There are also computational considerations to be made about the number α appearing in the definition of α-PNE. In our results (e.g., Theorems 2 and 3), we will prove NP-hardness of determining whether games have α-PNE for any arbitrary real α below the nonexistence bound,

[2] Alternatively, we could have simply assumed succinct representability of the strategies. A prominent such case is that of *network* congestion games, where each player's strategies are all feasible paths between two specific nodes of an underlying graph. Notice however that, since in this paper we are proving hardness results, insisting on explicit representation only makes our results even stronger.

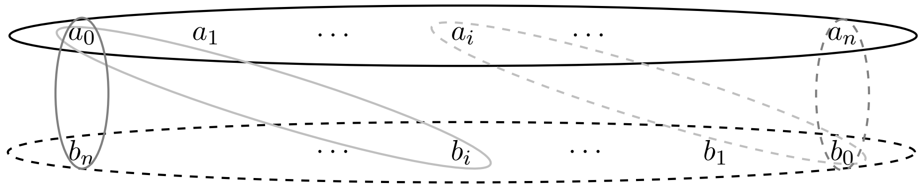

Figure 1: Strategies of the game $\mathcal{G}^d_{(n,k,w,\beta)}$. Resources contained in the two ellipses of the same colour correspond to the two strategies of a player. The strategies of the heavy player and light players n and i are depicted in black, grey and light grey, respectively.

regardless of whether α is rational or irrational, computable or uncomputable. However, to prove NP-completeness, i.e. to prove that the decision problem belongs in NP (as in Theorem 3), we need to be able to verify, given a strategy profile and a deviation of some player, whether this deviation is an α-improving move. This can be achieved by additionally assuming that the *upper Dedekind cut* of α, $R_\alpha = \{q \in \mathbb{Q} \mid q > \alpha\}$, is a language decidable in polynomial time. In this paper we will refer to such an α as a *polynomial-time computable* real number. In particular, notice that rationals are polynomial-time computable; thus the NP-completeness of the α-PNE problem does hold for α rational. We refer the interested reader to Ko [22] for a detailed discussion on polynomial-time computable numbers (which is beyond the scope of our paper), as well as for a comparison with other axiomatizations using binary digits representations or convergent sequences. If, more generally, $\alpha : \mathbb{N} \longrightarrow \mathbb{R}$ is a sequence of reals (as in Theorem 6), we say that α is a *polynomial-time computable* real sequence if $R_\alpha = \{(n,q) \in \mathbb{N} \times \mathbb{Q} \mid q > \alpha(n)\}$ is a language decidable in polynomial time.

3 The Nonexistence Gadget

In this section we give examples of polynomial congestion games of degree d, that do *not* have $\alpha(d)$-approximate equilibria; $\alpha(d)$ grows as $\Omega\left(\frac{\sqrt{d}}{\ln d}\right)$. Fixing a degree $d \geq 2$, we construct a family of games $\mathcal{G}^d_{(n,k,w,\beta)}$, specified by parameters $n \in \mathbb{N}, k \in \{1,\dots,d\}, w \in [0,1]$, and $\beta \in [0,1]$. In $\mathcal{G}^d_{(n,k,w,\beta)}$ there are $n+1$ players: a *heavy player* of weight 1 and n *light players* $1,\dots,n$ of equal weights w. There are $2(n+1)$ resources $a_0, a_1, \dots, a_n, b_0, b_1, \dots, b_n$ where a_0 and b_0 have the same cost function c_0 and all other resources $a_1,\dots,a_n, b_1,\dots,b_n$ have the same cost function c_1 given by

$$c_0(x) = x^k \quad \text{and} \quad c_1(x) = \beta x^d.$$

Each player has exactly two strategies, and the strategy sets are given by

$$S_0 = \{\{a_0,\dots,a_n\}, \{b_0,\dots,b_n\}\} \quad \text{and} \quad S_i = \{\{a_0,b_i\},\{b_0,a_i\}\} \quad \text{for } i = 1,\dots,n.$$

The structure of the strategies is visualized in Figure 1.

In the following theorem we give a lower bound on α, depending on parameters (n,k,w,β), such that games $\mathcal{G}^d_{(n,k,w,\beta)}$ do not admit an α-PNE. Maximizing this lower bound over all games in the family, we obtain a general lower bound $\alpha(d)$ on the inapproximability for polynomial congestion games of degree d (see (3) and its plot in Fig. 2). Finally, choosing specific values for the parameters (n,k,w,β), we prove that $\alpha(d)$ is asymptotically lower bounded by $\Omega(\frac{\sqrt{d}}{\ln d})$.

Theorem 1. *For any integer $d \geq 2$, there exist (weighted) polynomial congestion games of degree d that do not have α-approximate PNE for any $\alpha < \alpha(d)$, where*

$$\alpha(d) = \sup_{n,k,w,\beta} \min\left\{\frac{1+n\beta(1+w)^d}{(1+nw)^k+n\beta}, \frac{(1+w)^k+\beta w^d}{(nw)^k+\beta(1+w)^d}\right\} \tag{3}$$
$$s.t. \quad n \in \mathbb{N}, k \in \{1,\dots,d\}, w \in [0,1], \beta \in [0,1].$$

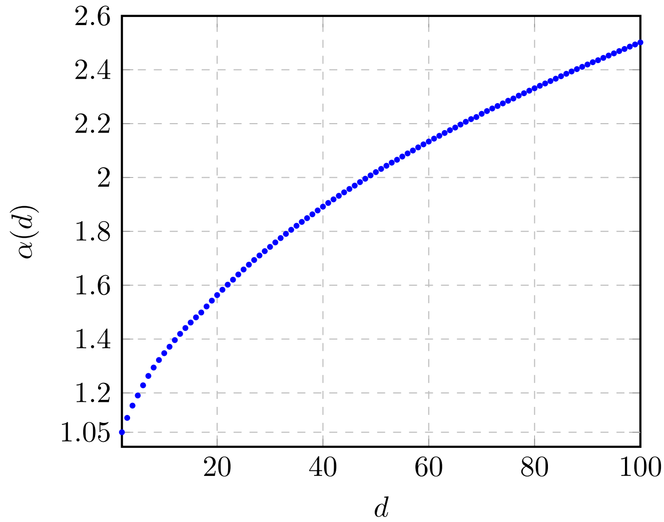

Figure 2: Nonexistence of $\alpha(d)$-PNE for weighted polynomial congestion games of degree d, as given by (3) in Theorem 1, for $d = 2, 3, \ldots, 100$.

In particular, we have the asymptotics $\alpha(d) = \Omega\left(\frac{\sqrt{d}}{\ln d}\right)$ *and the bound* $\alpha(d) \geq \frac{\sqrt{d}}{2\ln d}$*, valid for large enough* d*. A plot of the exact values of* $\alpha(d)$ *(given by (3)) for small degrees can be found in Figure 2.*

Proof. Due to symmetries, it is enough to just consider the following two cases for the strategy profiles in game $\mathcal{G}^d_{(n,k,w,\beta)}$ described above:

Case 1: *The heavy player is alone on resource* a_0. This means that every light player $i \in \{1, \ldots, n\}$ must have chosen strategy $\{b_0, a_i\}$. Thus the heavy player incurs a cost of $c_0(1) + nc_1(1+w)$; while, deviating to strategy $\{b_0, \ldots, b_n\}$, she would incur a cost of $c_0(1+nw)+nc_1(1)$. The improvement factor can then be lower bounded by

$$\frac{c_0(1) + nc_1(1 + w)}{c_0(1 + nw) + nc_1(1)} = \frac{1 + n\beta(1 + w)^d}{(1 + nw)^k + n\beta}.$$

Case 2: *The heavy player shares resource* a_0 *with at least one light player* $i \in \{1, \ldots, n\}$. Thus player i incurs a cost of at least $c_0(1 + w) + c_1(w)$; while, deviating to strategy $\{b_0, a_i\}$, she would incur a cost of at most $c_0(nw) + c_1(1 + w)$. The improvement factor can then be lower bounded by

$$\frac{c_0(1 + w) + c_1(w)}{c_0(nw) + c_1(1 + w)} = \frac{(1 + w)^k + \beta w^d}{(nw)^k + \beta(1 + w)^d}.$$

In order for the game to not have an α-PNE, it is enough to guarantee that both ratios are greater than α. Maximizing these ratios over all games in the family, yields the lower bound in the statement of the theorem,

$$\alpha(d) = \sup_{n,k,w,\beta} \min\left\{\frac{1 + n\beta(1 + w)^d}{(1 + nw)^k + n\beta}, \frac{(1 + w)^k + \beta w^d}{(nw)^k + \beta(1 + w)^d}\right\}$$
$$\text{s.t.} \quad n \in \mathbb{N}, k \in \{1, \ldots, d\}, w \in [0, 1], \beta \in [0, 1].$$

For small values of d the above quantity can be computed numerically (see Fig. 2); in particular, for $d = 2, 3, 4$ this yields the same lower bounds as in Hansknecht et al. [18], since $n = 1$ is the optimal choice.

Next we prove the asymptotics $\alpha(d) = \Omega\left(\frac{\sqrt{d}}{\ln d}\right)$. To that end, we take the following choice of parameters:

$$w = \frac{\ln d}{2d}\ ,\ k = \left\lceil \frac{\ln d}{2\ln\ln d} \right\rceil\ ,\ \beta = \frac{1}{d^{\frac{k}{2(k+1)}}(1+w)^d}\ ,\ n = \left\lfloor \frac{1}{d^{\frac{1}{2(k+1)}} w} \right\rfloor.$$

One can check that this choice satisfies $k \in \{1, \ldots, d\}$ (for $d \geq 4$) and $w, \beta \in [0, 1]$. We can bound the expressions appearing in (3) as follows.

$$\begin{aligned}
1 + n\beta(1+w)^d &\geq 1 + \left(\frac{1}{d^{\frac{1}{2(k+1)}} w} - 1\right) \frac{1}{d^{\frac{k}{2(k+1)}}(1+w)^d}(1+w)^d \\
&= 1 + \left(\frac{2d}{d^{\frac{1}{2(k+1)}} \ln d} - 1\right) \frac{1}{d^{\frac{k}{2(k+1)}}} \\
&= \frac{2d}{d^{\frac{1}{2(k+1)} + \frac{k}{2(k+1)}} \ln d} + 1 - \frac{1}{d^{\frac{k}{2(k+1)}}} \\
&\geq \frac{2d}{d^{1/2} \ln d} && \text{(since } d \geq 1\text{)} \\
&= \frac{2\sqrt{d}}{\ln d}; && (4) \\
(1+nw)^k + n\beta &\leq \left(1 + \frac{1}{d^{\frac{1}{2(k+1)}} w} w\right)^k + \frac{1}{d^{\frac{1}{2(k+1)}} w d^{\frac{k}{2(k+1)}} (1+w)^d} \\
&= \left(1 + d^{-\frac{1}{2(k+1)}}\right)^k + \frac{1}{d^{1/2} \frac{\ln d}{2d} \left(1 + \frac{\ln d}{2d}\right)^d} \\
&= \left(1 + d^{-\frac{1}{2(k+1)}}\right)^k + \frac{2\sqrt{d}}{\ln d \left(1 + \frac{\ln d}{2d}\right)^d}; && (5) \\
(1+w)^k + \beta w^d &\geq 1; && (6) \\
(nw)^k + \beta(1+w)^d &\leq \left(\frac{1}{d^{\frac{1}{2(k+1)}} w} w\right)^k + \frac{1}{d^{\frac{k}{2(k+1)}} (1+w)^d}(1+w)^d \\
&= 2 \cdot \frac{1}{d^{\frac{k}{2(k+1)}}} = \frac{2d^{\frac{1}{2(k+1)}}}{\sqrt{d}} \leq \frac{2d^{\frac{\ln\ln d}{\ln d}}}{\sqrt{d}} = \frac{2\ln d}{\sqrt{d}}. && (7)
\end{aligned}$$

In the Appendix, we prove (Lemma 3) that the final quantity in (5) converges to 1 as $d \to \infty$; in particular, it is upper bounded by 4 for d large enough. Thus, we can lower bound the ratios of (3) as

$$\begin{aligned}
\frac{1 + n\beta(1+w)^d}{(1+nw)^k + n\beta} &\geq \frac{\frac{2\sqrt{d}}{\ln d}}{4} = \frac{\sqrt{d}}{2\ln d} = \Omega\left(\frac{\sqrt{d}}{\ln d}\right), && \text{(from (4), (5) and large } d\text{)} \\
\frac{(1+w)^k + \beta w^d}{(nw)^k + \beta(1+w)^d} &\geq \frac{1}{\frac{2\ln d}{\sqrt{d}}} = \frac{\sqrt{d}}{2\ln d} = \Omega\left(\frac{\sqrt{d}}{\ln d}\right). && \text{(from (6) and (7))}
\end{aligned}$$

This proves the asymptotics and the bound $\alpha(d) \geq \frac{\sqrt{d}}{2\ln d}$ for large d. □

4 The Hardness Gadget

In this section we construct an unweighted polynomial congestion game from a Boolean circuit. In the α-PNE of this game the players emulate the computation of the circuit. This gadget will

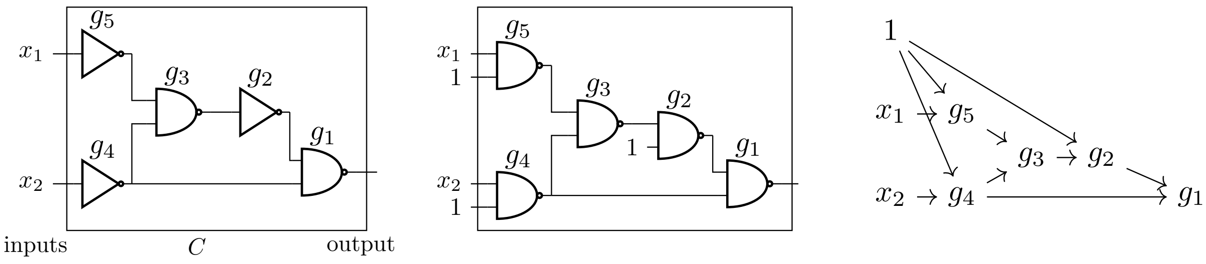

(a) valid circuit C (b) canonical form of C (c) directed acyclic graph

Figure 3: Example of a valid circuit C (having both NOT and NAND gates), its canonical form (having only NAND gates), and the directed acyclic graph corresponding to C.

be used in reductions from Circuit Satisfiability to show NP-hardness of several problems related to the existence of approximate equilibria with some additional properties. For example, deciding whether a congestion game has an α-PNE where a certain set of players choose a specific strategy profile (Theorem 2).

Circuit Model We consider Boolean circuits consisting of NOT gates and 2-input NAND gates only. We assume that the two inputs to every NAND gate are different. Otherwise we replace the NAND gate by a NOT gate, without changing the semantics of the circuit. We further assume that every input bit is connected to exactly one gate and this gate is a NOT gate. See Figure 3a for a *valid* circuit. In a valid circuit we replace every NOT gate by an equivalent NAND gate, where one of the inputs is fixed to 1. See the replacement of gates g_5, g_4 and g_2 in the example in Figure 3b. Thus, we look at circuits of 2-input NAND gates where both inputs to a NAND gate are different and every input bit of the circuit is connected to exactly one NAND gate where the other input is fixed to 1. A circuit of this form is said to be in *canonical form*. For a circuit C and a vector $x \in \{0,1\}^n$ we denote by $C(x)$ the output of the circuit on input x.

We model a circuit C in canonical form as a *directed acyclic graph*. The nodes of this graph correspond to the input bits $x_1, \dots, x_n$, the gates $g_1, \dots, g_K$ and a node 1 for all fixed inputs. There is an arc from a gate g to a gate g' if the output of g is input to gate g' and there are arcs from the fixed input and all input bits to the connected gates. We index the gates in reverse topological order, so that all successors of a gate g_k have a smaller index and the output of gate g_1 is the output of the circuit. Denote by $\delta^+(v)$ the set of the direct successors of node v. Then we have $|\delta^+(x_i)| = 1$ for all input bits x_i and $\delta^+(g_k) \subseteq \{g_{k'} \mid k' < k\}$ for every gate g_k. See Figure 3 for an example of a valid circuit, its canonical form and the corresponding directed acyclic graph.

Translation to Congestion Game Fix some integer $d \geq 2$ and a parameter $\mu > 1 + 3^{d+d/2}$. From a valid circuit in canonical form with input bits $x_1, \dots, x_n$, gates $g_1, \dots, g_K$ and the extra input fixed to 1, we construct a polynomial congestion game $\mathcal{G}_\mu^d$ of degree d. There are n *input players* $X_1, \dots, X_n$ for every input bit, a *static player* P for the input fixed to 1, and K *gate players* $G_1, \dots, G_K$ for the output bit of every gate. G_1 is sometimes called *output player* as g_1 corresponds to the output $C(x)$.

The idea is that every input and every gate player has a *zero* and a *one strategy*, corresponding to the respective bit being 0 or 1. In every α-PNE we want the players to emulate the computation of the circuit, i.e. the NAND semantics of the gates should be respected. For every gate g_k, we introduce two *resources* 0_k and 1_k. The zero (one) strategy of a player consists of the $0_{k'}$ ($1_{k'}$) resources of the direct successors in the directed acyclic graph corresponding to the circuit and its own 0_k (1_k) resource (for gate players). The static player has only one strategy playing all $1_{k'}$ resources of the gates where one input is fixed to 1.

Formally, we have

$$s^0_{X_i} = \left\{0_k \mid g_k \in \delta^+(x_i)\right\} \text{ and } s^1_{X_i} = \left\{1_k \mid g_k \in \delta^+(x_i)\right\}$$

for the zero and one strategy of an input player X_i. Recall that $\delta^+(x_i)$ is the set of direct successors of x_i, thus every strategy of an input player consists of exactly one resource. For a gate player G_k we have the two strategies

$$s^0_{G_k} = \{0_k\} \cup \left\{0_{k'} \mid g_{k'} \in \delta^+(g_k)\right\} \text{ and } s^1_{G_k} = \{1_k\} \cup \left\{1_{k'} \mid g_{k'} \in \delta^+(g_k)\right\}$$

consisting of at most k resources each. The strategy of the static player is $s_P = \{1_k \mid g_k \in \delta^+(1)\}$. Notice that all 3 players related to a gate g_k (gate player G_k and the two players corresponding to the input bits) are different and observe that every resource 0_k and 1_k can be played by exactly those 3 players.

We define the cost functions of the resources using parameter μ. The cost functions for resources 1_k are given c_{1_k} and for resources 0_k by c_{0_k}, where

$$c_{1_k}(x) = \mu^k x^d \qquad \text{and} \qquad c_{0_k}(x) = \lambda\mu^k x^d, \text{ with } \lambda = 3^{d/2}. \tag{8}$$

This construction is inspired by the lockable circuit games in Skopalik and Vöcking [33]. Our main contribution is the use of polynomial cost functions and a simplification of introducing only 2 resources per gate (and not 3). While Skopalik and Vöcking use these games in a PLS-reduction from Circuit/FLIP, we are interested in this gadget on its own.

Properties of the Gadget For a valid circuit C in canonical form consider the game $\mathcal{G}^d_\mu$ as defined above. We interpret any strategy profile $\mathbf{s}$ of the input players as a bit vector $x \in \{0,1\}^n$ by setting $x_i = 0$ if $s_{X_i} = s^0_{X_i}$ and $x_i = 1$ otherwise. The gate players are said to *follow the NAND semantics* in a strategy profile, if for every gate g_k the following holds:

- if both players corresponding to the input bits of g_k play their one strategy, then the gate player G_k plays her zero strategy;
- if at least one of the players corresponding to the input bits of g_k plays her zero strategy, then the gate player G_k plays her one strategy.

We show that for the right choice of α, the set of α-PNE in $\mathcal{G}^d_\mu$ is the same as the set of all strategy profiles where the gate players follow the NAND semantics.

Define

$$\varepsilon(\mu) = \frac{3^{d+d/2}}{\mu - 1}. \tag{9}$$

From our choice of μ, we obtain $3^{d/2} - \varepsilon(\mu) > 1$. For any valid circuit C in canonical form and a valid choice of μ the following lemma holds for $\mathcal{G}^d_\mu$.

Lemma 1. *Let $\mathbf{s}_X$ be any strategy profile for the input players $X_1, \dots, X_n$ and let $x \in \{0,1\}^n$ be the bit vector represented by $\mathbf{s}_X$. For any $\mu > 1 + 3^{d+d/2}$ and any $1 \le \alpha < 3^{d/2} - \varepsilon(\mu)$, there is a unique α-approximate PNE[3] in $\mathcal{G}^d_\mu$ where the input players play according to $\mathbf{s}_X$. In particular, in this α-PNE the gate players follow the NAND semantics, and the output player G_1 plays according to $C(x)$.*

Proof. Let $\mu > 1 + 3^{d+d/2}$ and $\alpha < 3^{d/2} - \varepsilon(\mu)$. First, we fix the input players to the strategies given by $\mathbf{s}_X$ and show that in any α-PNE every gate player follows the NAND semantics, as otherwise changing to the strategy corresponding to the NAND of its input bits is an α-improving

[3] Which, as a matter of fact, is actually also an *exact* PNE.

move. Second, we show that in any α-PNE where the gate players follow the NAND semantics, the input players have no incentive to change their strategy. In total we get that every strategy profile for the input players can be extended to an α-PNE, where the gate players emulate the circuit. Hence this α-PNE is unique.

Let $\mathbf{s}_X$ be any strategy profile for the input players $X_1, \ldots, X_n$ and let $\mathbf{s}$ be an α-PNE of $\mathcal{G}_\mu^d$ where the input players play according to $\mathbf{s}_X$. Take G_k to be any of the gate players and let P_a and P_b be the players corresponding to the input bits of gate g_k. Note that P_a and P_b can be other gate players or input players, and one of them can be the static player. To show that G_k follows the NAND semantics we consider two cases.

<u>Case 1:</u> *Both P_a and P_b play their one strategy in* $\mathbf{s}$. As both P_a and P_b play resource 1_k and all three players P_a, P_b and G_k are different, the cost of G_k's one strategy is at least $c_{1_k}(3)$. The cost of G_k's zero strategy is at most $c_{0_k}(1) + \sum_{k'=1}^{k-1} c_{0_{k'}}(3)$. Thus, we have

$$\frac{C_{G_k}(s_{G_k}^1, \mathbf{s}_{-G_k})}{C_{G_k}(s_{G_k}^0, \mathbf{s}_{-G_k})} \geq \frac{c_{1_k}(3)}{c_{0_k}(1) + \sum_{k'=1}^{k-1} c_{0_{k'}}(3)} = \frac{\mu^k 3^d}{\lambda\mu^k + \sum_{k'=1}^{k-1} \lambda\mu^{k'} 3^d} > \frac{3^d}{\lambda}\left(\frac{1}{1 + \frac{1}{\mu-1}3^d}\right),$$

where we used that $\frac{1}{\mu^k}\sum_{k'=1}^{k-1} \mu^{k'} = \frac{1}{\mu^k}\left(\frac{\mu^k-\mu}{\mu-1}\right) < \frac{1}{\mu-1}$. By the definition of λ (see (8)) and $\varepsilon(\mu)$ (see (9)), we obtain

$$\frac{3^d}{\lambda}\left(\frac{1}{1 + \frac{1}{\mu-1}3^d}\right) = 3^{d/2}\left(\frac{1}{1 + \frac{1}{\mu-1}3^d}\right) > 3^{d/2}\left(1 - \frac{1}{\mu-1}3^d\right) = 3^{d/2} - \varepsilon(\mu) > \alpha. \tag{10}$$

Hence, changing from the one to the zero strategy would be an α-improving move for G_k. Thus, G_k must follow the NAND semantics and play her zero strategy in $\mathbf{s}$.

<u>Case 2:</u> *At least one of P_a or P_b is playing her zero strategy in* $\mathbf{s}$. By similar arguments to the previous case, we obtain that the cost of G_k's zero strategy is at least $c_{0_k}(2)$ and the cost of the one strategy is at most $c_{1_k}(2) + \sum_{k'=1}^{k-1} c_{1_{k'}}(3)$. Then, we get that

$$\frac{C_{G_k}(s_{G_k}^0, \mathbf{s}_{-G_k})}{C_{G_k}(s_{G_k}^1, \mathbf{s}_{-G_k})} \geq \frac{c_{0_k}(2)}{c_{1_k}(2) + \sum_{k'=1}^{k-1} c_{1_{k'}}(3)} = \frac{\lambda\mu^k 2^d}{\mu^k 2^d + \sum_{k'=1}^{k-1} \mu^{k'} 3^d} > \lambda\left(\frac{1}{1 + \frac{1}{\mu-1}\left(\frac{3}{2}\right)^d}\right).$$

By the definition of λ and $\varepsilon(\mu)$, we obtain

$$\lambda\left(\frac{1}{1 + \frac{1}{\mu-1}\left(\frac{3}{2}\right)^d}\right) > 3^{d/2}\left(\frac{1}{1 + \frac{1}{\mu-1}3^d}\right) > 3^{d/2}\left(1 - \frac{1}{\mu-1}3^d\right) = 3^{d/2} - \varepsilon(\mu) > \alpha. \tag{11}$$

Hence, changing from the zero to the one strategy would be an α-improving move for G_k. Thus, G_k must follow the NAND semantics and play her one strategy in $\mathbf{s}$.

We just showed that, in an α-PNE, every gate player must follow the NAND semantics. This implies that there is *at most one* α-PNE where the input players play according to $\mathbf{s}_X$, since the NAND semantics uniquely define the strategy of the remaining players. To conclude the proof, we must argue that this yields in fact an α-PNE, meaning that the input players are also 'locked' to their strategies in $\mathbf{s}_X$ and have no incentive to deviate. To that end, let $\mathbf{s}$ be a strategy profile PNE of $\mathcal{G}_\mu^d$ where the gate players follow the NAND semantics and let X_i be any of the input players. Recall that every input bit x_i is connected to exactly one gate, say $g_k = g_{k(i)}$, while the other input is fixed to 1. To show that X_i does not have an incentive to change her strategy, we consider two cases.

<u>Case 1:</u> *X_i plays her zero strategy in* $\mathbf{s}$. As G_k follows the NAND semantics in $\mathbf{s}$ and the other input of g_k is fixed to 1, we know that G_k must be playing her one strategy. This incurs a cost of $c_{0_k}(1) = \lambda\mu^k$ to X_i. On the other hand, if X_i changed to her one strategy this would incur a cost of $c_{1_k}(3) = \mu^k 3^d$.

Case 2: X_i *plays her one strategy in* $\mathbf{s}$. As G_k follows the NAND semantics in $\mathbf{s}$ and the other input of g_k is fixed to 1, we know that G_k must be playing her zero strategy. Thus, incurring a cost of $c_{1_k}(2) = \mu^k 2^d$ for X_i. On the other hand, if X_i changed to her zero strategy this would incur a cost of $c_{0_k}(2) = \lambda\mu^k 2^d$.

In both cases it is α-dominating for X_i not to change her strategy, since $\alpha < 3^{d/2} - \varepsilon(\mu) < 3^{d/2} = 3^d/\lambda = \lambda$. □

We are now ready to show our main result of this section; using the circuit game described above, we show NP-hardness of deciding whether approximate equilibria with additional properties exist.

Theorem 2. *The following problems are NP-hard, even for* unweighted *polynomial congestion games of degree* $d \geq 2$*, for all* $\alpha \in [1, 3^{d/2})$ *and all* $z > 0$*:*

- *"Does there exist an* α*-approximate PNE in which a certain subset of players are playing a specific strategy profile?"*
- *"Does there exist an* α*-approximate PNE in which a certain resource is used by at least one player?"*
- *"Does there exist an* α*-approximate PNE in which a certain player has cost at most* z*?"*

Proof. For the first problem we reduce from Circuit Satisfiability: given a Boolean circuit with n input bits and one output bit, is there an assignment of the input bits where the output of the circuit is 1? This problem is NP-hard even for circuits consisting only of 2-input NAND gates [28]. Let C' be a Boolean circuit of 2-input NAND gates. We transform C' into a valid circuit C by connecting every input bit to a NOT gate and the output of this NOT gate to all gates connected to the input bit in C'. Thus, $C'(x) = C(\bar{x})$, where $\bar{x}$ denotes the vector obtained from $x \in \{0,1\}^n$ by flipping every bit. Hence, we have that C' is a YES-instance to Circuit Satisfiability if and only if C is a YES-instance.

Let $\alpha \in [1, 3^{d/2})$, then there is an $\varepsilon > 0$ with $\alpha < 3^{d/2} - \varepsilon$. We set $\mu = 1 + \frac{3^{d+d/2}}{\min\{\varepsilon,1\}}$. For this choice of μ, we obtain $\varepsilon(\mu) \leq \varepsilon$ and thus $3^{d/2} - \varepsilon(\mu) \geq 3^{d/2} - \varepsilon > \alpha$. From the canonical form of C' we construct[4] the game $\mathcal{G}^d_\mu$. The subset of players we are looking at is the output player G_1 and the specific strategy for G_1 is her one strategy $s^1_{G_1}$. We show that there is an α-PNE where G_1 plays $s^1_{G_1}$ if and only if C is a YES-instance to Circuit Satisfiability.

Suppose there is a bit vector $x \in \{0,1\}^n$ such that $C(x) = 1$. Let $\mathbf{s}_X$ be the strategy profile for the input players of $\mathcal{G}^d_\mu$ corresponding to x. Since $3^{d/2} - \varepsilon(\mu) > \alpha$, Lemma 1 holds for $\mathcal{G}^d_\mu$ and α. Hence, the profile $\mathbf{s}_X$ can be extended to an α-PNE where G_1 plays according to $C(x)$. Thus, there is an α-PNE where G_1 plays $s^1_{G_1}$.

On the other hand, suppose for all bit vectors $x \in \{0,1\}^n$ it holds $C(x) = 0$. Again, by Lemma 1 we know that for any choice of strategies for the input players, the only α-PNE is a profile where the gate players follow the NAND semantics. Thus in this case, G_1 is playing $s^0_{G_1}$ in any α-PNE.

To show NP-hardness of the second problem we reduce from the first part of this theorem. From the proof above we know that this problem is even NP-hard if the given subset of players consists only of one player. Let $(\mathcal{G}, P, s_P)$ be an instance of this problem, i.e. $\mathcal{G}$ is an unweighted polynomial congestion game of degree d, P is one of its players and s_P is a specific strategy for P. We construct a new game $\mathcal{G}'$ from $\mathcal{G}$ by adding a new resource r with cost 0 to the strategy

[4] To be precise, the description of the game $\mathcal{G}^d_\mu$ involves the quantities $\mu = 1 + \frac{3^{d+d/2}}{\min\{\varepsilon,1\}}$ and $\lambda = 3^{d/2}$, which in general might be irrational. In order to incorporate this game into our reduction, it is enough to choose a rational μ such that $\mu > 1 + \frac{3^{d+d/2}}{\min\{\varepsilon,1\}}$, and a rational λ such that $\alpha\left(1 + \frac{1}{\mu-1}3^d\right) < \lambda < \frac{3^d}{\alpha\left(1+\frac{1}{\mu-1}3^d\right)}$. In this way, $\mathcal{G}^d_\mu$ is described entirely via rational numbers, while preserving the inequalities in (10) and (11).

s_P. The rest of the game $\mathcal{G}$ stays unchanged. As r does not incur any cost, the set of α-PNE in $\mathcal{G}'$ and $\mathcal{G}$ are the same. Thus, there exists an α-PNE in $\mathcal{G}'$ where the resource r is used by at least one player if and only if there is an α-PNE in $\mathcal{G}$ where P plays s_P.

The hardness of the third problem is shown by a reduction from CIRCUIT SATISFIABILITY, similar to the proof for the first problem. For $\alpha \in [1, 3^{d/2})$ we choose μ as before, so that Lemma 1 holds for α and the game $\mathcal{G}_\mu^d$ for a suitable circuit. Let C' be an instance of CIRCUIT SATISFIABILITY. By negating the output of C', we obtain a circuit $\overline{C'}$. As before we transform $\overline{C'}$ to a valid circuit $\overline{C}$, so that $C'(x) = \neg\overline{C}(\bar{x})$ holds. From the canonical form of $\overline{C}$ we construct the game $\mathcal{G}_\mu^d$. Note that the output player G_1 of this game is the output of a gate, where one of the inputs is fixed to 1, as we negated the output of C' by connecting the output of C' to a NOT gate. We show that there is an α-PNE in $\mathcal{G}_\mu^d$, where G_1 has cost at most $\lambda\mu$, if and only if C' is a YES-instance to CIRCUIT SATISFIABILITY.

Suppose there is a bit vector $x \in \{0,1\}^n$ with $C'(x) = 1$, then there is a vector $\bar{x}$ with $\overline{C}(\bar{x}) = 0$. Let $\mathbf{s}_X$ be the strategy profile for the input players of $\mathcal{G}_\mu^d$ corresponding to $\bar{x}$. By Lemma 1 this profile can be extended to an α-PNE, where G_1 is playing her zero strategy. As the gate players follow the NAND semantics in this PNE, the cost of player G_1 is exactly $c_{0_1}(1) = \lambda\mu$.

If, on the other hand, for all bit vectors $x \in \{0,1\}^n$ we have $C'(x) = 0$, then for all $\bar{x} \in \{0,1\}^n$ we have $\overline{C}(\bar{x}) = 1$. Thus, using Lemma 1 we know that in every α-PNE G_1 plays her one strategy. As G_1 follows the NAND semantics in any α-PNE and the player corresponding to one of the inputs of g_1 is the static player, we obtain that the cost of G_1 is exactly $c_{1_1}(2) = \mu 2^d$. Noticing that $\lambda = 3^{d/2} < 2^d$, we have deduced the following: either C' is a YES-instance, and $\mathcal{G}_\mu^d$ has an α-PNE where G_1 has a cost of (at most) $\lambda\mu$; or C' is a NO-instance, and for every α-PNE of $\mathcal{G}_\mu^d$, G_1 has a cost of (at least) $2^d\mu$. This immediately implies that determining whether an α-PNE exists in which a certain player has cost at most z is NP-hard for $\lambda\mu < z < 2^d\mu$. To prove that the problem remains NP-hard for an arbitrary $z > 0$, simply take a rational c such that $c\lambda\mu < z < c2^d\mu$ and rescale all costs of the resources in $\mathcal{G}_\mu^d$ by c.

□

5 Hardness of Existence

In this section we show that it is NP-hard to decide whether a polynomial congestion game has an α-PNE. For this we use a black-box reduction: our hard instance is obtained by combining any (weighted) polynomial congestion game $\mathcal{G}$ without α-PNE (i.e., the game from Section 3) with the circuit gadget of the previous section. To achieve this, it would be convenient to make some assumptions on the game $\mathcal{G}$, which however do not influence the existence or nonexistence of approximate equilibria:

Structural Properties of $\mathcal{G}$ Without loss of generality, we assume that a weighted polynomial congestion game of degree d has the following structural properties.

- *No player has an empty strategy.* If, for some player i, $\emptyset \in S_i$, then this strategy would be α-dominating for i. Removing i from the game description would not affect the (non)existence of (approximate) equilibria[5].

- *No player has zero weight.* If a player i had zero weight, her strategy would not influence the costs of the strategies of the other players. Again, removing i from the game description would not affect the (non)existence of equilibria.

[5] By this we mean, if $\mathcal{G}$ has (resp. does not have) α-PNE, then $\tilde{\mathcal{G}}$, obtained by removing player i from the game, still has (resp. still does not have) α-PNE.

- *Each resource e has a monomial cost function with a strictly positive coefficient*, i.e. $c_e(x) = a_e x^{k_e}$ where $a_e > 0$ and $k_e \in \{0, \dots, d\}$. If a resource had a more general cost function $c_e(x) = a_{e,0} + a_{e,1}x + \dots + a_{e,d}x^d$, we could split it into at most $d+1$ resources with (positive) monomial costs, $c_{e,0}(x) = a_{e,0}$, $c_{e,1}(x) = a_{e,1}x$, ..., $c_{e,d}(x) = a_{e,d}x^d$. These monomial cost resources replace the original resource, appearing on every strategy that included e.

- *No resource e has a constant cost function.* If a resource e had a constant cost function $c_e(x) = a_{e,0}$, we could replace it by new resources having monomial cost. For each player i of weight w_i, replace resource e by a resource e_i with monomial cost $c_{e_i}(x) = \frac{a_{e,0}}{w_i}x$, that is used exclusively by player i on her strategies that originally had resource e. Note that $c_{e_i}(w_i) = a_{e,0}$, so that this modification does not change the player's costs, neither has an effect on the (non)existence of approximate equilibria. If a resource has cost function constantly equal to zero, we can simply remove it from the description of the game.

For a game having the above properties, we define the (strictly positive) quantities

$$a_{\min} = \min_{e\in E} a_e, \quad W = \sum_{i\in N} w_i, \quad c_{\max} = \sum_{e\in E} c_e(W). \tag{12}$$

Note that $c_{\max}$ is an upper bound on the cost of any player on any strategy profile.

Rescaling of $\mathcal{G}$ In our construction of the combined game we have to make sure that the weights of the players in $\mathcal{G}$ are smaller than the weights of the players in the circuit gadget. We introduce the following rescaling argument.

For any $\gamma \in (0, 1]$ define the game $\tilde{\mathcal{G}}_\gamma$, where we rescale the player weights and resource cost coefficients in $\mathcal{G}$ as

$$\tilde{w}_i = \gamma w_i, \quad \tilde{a}_e = \gamma^{d+1-k_e} a_e, \quad \tilde{c}_e(x) = \tilde{a}_e x^{k_e}. \tag{13}$$

This changes the quantities in (12) for $\tilde{\mathcal{G}}_\gamma$ as

$$\tilde{a}_{\min} \geq \gamma^d a_{\min}, \quad \tilde{W} = \gamma W, \quad \tilde{c}_{\max} = \gamma^{d+1} c_{\max}.$$

In $\tilde{\mathcal{G}}_\gamma$ the player costs are all uniformly scaled as $\tilde{C}_i(\mathbf{s}) = \gamma^{d+1} C_i(\mathbf{s})$, so that the Nash dynamics and the (non)existence of equilibria are preserved.

The next lemma formalizes the combination of both game gadgets and, furthermore, establishes the gap-introduction in the equilibrium factor. Using it, we will derive our key hardness tool of Theorem 3.

Lemma 2. *Fix any integer $d \geq 2$ and real $\alpha \geq 1$. Suppose there exists a weighted polynomial congestion game $\mathcal{G}$ of degree d that does* not *have an α-approximate PNE. Then, for any circuit C there exists a game $\tilde{\mathcal{G}}_C$ with the following property: the sets of α-approximate PNE and exact PNE of $\tilde{\mathcal{G}}_C$ coincide and are in one-to-one correspondence with the set of satisfying assignments of C. In particular, one of the following holds: either*

1. *C has a satisfying assignment, in which case $\tilde{\mathcal{G}}_C$ has an exact PNE (and thus, also an α-approximate PNE); or*

2. *C has no satisfying assignments, in which case $\tilde{\mathcal{G}}_C$ has no α-approximate PNE (and thus, also no exact PNE).*

Proof. Let $\mathcal{G}$ be a congestion game as in the statement of the theorem having the above mentioned structural properties. Recalling that weighted polynomial congestion games of degree d have d-PNE [3], this implies that $\alpha < d < 3^{d/2}$. Fix some $0 < \varepsilon < 3^{d/2} - \alpha$ and take $\mu \geq 1 + \frac{3^{d+d/2}}{\min\{\varepsilon,1\}}$; in this way $\alpha < 3^{d/2} - \varepsilon \leq 3^{d/2} - \varepsilon(\mu)$.

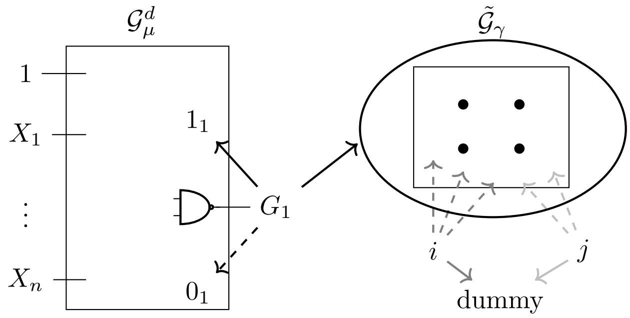

Figure 4: Merging a circuit game (on the left) and a game without approximate equilibria (on the right). Changes to the subgames are indicated by solid arrows. The new one strategy of G_1 consists of 1_1 and all resources in $\tilde{\mathcal{G}}_\gamma$, while the zero strategy stays unchanged. The players of $\tilde{\mathcal{G}}_\gamma$ get a new strategy (the dummy resource), and keep their old strategies playing in $\tilde{\mathcal{G}}_\gamma$.

Given a circuit C we construct the game $\tilde{\mathcal{G}}_C$ as follows. We combine the game $\mathcal{G}_\mu^d$ whose Nash dynamics model the NAND semantics of C, as described in Section 4, with the game $\tilde{\mathcal{G}}_\gamma$ obtained from $\mathcal{G}$ via the aforementioned rescaling. We choose $\gamma \in (0,1]$ sufficiently small such that the following three inequalities hold for the quantities in (12) for $\mathcal{G}$:

$$\gamma W < 1, \quad \gamma \sum_{e \in E} a_e < \frac{\mu}{\mu-1}\left(\frac{3}{2}\right)^d, \quad \gamma \alpha^2 < \frac{a_{\min}}{c_{\max}}. \tag{14}$$

Thus, the set of players in $\tilde{\mathcal{G}}_C$ corresponds to the (disjoint) union of the static, input and gate players in $\mathcal{G}_\mu^d$ (which all have weights 1) and the players in $\tilde{\mathcal{G}}_\gamma$ (with weights $\tilde{w}_i$). We also consider a new dummy resource with constant cost $c_{\text{dummy}}(x) = \frac{\tilde{a}_{\min}}{\alpha}$. Thus, the set of resources corresponds to the (disjoint) union of the gate resources $0_k, 1_k$ in $\mathcal{G}_\mu^d$, the resources in $\tilde{\mathcal{G}}_\gamma$, and the dummy resource. We augment the strategy space of the players as follows:

- each input player or gate player of $\mathcal{G}_\mu^d$ that is *not* the output player G_1 has the same strategies as in $\mathcal{G}_\mu^d$ (i.e. either the zero or the one strategy);
- the zero strategy of the output player G_1 is the same as in $\mathcal{G}_\mu^d$, but her one strategy is augmented with *every* resource in $\tilde{\mathcal{G}}_\gamma$; that is, $s_{G_1}^1 = \{1_1\} \cup E(\tilde{\mathcal{G}}_\gamma)$;
- each player i in $\tilde{\mathcal{G}}_\gamma$ keeps her original strategies as in $\tilde{\mathcal{G}}_\gamma$, and gets a new dummy strategy $s_{i,\text{dummy}} = \{\text{dummy}\}$.

A graphical representation of the game $\tilde{\mathcal{G}}_C$ can be seen in Fig. 4.

To finish the proof, we need to show that every α-PNE of $\tilde{\mathcal{G}}_C$ is an exact PNE and corresponds to a satisfying assignment of C; and, conversely, that every satisfying assignment of C gives rise to an exact PNE of $\tilde{\mathcal{G}}_C$ (and thus, an α-PNE as well).

Suppose that $\mathbf{s}$ is an α-PNE of $\tilde{\mathcal{G}}_C$, and let $\mathbf{s}_X$ denote the strategy profile restricted to the input players of $\mathcal{G}_\mu^d$. Then, as in the proof of Lemma 1, every gate player that is not the output player must respect the NAND semantics, and this is an α-dominating strategy. For the output player, either $\mathbf{s}_X$ is a non-satisfying assignment, in which case the zero strategy of G_1 was α-dominating, and this remains α-dominating in the game $\tilde{\mathcal{G}}_C$ (since only the cost of the one strategy increased for the output player); or $\mathbf{s}_X$ is a satisfying assignment. In the second case, we now argue that the one strategy of G_1 remains α-dominating. The cost of the output player on the zero strategy is at least $c_{0_1}(2) = \lambda\mu 2^d$, and the cost on the one strategy is at most

$$c_{1_1}(2) + \sum_{e \in E} \tilde{c}_e(1+\gamma W) = \mu 2^d + \sum_{e \in E} \gamma^{d+1-k_e} a_e (1+\gamma W)^{k_e} < \mu 2^d + \gamma \sum_{e \in E} a_e 2^d < \mu 2^d + \frac{\mu}{\mu-1} 3^d,$$

where we used the first and second bounds from (14). Thus, the ratio between the costs is at least

$$\frac{\lambda\mu 2^d}{\mu 2^d + \frac{\mu}{\mu-1}3^d} = \lambda\left(\frac{1}{1+\frac{1}{\mu-1}\left(\frac{3}{2}\right)^d}\right) > 3^{d/2}\left(\frac{1}{1+\frac{1}{\mu-1}3^d}\right) > 3^{d/2} - \varepsilon(\mu) > \alpha.$$

Given that the gate players must follow the NAND semantics, the input players are also locked to their strategies (i.e. they have no incentive to change) due to the proof of Lemma 1. The only players left to consider are the players from $\tilde{\mathcal{G}}_\gamma$. First we show that, since $\mathbf{s}$ is an α-PNE, the output player must be playing her one strategy. If this was not the case, then each dummy strategy of a player in $\tilde{\mathcal{G}}_\gamma$ is α-dominated by any other strategy: the dummy strategy incurs a cost of $\frac{\tilde{a}_{\min}}{\alpha} \geq \gamma^d \frac{a_{\min}}{\alpha}$, whereas any other strategy would give a cost of at most $\tilde{c}_{\max} = \gamma^{d+1} c_{\max}$ (this is because the output player is not playing any of the resources in $\tilde{\mathcal{G}}_\gamma$). The ratio between the costs is thus at least

$$\frac{\gamma^d a_{\min}}{\gamma^{d+1} c_{\max}\alpha} = \frac{a_{\min}}{\gamma c_{\max}\alpha} > \alpha.$$

Since the dummy strategies are α-dominated, the players in $\tilde{\mathcal{G}}_\gamma$ must be playing on their original sets of strategies. The only way for $\mathbf{s}$ to be an α-PNE would be if $\mathcal{G}$ had an α-PNE to begin with, which yields a contradiction. Thus, the output player is playing the one strategy (and hence, is present in every resource in $\tilde{\mathcal{G}}_\gamma$). In such a case, we can conclude that each dummy strategy is now α-dominating. If a player i in $\tilde{\mathcal{G}}_\gamma$ is not playing a dummy strategy, she is playing at least one resource in $\tilde{\mathcal{G}}_\gamma$, say resource e. Her cost is at least $\tilde{c}_e(1+\tilde{w}_i) = \tilde{a}_e(1+\tilde{w}_i)^{k_e} > \tilde{a}_e \geq \tilde{a}_{\min}$ (the strict inequality holds since, by the structural properties of our game, all of $\tilde{a}_e$, $\tilde{w}_i$ and k_e are strictly positive quantities). On the other hand, the cost of playing the dummy strategy is $\frac{\tilde{a}_{\min}}{\alpha}$. Thus, the ratio between the costs is greater than α.

We have concluded that, if $\mathbf{s}$ is an α-PNE of $\tilde{\mathcal{G}}_C$, then $\mathbf{s}_X$ corresponds to a satisfying assignment of C, all the gate players are playing according to the NAND semantics, the output player is playing the one strategy, and all players of $\tilde{\mathcal{G}}_\gamma$ are playing the dummy strategies. In this case, we also have observed that each player's current strategy is α-dominating, so the strategy profile is an exact PNE. To finish the proof, we need to argue that every satisfying assignment gives rise to a unique α-PNE. Let $\mathbf{s}_X$ be the strategy profile corresponding to this assignment for the input players in $\mathcal{G}^d_\mu$. Then, as before, there is one and exactly one α-PNE $\mathbf{s}$ in $\tilde{\mathcal{G}}_C$ that agrees with $\mathbf{s}_X$; namely, each gate player follows the NAND semantics, the output player plays the one strategy, and the players in $\tilde{\mathcal{G}}_\gamma$ play the dummy strategies. □

Theorem 3. *For any integer $d \geq 2$ and real $\alpha \geq 1$, suppose there exists a weighted polynomial congestion game which does not have an α-approximate PNE. Then it is NP-hard to decide whether (weighted) polynomial congestion games of degree d have an α-approximate PNE. If in addition α is polynomial-time computable,*[6] *the aforementioned problem is NP-complete.*

Proof. Let $d \geq 2$ and $\alpha \geq 1$. Let $\mathcal{G}$ be a weighted polynomial congestion game of degree d that has no α-PNE; this means that for every strategy profile $\mathbf{s}$ there exists a player i and a strategy $s'_i \neq s_i$ such that $C_i(s_i, \mathbf{s}_{-i}) > \alpha \cdot C_i(s'_i, \mathbf{s}_{-i})$. Note that the functions C_i are polynomials of degree d and hence they are continuous on the weights w_i and the coefficients a_e appearing on the cost functions. Hence, any arbitrarily small perturbation of the w_i, a_e does not change the sign of the above inequality. Thus, without loss of generality, we can assume that all w_i, a_e are rational numbers. By a similar reasoning, we can let $\bar{\alpha} > \alpha$ be a rational number sufficiently close to α such that $\mathcal{G}$ still does not have an $\bar{\alpha}$-PNE.

Next, we consider the game $\tilde{\mathcal{G}}_\gamma$ obtained from $\mathcal{G}$ by rescaling, as in the proof of Lemma 2, but with $\bar{\alpha}$ playing the role of α. Notice that the rescaling is done via the choice of a sufficiently small γ, according to (14), and hence in particular we can take γ to be a sufficiently small

[6] Recall the definition of polynomial-time computable real number at the end of Section 2.

rational. In this way, all the player weights and coefficients in the cost of resources are rational numbers scaled by a rational number and hence rationals.

Finally, we are able to provide the desired NP reduction from Circuit Satisfiability. Given a Boolean circuit C' built with 2-input NAND gates, transform it into a valid circuit C in canonical form. From C we can construct in polynomial time the game $\tilde{\mathcal{G}}_C$ as described in the proof of Lemma 2. The 'circuit part', i.e. the game $\mathcal{G}_\mu^d$, is obtained in polynomial time from C, as in the proof of Theorem 2; the description of the game $\tilde{\mathcal{G}}_\gamma$ involves only rational numbers, and hence the game can be represented by a constant number of bits (i.e. independent of the circuit C). Similarly, the additional dummy strategy has a constant delay of $\tilde{a}_{\min}/\bar{\alpha}$, and can be represented with a single rational number. Merging both $\mathcal{G}_\mu^d$ and $\tilde{\mathcal{G}}_\gamma$ into a single game $\tilde{\mathcal{G}}_C$ can be done in linear time. Since C has a satisfying assignment iff $\tilde{\mathcal{G}}_C$ has an α-PNE (or $\bar{\alpha}$-PNE), this concludes that the problem described is NP-hard.

If α is polynomial-time computable, the problem is clearly in NP: given a weighted polynomial congestion game of degree d and a strategy profile $\mathbf{s}$, one can check if $\mathbf{s}$ is an α-PNE by computing the ratios between the cost of each player in $\mathbf{s}$ and their cost for each possible deviation, and comparing these ratios with α. □

Combining the hardness result of Theorem 3 together with the nonexistence result of Theorem 1 we get the following corollary, which is the main result of this section.

Corollary 1. *For any integer $d \geq 2$ and real $\alpha \in [1, \alpha(d))$, it is NP-hard to decide whether (weighted) polynomial congestion games of degree d have an α-approximate PNE, where $\alpha(d) = \tilde{\Omega}(\sqrt{d})$ is the same as in Theorem 1. If in addition α is polynomial-time computable, the aforementioned problem is NP-complete.*

Notice that, in the proof of Lemma 2 and Theorem 3, we constructed a polynomial-time reduction from Circuit Satisfiability to the problem of determining whether a given congestion game has an α-PNE. Not only does this reduction map YES-instances of one problem to YES-instances of the other, but it also induces a bijection between the sets of satisfying assignments of a circuit C and α-PNE of the corresponding game $\tilde{\mathcal{G}}_C$. That is, this reduction is *parsimonious*. As a consequence, we can directly lift hardness of problems associated with counting satisfying assignments to Circuit Satisfiability into problems associated with counting α-PNE of congestion games:

Corollary 2. *Let $k \geq 1$ and $d \geq 2$ be integers and $\alpha \in [1, \alpha(d))$ where $\alpha(d) = \tilde{\Omega}(\sqrt{d})$ is the same as in Theorem 1. Then*

- *it is #P-hard to count the number of α-approximate PNE of (weighted) polynomial congestion games of degree d;*
- *it is NP-hard to decide whether a (weighted) polynomial congestion game of degree d has at least k distinct α-approximate PNE.*

Proof. The hardness of the first problem comes from the #P-hardness of the counting version of Circuit Satisfiability (see, e.g., [28, Ch. 18]). For the hardness of the second problem, it is immediate to see that the following problem is NP-complete, for any fixed integer $k \geq 1$: given a circuit C, decide whether there are at least k distinct satisfying assignments for C (simply add "dummy" variables to the description of the circuit). □

6 General Cost Functions

In this final section we leave the domain of polynomial latencies and study the existence of approximate equilibria in general congestion games having arbitrary (nondecreasing) cost functions. Our parameter of interest, with respect to which both our positive and negative results are going to be stated, is the number of players n. We start by showing that n-PNE always exist:

Theorem 4. *Every weighted congestion game with n players and arbitrary (nondecreasing) cost functions has an n-approximate PNE.*

Proof. Fix a weighted congestion game with $n \geq 2$ players, some strategy profile $\mathbf{s}$, and a possible deviation s'_i of player i. First notice that we can bound the change in the cost of any other player $j \neq i$ as

$$\begin{aligned} C_j(s'_i, \mathbf{s}_{-i}) - C_j(\mathbf{s}) &= \sum_{e \in s_j} c_e(x_e(s'_i, \mathbf{s}_{-i})) - \sum_{e \in s_j} c_e(x_e(\mathbf{s})) \\ &= \sum_{e \in s_j \cap (s'_i \setminus s_i)} \left[c_e(x_e(s'_i, \mathbf{s}_{-i})) - c_e(x_e(\mathbf{s}))\right] \\ &\quad + \sum_{e \in s_j \cap (s_i \setminus s'_i)} \left[c_e(x_e(s'_i, \mathbf{s}_{-i})) - c_e(x_e(\mathbf{s}))\right] \qquad (15) \\ &\leq \sum_{e \in s_j \cap (s'_i \setminus s_i)} \left[c_e(x_e(s'_i, \mathbf{s}_{-i})) - c_e(x_e(\mathbf{s}))\right] \\ &\leq \sum_{e \in s'_i} c_e(x_e(s'_i, \mathbf{s}_{-i})) \\ &= C_i(s'_i, \mathbf{s}_{-i}), \qquad (16) \end{aligned}$$

the first inequality holding due to the fact that the second sum in (15) contains only nonpositive terms (since the latency functions are nondecreasing).

Next, define the social cost $C(\mathbf{s}) = \sum_{i \in N} C_i(\mathbf{s})$. Adding the above inequality over all players $j \neq i$ (of which there are $n-1$) and rearranging, we successively derive:

$$\begin{aligned} \sum_{j \neq i} C_j(s'_i, \mathbf{s}_{-i}) - \sum_{j \neq i} C_j(\mathbf{s}) &\leq (n-1) C_i(s'_i, \mathbf{s}_{-i}) \\ \left(C(s'_i, \mathbf{s}_{-i}) - C_i(s'_i, \mathbf{s}_{-i})\right) - \left(C(\mathbf{s}) - C_i(\mathbf{s})\right) &\leq (n-1) C_i(s'_i, \mathbf{s}_{-i}) \\ C(s'_i, \mathbf{s}_{-i}) - C(\mathbf{s}) &\leq n C_i(s'_i, \mathbf{s}_{-i}) - C_i(\mathbf{s}). \qquad (17) \end{aligned}$$

We conclude that, if s'_i is an n-improving deviation for player i (i.e., $nC_i(s'_i, \mathbf{s}_{-i}) < C_i(\mathbf{s})$), then the social cost must strictly decrease after this move. Thus, any (global or local) minimizer of the social cost must be an n-PNE (the existence of such a minimizer is guaranteed by the fact that the strategy spaces are finite). □

The above proof not only establishes the existence of n-approximate equilibria in general congestion games, but also highlights a few additional interesting features. First, due to the key inequality (17), n-PNE are reachable via sequences of n-improving moves, in addition to arising also as minimizers of the social cost function. These attributes give a nice "constructive" flavour to Theorem 4. Secondly, exactly because social cost optima are n-PNE, the *Price of Stability*[7] of n-PNE is optimal (i.e., equal to 1) as well. Another, more succinct way, to interpret these observations is within the context of *approximate potentials* (see, e.g., [6, 8, 9]); (17) establishes that the social cost itself is always an n-approximate potential of any congestion game.

Next, we design a family of games that do not admit $\Theta\left(\frac{n}{\ln n}\right)$-PNE, thus nearly matching the upper bound Theorem 4.

[7]The Price of Stability (PoS) is a well-established and extensively studied notion in algorithmic game theory, originally studied in [2, 10]. It captures the minimum approximation ratio of the social cost between equilibria and the optimal solution (see, e.g., [7, 9]); in other words, it is the best-case analogue of the the Price of Anarchy (PoA) notion of Koutsoupias and Papadimitriou [23].

Theorem 5. *For any integer $n \geq 2$, there exist weighted congestion games with n players and general cost functions that do not have α-approximate PNE for any $\alpha < \Phi_{n-1}$, where $\Phi_m \sim \frac{m}{\ln m}$ is the unique positive solution of $(x+1)^m = x^{m+1}$.*

Proof. For any integer $n \geq 2$, let $\xi = \Phi_{n-1}$ be the positive solution of $(x+1)^{n-1} = x^n$. Then, equivalently,

$$\left(1 + \frac{1}{\xi}\right)^{n-1} = \xi. \tag{18}$$

Furthermore, as we mentioned in Section 2, $\xi > 1$ and asymptotically $\Phi_{n-1} \sim \frac{n}{\ln n}$.

Consider the following congestion game $\mathcal{G}_n$. There are $n = m+1$ players $0, 1, \ldots, m$, where player i has weight $w_i = 1/2^i$. In particular, this means that for any $i \in \{1, \ldots, m\}$: $\sum_{k=i}^m w_k < w_{i-1} \leq w_0$. Furthermore, there are $2(m+1)$ resources $a_0, a_1, \ldots, a_m, b_0, b_1, \ldots, b_m$, where resources a_i and b_i have the same cost function c_i given by

$$c_{a_0}(x) = c_{b_0}(x) = c_0(x) = \begin{cases} 1, & \text{if } x \geq w_0, \\ 0, & \textit{otherwise}; \end{cases}$$

and for all $i \in \{1, \ldots, m\}$,

$$c_{a_i}(x) = c_{b_i}(x) = c_i(x) = \begin{cases} \frac{1}{\xi}\left(1 + \frac{1}{\xi}\right)^{i-1}, & \text{if } x \geq w_0 + w_i, \\ 0, & \textit{otherwise.} \end{cases}$$

The strategy set of player 0 and of all players $i \in \{1, \ldots, m\}$ are, respectively,

$$S_0 = \{\{a_0, \ldots, a_m\}, \{b_0, \ldots, b_m\}\}, \qquad \text{and} \qquad S_i = \{\{a_0, \ldots, a_{i-1}, b_i\}, \{b_0, \ldots, b_{i-1}, a_i\}\}.$$

We show that this game has no α-PNE, for any $\alpha < \xi$, by proving that in any outcome there is at least one player that can deviate and improve her cost by a factor of at least ξ. Due to symmetry it is sufficient to consider the following two kinds of outcomes:

<u>Case 1:</u> *Player 0 is alone on resource a_0.*

Then player 0 must have chosen $\{a_0, \ldots, a_m\}$, and all other players $i \in \{1, \ldots, m\}$ must have chosen strategy $\{b_0, \ldots, b_{i-1}, a_i\}$. In this outcome, player 0 has a cost of

$$c_0(w_0) + \sum_{i=1}^m c_i(w_0 + w_i) = 1 + \frac{1}{\xi}\sum_{i=1}^m \left(1 + \frac{1}{\xi}\right)^{i-1} = \left(1 + \frac{1}{\xi}\right)^m = \xi,$$

where the last equality follows by the fact that $m = n-1$ and (18). Deviating to $\{b_0, \ldots, b_m\}$, player 0 would get a cost of

$$c_0(w_0 + \ldots + w_m) + \sum_{i=1}^m c_i\left(w_0 + \sum_{j=i+1}^m w_j\right) = 1 + 0.$$

Thus player 0 can improve by a factor of at least ξ.

<u>Case 2:</u> *Player 0 is sharing resource a_0 with at least one other player.*

Let j be the smallest index of such a player, i.e., player j plays $\{a_0, \ldots, a_{j-1}, b_j\}$ and all players $i \in \{1, \ldots, j-1\}$ have chosen strategy $\{b_0, \ldots, b_{i-1}, a_i\}$. In such a profile the cost of player j is at least

$$c_0(w_0 + w_j) + \sum_{i=1}^{j-1} c_i(w_0 + w_i + w_j) = 1 + \frac{1}{\xi}\sum_{i=1}^{j-1}\left(1 + \frac{1}{\xi}\right)^{i-1} = \left(1 + \frac{1}{\xi}\right)^{j-1},$$

while deviating to j's other strategy would result in a cost of at most

$$c_0\left(\sum_{i=1}^m w_i\right) + \sum_{i=1}^{j-1} c_i\left(\sum_{k=i+1}^m w_k\right) + c_j\left(w_0 + \sum_{k=j}^m w_k\right) = 0 + 0 + \frac{1}{\xi}\left(1 + \frac{1}{\xi}\right)^{j-1}.$$

Thus player j can improve by a factor of at least ξ.

$\square$

Similar to the spirit of the rest of our paper so far, we'd like to show an NP-hardness result for deciding existence of α-PNE for general games as well. We do exactly that in the following theorem, where now α grows as $\tilde{\Theta}(n)$. Again, we use the circuit gadget and combine it with the game from the previous nonexistence Theorem 5. The main difference to the previous reductions is that now n is part of the input. On the other hand we are not restricted to polynomial latencies, so we use step functions having a single breakpoint.

Theorem 6. *Let $\varepsilon > 0$, and let $\tilde{\alpha} : \mathbb{N}_{\geq 2} \longrightarrow \mathbb{R}$ be any sequence of reals such that $1 \leq \tilde{\alpha}(n) < \frac{\Phi_{n-1}}{1+\varepsilon} = \tilde{\Theta}(n)$, where $\Phi_m \sim \frac{m}{\ln m}$ is the unique positive solution of $(x+1)^m = x^{m+1}$. Then, it is NP-hard to decide whether a (weighted) congestion game with n players has an $\tilde{\alpha}(n)$-approximate PNE. If in addition $\tilde{\alpha}$ is a polynomial-time computable real sequence (as defined in Section 2), the aforementioned problem is NP-complete.*

Proof. Recall that we have $\Phi_{n-1} \sim \frac{n}{\ln n}$. Given $\varepsilon > 0$, without loss of generality assume $\varepsilon < 1$, so that $1 + \varepsilon/3 < (1+\varepsilon)(1 - \varepsilon/3)$. Let $n_0, \ell \in \mathbb{N}$ be large enough natural numbers such that

$$1 + \frac{1}{\ell} < \frac{(1+\varepsilon)(1-\frac{\varepsilon}{3})}{1+\frac{\varepsilon}{3}} \qquad \text{and} \qquad \left(1 - \frac{\varepsilon}{3}\right)\frac{n}{\ln n} \leq \Phi_{n-1} \leq \left(1 + \frac{\varepsilon}{3}\right)\frac{n}{\ln n} \quad \text{for all } n \geq n_0. \tag{19}$$

We will again reduce from Circuit Satisfiability: given a circuit C, we must construct (in polynomial time) a game $\tilde{\mathcal{G}}$, say with $\tilde{n}$ players, that has an $\tilde{\alpha}(\tilde{n})$-PNE if and only if C has a satisfying assignment. Without loss of generality assume that C is in canonical form (as described in Section 4); add also one extra gate that negates the output of C, making this the new output of a circuit $\bar{C}$, say with m inputs and K NAND gates. Let $s = m + K + 1$, $n = \ell s$, and take a large enough integer d such that $3^{d/2} > \Phi_{n-1}$. Note that s, n and a suitable d can all be found in time polynomial in the description of C. To conclude the preliminaries of this proof, assume also without loss of generality that $s \geq n_0$; if s is bounded by a constant, determining whether C has a satisfying assignment can be done in constant time.

Next, given $\bar{C}$ and d, construct the game $\mathcal{G}_\mu^d$ where μ is such that $3^{d/2} - \varepsilon(\mu) > \Phi_{n-1}$, as in Section 4. Notice that $\mathcal{G}_\mu^d$ can be computed in polynomial time from C, and that the Φ_{n-1}-improving Nash dynamics of this game emulate the computation of the circuit. Consider also the game $\mathcal{G}_n$ with n players from Theorem 5 that does not have α-PNE for any $\alpha < \Phi_{n-1}$.

We would like to merge $\mathcal{G}_\mu^d$ and $\mathcal{G}_n$ into a single game $\tilde{\mathcal{G}}$, in such a way that $\tilde{\mathcal{G}}$ has an approximate PNE if and only if C has a satisfying assignment. Following the same technique as in Lemma 2, we would like to extend the strategies of the output player of $\mathcal{G}_\mu^d$ to include resources that are used by players in $\mathcal{G}_n$. For this technique to work, we must rescale the weights and cost functions in $\mathcal{G}_n$. In particular, we divide all weights of the players in $\mathcal{G}_n$ by 2 (so that the sum of the weights of all the players is less than 1) and halve the breakpoints of the cost functions accordingly. We also add a new dummy resource with cost function

$$c_{\text{dummy}}(x) = \begin{cases} \Phi_{n-1}^2, & \text{if } x \geq 1, \\ 0, & \text{otherwise;} \end{cases}$$

We are now ready to describe the congestion game $\tilde{\mathcal{G}}$ that is obtained by merging the circuit game $\mathcal{G}_\mu^d$ with the (rescaling of) game $\mathcal{G}_n$. Note that this game has $n + s = (\ell+1)s$ players: s from the circuit game (which all have weight 1) and n from the nonexistence gadget. The set of resources corresponds to the union of the gate resources of $\mathcal{G}_\mu^d$, the resources in $\mathcal{G}_n$, and the dummy resource. Similarly to the proof of Lemma 2,

- we do not change the strategies of the players in $\mathcal{G}_\mu^d$, with the exception of the output player G_1;
- the zero strategy of the output player G_1 remains the same as in $\mathcal{G}_\mu^d$, but her one strategy is augmented with the dummy resource; that is, $s_{G_1}^1 = \{1_1, \text{dummy}\}$;
- each player i in $\mathcal{G}_n$ keeps her original strategies, and gets a new dummy strategy $s_{i,\text{dummy}} = \{\text{dummy}\}$.

With the above description,[8] the only thing left to prove NP-hardness is that C has a satisfying assignment if and only if $\tilde{\mathcal{G}}$ has an $\tilde{\alpha}(n+s)$-PNE. The proof follows the same approach as in Lemma 2. Letting $\alpha < \Phi_{n-1}$, we suppose that $\tilde{\mathcal{G}}$ has an α-PNE, say $\mathbf{s}$, and proceed to prove that C has a satisfying assignment.

As before, if $\mathbf{s}$ is an α-PNE, then every gate player that is not the output player must respect the NAND semantics, and this strategy is α-dominating. For the output player, the cost of her zero strategy remains the same, and the cost of her one strategy increases by exactly $\Phi_{n-1}^2 < 3^d < \frac{\mu}{\mu-1}3^d$. Hence, if $\mathbf{s}_X$ is a satisfying assignment, then the zero strategy of the output player (which negates the output of the original circuit C) remains α-dominating; on the other hand, if $\mathbf{s}_X$ is not a satisfying assignment, then the ratio between the costs of the zero strategy and the one strategy of the output player is at least

$$\frac{c_{0_1}(2)}{c_{1_1}(2)+\Phi_{n-1}^2} > \frac{\lambda\mu 2^d}{\mu 2^d + \frac{\mu}{\mu-1}3^d} = \lambda\left(\frac{1}{1+\frac{1}{\mu-1}\left(\frac{3}{2}\right)^d}\right) > 3^{d/2}\left(\frac{1}{1+\frac{1}{\mu-1}3^d}\right) > 3^{d/2} - \varepsilon(\mu) > \alpha.$$

Hence, respecting the NAND semantics remains α-dominating for the output player as well. As a consequence, the input players are also locked to their strategies (i.e. they have no incentive to change).

Now, if the output player happened to be playing her one strategy, this could not be an α-PNE. For each of the players in $\mathcal{G}_n$, the dummy strategy would incur a cost of Φ_{n-1}^2, whereas any other strategy would give a cost of at most Φ_{n-1}. Thus the dummy strategy would be Φ_{n-1}-dominated, and the players in $\mathcal{G}_n$ must be playing on their original sets of strategies, for which we know that α-PNE do not occur.

The above argument proves that, in an α-PNE, the output player must be playing her zero strategy. Since the output player, by construction, negates the output of C, this implies that C must have a satisfying assignment. This also implies that the congestion on the dummy resource cannot reach the breakpoint of 1, and hence it would be α-dominating for each of the players in $\mathcal{G}_n$ to play her dummy strategy (and incur a cost of 0). Thus, $\mathbf{s}$ is an exact PNE as well.

For the converse direction, suppose C has a satisfying assignment $\mathbf{s}_X$. Then this can be extended to an α-PNE of $\tilde{\mathcal{G}}$ in which the input players play according to $\mathbf{s}_X$, the gate players play according to the NAND semantics, the output player of $\mathcal{G}_\mu^d$ plays the zero strategy, and each player in $\mathcal{G}_n$ plays the dummy strategy.

We have proven that, for any $\alpha < \Phi_{n-1}$, C has a satisfying assignment iff $\tilde{\mathcal{G}}$ has an α-PNE. To conclude the proof, we verify that $\tilde{\alpha}(n+s) < \Phi_{n-1}$:

$$\begin{aligned}\tilde{\alpha}(n+s) &< \frac{\Phi_{n+s-1}}{1+\varepsilon} \\ &\leq \frac{1+\frac{\varepsilon}{3}}{1+\varepsilon}\frac{n+s}{\ln(n+s)}\end{aligned}$$

[8]This almost concludes the description of the game – the only problem is that some of the cost functions of the game are defined in terms of Φ_{n-1}, which is not a rational number. To make the proof formally correct, one can approximate Φ_{n-1} sufficiently close by a rational $\bar{\Phi}_{n-1} < \Phi_{n-1}$. We omit the details.

$$\leq \frac{1+\frac{\varepsilon}{3}}{(1+\varepsilon)(1-\frac{\varepsilon}{3})} \frac{(n+s)\ln n}{n\ln(n+s)} \Phi_{n-1}$$
$$< \frac{(1+\frac{s}{n})(1+\frac{\varepsilon}{3})}{(1+\varepsilon)(1-\frac{\varepsilon}{3})} \Phi_{n-1}$$
$$= \frac{(1+\frac{1}{\ell})(1+\frac{\varepsilon}{3})}{(1+\varepsilon)(1-\frac{\varepsilon}{3})} \Phi_{n-1} < \Phi_{n-1}.$$

The first inequality comes from the assumption on $\tilde{\alpha}$, the second and third come from the upper and lower bounds on Φ_n from (19) and the fact that $n+s \geq n \geq s \geq n_0$, the fourth comes from the trivial bound $\ln n < \ln(n+s)$, the equality comes from the definition of $n = \ell s$, and the final step comes from the choice of ℓ in (19).

We conclude that the problem of deciding whether a (weighted) congestion game with n players has an $\tilde{\alpha}(n)$-PNE is NP-hard. If in addition $\tilde{\alpha}$ is a polynomial-time computable real sequence, the problem is also in NP; given a game with n players and a (candidate) strategy profile, verify that this is an $\tilde{\alpha}(n)$-PNE by iterating over all possible moves of all players and verifying that none of these are $\tilde{\alpha}(n)$-improving. □

7 Discussion and Future Directions

In this paper we showed that weighted congestion games with polynomial latencies of degree d do not have α-PNE for $\alpha < \alpha(d) = \Omega\left(\frac{\sqrt{d}}{\ln d}\right)$. For general cost functions, we proved that n-PNE always exist whereas α-PNE in general do not, where n is the number of players and $\alpha < \Phi_{n-1} = \Theta\left(\frac{n}{\ln n}\right)$. We also transformed the nonexistence results into complexity-theoretic results, establishing that deciding whether such α-PNE exist is itself an NP-hard problem.

We now identify two possible directions for follow-up work. A first obvious question would be to reduce the nonexistence gap between $\Omega\left(\frac{\sqrt{d}}{\ln d}\right)$ (derived in Theorem 1 of this paper) and d (shown in [3]) for polynomials of degree d; similarly for the gap between $\Theta\left(\frac{n}{\ln n}\right)$ (Theorem 5) and n (Theorem 4) for general cost functions and n players. Notice that all current methods for proving upper bounds (i.e., existence) are essentially based on potential function arguments; thus it might be necessary to come up with novel ideas and techniques to overcome the current gaps.

A second direction would be to study the complexity of *finding* α-PNE, when they are guaranteed to exist. For example, for polynomials of degree d, we know that d-improving dynamics eventually reach a d-PNE [3], and so finding such an approximate equilibrium lies in the complexity class PLS of local search problems (see, e.g., [21, 32]). However, from a complexity theory perspective the only known lower bound is the PLS-completeness of finding an *exact* equilibrium for *unweighted* congestion games [12] (and this is true even for $d = 1$, i.e., affine cost functions; see [1]). On the other hand, we know that $d^{O(d)}$-PNE can be computed in polynomial time (see, e.g., [5, 13, 16]). It would be then very interesting to establish a "gradation" in complexity (e.g., from NP-hardness to PLS-hardness to P) as the parameter α increases from 1 to $d^{O(d)}$.

References

[1] H. Ackermann, H. Röglin, and B. Vöcking. On the impact of combinatorial structure on congestion games. *Journal of the ACM*, 55(6):1–22, Dec. 2008. doi:10.1145/1455248.1455249.

[2] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008. doi:10.1137/070680096.

[3] I. Caragiannis and A. Fanelli. On approximate pure Nash equilibria in weighted congestion games with polynomial latencies. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 133:1–133:12, 2019. doi:10.4230/LIPIcs.ICALP.2019.133.

[4] I. Caragiannis, A. Fanelli, N. Gravin, and A. Skopalik. Efficient computation of approximate pure Nash equilibria in congestion games. In *Proceedings of the 52nd IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 532–541, 2011. doi:10.1109/focs.2011.50.

[5] I. Caragiannis, A. Fanelli, N. Gravin, and A. Skopalik. Approximate pure nash equilibria in weighted congestion games: Existence, efficient computation, and structure. *ACM Trans. Econ. Comput.*, 3(1):2:1–2:32, Mar. 2015. doi:10.1145/2614687.

[6] H.-L. Chen and T. Roughgarden. Network design with weighted players. *Theory of Computing Systems*, 45(2):302, 2008. doi:10.1007/s00224-008-9128-8.

[7] G. Christodoulou and M. Gairing. Price of stability in polynomial congestion games. *ACM Transactions on Economics and Computation*, 4(2):1–17, 2015. doi:10.1145/2841229.

[8] G. Christodoulou, E. Koutsoupias, and P. G. Spirakis. On the performance of approximate equilibria in congestion games. *Algorithmica*, 61(1):116–140, 2011. doi:10.1007/s00453-010-9449-2.

[9] G. Christodoulou, M. Gairing, Y. Giannakopoulos, and P. G. Spirakis. The price of stability of weighted congestion games. *SIAM Journal on Computing*, 48(5):1544–1582, 2019. doi:10.1137/18M1207880.

[10] J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):961–976, 2004. doi:10.1287/moor.1040.0098.

[11] J. Dunkel and A. S. Schulz. On the complexity of pure-strategy Nash equilibria in congestion and local-effect games. *Mathematics of Operations Research*, 33(4):851–868, 2008. doi:10.1287/moor.1080.0322.

[12] A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure nash equilibria. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*, pages 604–612, 2004. doi:10.1145/1007352.1007445.

[13] M. Feldotto, M. Gairing, G. Kotsialou, and A. Skopalik. Computing approximate pure nash equilibria in shapley value weighted congestion games. In *Web and Internet Economics*, pages 191–204. 2017. doi:10.1007/978-3-319-71924-5_14.

[14] D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish unsplittable flows. *Theoretical Computer Science*, 348(2):226–239, 2005. doi:10.1016/j.tcs.2005.09.024.

[15] D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The structure and complexity of Nash equilibria for a selfish routing game. *Theoretical Computer Science*, 410(36):3305–3326, 2009. doi:10.1016/j.tcs.2008.01.004.

[16] Y. Giannakopoulos, G. Noarov, and A. S. Schulz. An improved algorithm for computing approximate equilibria in weighted congestion games. *CoRR*, abs/1810.12806, Oct. 2018. URL `https://arxiv.org/abs/1810.12806`.

[17] M. Goemans, V. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 142–151, 2005. doi:10.1109/SFCS.2005.68.

[18] C. Hansknecht, M. Klimm, and A. Skopalik. Approximate pure Nash equilibria in weighted congestion games. In *Proceedings of APPROX/RANDOM*, pages 242–257, 2014. doi:10.4230/LIPIcs.APPROX-RANDOM.2014.242.

[19] T. Harks and M. Klimm. On the existence of pure Nash equilibria in weighted congestion games. *Mathematics of Operations Research*, 37(3):419–436, 2012. doi:10.1287/moor.1120.0543.

[20] T. Harks, M. Klimm, and R. H. Möhring. Strong equilibria in games with the lexicographical improvement property. *International Journal of Game Theory*, 42(2):461–482, 2012. doi:10.1007/s00182-012-0322-1.

[21] D. S. Johnson, C. H. Papadimitriou, and M. Yannakakis. How easy is local search? *Journal of Computer and System Sciences*, 37(1):79–100, 1988. doi:10.1016/0022-0000(88)90046-3.

[22] K.-I. Ko. On the definitions of some complexity classes of real numbers. *Mathematical systems theory*, 16(1):95–109, 1983. doi:10.1007/BF01744572.

[23] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3 (2):65–69, 2009. doi:10.1016/j.cosrev.2009.04.003.

[24] L. Libman and A. Orda. Atomic resource sharing in noncooperative networks. *Telecommunication Systems*, 17(4):385–409, Aug. 2001. doi:10.1023/A:1016770831869.

[25] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[26] F. W. J. Olver, D. W. Lozier, R. F. Boisvert, and C. W. Clark. *NIST Handbook of Mathematical Functions*. Cambridge University Press, 2010. URL `http://dlmf.nist.gov/4.5.E13`.

[27] P. N. Panagopoulou and P. G. Spirakis. Algorithms for pure Nash equilibria in weighted congestion games. *Journal of Experimental Algorithmics*, 11:27, Feb. 2007. doi:10.1145/1187436.1216584.

[28] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994. ISBN 0201530821.

[29] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973. doi:10.1007/BF01737559.

[30] R. W. Rosenthal. The network equilibrium problem in integers. *Networks*, 3(1):53–59, 1973. doi:10.1002/net.3230030104.

[31] T. Roughgarden. *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press, 2016.

[32] A. A. Schäffer and M. Yannakakis. Simple local search problems that are hard to solve. *SIAM Journal on Computing*, 20(1):56–87, 1991. doi:10.1137/0220004.

[33] A. Skopalik and B. Vöcking. Inapproximability of pure Nash equilibria. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 355–364, 2008. doi:10.1145/1374376.1374428.

A Technical Lemmas

Lemma 3. *For any integer $d \geq 2$ define the sequence*

$$g(d) = \left(1 + d^{-\frac{1}{2(k_d+1)}}\right)^{k_d} + \frac{2\sqrt{d}}{\ln d \left(1 + \frac{\ln d}{2d}\right)^d} \qquad \textit{where } k_d = \left\lceil \frac{\ln d}{2 \ln \ln d} \right\rceil.$$

Then $\lim_{d\to\infty} g(d) = 1$.

Proof. Define

$$g_1(d) = \left(1 + d^{-\frac{1}{2(k_d+1)}}\right)^{k_d} \quad \text{and} \quad g_2(d) = \frac{2\sqrt{d}}{\left(1 + \frac{\ln d}{2d}\right)^d \ln d}$$

so that $g(d) = g_1(d) + g_2(d)$. We will show the desired convergence by establishing that $\lim_{d\to\infty} g_1(d) = 1$ and $\lim_{d\to\infty} g_2(d) = 0$. We will make use of the following inequalities (see, e.g., [26, Eq. 4.5.13]):

$$\exp\left(\frac{xy}{x+y}\right) < \left(1 + \frac{x}{y}\right)^y < \exp(x), \quad \text{for all } x, y > 0. \tag{20}$$

First, we show $\lim_{d\to\infty} g_1(d) = 1$. As d and k_d are positive, we have $g_1(d) > 1$ for every d. Furthermore, g_1 is increasing in k_d, and $k_d < \frac{\ln d}{2 \ln \ln d} + 1$. Thus,

$$g_1(d) < \left(1 + d^{-\frac{1}{\frac{\ln d}{\ln \ln d} + 4}}\right)^{\frac{\ln d}{2 \ln \ln d} + 1}.$$

Using the second inequality of (20) with $y = \frac{\ln d}{2 \ln \ln d} + 1$ and $x = y d^{-\frac{\ln \ln d}{\ln d + 4 \ln \ln d}}$, we can further bound

$$g_1(d) < \exp\left(\frac{\frac{\ln d}{2 \ln \ln d} + 1}{d^{\frac{\ln \ln d}{\ln d + 4 \ln \ln d}}}\right). \tag{21}$$

We will show that the argument of the exponential function on the r.h.s. of (21) goes to 0 for $d \to \infty$, thus proving the claim. Replacing $\ln d = \exp(\ln \ln d)$ in the numerator and $d = \exp(\ln d)$ in the denominator, that argument can be written as

$$\frac{\exp(\ln \ln d)\left(\frac{1}{2 \ln \ln d} + \frac{1}{\ln d}\right)}{\exp\left(\frac{\ln d \ln \ln d}{\ln d + 4 \ln \ln d}\right)} = \left(\frac{1}{2 \ln \ln d} + \frac{1}{\ln d}\right) \exp\left(\frac{4(\ln \ln d)^2}{\ln d + \ln \ln d}\right). \tag{22}$$

The argument of the exponential function on the r.h.s. of (22) goes to 0, as $\ln d$ is the dominating term in the denominator for $d \to \infty$. Thus, the whole expression in (22) goes to 0.

Next, we show that $\lim_{d\to\infty} g_2(d) = 0$. As $d \geq 2$, we have that $g_2(d) > 0$ for every d. Using the first inequality of (20) with $x = \frac{\ln d}{2}$ and $y = d$, we have

$$\left(1 + \frac{\ln d}{2d}\right)^d > \exp\left(\frac{d \ln d}{\ln d + 2d}\right).$$

Thus, we obtain an upper bound on g_2 by writing $\sqrt{d} = \exp\left(\frac{\ln d}{2}\right)$:

$$g_2(d) < \frac{2 \exp\left(\frac{\ln d}{2}\right)}{\ln d \exp\left(\frac{d \ln d}{\ln d + 2d}\right)} = \frac{2}{\ln d} \exp\left(\frac{(\ln d)^2}{2 \ln d + 4d}\right). \tag{23}$$

As $4d$ is the dominating term in the denominator of the argument of the exponential function, the argument goes to 0 for $d \to \infty$, and thus the r.h.s. of (23) goes to 0, showing the claim. □