

Congestion Games with Multisets of Resources and Applications in Synthesis

Guy Avni¹, Orna Kupferman¹, and Tami Tamir²

1 School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel

2 School of Computer Science, The Interdisciplinary Center, Israel

Abstract

In classical congestion games, players' strategies are subsets of resources. We introduce and study *multiset congestion games*, where players' strategies are multisets of resources. Thus, in each strategy a player may need to use each resource a different number of times, and his cost for using the resource depends on the load that he and the other players generate on the resource.

Beyond the theoretical interest in examining the effect of a repeated use of resources, our study enables better understanding of non-cooperative systems and environments whose behavior is not covered by previously studied models. Indeed, congestion games with multiset-strategies arise, for example, in production planning and network formation with tasks that are more involved than reachability. We study in detail the application of synthesis from component libraries: different users synthesize systems by gluing together components from a component library. A component may be used in several systems and may be used several times in a system. The performance of a component and hence the system's quality depends on the load on it.

Our results reveal how the richer setting of multisets congestion games affects the stability and equilibrium efficiency compared to standard congestion games. In particular, while we present very simple instances with no pure Nash equilibrium and prove tighter and simpler lower bounds for equilibrium inefficiency, we are also able to show that some of the positive results known for affine and weighted congestion games apply to the richer setting of multisets.

Keywords and phrases Congestion games, Multiset strategies, Equilibrium existence and computation, Equilibrium inefficiency

Digital Object Identifier 10.4230/LIPIcs.xxx.yyy.p

1 Introduction

Congestion games model non-cooperative resource sharing among selfish players. Resources may be shared by the players and the cost of using a resource increases with the load on it. Such a cost paradigm models settings where high congestion corresponds to lower quality of service or higher delay. Formally, each resource e is associated with an increasing latency function $f_e : \mathbb{N} \rightarrow \mathbb{R}$, where $f_e(\ell)$ is the cost of a single use of e when it has load ℓ .

Previous work on congestion games assumes that players' strategies are subsets of resources, as is the case in many applications, most notably routing and network design. For example, in the setting of networks, players have reachability objectives and strategies are subsets of edges, each inducing a simple path from the source to the target [29, 3, 19]. We introduce and study multiset games, where players' strategies are multisets of resources. Thus, a player may need a resource multiple times – depending on the specific resource and strategy, and his cost for using the resource depends on the load that he and the other players generate on it. Formally, in *multiset congestion games* (MCGs, for short), a player that uses j times a resource e that is used ℓ times by all players together, pays $j \cdot f_e(\ell)$ for these uses.

Beyond the theoretical interest in examining the effect of multisets on the extensively-studied classical games, multiset congestion games arise naturally in many applications and environments. The use of multisets enables the specification of rich settings that cannot be specified by means of subsets. We give here several examples.

As a first example, consider *network formation*. In addition to reachability tasks, which involve simple paths (and hence, subsets of resources), researchers have studied tasks whose satisfaction may involve paths that are not simple. For example, a user may want to specify that each traversal of a low-security channel is followed by a visit to a check-sum node. A well-studied class of tasks that involve paths that need not be simple are these associated with a specific length, such as patrols in a geographical region. Several communication protocols are based on the fact that a message must pass a pre-defined length before reaching its destination, either for security reasons (e.g., in *Onion routing*, where the message is encrypted in layers [27]) or for marketing purposes (e.g., advertisement spread in social networks). In addition, tasks of a pre-defined length are the components of *proof-of-work* protocols that are used to deter denial of service attacks and other service abuses such as spam (e.g., [15]), and of several protocols for sensor networks [7]. The introduction of multiset corresponds to strategies that are not necessarily simple paths [5].

In *production systems* or in *planning*, a system is modeled by a network whose nodes correspond to configurations and whose edges correspond to actions performed by resources. Users have tasks, that need to be fulfilled by taking sequences of actions. This setting corresponds to an MCG in which the strategies of the players are multisets of actions that fulfill their tasks, which indeed often involve repeated execution of actions [13]; for example “once the arm is up, do not put it down until the block is placed”. Also, multiset games can model *preemptive scheduling*, where the processing of a job may split in several feasible ways among a set of machines.

Our last example, which we are going to study in detail, is *synthesis from component libraries*. A central problem in formal methods is synthesis [26], namely the automated construction of a system from its specification. In real life, hardware and software systems are rarely constructed from scratch. Rather, a system is typically constructed from a library of components by *gluing* components from a library (allowing multiple uses) [23]. For example, when designing an internet browser, a designer does not implement the TCP protocol but uses existing implementations as black boxes. The library of components is used by multiple users simultaneously, and the usages are associated with costs. The usage cost can either decrease with load (e.g., when the cost of a component represents its construction price, the users of a component share this price) as was studied in [4], or increase with load (e.g., when the components are processors and a higher load means slower performance). The

later scenario induces an instance of an MCG.

Let us demonstrate the intricacy of the multiset setting with the question of the existence of a *pure Nash equilibrium* (PNE). That is, whether each instance of the game has a profile of pure strategies that constitutes a PNE – a profile such that no player can decrease his cost by unilaterally deviating from his current strategy. By [28], classical congestion games are potential games and thus always have a PNE. Moreover, by [19], in a symmetric congestion game, a PNE can be found in polynomial time. As we show in Example 1 below, a PNE might not exist in an MCG even in a symmetric two-player game over identical resources.

Example 1: Consider the following symmetric MCG with two players and three resources: a , b , and c . The players' strategy space is $\{a, a, b\}$ or $\{b, b, c\}$ or $\{c, c, a\}$. That is, a player needs to access some resource twice and the (cyclically) consequent resource once. The latency function of all three resources is the same, specifically, $f_a(\ell) = f_b(\ell) = f_c(\ell) = \ell^2$. The players' costs in all possible profiles are given in Table 1. We show that no PNE exists in this game. Assume first that the two players select distinct strategies, w.l.o.g. $\{a, a, b\}$ and $\{b, b, c\}$. In this profile, a is accessed twice, b is accessed three times, and c is accessed once. Thus, every access of a , b and c costs 4, 9 and 1 respectively. The cost of Player 1 is $8 + 9 = 17$, while the cost of Player 2 is $18 + 1 = 19$. By deviating to $\{c, c, a\}$, the cost of Player 2 will reduce to 17 (while the cost of Player 1 will increase to 19). Thus, no PNE in which the players select different strategies exists. If the player select the the same strategy, then one resource is accessed 4 times, and one resource is accessed twice, implying that the cost of both players is $2 \cdot 16 + 1 \cdot 4 = 36$, and any deviation is profitable. We conclude that no PNE exists in the game. ◀

	$\{a, a, b\}$	$\{b, b, c\}$	$\{c, c, a\}$
$\{a, a, b\}$	36, 36	19, 17	17, 19
$\{b, b, c\}$	17, 19	36, 36	19, 17
$\{c, c, a\}$	19, 17	17, 19	36, 36

■ **Table 1** Players costs. Each entry describes the cost of Player 1 followed by the cost of Player 2.

We study and answer the following questions in general and for various classes of multiset congestion games (for formal definitions, see Section 2): (i) Existence of a PNE. (ii) An analysis of *equilibrium inefficiency*. A *social optimum* (SO) of the game is a profile that minimizes the total cost of the players; thus, the one obtained when the players obey some centralized authority. It is well known that decentralized decision-making may lead to solutions that are sub-optimal from the point of view of society as a whole. We quantify the inefficiency incurred due to selfish behavior according to the *price of anarchy* (PoA) [22] and *price of stability* (PoS) [3] measures. The PoA is the worst-case inefficiency of a PNE (that is, the ratio between the cost of a worst PNE and the SO). The PoS is the best-case inefficiency of a Nash equilibrium (that is, the ratio between the cost of a best PNE and the SO). (iii) *Computational complexity* of finding a PNE.

Before we turn to describe our results, let us review related work. *Weighted* congestion games (WCGs, for short), introduced in [25], are congestion games in which each player i has a *weight* $w_i \in \mathbb{N}$, and his contribution to the load of the resources he uses as well as his payments are multiplied by w_i . WCGs can be viewed as a special case of MCGs, where each resource in a strategy for Player i repeats w_i times. A different extension of WCGs in which players may use a resource more than once is *integer splittable WCGs* [24, 30]. These games model the setting in which a player has a number (integer) of tasks he needs to perform and can split them among the resources. For example, in the network setting, a player might need to send $\ell \in \mathbb{N}$ packets from vertex s to t . He can send the packets on different paths, but a packet cannot be split. MCGs are clearly more general than WCGs and integer splittable WCGs – the ability to repeat each resource a different number of

times lead to a much more complex setting. Thus, it is interesting to compare our results with these known for these games.

It is shown in [17, 21] that the existence of a PNE in WCGs depends on the latency function: when the latency functions are either affine or exponential, WCGs are guaranteed to admit a PNE, whereas WCGs with a polynomial latency function need not have a PNE. In [24], the author shows that PNE always exists when the latency functions are linear using a *potential function* argument. This argument fails when the latency functions are convex, but [30] are still able to show that there is always a PNE in these games. We are able to show that the exact potential function of [17] applies also to (the much richer) affine MCGs (that is, MCGs with a affine latency function), and thus they always admit a PNE. As demonstrated in Example 1, very simple MCGs with quadratic latency functions might have no PNE.

We turn on to results in the front of equilibrium inefficiency. In congestion games with affine latency functions, both the PoA and PoS measures are well understood. It was shown in [12] that $\text{PoS} \geq 1 + \frac{1}{\sqrt{3}} \approx 1.577$ and is at most 1.6. A tight upper bound was later shown in [10]. Also, $\text{PoA} = \frac{5}{2}$ [12]. Going one step towards our setting to the study of affine WCGs, [6] shows that $\text{PoA} = 1 + \phi$, where $\phi \approx 1.618$ is the golden ratio. The PoS question is far from being settled. Only recently, [9] shows a first upper bound of 2 for PoS in linear WCGs, which is a subclass of affine WCGs. As far as we know, the only lower bound that is known for affine WCGs is the lower bound from the unweighted setting. So there is a relatively large gap between the upper- and lower-bounds for the PoS in these games.

We bound the potential function in order to show that every affine MCG G has $\text{PoS}(G) < 2$. This improves and generalizes the result in [9]. Our most technically-challenging result is the PoS lower-bound proof, which involves the construction of a family \mathcal{G} of linear MCGs. Essentially, the game $G_k \in \mathcal{G}$ is parameterized by the number of players and defined recursively. The use of multisets enables us to define a game in which, although the sharing of resources dramatically changes between its profiles, the cost a player pays is equal in all of them. For $k = 17$ we obtain that $\text{PoS}(G_{17}) > 1.631$. This is the first lower bound in these models that exceeds the 1.577 lower bound in congestion games. Finally, the PNE in \mathcal{G} is achieved with dominant strategies, so our bound holds for stronger equilibrium concepts.

As for the PoA, we show that MCGs with latency functions that are polynomials of degree at most d have $\text{PoA} = \Phi_d^{d+1}$, where Φ_d is the unique nonnegative real solution to $(x+1)^d = x^{d+1}$. Observe that Φ_d is a natural generalization of the golden ratio to higher degrees. Specifically, $\Phi_1 = \phi$. For the upper bound, we adjust the upper-bound proof of [2] to our setting. We show a simplified matching lower bound; a simple two-player MCG with only two resources and latency functions of the form $f(\ell) = \ell^d$. For general latency functions we show that the PoA can be arbitrarily high.

We turn to study the application of synthesis from component libraries by multiple players. Recall that in this application, different users synthesize systems from components. A component may be used in several systems and may be used several times in a system. The quality of a system depends on the load on its components. This gives rise to an MCG, which we term a *component library game* (CLG, for short). On the one hand, a CLG is a special case of MCG, so one could expect positive results about MCGs to apply to CLGs. On the other hand, while in MCGs the strategies of the players are given explicitly by means of multisets of resources, in CLGs the strategies of the players are given symbolically by means of a specification deterministic finite automaton – the one whose language has to be composed from the library's components.

We prove that every MCG has a corresponding CLG, implying that negative results for MCGs apply to CLGs. Moreover, we show that the succinctness of the presentation of the strategies makes decision problems about MCGs more complex in the setting of CLGs. We demonstrate this by studying the complexity of the *best-response* problem – deciding whether a player can benefit from a

unilateral deviation from his strategy, and the problem of deciding whether a PNE exists in a given game. For the best-response problem, which is in P for MCGs, we prove NP-completeness for CLGs. The problem of deciding the existence of a PNE is known to be strongly NP-complete for weighted symmetric congestion games. For network congestion games with player specific cost functions, this problem is NP-complete for arbitrary networks, while a PNE can be found efficiently for constant size networks [1]. We provide a simpler hardness proof for MCGs, which is valid also for a constant number of resources, and we show that for CLGs the problem is Σ_2^P -complete. As good news, we are able to prove a “small-design property” for CLGs, which bounds the number of strategies that one needs to consider and enables us to lift to CLGs the positive results for MCGs with linear latency functions. Thus, such CLGs always have a PNE and their PoS is at most 2.

Due to space constraints, some proofs and examples are given in the appendix.

2 Preliminaries

A *multiset* over a set E of elements is a generalization of a subset of E in which each element may appear more than once. For a multiset A over E and an element $e \in E$, we use $A(e)$ to denote the number of times e appears in A , and use $e \in A$ to indicate that $A(e) \geq 1$. When describing multisets, we use e^m , for $m \in \mathbb{N}$, to denote m occurrences of e .

A *multiset congestion game* (MCG) is a tuple $G = \langle K, E, \{\Sigma_i\}_{i \in K}, \{f_e\}_{e \in E} \rangle$, where $K = \{1, \dots, k\}$ is a set of players, E is a set of *resources*, for every $1 \leq i \leq k$, the *strategy space* Σ_i of Player i is a collection of multisets over E , and for every resource $e \in E$, the *latency function* $f_e : \mathbb{N} \rightarrow \mathbb{R}$ is a non-decreasing function. The MCG G is an *affine* MCG if for every $e \in E$, the latency function f_e is affine, i.e., $f_e(x) = a_e x + b_e$, for non-negative constants a_e and b_e . Similarly, we say that G is a *linear* MCG if it is affine and for $e \in E$ we have $b_e = 0$. We assume w.l.o.g. that for $e \in E$ we have $a_e \geq 1$. Classical congestion games are a special case of MCGs where the players’ strategies are sets of resources. Weighted congestion games can be viewed as a special case of MCGs, where for every $1 \leq i \leq k$, multiset $s_i \in \Sigma$ and $e \in s_i$, we have $s_i(e) = w_i$.

A profile of a game G is a tuple $P = \langle s_1, s_2, \dots, s_k \rangle \in (\Sigma_1 \times \Sigma_2 \times \dots \times \Sigma_k)$ of strategies selected by the players. For a resource $e \in E$, we use $L_{e,i}(P)$ to denote the number of times e is used in P by Player i . Note that $L_{e,i}(P) = s_i(e)$. We define the *load* on e in P , denoted $L_e(P)$, as the number of times it is used by all players, thus $L_e(P) = \sum_{1 \leq i \leq k} L_{e,i}(P)$ ¹.

In classical congestion games, all players that use a resource e pay $f_e(\ell)$, where ℓ is the number of players that use e . As we formalize below, in MCGs, the payment of a player for using a resource e depends on the number of times he uses it. Given a profile P , a resource $e \in E$, and $1 \leq i \leq k$, the *cost of e for Player i in P* is $cost_{e,i}(P) = L_{e,i}(P) \cdot f_e(L_e(P))$. That is, for each of the $L_{e,i}(P)$ uses of e , Player i pays $f_e(L_e(P))$. The cost of Player i in the profile P is then $cost_i(P) = \sum_{e \in E} cost_{e,i}(P)$ and the cost of the profile P is $cost(P) = \sum_{1 \leq i \leq k} cost_i(P)$. We also refer to the cost of a resource e in P , namely $cost_e(P) = \sum_{i \in K} cost_{e,i}(P)$.

Consider a game G . For a profile P , player $i \in K$, and a strategy $s'_i \in \Sigma$ for Player i , let $P[i \leftarrow s'_i]$ denote the profile obtained from P by replacing the strategy for Player i by s'_i . A profile P is a *pure Nash equilibrium* (PNE) if no Player i can benefit from unilaterally deviating from his strategy in P to another strategy; i.e., for every player i and every strategy $s'_i \in \Sigma$ it holds that $cost_i(P[i \leftarrow s'_i]) \geq cost_i(P)$.

¹ Since our strategies are multisets, we have that $s_i(e)$, for all i and e , is an integer. Our considerations, however, are independent of this, thus all our results are valid also for games in which strategies might include fractional demands for resources. In non-splittable (atomic) games, the players must select a single strategy, even if fractional demands are allowed.

We denote by OPT the cost of a social-optimal solution; i.e., $OPT = \min_P \text{cost}(P)$. It is well known that decentralized decision-making may lead to sub-optimal solutions from the point of view of society as a whole. We quantify the inefficiency incurred due to self-interested behavior according to the *price of anarchy* (PoA) [22] and *price of stability* (PoS) [3] measures. The PoA is the worst-case inefficiency of a Nash equilibrium, while the PoS measures the best-case inefficiency of a Nash equilibrium. Formally,

► **Definition 1.** Let \mathcal{G} be a family of games, and let G be a game in \mathcal{G} . Let $\Upsilon(G)$ be the set of Nash equilibria of the game G . Assume that $\Upsilon(G) \neq \emptyset$.

- The *price of anarchy* of G is the ratio between the *maximal* cost of a PNE and the social optimum of G . That is, $\text{PoA}(G) = \max_{P \in \Upsilon(G)} \text{cost}(P) / OPT(G)$. The *price of anarchy* of the family of games \mathcal{G} is $\text{PoA}(\mathcal{G}) = \sup_{G \in \mathcal{G}} \text{PoA}(G)$.
- The *price of stability* of G is the ratio between the *minimal* cost of a PNE and the social optimum of G . That is, $\text{PoS}(G) = \min_{P \in \Upsilon(G)} \text{cost}(P) / OPT(G)$. The *price of stability* of the family of games \mathcal{G} is $\text{PoS}(\mathcal{G}) = \sup_{G \in \mathcal{G}} \text{PoS}(G)$.

3 Existence of a Pure Nash Equilibrium

As demonstrated in Example 1, MCGs are less stable than weighted congestion games:

► **Theorem 2.** *There exists a symmetric two-player MCG with identical resources and quadratic latency function that has no PNE.*

On the positive side, we show that a PNE exists in all MCGs with affine latency functions. We do so by showing that an exact potential function exists, which is a generalization of the one in [9, 18].

► **Theorem 3.** *Affine MCGs are potential games.*

Proof. For a profile P and a resource $e \in E$, define

$$\Phi_e(P) = a_e \cdot \left(\sum_{i=1}^k \sum_{j=i}^k L_{e,i}(P) \cdot L_{e,j}(P) \right) + \left(b_e \cdot \sum_{i=1}^k L_{e,i}(P) \right).$$

Also, $\Phi(P) = \sum_{e \in E} \Phi_e(P)$. In the appendix, we prove that Φ is an exact potential function. ◀

The negative result in Theorem 2 gives rise to the decision problem $\exists\text{PNE}$; given an MCG, decide whether it has a PNE. Being a generalization of WCGs, the hardness results known for WCGs imply that $\exists\text{PNE}$ is NP-hard [14]. Using the richer definition of MCGs, we show below a much simpler hardness proof. We also show hardness for games with a constant number of resources, unlike congestion games with user-specific cost functions [1].

► **Theorem 4.** *Given an instance of an MCG, it is strongly NP-complete to decide whether the game has a PNE, as well as to find a PNE given that one exists. For games with a constant number of resources, the problems are NP-Complete.*

Remark 4.1: In *splittable (non-atomic) games*, each player can split his task among several strategies. This can be seen as if each player is replaced by $M \rightarrow \infty$ identical players all having the same strategy space scaled by $1/M$. This model suits several applications, in particular planning of preemptive production. Splittable games are well-understood in classical and weighted congestion games [29, 8]. In Appendix B we define the corresponding MCG and show that the positive PNE-existence result, known for weighted congestion games, carry over to games with multisets of resources. ◀

4 Equilibrium Inefficiency in MCGs

4.1 The Price of Stability

The PoS problem in affine congestion games is settled: [12, 10] show that $\text{PoS} = 1 + \frac{1}{\sqrt{3}} \approx 1.577$. For affine WCGs, the problem was open for a long time, and only recently progress was made by [9], who showed that $\text{PoS} \leq 2$ for linear WCGs. As far as we know, there is no known lower bound for linear WCGs that exceeds the 1.577 bound for unweighted games. We show that every affine MCG G has $\text{PoS}(G) < 2$. Thus, we both improve the result to include affine functions, tighten the bound, and generalize it. For the lower bound, we show a family of linear MCGs \mathcal{G} that has $\text{PoS}(\mathcal{G}) > 1.631$. We start with the upper bound.

► **Theorem 5.** *Every affine MCG G has $\text{PoS}(G) < 2$.*

Proof. Consider an affine MCG G and a profile P . It is not hard to see that for the potential function Φ that is presented in Theorem 3 we have $\Phi(P) \leq \text{cost}(P)$. Moreover, for $e \in E$ we have $2\Phi_e(P) = \text{cost}_e(P) - a_e \sum_{1 \leq i \leq k} L_{e,i}^2(P) - b_e \sum_{1 \leq i \leq k} L_{e,i}(P)$. Thus, $\Phi(P) > \frac{1}{2}\text{cost}(P)$. The theorem follows using standard techniques: $\text{cost}(O) \geq \Phi(O) \geq \Phi(N) > \frac{1}{2}\text{cost}(N)$, where O is the social optimum and N is a PNE that is reached from O by a sequence of best-respond moves of the players. Then, $\text{PoS}(G) \leq \frac{\text{cost}(N)}{\text{cost}(O)} < 2$. The details of the proof can be found in the appendix. ◀

Note that while the PoS can get arbitrarily close to 2, it is strictly smaller than 2 for every game instance. The proof in [9], on the other hand, only shows $\text{PoS} \leq 2$ for the family of affine MCGs, and our result does not improve this bound.

For the lower bound, we show a family of linear MCG $\mathcal{G} = \{G_k\}_{k \geq 2}$ that are parameterized by the number of players. Using a computerized simulation, we obtain that for the game with 17 players, we have $\text{PoS}(G_{17}) > 1.631$. We leave open the problem of calculating the exact value the PoS tends to as the number of players increases. The graph depicted in Figure 3 in the appendix, of $\text{PoS}(G_k)$ as a function of k , hints that the answer is only slightly higher than 1.631.

The PNE in the games in the family is achieved with dominant strategies, and thus it is resistant to stronger types of equilibria.

► **Theorem 6.** *There is a linear MCG G with $\text{PoS}(G) > 1.631$.*

Proof. We define a family of games $\{G_k\}_{k \geq 2}$ as follows. The game G_k is played by k players, thus $K_k = \{1, \dots, k\}$. For Player 1, all strategies $\Sigma_1^k = \{O_1^k\}$ consists of a single multiset. For ease of presentation we sometimes refer to O_1^k as N_1^k . For $i \geq 2$, the strategy space of Player i consists of two multisets, $\Sigma_i^k = \{O_i^k, N_i^k\}$. We define G_k so that for all $k \geq 2$, the profile $\bar{O}_k = \langle O_1^k, \dots, O_k^k \rangle$ is the social optimum and the profile $\bar{N}_k = \langle N_1^k, \dots, N_k^k \rangle$ is the only PNE.

When describing the games in the family, we partition the resources into *types* and describe a multiset as a collection of triples. A triple $\langle t, y, l \rangle$ stands for y different resources of type t , each appearing l times. For example, $\{\langle a, 2, 1 \rangle, \langle b, 1, 3 \rangle, \langle c, 2, 2 \rangle\}$ stands for the multiset $\{a_1, a_2, b_1, b_1, b_1, c_1, c_1, c_2, c_2\}$. In all games and resources, there are two types of latency functions; the identity function, or identity plus epsilon, where the second type of function are linear functions of the form $f(x) = (1 + \epsilon) \cdot x$, for some $\epsilon > 0$. The latency function of resources of the same type is the same, and we use the terms “ a has identity latency” and “ b has identity plus ϵ latency” to indicate that all the resources a' of type a have $f_{a'}(j) = j$ and all the resources b' of type b have $f_{b'}(j) = (1 + \epsilon) \cdot j$, for all numbers j of uses.

The definition of G_k is complicated and we start by describing the idea in the construction of G_2 and G_3 . In the appendix we also describe G_4 . We start by describing G_2 . The game G_2 is defined with respect to two types of resources, a and b , with identity and identity plus ϵ latency, respectively. We define Player 1’s strategy space $\Sigma_1^2 = \{O_1^2\}$ and Player 2’s strategy space $\Sigma_2^2 = \{O_2^2, N_2^2\}$, with

$O_1^2 = N_2^2 = \langle a, 2, 1 \rangle$ and $O_2^2 = \langle b, 1, 2 \rangle$. That is, $\Sigma_1^2 = \{\{a_1, a_2\}\}$ and $\Sigma_2^2 = \{\{a_1, a_2\}, \{b_1, b_1\}\}$. Clearly, the profile $\bar{N}_2 = \langle O_1^2, N_2^2 \rangle$ is the only PNE in G_2 .

We continue to describe G_3 . The game G_3 is defined with respect to four types of resources, a, b, c^1 and c^2 , where b has identity plus ϵ latency, c^1 has identity plus ϵ' latency, and the other resources have identity latency. Let $x_3 = 3! = 6$. We define $\Sigma_1^3 = \{O_1^3\}$, $\Sigma_2^3 = \{O_2^3, N_2^3\}$, and $\Sigma_3^3 = \{O_3^3, N_3^3\}$, with $O_1^3 = N_2^3 = \langle a, x_3, 1 \rangle$, $O_2^3 = \langle b, \frac{x_3}{2}, 2 \rangle$, $O_3^3 = \{\langle c^1, \frac{x_3}{3}, 3 \rangle, \langle c^2, \frac{x_3}{2}, 1 \rangle\}$, and $N_3^3 = \{\langle b, \frac{x_3}{2}, 1 \rangle, \langle a, x_3, 1 \rangle\}$. We claim that $\bar{N}_3 = \langle O_1^3, N_2^3, N_3^3 \rangle$ is the only PNE. Our goal here is not to show a complete proof, but to demonstrate the idea of the construction. It is not hard to see that Player 2 deviates to N_2^3 from the profile $\bar{O}_3 = \langle O_1^3, O_2^3, O_3^3 \rangle$, Player 3 deviates from the resulting profile $\bar{N}_3 = \langle O_1^3, N_2^3, N_3^3 \rangle$. The crux of the construction is to keep Player 2 from deviating back from \bar{N}_3 . Note that since Player 3 uses the b -type resources once in \bar{N}_3 , when Player 2 deviates from N_2^3 to O_2^3 , their load increases to 3. Thus, $cost_2(\bar{N}_3[2 \leftarrow O_2^3]) = 3(3 \cdot 2 \cdot (1 + \epsilon)) > 6(3 \cdot 1) = cost_2(\bar{N}_3)$ and the deviation is not beneficial.

We define the game G_k , for $k \geq 2$, as follows. Let $x_k = k!$. Player 1's strategy space consists of a single multiset $O_1^k = \langle e_{1,1}, x_k, 1 \rangle$. For $2 \leq i \leq k$, assume we have defined the strategies and resources for players $1, \dots, i-1$. We define Player i 's strategies as follows. We start with the multiset N_i^k , which does not introduce new resources. We define $N_i^k = \cup_{1 \leq j \leq i-1} \{\langle t, x, 1 \rangle : \langle t, x, l \rangle \subseteq O_j^k\}$. The definition of O_i^k is more involved, but the idea is simple. We define O_i^k so that it satisfies two properties. First, O_i^k uses new resources. That is, for every $1 \leq j \leq i-1$, both $O_i^k \cap O_j^k = \emptyset$ and $O_i^k \cap N_j^k = \emptyset$. Consider the profile P_i in which, for every $1 \leq j < i$, Player j uses N_j^k and, for every $i \leq l \leq k$, Player l uses O_l^k . We define O_i^k so that when all resources have identity latency, $cost_i(P_i) = cost_i(P_i[i \leftarrow N_i^k])$. For every multiset $\langle e_{j,a}, x_{j,a}, 1 \rangle$ in N_i^k , which we have just defined, we introduce a multiset $\langle e_{i,b}, x_{i,b}, l_{i,b} \rangle$ in O_i^k that uses new resources, where b is a unique index that is arbitrarily chosen, and $x_{i,b}$ and $l_{i,b}$ are defined as follows. Let $l = |\{j : e_{j,a} \in N_j^k\}|$. We define $l_{i,b} = l + 1$ and $x_{i,b} = x_{j,a}/l_{i,b}$. Since O_i^k uses new resources, showing the first property is easy. In the appendix we show it satisfies a much stronger property.

Claim 6.1: Consider $k \in \mathbb{N}$, a profile P in G_k , and $1 < i \leq k$. Assume Player i plays O_i^k in P . When the latency functions are identity, we have $cost_i(P) = cost_i(P[i \leftarrow N_i^k])$.

To complete the construction, we define the latency functions so that for every $2 \leq i \leq k$, we have that $e_{i,1}$ -type resources have identity plus ϵ_i latency for $0 < \epsilon_2 < \dots < \epsilon_k$. By Claim 6.1 there are such values that make N_i^k a dominant strategy for Player i . Thus, the only PNE in G_k , for $k \geq 2$, is the profile $\bar{N}_k = \langle O_1^k, N_2^k, \dots, N_k^k \rangle$. Next, we identify the social optimum.

Claim 6.2: The profile $\bar{O}_k = \langle O_1^k, \dots, O_k^k \rangle$ is the social optimum.

Once we identify \bar{O}_k as the social optimum and \bar{N}_k as the only PNE, the calculation of the PoS boils down to calculating their costs, which we do using a computer. Specifically, we have $PoS(G_{17}) = 1.6316$, and we depict the values of G_k , for $2 \leq k \leq 17$, in Figure 3 in the appendix. ◀

Remark 6.1: We conjecture that the correct value for the PoS is closer to our lower bound of 1.631 rather than to the upper bound of 2. In the appendix we show a more careful analysis of the potential function than the one in Theorem 5 that shows that for every linear MCG G we have $PoS(G) \leq 2 - \frac{\sum_{e \in E} \sqrt{cost_e(N_G)}}{cost(O_G)}$, where N_G and O_G denote the cheapest PNE and the social optimum of G , respectively. Also, we show that for every $n \geq 2$, for the MCG G_n that is described in Theorem 6, the inequality in the expression is essentially an equality. ◀

Remark 6.2: We can alter the family in Theorem 6 to have quadratic latency functions instead of identity functions. Although Claim 6.1 does not hold in the altered family, a computerized simulation shows that the N strategies are still dominant strategies. Also, using a computerized simulation, we show that the PoS for G_{15} is 2.399, higher than the upper bound of 2.362 for congestion games,

which is shown in [9, 11]. ◀

4.2 The Price of Anarchy

In this section we study the PoA for MCGs. We start with MCGs with polynomial latency functions and show that the upper bound proven in [2] for WCGs can be adjusted to our setting. Being a special case of MCGs, the matching lower bound for WCGs applies too. Still, we present a different and much simpler lower-bound example, which uses a two-player singleton MCG. In a singleton game, each strategy consists of (multiple accesses to) a single resource. Finally, when the latency functions are not restricted to be polynomials, we show that the PoA is unbounded, and it is unbounded already in a singleton MCG with only two players.

We start by showing that the PoA in polynomial MCGs is not higher than in polynomial WCGs. The proof adjusts the one known for WCGs [2] to our setting. For $d \in \mathbb{N}$, we denote by \mathcal{P}_d the set of polynomials of degree at most d .

► **Theorem 7.** *The PoA in MCGs with latency functions in \mathcal{P}_d is at most Φ_d^{d+1} , where Φ_d is the unique nonnegative real solution to $(x + 1)^d = x^{d+1}$.*

Next, we show a matching lower bound that is stronger and simpler than the one in [2].

► **Theorem 8.** *For $d \in \mathbb{N}$, the PoA in two-player singleton MCG with latency functions in \mathcal{P}_d is at least Φ_d^{d+1} .*

Proof. Let $d \in \mathbb{N}$. Consider the two-player singleton MCG G with resources $E = \{e_1, e_2\}$, strategy spaces $\Sigma_1 = \{e_1^x, e_2^y\}$ and $\Sigma_2 = \{e_1^y, e_2^x\}$, and for $\ell \in \mathbb{R}$, we define the latency functions $f_{e_1}(\ell) = f_{e_2}(\ell) = \ell^d$. We define $x = \Phi_d$ and $y = 1$. Since $x > y$ the social optimum is attained in the profile $\langle e_1^y, e_2^y \rangle$ and its cost is $2y^d = 2$. Recall that in MCGs, the players' strategies are multisets. In particular, x should be a natural number. To fix this, we consider a family of MCGs in which the ratio between x and y tends to the ratio above.

We claim that the profile $N = \langle e_1^x, e_2^x \rangle$ is a PNE. This would imply that $\text{PoA}(G) = \frac{2x^{d+1}}{2} = \Phi_d^{d+1}$, which would conclude the proof. We continue to prove the claim. The cost of a player in N is $x \cdot x^d = x^{d+1}$ and by deviating, the cost changes to $y \cdot (x + y)^d = (x + 1)^d$. Our definition of x implies that $x^{d+1} = (x + 1)^d$. Thus, the cost does not change after deviating. Since the players are symmetric, we conclude that the profile N is a PNE, and we are done. ◀

Finally, by taking variants with factorial latency functions to the game described in Theorem 8, we are able to increase the PoA in an unbounded manner.

► **Theorem 9.** *The PoA in two-player MCGs is unbounded.*

5 Synthesis from Component Libraries

In this section we describe the application of MCGs in synthesis from component libraries. As briefly explained in Section 1, in this application, different users synthesize systems by gluing together components from a component library. A component may be used in several systems and may be used several times in a system. The performance of a component and hence the system's quality depends on the load on it. We describe the setting in more detail, formalize it by means of MCGs, and relate to the results studied in earlier sections.

Today's rapid development of complex and safety-critical systems requires reliable verification methods. In *formal methods*, we reason about systems and their specifications by solving mathematical questions about their models. A central problem in formal methods is synthesis, namely the automated construction of a system from its specification. In real life, systems are rarely constructed from

scratch. Rather, a system is typically constructed from a library of components by *gluing* components from the library [23]. In this setting, the input to the synthesis problem is a specification and a library of components, and the goal is to construct from the components a system that exhibits exactly the behaviors specified in the specification.

Remark 9.1: The above setting corresponds to *closed systems*, whose behavior is independent of their environment. It is possible to generalize the definitions to *open systems*, which interact with their environment. In [4], we studied both the closed and open settings in the context of cost-sharing (rather than congestion) games. The technical challenges that have to do with the system being open are orthogonal to these that arise from the congestion effects, and on which we focus in this work. ◀

In our setting, we use deterministic finite automata (DFAs, for short) to model the specification and use *box-DFAs* to model the components in the library. Formally, a DFA is $\mathcal{A} = \langle \Sigma, Q, \delta, q_0, F \rangle$, where Σ is an alphabet, Q is a set of states, $\delta : Q \times \Sigma \rightarrow Q$ is a partial transition function, $q_0 \in Q$ is an initial states, and $F \subseteq Q$ is a set of accepting states. The *run* of \mathcal{A} on a word $w = w_1 \dots w_n \in \Sigma^*$ is the sequence of states $r = r_0, r_1, \dots, r_n$ such that $r_0 = q_0$ and for every $0 \leq i \leq n-1$, we have $r_{i+1} = \delta(r_i, w_{i+1})$. Now, a box-DFA \mathcal{B} is a DFA augmented with a set of *exit states*. When a run of \mathcal{B} reaches an exit state, it moves to another box-DFA, as we formalize below.

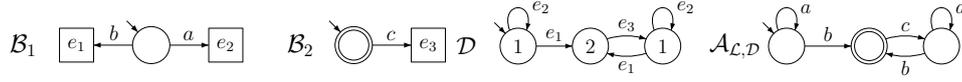
The input to the synthesis from component libraries problem is a specification DFA \mathcal{S} over an alphabet Σ and a library of box-DFAs components $\mathcal{L} = \{\mathcal{B}_1, \dots, \mathcal{B}_n\}$. The goal is to produce a *design*, which is a recipe to compose the components from \mathcal{L} to a DFA. A design is correct if the language of the system it induces coincides with that of the specification.

Intuitively, the design can be thought of as a scheduler; it passes control between the different components in \mathcal{L} . When a component \mathcal{B}_i is in *control*, it reads letters in Σ , visits the states of \mathcal{B}_i , follows its transition function, and if the run terminates, it is accepting iff it terminates in one of \mathcal{B}_i 's accepting states. A component relinquishes control when the run reaches one of its exit states. It is then the design's duty to choose the next component, which gains control through its initial state.

Formally (see an example in Figure 1), a *transducer* is a DFA that has, in addition to the input alphabet that labels the transitions, also an output alphabet that labels the states. Also, a transducer has no rejecting states. Let $[n] = \{1, \dots, n\}$. A design is a transducer \mathcal{D} whose input alphabet is the set \mathcal{E} of all exit states of all the components in \mathcal{L} and whose output alphabet is $[n]$. We can think of \mathcal{D} as running beside the components. When a component reaches an exit state e , then \mathcal{D} reads the input letter e , proceeds to its next state, and outputs the index of the component to gain control next. Note that the components in the library are black boxes: the design \mathcal{D} does not read the alphabet Σ of the components and has no information about the states that the component visits. It only sees which exit state have been reached. Given a library \mathcal{L} and a design \mathcal{D} , their *composition* is a DFA $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$ obtained by composing the components in \mathcal{L} according to \mathcal{D} . We say that a design \mathcal{D} is *correct* with respect to a specification DFA \mathcal{S} iff $L(\mathcal{A}_{\mathcal{L}, \mathcal{D}}) = L(\mathcal{S})$. In the appendix we construct $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$ formally.

For example, consider the library $\mathcal{L} = \{\mathcal{B}_1, \mathcal{B}_2\}$ over the alphabet $\Sigma = \{a, b, c\}$, and the design \mathcal{D} that are depicted in Figure 1. We describe the run on the word bc . The component that gains initial control is \mathcal{B}_1 as the initial state of \mathcal{D} outputs 1. The run in \mathcal{B}_1 proceeds with the letter b to the exit state e_1 and relinquishes control. Intuitively, control is passed to the design that advances with the letter e_1 to the state that outputs 2. Thus, the component \mathcal{B}_2 gains control, and it gains it through its initial state. Then, the letter c is read, \mathcal{B}_2 proceeds to the exit state e_3 and relinquishes control. The design advances with the letter e_3 to a state that outputs 1, and control is assigned to \mathcal{B}_1 . Since the initial state of \mathcal{B}_1 is rejecting, the word ab is rejected. As a second example, consider the word ab . Again, \mathcal{B}_1 gains initial control. After visiting the exit state e_2 , control is reassigned to \mathcal{B}_1 . Finally, after visiting the state e_1 , control is assigned to \mathcal{B}_2 , where the run ends. Since the initial state of \mathcal{B}_2 is accepting, the run is accepting.

The synthesis problem defined above is aimed at synthesizing correct designs. We now add costs to



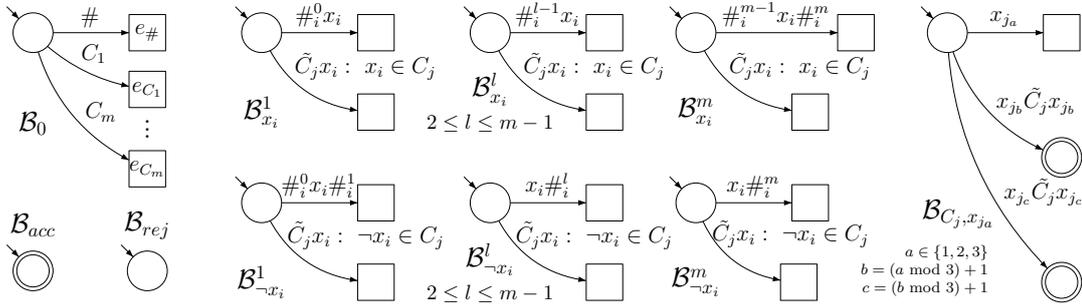
■ **Figure 1** An example of a library $\mathcal{L} = \{\mathcal{B}_1, \mathcal{B}_2\}$, a design \mathcal{D} , and the resulting composition $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$.

the setting. A *component library game* (CLG, for short) is a tuple $\langle K, \mathcal{L}, \{\mathcal{S}_i\}_{i \in K}, \{f_{\mathcal{B}}\}_{\mathcal{B} \in \mathcal{L}} \rangle$, where $K = \{1, \dots, k\}$ is a set of players, \mathcal{L} is a collection of box-DFAs, the objective of Player $i \in K$ is given by means of a specification DFA \mathcal{S}_i , and, as in MCGs, the latency function $f_{\mathcal{B}}$ for a component $\mathcal{B} \in \mathcal{L}$ maps the load on \mathcal{B} to its cost with this load. For $i \in K$, the set of strategies for Player i is the set of designs that are correct with respect to \mathcal{S}_i . A CLG corresponds to an MCG with a slight difference that there might be infinitely many correct designs. Consider a profile $P = \langle \mathcal{D}_1, \dots, \mathcal{D}_k \rangle$. For a component $\mathcal{B} \in \mathcal{L}$, we use $L_{\mathcal{B}, i}(P)$ to denote the number of times Player i uses \mathcal{B} in P . Recall that each state in the transducer \mathcal{D}_i is labeled by a component in \mathcal{L} . We define $L_{\mathcal{B}, i}(P)$ to be the number of states in \mathcal{D}_i that are labeled with \mathcal{B} . The rest of the definitions are as in MCGs.

We first show that every MCG can be translated to a CLG:

► **Theorem 10.** *Consider a k -player MCG G . There is a CLG G' between k players with corresponding profiles. Formally, there is a one-to-one and onto function f from profiles of G to profiles of G' such that for every profiles P in G and Player $i \in [k]$, we have that $\text{cost}_i(P) = \text{cost}_i(f(P))$.*

Proof. Consider an MCG $\langle K, E, \{\Sigma_i\}_{i \in K}, \{f_e\}_{e \in E} \rangle$. Recall that Σ_i is the set of strategies for Player i that consists of multisets over E . We construct a CLG with alphabet $E \cup \bigcup_{i \in K} \Sigma_i$. For $i \in K$, the specification \mathcal{S}_i for Player i consists of $|\Sigma_i|$ words. Every strategy $s = \{e_1, \dots, e_n\}$ (allowing duplicates) in Σ_i contributes to $L(S)$ the word $s \cdot e_1 \cdot e_2 \cdot \dots \cdot e_n$. We construct a library \mathcal{L} with $|E| + \sum_{i \in K} |\Sigma_i|$ components of two types: a *strategy component* \mathcal{B}_s for each $s \in \Sigma_i$ and a *resource component* \mathcal{B}_e for each $e \in E$. In addition, \mathcal{L} contains the component \mathcal{B}_{acc} that is depicted in Figure 2. Intuitively, a correct design must choose one strategy component \mathcal{B}_s and then use the component \mathcal{B}_e the same number of times e appears in s . We continue to describe the components. For $s \in \Sigma_i$, the component \mathcal{B}_s relinquishes control only if the letter s is read. It accepts every word in $L(\mathcal{S}_i)$ that does not start with s . For $e \in E$, the resource component \mathcal{B}_e has an initial state with an e -labeled transition to an exit state. Finally, the latency function for the resource components coincides with latency functions of the resources in the MCG, thus for $e \in E$, we have $f_{\mathcal{B}_e} = f_e$. The other latency functions are $f \equiv 0$. In the appendix we prove that there is a cost-preserving one-to-one and onto correspondence between correct designs with respect to \mathcal{S}_i and strategies in Σ_i , implying the existence of the required function between the profiles. ◀



■ **Figure 2** The components in the library \mathcal{L} .

Theorem 10 implies that the negative results we show for MCGs apply to CLGs:

► **Corollary 11.** *There is a CLG with quadratic latency functions with no PNE; for CLGs with affine latency functions, we have $PoS(CLG) > 1.631$; for $d \in \mathbb{N}$, the PoA in a two-player singleton MCG with latency functions in \mathcal{P}_d is at least Φ_d^{d+1} .*

Remark 11.1: We note that the positive results for CLGs with linear latency functions, namely existence of PNE and $PoS(CLG) \leq 2$, do not follow immediately from Theorem 3, as its proof relies on the fact that an MCG has only finitely many profiles. Since the strategy space of a player might have infinitely many strategies, a CLG might have infinitely many profiles. In order to show that CLGs with linear latency functions have a PNE we need Lemma 12 below, which implies that even in games with infinitely many profiles, there is a best response dynamics that only traverses profiles with “small” designs. Such a traversal is guaranteed to reach a PNE as there are only finitely many such profiles. ◀

Computational complexity We turn to study two computational problems for CLGs: finding a *best-response* and deciding the existence of a PNE. We show that the succinctness of the representation of the objectives of the players in CLGs makes these problems much harder than for MCGs. Our upper bounds rely on the following lemma. The lemma is proven in [4] for cost-sharing games, and the considerations in the proof there applies also for congestion games.

► **Lemma 12.** *Consider a library \mathcal{L} , a specification \mathcal{S} , and a correct design \mathcal{D} . There is a correct design \mathcal{D}' with at most $|\mathcal{S}| \cdot |\mathcal{L}|$ states, where $|\mathcal{L}|$ is the number of states in the components of \mathcal{L} , such that for every component $\mathcal{B} \in \mathcal{L}$, the number of times \mathcal{D}' uses \mathcal{B} is at most the number of times \mathcal{D} uses \mathcal{B} .*

We start with the best-response problem (BR problem, for short): Given an MCG \mathcal{G} between k players, a profile P , an index $i \in K$, and $\mu \in \mathbb{R}$, decide whether Player i has a strategy S'_i such that $cost_i(P[i \leftarrow S'_i]) \leq \mu$.

► **Theorem 13.** *The BR problem for MCGs is in P. For CLGs it is NP-complete, and NP-hardness holds already for games with one player and linear latency functions.*

Proof. Showing that the BR problem is in P for MCGs follows easily from the fact the set of strategies for Player i is given implicitly and calculating the cost for a player in a profile can be done in polynomial time.

The upper bound for CLGs follows from Lemma 12, which implies an upper bound on the size of the cheapest correct designs. Since checking whether a design is correct and calculating its cost can both be done in polynomial time, membership in NP follows.

We continue to the lower bound. We describe the intuition of the reduction and the formal definition along with the correctness proof can be found in the appendix. Given a 3SAT formula φ with clauses C_1, \dots, C_m and variables x_1, \dots, x_n , we construct a library \mathcal{L} and a specification \mathcal{S} such that there is a design \mathcal{D} that costs at most $\mu = nm + m$ iff φ is satisfiable. The library \mathcal{L} consists of an *initial component* \mathcal{B}_0 , *variable components* $\mathcal{B}_{x_i}^j$ and $\mathcal{B}_{\neg x_i}^j$ for $j \in [m]$ and $i \in [n]$, *clause components* $\mathcal{B}_{C_j, x_{j_k}}$ for $j \in [m]$ and $k \in \{1, 2, 3\}$, and component \mathcal{B}_{acc} and \mathcal{B}_{rej} . The components of the library are depicted in Figure 2. The latency function of the variable components is the identity function $f(x) = x$, thus using such a component once costs 1. The latency functions of the other components is the constant function $f \equiv 0$, thus using such components any number of times is free.

Intuitively, a correct design corresponds to an assignment to the variable and must use nm variable components as follows. For $i \in [n]$, either use all the components $\mathcal{B}_{x_i}^1, \dots, \mathcal{B}_{x_i}^m$ or all the components $\mathcal{B}_{\neg x_i}^1, \dots, \mathcal{B}_{\neg x_i}^m$ with a single use each. Thus, a correct design implies an assignment $\eta : \{x_1, \dots, x_n\} \rightarrow \{T, F\}$. Choosing $\mathcal{B}_{x_i}^1, \dots, \mathcal{B}_{x_i}^m$ corresponds to $\eta(x_i) = F$ and choosing $\mathcal{B}_{\neg x_i}^1, \dots, \mathcal{B}_{\neg x_i}^m$ corresponds to $\eta(x_i) = F$.

Additionally, in order to verify that a correct design corresponds to a satisfying assignment, it must use m clause components and m more variable components as follows. Consider a correct design \mathcal{D} , and let $\eta : \{x_1, \dots, x_n\} \rightarrow \{T, F\}$ be the corresponding assignment as described above. For every $j \in [m]$, \mathcal{D} must use a clause component \mathcal{B}_{C_j, x_i} , where recall that the clause C_j includes a literal $\ell \in \{x_i, \neg x_i\}$. Using the component \mathcal{B}_{C_j, x_i} requires \mathcal{D} to use a variable component \mathcal{B}_ℓ^t , for some $t \in [m]$. So, a correct design uses a total of $nm + m$ components with identity latency. If $\eta(\ell) = F$, then \mathcal{B}_ℓ^t is already in use and a second use will cost more than 1, implying that the design costs more than $nm + m$. ◀

The next problem we study is deciding the existence of a PNE. As we show in Theorem 4, the problem is NP-complete for MCGs. As we show below, the succinctness of the representation makes this problem harder for CLGs.

► **Theorem 14.** *The \exists PNE problem for CLGs is Σ_2^P -complete.*

Proof. The upper bound is easy and follows from Lemma 12. For the lower bound we show a reduction from the complement of *not all equal* $\forall \exists$ 3SAT (NAE, for short), which is known to be Σ_2^P -complete [16]. An input to NAE is a 3CNF formula φ over variables $x_1, \dots, x_n, y_1, \dots, y_n$. It is in NAE if for every assignment $\eta : \{x_1, \dots, x_n\} \rightarrow \{T, F\}$ there is an assignment $\rho : \{y_1, \dots, y_n\} \rightarrow \{T, F\}$ such that every clause in φ has a literal that gets value truth and a literal that gets value false (in η or ρ , according to whether the variable is an x or a y variable). We say that such a pair of assignments $\langle \eta, \rho \rangle$ is *legal* for φ .

Given a 3CNF formula φ , we construct a CLG G with three players such that $\varphi \in \text{NAE}$ iff G does not have a PNE. We describe the intuition of the reduction. The details can be found in the appendix. There is a one-to-one correspondence between Player 3 correct designs and assignments to the variables $\{x_1, \dots, x_n\}$. For an assignment $\eta : \{x_1, \dots, x_n\} \rightarrow \{T, F\}$ we refer to the corresponding correct design by \mathcal{D}_η . Consider a legal pair of assignments $\langle \eta, \rho \rangle$, and assume Player 3 chooses the design \mathcal{D}_η . Similarly to the proof of Theorem 13, the library contains variable components with identity latency function. We construct the library and the players' objectives so that there is a correct design \mathcal{D}_ρ for Player 1 that uses $mn + 2m$ variable components each with load 1 iff $\langle \eta, \rho \rangle$ is a legal pair for φ . More technically, both \mathcal{D}_η and \mathcal{D}_ρ use mn variable components that correspond to the variables x_1, \dots, x_n and y_1, \dots, y_n , respectively. For every $j \in [m]$, assuming the j -th clause is $\ell_j^1 \vee \ell_j^2 \vee \ell_j^3$, the design \mathcal{D}_ρ must use two additional variable components $\mathcal{B}_{\ell_j^a}^{t_1}$ and $\mathcal{B}_{\ell_j^b}^{t_2}$, for $a \neq b \in \{1, 2, 3\}$ and $t_1, t_2 \in [m]$, which corresponds to η or ρ assigning value true to ℓ_j^a and value false to ℓ_j^b .

Player 1 has an additional correct design \mathcal{D}_{ALL} in which he does not share any components regardless of the other players' choices. Player 2 has two possible designs \mathcal{D}_A and \mathcal{D}_B . Assume Player 3 chooses a design \mathcal{D}_η . We describe the interaction between Player 1 and Player 2. We define the library and the players' objectives so that when Player 1 chooses some design \mathcal{D}_ρ , Player 2 prefers \mathcal{D}_B over \mathcal{D}_A , thus $\text{cost}_2(\langle \mathcal{D}_\rho, \mathcal{D}_A, \mathcal{D}_\eta \rangle) > \text{cost}_2(\langle \mathcal{D}_\rho, \mathcal{D}_B, \mathcal{D}_\eta \rangle)$. When Player 2 plays \mathcal{D}_B , Player 1 prefers \mathcal{D}_{ALL} over every design \mathcal{D}_ρ , thus $\text{cost}_1(\langle \mathcal{D}_\rho, \mathcal{D}_B, \mathcal{D}_\eta \rangle) > \text{cost}_1(\langle \mathcal{D}_{ALL}, \mathcal{D}_B, \mathcal{D}_\eta \rangle)$. When Player 1 chooses \mathcal{D}_{ALL} , Player 2 prefers \mathcal{D}_A over \mathcal{D}_B , thus $\text{cost}_2(\langle \mathcal{D}_{ALL}, \mathcal{D}_B, \mathcal{D}_\eta \rangle) > \text{cost}_2(\langle \mathcal{D}_{ALL}, \mathcal{D}_A, \mathcal{D}_\eta \rangle)$. Finally, when Player 2 chooses \mathcal{D}_A , Player 1 prefers the design \mathcal{D}_ρ iff the pair $\langle \eta, \rho \rangle$ is legal for φ , thus $\text{cost}_1(\langle \mathcal{D}_{ALL}, \mathcal{D}_A, \mathcal{D}_\eta \rangle) > \text{cost}_1(\langle \mathcal{D}_\rho, \mathcal{D}_A, \mathcal{D}_\eta \rangle)$, for a legal pair $\langle \eta, \rho \rangle$.

Thus, if $\varphi \in \text{NAE}$, then for every assignment η , there is an assignment ρ such that $\langle \eta, \rho \rangle$ is a legal pair. Then, assuming Player 3 chooses a design \mathcal{D}_η , Player 1 prefers either choosing \mathcal{D}_{ALL} or \mathcal{D}_ρ over every other design, where $\langle \eta, \rho \rangle$ is a legal pair. By the above, there is no PNE in the game. If $\varphi \notin \text{NAE}$, then there is an assignment η such that for every assignment ρ , the pair $\langle \eta, \rho \rangle$ is illegal. Then, the profile $\langle \mathcal{D}_{ALL}, \mathcal{D}_A, \mathcal{D}_\eta \rangle$ is a PNE, and we are done. ◀

References

- 1 H. Ackermann and A. Skopalik. Complexity of pure Nash equilibria in player-specific network congestion games. *Internet Mathematics*, 5(4):321–515, 2008.
- 2 S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact price of anarchy for polynomial congestion games. *SIAM J. Comput.*, 40(5):1211–1233, 2011.
- 3 E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM J. Comput.*, 38(4):1602–1623, 2008.
- 4 G. Avni and O. Kupferman. Synthesis from component libraries with costs. In *Proc. 25th CONCUR*, LNCS 8704, pages 156–172. Springer, 2014.
- 5 G. Avni, O. Kupferman, and T. Tamir. Network-formation games with regular objectives. In *Proc. 17th FoSSaCS*, LNCS 8412, pages 119–133. Springer, 2014.
- 6 B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. *SIAM J. Comput.*, 42(1):160–177, 2013.
- 7 N. Basilico, N. Gatti, and F. Amigoni. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *Proc. 8th AAMAS*, 2009.
- 8 K. Bhawalkar, M. Gairing, and T. Roughgarden. Weighted congestion games: Price of anarchy, universal worst-case examples, and tightness. In *ESA (2)*, pages 17–28, 2010.
- 9 V. Bilò. A unifying tool for bounding the quality of non-cooperative solutions in weighted congestion games. In *WAOA*, pages 215–228, 2012.
- 10 I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli. Tight bounds for selfish and greedy load balancing. *Algorithmica*, 61(3):606–637, 2011.
- 11 G. Christodoulou and M. Gairing. Price of stability in polynomial congestion games. In *Proc. 40th ICALP*, pages 496–507, 2013.
- 12 G. Christodoulou and E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *ESA*, pages 59–70, 2005.
- 13 N. Daniele, F. Guinchiglia, and M.Y. Vardi. Improved automata generation for linear temporal logic. In *Proc. 11th CAV*, LNCS 1633, pages 249–260. Springer, 1999.
- 14 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.
- 15 C. Dwork and M. Naor. Pricing via processing, or, combatting junk mail. In *Proc. CRYPTO*, pages 139–177, 2009.
- 16 T. Eiter and G. Gottlob. Note on the complexity of some eigenvector problems. Technical Report CD-TR 95/89, Christian Doppler Laboratory for Expert Systems, TU Vienna, 1995.
- 17 D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish unsplittable flows. *Theoretical Computer Science*, 348(2-3):226–239, 2005.
- 18 D. Fotakis, S. Kontogiannis, and P. Spirakis. Symmetry in Network Congestion Games: Pure Equilibria and Anarchy Cost. In *Proc. WAOA*, pages 161-175, 2005.
- 19 A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *Proc. 36th STOC*, pages 604–612, 2004.
- 20 M. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*. W. Freeman and Co., 1979.
- 21 T. Harks and M. Klimm. On the existence of pure Nash equilibria in weighted congestion games. *Math. Oper. Res.*, 37(3):419–436, 2012.
- 22 E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009.
- 23 Y. Lustig and M.Y. Vardi. Synthesis from component libraries. *STTT*, 15(5-6):603–618, 2013.
- 24 C. Meyers. *Network flow problems and congestion games: complexity and approximation results*. PhD thesis, MIT, 2006.
- 25 I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111 – 124, 1996.

- 26 A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190, 1989.
- 27 M.G. Reed, P.F. Syverson, and D.M. Goldschlag. Anonymous connections and onion routing. *IEEE J. on Selected Areas in Communication*, 1998. Issue on Copyright and Privacy Protection.
- 28 R.W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
- 29 T. Roughgarden and E. Tardos. How bad is selfish routing? *JACM*, 49(2):236–259, 2002.
- 30 L. Tran-Thanh, M. Polukarov, A. C. Chapman, A. Rogers, and N. R. Jennings. On the existence of pure strategy nash equilibria in integer-splittable weighted congestion games. In *SAGT*, pages 236–253, 2011.

A Proofs

The function Φ in Theorem 3 is an exact potential function: Consider a profile P , and assume Player i deviates from his strategy s_i in P to a strategy s'_i , and let P' be the resulting profile. We claim that $\Phi(P) > \Phi(P')$. Moreover, $\Phi(P) - \Phi(P')$ is exactly the gain of Player i from the deviation.

Consider a transition $e \in E$. Note that for every $j \neq i$, we have $L_{e,j}(P) = L_{e,j}(P')$. Thus, for $j, j' \neq i$, the values $a_e \cdot L_{e,j}(P) \cdot L_{e,j'}(P)$ and $b_e \cdot L_{e,j}(P)$ appear both in $\Phi_e(P)$ and $\Phi_e(P')$. Thus,

$$\begin{aligned} \Phi_e(P) - \Phi_e(P') &= \\ &= a_e \cdot \sum_{j=1}^k L_{e,i}(P) \cdot L_{e,j}(P) + b_e \cdot L_{e,i}(P) - \left(a_e \cdot \left(\sum_{j=1}^k L_{e,i}(P') \cdot L_{e,j}(P') \right) + b_e \cdot L_{e,i}(P') \right) = \\ &= L_{e,i}(P) \cdot (a_e \cdot L_e(P) + b_e) - L_{e,i}(P') \cdot (a_e \cdot L_e(P') + b_e) = \text{cost}_{e,i}(P) - \text{cost}_{e,i}(P'). \end{aligned}$$

Note that if $L_{e,i}(P) = L_{e,i}(P') = 0$, then $\Phi_e(P) - \Phi_e(P') = 0$. Thus,

$$\Phi(P) - \Phi(P') = \sum_{e \in E} \Phi_e(P) - \Phi_e(P') = \sum_{e \in E} \text{cost}_{e,i}(P) - \text{cost}_{e,i}(P') = \text{cost}_i(P) - \text{cost}_i(P').$$

Since we assume Player i improves his cost, we have $\text{cost}_i(P) > \text{cost}_i(P')$, and we are done. \blacktriangleleft

Proof of Theorem 4: The problem is clearly in NP. For the hardness proof we present a reduction from the *3-Partition Problem*, which is known to be strongly NP-hard [20]. Given a set of $3m$ integers $\mathcal{A} = \{s_1, s_2, \dots, s_{3m}\}$ whose total sum is mS , and for every $1 \leq i \leq 3m$ it holds that $S/4 < s_i < S/2$, the goal is to decide whether it is possible to partition \mathcal{A} into sets of size 3, each with total sum S .

Given \mathcal{A} , define the following MCG G over $3m+2$ players and $m+4$ resources $\{e_1, e_2, \dots, e_m, z, a, b, c\}$. The latency functions of the resources are $f_{e_1}(x) = \dots = f_{e_m}(x) = x$; $f_z(x) = \lceil 2x/S \rceil \cdot S$; $f_a(x) = \max(0, (x-1)^2)$; $f_b(x) = f_c(x) = x^2$. For every $1 \leq i \leq 3m$, the strategy space of Player i is $\{\{e_1^{s_i}\}, \dots, \{e_m^{s_i}\}, \{z^{s_i}, a\}\}$. The additional two players play the no-PNE game introduced in Example 1, thus we have the same strategy space $\{\{a, a, b\}, \{b, b, c\}, \{c, c, a\}\}$. We show that a PNE exists in G if and only if a 3-partition of \mathcal{A} exists.

If a 3-partition of \mathcal{A} exists, then a PNE exists: the first $3m$ players can split between the resources e_1, \dots, e_m , causing load S on each resource, thus, cost $s_i \cdot S$ for each player $1 \leq i \leq 3m$. A deviation to $\{z, a\}$ will cause cost $s_i \cdot S + 1$ and is therefore not beneficial. Since the latency function of a is lower than the latency functions of b and c , and the first $3m$ players do not cause any load on a , it is easy to verify that the profile $\{\{a, a, b\}$ and $\{c, c, a\}\}$ is stable for the two last players (achieving costs 9 and 12).

If a 3-partition does not exist, then we claim that in every profile, exactly one out of the first $3m$ players selects $\{z^{s_i}, a\}$. If no players selects this strategy, then at least one of the resources

e_1, \dots, e_m has load larger than S and a deviation to $\{z^{s_i}, a\}$ is beneficial for any of the players using this resource. If more than one player selects this strategy, then the load on z is $L_z \geq 2s_{min} > S/2$. Combining the facts that $f_z(x) = \lceil 2x/S \rceil \cdot S$ and $L_z > S/2$, we get that $f_z(L_z) \geq 2S$, thus, deviating to a resource e_1, \dots, e_m with load less than $1.5S$ (by averaging argument, at least one such resource exists) is beneficial.

Given that exactly one player selects $\{z, a\}$, we have that the load on a is 1. Thus, the two last players face exactly the game introduced in Example 1 - that has no PNE. We conclude that a PNE exists if and only if a 3-partition exists. Moreover, even if a 3-partition exists, it is strongly NP-hard to find it, and therefore, finding a PNE is strongly NP-hard as well.

In a similar way, it is possible to define a reduction from *Equal-Partition*, in which $2n$ integers should split into two sets of the same cardinality and the same total sum. The induced game will be over six resources: the first two will be assigned the partition items, and the four additional items will have the same role as z, a, b and c in the above reduction. Thus, the hardness proof is valid also for games with a constant number of resources, unlike congestion games with user-specific cost functions [1]. ◀

Proof of Theorem 5: Our proof is based on the potential function defined in [9, 18] (see Theorem 3). Consider a affine MCG G . We claim that for every profile P in G we have:

$$\frac{1}{2} \cdot \text{cost}(P) < \Phi(P) \leq \text{cost}(P).$$

In order to prove the claim, we prove that for every $e \in E$:

$$\frac{1}{2} \cdot \text{cost}_e(P) < \Phi_e(P) \leq \text{cost}_e(P).$$

For the second inequality, recall that $\text{cost}_e(P) = \sum_{1 \leq i \leq k} L_{e,i}(P) \cdot (a_e \cdot L_e(P) + b_e)$ and $L_e(P) = \sum_{1 \leq i \leq k} L_{e,i}(P)$. Thus,

$$\begin{aligned} \text{cost}_e(P) &= \sum_{i=1}^k L_{e,i}(P) \cdot (a_e \sum_{j=1}^k L_{e,j}(P) + b_e) = a_e \cdot \sum_{i=1}^k \sum_{j=1}^k L_{e,j}(P) \cdot L_{e,i}(P) + b_e \cdot \sum_{i=1}^k L_{e,i}(P) \geq \\ &\geq a_e \cdot \sum_{i=1}^k \sum_{j=i}^k L_{e,j}(P) \cdot L_{e,i}(P) + b_e \cdot \sum_{i=1}^k L_{e,i}(P) = \Phi_e(P). \end{aligned}$$

We continue to prove the first inequality. Using the same calculations as above, we have:

$$\begin{aligned} \text{cost}_e(P) - \Phi_e(P) &= a_e \cdot \sum_{i=2}^k \sum_{j=1}^{i-1} L_{e,j}(P) \cdot L_{e,i}(P) = \\ &= a_e \cdot \sum_{i=1}^k \sum_{j=i+1}^k L_{e,j}(P) \cdot L_{e,i}(P) = \Phi_e(P) - a_e \sum_{i=1}^k L_{e,i}(P)^2 - b_e \sum_{i=1}^k L_{e,i}(P). \end{aligned}$$

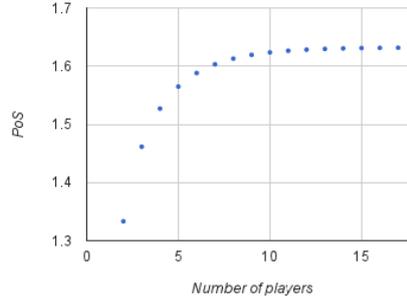
Thus,

$$\text{cost}_e(P) = 2 \cdot \Phi_e(P) - a_e \sum_{i=1}^k L_{e,i}(P)^2 - b_e \sum_{i=1}^k L_{e,i}(P).$$

Since a_e, b_e , and $L_{e,i}(P)$, for $1 \leq i \leq k$, are all positive, we get that $\text{cost}_e(P) < 2 \cdot \Phi_e(P)$, and we are done. ◀

A description of G_4 from Theorem 6:

The game G_4 is defined with respect to resources $a, b, c^1, c^2, d^1, d^2, d^3$, and d^4 , where the latency of a, b, c^1 , and c^2 are as in G_3 , d^1 has latency identity plus ϵ'' and d^2, d^3 , and d^4 have identity latency. Let $x_4 = 4!$. The strategies of players 1, 2 and 3, are the same as in G_3 only that we use x_4 instead of x_3 . We define $O_4^4 = \{\langle d^1, \frac{x_4}{4}, 4 \rangle, \langle d^2, \frac{x_4}{4}, 2 \rangle, \langle d^3, \frac{x_4}{3}, 1 \rangle, \langle d^4, \frac{x_4}{2}, 1 \rangle\}$ and $N_4^4 = \{\langle a, x_4, 1 \rangle, \langle b, \frac{x_4}{2}, 1 \rangle, \langle c^1, \frac{x_4}{3}, 1 \rangle, \langle c^2, \frac{x_4}{2}, 1 \rangle\}$. Again, we claim that \bar{N}_4 is the only PNE in the game. Note that with a slight adjustment due to the transition from x_3 to x_4 , our calculations in G_3 are valid here. So, to illustrate the idea why \bar{N}_4 is the only PNE, we show that Player 4 deviates from $\bar{N}_4[4 \leftarrow O_4^4]$ to N_4^4 , and that players 2 and 3 do not benefit by deviating from \bar{N}_4 to O_2^4 and O_3^4 , respectively. Thus, $cost_4(\bar{N}_4[4 \leftarrow O_4^4]) = \frac{x_4}{4}(4 \cdot 4 \cdot (1 + \epsilon'')) + \frac{x_4}{2 \cdot 2}(2 \cdot 2) + \frac{x_4}{3} + \frac{x_4}{2}$. In the profile $\bar{N}_4[4 \leftarrow O_4^4]$, the load on the a -type resources is 3, the load on the b -type resources is 1, and the load on the c^1 - and c^2 -type resources is 0. Thus, in \bar{N}_4 , the loads increase to 4, 2, 1, and 1, respectively, and $cost_4(\bar{N}_4) = x_4(4 \cdot 1) + \frac{x_4}{2}(2 \cdot 1 \cdot (1 + \epsilon)) + \frac{x_4}{3}(1 \cdot 1 \cdot (1 + \epsilon')) + \frac{x_4}{2}(1 \cdot 1)$. We choose $\epsilon > \epsilon'$ so that $cost_4(\bar{N}_4[4 \leftarrow O_4^4]) > cost_4(\bar{N}_4)$. We continue to show that players 2 and 3 do not benefit by deviating from \bar{N}_4 to O_2^4 . In \bar{N}_4 , the load on the b -type resources is 2, and in $\bar{N}_4[2 \leftarrow O_2^4]$ the load increases to 4 as Player 2 has two uses of every b -type resource. Thus, $cost_2(\bar{N}_4) = 24(4 \cdot 1) < 12(4 \cdot 2 \cdot (1 + \epsilon)) = cost_2(\bar{N}_4[2 \leftarrow O_2^4])$. Similarly, $cost_3(\bar{N}_4) = 24(4 \cdot 1) + 12(2 \cdot 1 \cdot (1 + \epsilon)) > 8(4 \cdot 3 \cdot (1 + \epsilon')) + 12(2 \cdot 1) = cost_3(\bar{N}_4[3 \leftarrow O_3^4])$. Recall that $x_4 = 4! = 24$. Thus, $cost(\bar{N}_4) = 24(4 \cdot 4) + 12(2 \cdot 2 \cdot (1 + \epsilon)) + 8 + 12$ and $cost(\bar{O}_4) = 24 + 12(2 \cdot 2(1 + \epsilon)) + 8(3 \cdot 3(1 + \epsilon')) + 12 + 6(4 \cdot 4(1 + \epsilon'')) + 6(2 \cdot 2) + 8 + 12$, and thus, $PoS(G_4) = 1.527$.



■ **Figure 3** A graph of $PoS(G_k)$ as a function of k . The value of $PoS(G_{17})$ is 1.6316.

Proof of Claim 6.1: We prove the claim by induction on k . For the base case, $k = 2$ and the correctness of the claim is shown in the description of G_2 above. We assume correctness for $k - 1$ and prove for k . Consider $1 \leq i \leq k$ and a profile $P = \langle P_1^k, \dots, P_k^k \rangle$ having $P_i^k = O_i^k$. Let $P' = P[i \leftarrow N_i^k]$. We claim that $cost_i(P) > cost_i(P')$.

We distinguish between two cases. In the first case, $i \neq k$. Let $Q = \langle P_1^{k-1}, \dots, P_{k-1}^{k-1} \rangle$, where for $2 \leq j \leq k$, if $P_j^k = N_j^k$, then $P_j^{k-1} = N_j^{k-1}$ and otherwise $P_j^k = O_j^k$ and $P_j^{k-1} = O_j^{k-1}$. Let $Q' = Q[i \leftarrow N_i^{k-1}]$. By the induction hypothesis, $cost_i(Q) = cost_i(Q')$. We distinguish between two cases. In the first case $P_k^k = O_k^k$, and by the construction, Player k does not share any of its resources. It is not hard to prove that $\frac{1}{x_k} \cdot cost_i(P) = \frac{1}{x_{k-1}} \cdot cost_i(Q)$ and $\frac{1}{x_k} \cdot cost_i(P') = \frac{1}{x_{k-1}} \cdot cost_i(Q')$, and the claim follows.

For the second case, $P_k^k = N_k^k$. Recall that N_k^k has a single use for every resource in N_j^k or O_j^k , for $1 \leq j \leq k - 1$. Specifically, for every $e \in O_i^k$ we have $L_e(P) = L_e(Q) + 1$, and for every

$e \in N_i^k$ we have $L_e(P') = L_e(Q') + 1$. Thus,

$$\text{cost}_i(P) = \sum_{e \in O_i^k} L_{e,i}(P) \cdot f_e(L_e(P)) = \sum_{e \in O_i^k} L_{e,i}(P) \cdot f_e(L_e(Q)+1) = \sum_{e \in O_i^k} L_{e,i}(P) \cdot f_e(L_e(Q)) + \sum_{e \in O_i^k} L_{e,i}(P).$$

Since $L_{e,i}(P) = L_{e,i}(Q)$, we have

$$\sum_{e \in O_i^k} L_{e,i}(P) \cdot f_e(L_e(Q)) = \sum_{e \in O_i^k} L_{e,i}(Q) \cdot f_e(L_e(Q)) = \text{cost}_i(Q).$$

By the induction hypothesis, $\text{cost}_i(Q) = \text{cost}_i(Q')$. Going the opposite direction, we have $\text{cost}_i(Q') = \sum_{e \in N_i^k} L_{e,i}(P') \cdot f_e(L_e(Q'))$. Note that $\sum_{e \in O_i^k} L_{e,i}(P) = |O_i^k|$. We show in the construction that $|O_i^k| = |N_i^k| = \sum_{e \in N_i^k} L_{e,i}(P')$. To conclude,

$$\text{cost}_i(P) = \sum_{e \in N_i^k} L_{e,i}(P') \cdot f_e(L_e(Q')) + \sum_{e \in N_i^k} L_{e,i}(P') = \sum_{e \in N_i^k} L_{e,i}(P') \cdot f_e(L_e(P')) = \text{cost}_i(P').$$

We continue to prove the second case in which $i = k$. Recall that we construct O_k^k and N_k^k so that $\text{cost}_k(O_k) = \text{cost}_k(\bar{N}_k)$. We claim that $\text{cost}_k(\bar{O}_k) = \text{cost}_k(P)$ and $\text{cost}_k(\bar{N}_k) = \text{cost}_k(P')$, which would conclude the proof. Recall that in the construction of O_k^k we use new resources. Thus, Player k does not share any resources when he plays O_k^k , and we have $\text{cost}_k(\bar{O}) = \text{cost}_k(P)$. We continue to show that $\text{cost}_k(\bar{N}_k) = \text{cost}_k(P')$. It is not hard to prove by induction on k , that for every two profiles S and S' in G_k , we have $\sum_{e \in E_k} L_e(S) = \sum_{e \in E_k} L_e(S')$, where E_k are the resources of the game G_k . When Player k plays N_k^k , he uses every resource in $E_k \setminus O_k^k$ exactly once. Since O_k^k uses new resources, when Player k plays N_k^k , the load on $e \in O_k^k$ is 0, and $\sum_{e \in E_k} L_e(\bar{N}_k) = \sum_{e \in (E_k \setminus O_k^k)} L_e(\bar{N}_k)$. Let $y = \sum_{2 \leq i \leq k-1} \frac{x_i}{i} \epsilon_i$. To conclude the proof,

$$\begin{aligned} \text{cost}_k(\bar{N}_k) &= \sum_{e \in N_i^k} L_{e,i}(\bar{N}_k) \cdot f_e(L_e(\bar{N}_k)) = y + \sum_{e \in N_i^k} L_e(\bar{N}_k) \cdot 1 = \\ &= y + \sum_{e \in E} L_e(\bar{N}_k) = y + \sum_{e \in E} L_e(P') \cdot 1 = \text{cost}_k(P'). \end{aligned}$$

◀

Proof of Claim 6.2: Assume towards contradiction that the social optimum profile P^* is not \bar{O} . Let $1 \leq i \leq k$ be the highest index of a player that uses his N_i^k strategy in P^* . We claim that the profile P' in which Player i plays O_i^k has $\text{cost}(P^*) > \text{cost}(P')$. For $i < j \leq k$, in P^* , Player j uses his O_j^k strategy, and thus does not share any resources with Player i in P' . It follows that $\text{cost}_j(P^*) = \text{cost}_j(P')$. For $1 \leq j < i$, since Player i shares resources with Player j in P^* and does not share resources in P' , we have that $\text{cost}_j(P^*) > \text{cost}_j(P')$. Moreover, since the e_1 -type resources have identity latency, we have $\text{cost}_j(P^*) \geq 1 + \text{cost}_j(P')$. By the construction, we have $\text{cost}_i(P^*) = \text{cost}_i(P') + \epsilon_i$, and it is possible to select ϵ_i so that $\text{cost}(P') < \text{cost}(P^*)$, which is a contradiction to the minimality of P^* , and we are done. ◀

Proof of Remark 6.1: For the upper bound, consider a linear MCG G and let N be a PNE profile that is reached by a sequence of best-response moves from the social optimum O . We tighten our analysis from Theorem 5. For $e \in E$, we have $\text{cost}_e(N) = 2\text{cost}(N) - a_e \cdot \sum_{1 \leq i \leq k} L_{e,i}^2(N) \geq \sqrt{a_e} \cdot \sum_{1 \leq i \leq k} L_{e,i}(N) = \sqrt{a_e} \cdot L_e(N) = \sqrt{\text{cost}_e(N)}$. Thus, $\text{cost}_e(N) \geq 2 - \sqrt{\text{cost}_e(N)}$ and $\text{cost}(N) \geq 2 - \sum_{e \in E} \sqrt{\text{cost}_e(N)}$. Let N' be the cheapest PNE profile. Then, $\sum_{e \in E} \sqrt{\text{cost}_e(N)} \geq \sum_{e \in E} \sqrt{\text{cost}_e(N')}$ and $\text{PoS}(G) \leq 2 - \text{cost}^{-1}(O) \cdot \sum_{e \in E} \sqrt{\text{cost}_e(N')}$.

For the lower bound, we show that the inequality is essentially an equality for the MCGs in the family described in Theorem 6. Consider $k \in \mathbb{N}$. Since in the social optimum in G_k the

resources are used by exactly one player, we have $\Phi(\bar{O}_k) = \text{cost}(\bar{O}_k)$. Claim 6.1 implies that, essentially, $\Phi(\bar{O}_k) = \Phi(\bar{N}_k)$. Finally, since in \bar{N}_k a resource is used by a player at most one time, we have $\Phi(\bar{N}_k) = \frac{1}{2} \text{cost}(\bar{N}_k) + \frac{1}{2} \sum_{e \in E} \sqrt{\text{cost}_e(\bar{N}_k)}$. Thus, $\text{PoS}(G_k) = 2 - \text{cost}^{-1}(O_k) \cdot \sum_{e \in E} \sqrt{\text{cost}_e(\bar{N}_k)}$. Since \bar{N}_k is the only PNE in G_k , we are done.

Proof of Theorem 7: Recall that \mathcal{P}_d is the set of polynomials of degree at most d . The proof relies on the following lemma [2]: Let $d \in \mathbb{N}$. Then,

$$\min_{(\lambda, \mu) \in \mathbb{R}^2} \left\{ \frac{\lambda}{1 - \mu} : \forall x, y \in \mathbb{R}, f \in \mathcal{P}_d \text{ we have } y \cdot f(x+y) \leq \lambda \cdot y \cdot f(y) + \mu \cdot x \cdot f(x) \right\} = \Phi_d^{d+1}. \quad (1)$$

Consider an MCG G , and let N be a PNE profile and O be the social optimum profile. Consider $1 \leq i \leq k$. Recall that for a profile P and resource $e \in E$, we use P_e to denote the load on e in P , the latency function on $e \in E$ is $f_e \in \mathcal{P}_d$, and for $1 \leq i \leq k$, we use P^i to denote the multiset that Player i chooses in P and $P^i(e)$ the number of times Player i uses e in P^i . By the definition of a PNE, we have $\text{cost}_i(N) \leq \text{cost}_i(N[i \leftarrow O^i])$. Next,

$$\text{cost}_i(N[i \leftarrow O^i]) = \sum_{e \in E} O^i(e) \cdot f_e(N_e - N_e^i + O_e^i) \leq \sum_{e \in E} O^i(e) \cdot f_e(N_e + O_e)$$

Where the last inequality follows from the fact that the latency functions we consider are monotonically increasing.

$$\text{cost}(N) = \sum_{1 \leq i \leq k} \text{cost}_i(N) \leq \sum_{1 \leq i \leq k} \sum_{e \in E} O^i(e) \cdot f_e(N_e + O_e) = \sum_{e \in E} O_e \cdot f_e(N_e + O_e)$$

Let $\lambda, \mu \in \mathbb{R}$ that minimize the expression in (1). Then,

$$\sum_{e \in E} O_e \cdot f_e(N_e + O_e) \leq \sum_{e \in E} \lambda \cdot O_e \cdot f(O_e) + \mu \cdot N_e \cdot f(N_e) = \lambda \cdot \text{cost}(O) + \mu \cdot \text{cost}(N).$$

Rearranging yields the theorem.

Proof of Theorem 9: We show a family of two-player congestion games in which the PoA is arbitrarily high. For $n \in \mathbb{N}$, we define the game G_n . The resources are $E = \{e_1, e_2\}$. The strategy spaces of the players are mirrored. Player 1's strategy space is $\Sigma_1 = \{e_1, e_2^n\}$ and Player 2 strategy space is $\Sigma_2 = \{e_1^n, e_2\}$. The latency functions are both the factorial function, i.e., $f_{e_1}(l) = f_{e_2}(l) = l!$.

We continue to calculate $\text{PoA}(G_n)$. Clearly, the social optimum is attained in the profile $\langle e_1, e_2 \rangle$ and its cost is 2. We claim that the profile $N = \langle e_2^n, e_1^n \rangle$ is a PNE. Indeed, $\text{cost}_1(N) = n \cdot n!$, and by deviating to e_1 , his cost increases to $1 \cdot (n+1)!$. The proof for Player 2 is dual. Since $\text{cost}(N) = 2n \cdot n!$, we have $\text{PoA}(G_n) = n \cdot n!$, and we are done. \blacktriangleleft

A formal construction of the DFA $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$: Recall that the library of components is $\mathcal{L} = \{\mathcal{B}_1, \dots, \mathcal{B}_n\}$. For $i \in [n]$, we have $\mathcal{B}_i = \langle \Sigma, B_i, \delta_i, b_i^0, F_i, E_i \rangle$, where $E_i \subseteq B_i$ is a set of exit states having $b_i^0 \notin E_i$ and $F_i \cap E_i = \emptyset$. The states of the components are disjoint, thus $B_i \cap B_j = \emptyset$, for $j \neq i$. We denote by \mathcal{B} , \mathcal{F} , and \mathcal{E} the union of all states, accepting states, and exit states, respectively, thus $\mathcal{B} = \bigcup_{i \in [n]} B_i$, $\mathcal{F} = \bigcup_{i \in [n]} F_i$, and $\mathcal{E} = \bigcup_{i \in [n]} E_i$. Recall that a design \mathcal{D} is a transducer over input alphabet \mathcal{E} and output alphabet $[n]$. Consider a design $\mathcal{D} = \langle \mathcal{E}, [n], D, \delta_{\mathcal{D}}, d^0, \nu \rangle$, where $\nu : D \rightarrow [n]$.

We construct the composition system $\mathcal{A}_{\mathcal{L}, \mathcal{D}} = \langle \Sigma, Q_{\mathcal{L}, \mathcal{D}}, \delta_{\mathcal{L}, \mathcal{D}}, q_{\mathcal{L}, \mathcal{D}}^0, F_{\mathcal{L}, \mathcal{D}} \rangle$ as follows. The set of states $Q_{\mathcal{L}, \mathcal{D}} \subseteq (\mathcal{B} \setminus \mathcal{E}) \times D$ consists of pairs of a *component state* from \mathcal{B} and an *design state* from

D . The component states are consistent with ν , thus $Q_{\mathcal{L},\mathcal{D}} = \bigcup_{i \in [n]} (B_i \setminus E_i) \times \{q : \nu(q) = i\}$. In exit states, the composition immediately moves to the initial state of the next component, which is why the component states of $\mathcal{A}_{\mathcal{L},\mathcal{D}}$ do not include \mathcal{E} . Consider a state $\langle b, q \rangle \in Q_{\mathcal{L},\mathcal{D}}$ and a letter $\sigma \in \Sigma$. Let $i \in [n]$ be such that $b \in B_i$. When a run of $\mathcal{A}_{\mathcal{L},\mathcal{D}}$ reaches the state $\langle b, q \rangle$, the component \mathcal{B}_i is in control. Recall that b is not an exit state. Let $b' = \delta_i(b, \sigma)$. If $b' \notin E_i$, then \mathcal{B}_i does not relinquish control after reading σ and $\delta_{\mathcal{L},\mathcal{D}}(\langle b, q \rangle, \sigma) = \langle b', q \rangle$. If $b' \in E_i$, then \mathcal{B}_i relinquishes control through b' , and it is the design's task to choose the next component to gain control. Let $q' = \delta(q, b')$ and let $j = \nu(q')$. Then, \mathcal{B}_j is the next component to gain control (possibly $j = i$). Accordingly, we advance \mathcal{D} to q' and continue to the initial state of \mathcal{B}_j . Formally, $\delta_{\mathcal{L},\mathcal{D}}(\langle b, q \rangle, \sigma) = \langle b_j^0, q' \rangle$. (Recall that $b_j^0 \notin E_j$, so the new state is in $Q_{\mathcal{L},\mathcal{D}}$.) Note also that a visit in b' is skipped. The component that gains initial control is chosen according to $\nu(d^0)$. Thus, $q_{\mathcal{L},\mathcal{D}}^0 = \langle b_j^0, d^0 \rangle$, where $j = \nu(d^0)$. Finally, the accepting states of $\mathcal{A}_{\mathcal{L},\mathcal{D}}$ are these in which the component state is accepting, thus $F_{\mathcal{L},\mathcal{D}} = \mathcal{F} \times D$. (Recall that $\mathcal{F} \cap \mathcal{E} = \emptyset$, so $F_{\mathcal{L},\mathcal{D}} \subseteq Q_{\mathcal{L},\mathcal{D}}$.)

Proof of Theorem 10: Consider an MCG $\langle K, E, \{\Sigma_i\}_{i \in K}, \{f_e\}_{e \in E} \rangle$. Recall that Σ_i is the set of strategies for Player i that consists of multisets over E . We construct a CLG with alphabet $E \cup \bigcup_{i \in K} \Sigma_i$. For $i \in K$, the specification \mathcal{S}_i for Player i consists of $|\Sigma_i|$ words. Every strategy $s = \{e_1, \dots, e_n\}$ (allowing duplicates) in Σ_i contributes to $L(\mathcal{S}_i)$ the word $s \cdot e_1 \cdot e_2 \cdot \dots \cdot e_n$. We construct a library \mathcal{L} with $|E| + \sum_{i \in K} |\Sigma_i|$ components of two types: a *strategy component* \mathcal{B}_s for each $s \in \Sigma_i$ and a *resource component* \mathcal{B}_e for each $e \in E$. In addition, \mathcal{L} contains the component \mathcal{B}_{acc} that is depicted in Figure 2. Intuitively, a correct design must choose one strategy component \mathcal{B}_s and then use the component \mathcal{B}_e the same number of times e appears in s . We continue to describe the components. For $s \in \Sigma_i$, the component \mathcal{B}_s relinquishes control only if the letter s is read. It accepts every word in $L(\mathcal{S}_i)$ that does not start with s . For $e \in E$, the resource component \mathcal{B}_e has an initial state with an e -labeled transition to an exit state. Finally, the latency function for the resource components coincides with latency functions of the resources in the MCG, thus for $e \in E$, we have $f_{\mathcal{B}_e} = f_e$. The other latency functions are $f \equiv 0$. We prove that there is a cost-preserving one-to-one and onto correspondence between correct designs with respect to \mathcal{S}_i and strategies in Σ_i , implying the existence of the required function between the profiles.

Consider $i \in K$. We claim that there is a one-to-one and onto correspondence between correct designs with respect to \mathcal{S}_i and strategies in Σ_i . Thus, there is a one-to-one and onto correspondence between the sets of profiles in the MCG and the CLG. We describe the design \mathcal{D}_s that corresponds to the strategy $s = \{e_1, \dots, e_n\} \in \Sigma_i$. The design \mathcal{D}_s assigns initial control to the strategy component \mathcal{B}_s . Recall that \mathcal{B}_s relinquishes control after reading the letter s . Then, \mathcal{D}_s assigns control to the components $\mathcal{B}_{e_1}, \dots, \mathcal{B}_{e_n}$ and finally control is assigned to \mathcal{B}_{acc} . Clearly the word $s \cdot e_1 \cdot e_2 \cdot \dots \cdot e_n \in L(\mathcal{S}_i)$ is accepted, no other word that starts with s is accepted, and since \mathcal{B}_s accepts every word in $L(\mathcal{S}_i)$ that does not start with s , we have that \mathcal{D}_s is correct. It is not hard to see that there are no other correct designs. To see that the correspondence is onto, note that since control has to be assigned to some component \mathcal{B}_s , each correct design corresponds to a strategy $s \in \Sigma_i$. Finally, consider a profile $P = \langle \mathcal{D}_{s_1}, \dots, \mathcal{D}_{s_k} \rangle$ that corresponds to the profile $P' = \langle s_1, \dots, s_k \rangle$. Clearly, for $i \in K$, we have $cost_i(P) = cost_i(P')$, and we are done.

Proof of Theorem 13: We describe the specification \mathcal{S} . The alphabet is $\Sigma = \{\#\} \cup \bigcup_{i \in [n]} \{\#_i^0, \dots, \#_i^m, x_i\} \cup \bigcup_{j \in [m]} \{C_j, \tilde{C}_j\}$. The language $L(\mathcal{S})$ includes $3m + 1$ words. A long word $\# \cdot w_1 \cdot \dots \cdot w_n$, where for $i \in [n]$, we have $w_i = \#_i^0 x_i \#_i^1 \dots x_i \#_i^m$, and $3m$ short words, which we describe next. For $j \in [m]$, let $C_j = \ell_j^1 \vee \ell_j^2 \vee \ell_j^3$. For $k \in \{1, 2, 3\}$, let j_k be the index of the variable in the literal ℓ_j^k . That is, $\ell_j^1 \in \{x_{j_1}, \neg x_{j_1}\}$, $\ell_j^2 \in \{x_{j_2}, \neg x_{j_2}\}$, and $\ell_j^3 \in \{x_{j_3}, \neg x_{j_3}\}$. For $k \in \{1, 2, 3\}$, we have $C_j x_{j_k} \tilde{C}_j x_{j_k} \in L(\mathcal{S})$. Clearly it is possible to construct \mathcal{S} so that its size is polynomial in n and m .

The components of the library \mathcal{L} are \mathcal{B}_0 , $\mathcal{B}_{x_i}^j$ and $\mathcal{B}_{\neg x_i}^j$ for $j \in [m]$ and $i \in [n]$, and $\mathcal{B}_{C_j, x_{j_k}}$ for

$j \in [m]$ and $k \in \{1, 2, 3\}$. The components of the library are depicted in Figure 2. The latency functions of the components \mathcal{B}_0 and $\mathcal{B}_{C_j, x_{j_k}}$, for $j \in [m]$ and $k \in \{1, 2, 3\}$ is the constant function $f \equiv 0$. The latency function of the components $\mathcal{B}_{x_i}^j$, for $j \in [m]$ and $i \in [n]$ is the identity function $f(x) = x$. Thus, using a component $\mathcal{B}_{x_i}^j$ once costs 1.

Assume that φ is satisfiable, and let $\eta : \{x_1, \dots, x_n\} \rightarrow \{T, F\}$ be a satisfying assignment to the variables. We describe a correct design \mathcal{D} with $\text{cost}(\mathcal{D}) = \mu = nm + m$. The component that gains control first is \mathcal{B}_0 . We distinguish between two types of words. First, assume \mathcal{B}_0 exists through $e_\#$ after reading $\#$. Thus, the only word that should be accepted from this point is the suffix $w_1 \cdot \dots \cdot w_m$ of the long word in $L(\mathcal{S})$, where recall that for $i \in [n]$ we have $w_i = \#_i^0 x_i \#_i^1 \dots x_i \#_i^m$. For $i \in [n]$, assuming the prefix $\#w_1 \cdot \dots \cdot w_{i-1}$ has been read. If $\eta(x_i) = T$, then control is assigned to $\mathcal{B}_{\neg x_i}^1$, and otherwise control is assigned to $\mathcal{B}_{x_i}^1$. Then, the next $m - 1$ assignments of control, assuming the prefix that is read is a prefix of w_i , are to the components $\mathcal{B}_{\neg x_i}^2, \dots, \mathcal{B}_{\neg x_i}^m$ if $f(x_i) = T$, and otherwise to $\mathcal{B}_{x_i}^2, \dots, \mathcal{B}_{x_i}^m$. If some other word is read, control is assigned to \mathcal{B}_{rej} . Finally, control is assigned to \mathcal{B}_{acc} . Note that so far we have used nm components with identity latency, each with one use, for a total cost of nm .

Recall that $L(\mathcal{S})$ has $3m$ short words of the form $C_j x_{j_1} \tilde{C}_j x_{j_k}$, where $j \in [m]$ and $k \in \{1, 2, 3\}$. When reading the prefix C_j of such a word, the component \mathcal{B}_0 relinquishes control through e_{C_j} . Assume, WLog, that $f(\ell_j^1) = T$ (there must be such a literal as f is a satisfying assignment). We refer to x_{j_1} as the *witness* for C_j . The component that gains control after \mathcal{B}_0 is $\mathcal{B}_{C_j, x_{j_1}}$, which relinquishes control only if x_{j_1} is read. Let $0 \leq j' < j$ be the number of clauses with index less than j that have x_{j_1} as their witness. The component that gains control after $\mathcal{B}_{C_j, x_{j_1}}$ is $\mathcal{B}_{x_{j_1}}^{j'+1}$. This component has many exit states. If it exists after reading $\tilde{C}_j x_{j_1}$, then control is assigned to \mathcal{B}_{acc} , and otherwise control is assigned to \mathcal{B}_{rej} . Note that we have used another m components with identity latency, where each component is used once. Thus, the total cost of \mathcal{D} is $nm + m$. It is not hard to see that \mathcal{D} is correct, and we are done.

For the other direction, assume there is a correct design \mathcal{D} that costs at most $nm + m$. We define an assignment $\eta : \{x_1, \dots, x_n\} \rightarrow \{T, F\}$. Let $i \in [n]$. Since the long word $\#w_1 \dots w_n$ is in $L(\mathcal{S})$, it must be in $L(\mathcal{A}_{\mathcal{L}, \mathcal{D}})$. Note that the only components that can process the letter $\#_i^0$ are $\mathcal{B}_{x_i}^1$ and $\mathcal{B}_{\neg x_i}^1$. Thus, after reading $\#w_1 \dots w_{i-1}$, one of these components must gain control. We define $\eta(x_i) = T$ if $\mathcal{B}_{\neg x_i}^1$ gains control and $\eta(x_i) = F$ if $\mathcal{B}_{x_i}^1$ gains control.

We claim that η is a satisfying assignment for φ . In order to do so, we make two observations. First, recall that, for $i \in [n]$, if $\eta(x_i) = T$, then $\mathcal{B}_{\neg x_i}^1$ gains control after $\#w_1 \dots w_{i-1}$ is read. If the prefix $\#_i^1 x_i$ of w_i is read, then $\mathcal{B}_{\neg x_i}^2$ must gain control as this is the only component that can process the next prefix $\#_i^2 x_i$ of w_i . Generalizing this observation, after reading $\#w_1 \dots w_{i-1}$, if $\eta(x_i) = T$, when reading the word w_i , the components that are in control are $\mathcal{B}_{\neg x_i}^1, \dots, \mathcal{B}_{\neg x_i}^m$, and if $\eta(x_i) = F$, then $\mathcal{B}_{x_i}^1, \dots, \mathcal{B}_{x_i}^m$ are in control. Thus, \mathcal{D} uses at least nm components with identity latency.

Next, let $j \in [m]$ with $C_j = \ell_j^1 \vee \ell_j^2 \vee \ell_j^3$ such that $\ell_j^k \in \{x_{j_k}, \neg x_{j_k}\}$, for $k \in \{1, 2, 3\}$. Note that after reading the letter C_j control must be assigned to a component $\mathcal{B}_{C_j, x_{j_k}}$ as $C_j x_{j_k} \tilde{C}_j x_{j_k} \in L(\mathcal{S})$ and these are the only components that can process words that start with one of the three letters x_{j_1}, x_{j_2} , or x_{j_3} . Denote by ℓ_j and x_j the literal and variable with $\ell_j \in \{x_j, \neg x_j\}$ such that \mathcal{B}_{C_j, x_j} gains control from \mathcal{D} after C_j is read. When x_j is read in \mathcal{B}_{C_j, x_j} , control is relinquished. Note that the component that gains control next must be a component $\mathcal{B}_{\ell_j^t}^t$, for $t \in [m]$, as these are the only components that can process $\tilde{C}_j x_j$. Let d be the state in \mathcal{D} that is reached after reading $C_j x_j^2$. We claim that there is no other word $w \in \Sigma^*$ such that the run of \mathcal{D} on w reaches d . Indeed, the word $\tilde{C}_j x_j$ is accepted from d and there is only one word in $L(\mathcal{S})$ that ends with this suffix.

² Recall that \mathcal{D} reads exit states and not words over Σ , so this notation is not formal and we mean the exit states that are read during the run on the word.

Since we have already counted nm uses of components with identity latency and \mathcal{D} costs at most $nm + m$, we conclude that \mathcal{D} uses exactly $nm + m$ such components each with a single use. Specifically, for $j \in [m]$, the component $\mathcal{B}_{\ell_j^t}$ that gains control after $C_j x_j$ is read is used once and it is used when reading the word $C_j x_j \tilde{C}_j x_j$. Thus, when reading the long word $\#w_1 \dots w_m$ it does not gain control, and by our definition of η , we have $\eta(\ell_j) = T$. We conclude that η is a satisfying assignment, and we are done.

Proof of Theorem 14: We now describe the reduction in detail. Let $\varphi = C_1 \wedge \dots \wedge C_m$, where for $j \in [t]$ we have $C_j = \ell_j^1 \vee \ell_j^2 \vee \ell_j^3$, we construct a CLG G with three players such that $\varphi \in \text{NAE}$ iff G does not have a PNE. We describe the specifications of the players. The alphabet is $\Sigma = \{\#, \&, T, F, a, b\} \cup \bigcup_{i \in [n]} \{\#_{x_i}^0, \dots, \#_{x_i}^m, x_i, \#_{y_i}^0, \dots, \#_{y_i}^m, y_i\} \cup \bigcup_{j \in [m]} \{C_j, \tilde{C}_j\} \cup \bigcup_{j \in [m]} \{\ell_j^1, \ell_j^2, \ell_j^3\}$. The specification \mathcal{S}_3 for Player 3 consists of a single word $w_1 \dots w_n$, where for $i \in [n]$, we have $w_i = \#_{x_i}^0 x_i \#_{x_i}^1 \dots x_i \#_{x_i}^m$. The specification \mathcal{S}_2 for Player 2 consists of two words a^3 and b^5 . The specification \mathcal{S}_1 for Player 1 consists of $6m + 1$ words. Similarly to Player 3, it has a long word $\# \cdot v_1 \dots v_n$, where for $i \in [n]$, we have $v_i = \#_{y_i}^0 y_i \#_{y_i}^1 \dots y_i \#_{y_i}^m$, and $6m$ short words, which we describe next. For $j \in [m]$, recall that $C_j = \ell_j^1 \vee \ell_j^2 \vee \ell_j^3$. For $k \in \{1, 2, 3\}$, we have $C_j \ell_j^k T \tilde{C}_j T \ell_j^k, C_j \ell_j^k F \tilde{C}_j F \ell_j^k \in L(\mathcal{S}_1)$. Clearly it is possible to construct $\mathcal{S}_1, \mathcal{S}_2$, and \mathcal{S}_3 so that their size is polynomial in n and m .

We continue to describe the library \mathcal{L} . For every $i \in [n]$ and $t \in [m]$, there are *variable components* $\mathcal{B}_{x_i}^t, \mathcal{B}_{-x_i}^t, \mathcal{B}_{y_i}^t$, and $\mathcal{B}_{-y_i}^t$, which have identity latency functions and are similar to these in the lower bound proof of Theorem 13. The *upper* part of the components is almost identical to these depicted in Figure 2. For example, $\mathcal{B}_{x_i}^1$ relinquishes control after reading $\#_{x_i}^0 x_i$. The *lower* part is slightly different. Consider $j \in [m]$ and a literal ℓ that appears in C_j . Intuitively, as in Theorem 13, designs of Player 1 and Player 3 correspond to assignments and must use either all the components $\mathcal{B}_{\ell}^1, \dots, \mathcal{B}_{\ell}^m$ or all the components $\mathcal{B}_{-\ell}^1, \dots, \mathcal{B}_{-\ell}^m$. Choosing the first corresponds to assigning value false to ℓ and choosing the second corresponds to assigning value true to ℓ . The components $\mathcal{B}_{\ell}^1, \dots, \mathcal{B}_{\ell}^m$ have a path labeled $\tilde{C}_j T \ell$ that leads to an exit state. The components $\mathcal{B}_{-\ell}^1, \dots, \mathcal{B}_{-\ell}^m$ have a path labeled $\tilde{C}_j F \ell$ that leads to an exit state.

For every $j \in [m]$, there are six components $\mathcal{B}_{C_j, \ell_j^a, \ell_j^b}$ for $a \neq b \in \{1, 2, 3\}$ with latency function $f \equiv 0$. The component $\mathcal{B}_{C_j, \ell_j^a, \ell_j^b}$ corresponds to assigning value true to ℓ_j^a and value false to ℓ_j^b . Let $c \in \{1, 2, 3\}$ such that $c \neq a$ and $c \neq b$. The component $\mathcal{B}_{C_j, \ell_j^a, \ell_j^b}$ relinquishes control after reading $\ell_j^a T$ and $\ell_j^b T$, after which a correct design must assign control to components $\mathcal{B}_{\ell_j^a}^{t_1}$ and $\mathcal{B}_{-\ell_j^b}^{t_2}$, for some $t_1, t_2 \in [m]$, in order to accept the suffixes $\tilde{C}_j T \ell_j^a$ and $\tilde{C}_j F \ell_j^b$, respectively. The other words that start with C_j are accepted *for free*. The component $\mathcal{B}_{C_j, \ell_j^a, \ell_j^b}$ accepts the words $\ell_j^a F \tilde{C}_j F \ell_j^a, \ell_j^b T \tilde{C}_j T \ell_j^b, \ell_j^c T \tilde{C}_j T \ell_j^c$, and $\ell_j^c F \tilde{C}_j F \ell_j^c$.

Similarly to the proof of Theorem 13, a correct design \mathcal{D}_η for Player 3 corresponds to an assignment $\eta : \{x_1, \dots, x_n\} \rightarrow \{T, F\}$ and a correct design \mathcal{D}_ρ corresponds to an assignment $\rho : \{y_1, \dots, y_n\} \rightarrow \{T, F\}$. Moreover, every correct design \mathcal{D}_ρ must use $mn + 2m$ variable components. Assuming Player 3 chooses the design \mathcal{D}_η , then there is a Player 1 correct design \mathcal{D}_ρ that uses only variable components with load 1 iff the pair $\langle \eta, \rho \rangle$ is legal for φ .

We describe the rest of the components in the library. There are components \mathcal{B}_{acc} and \mathcal{B}_{rej} that have constant latency function 0 and appear at the *end* of runs. We do not specify when they gain control below as it is immediate. There are components $\mathcal{B}_0^A, \mathcal{B}_0^B, A$, and B , with latency functions 0, 0, $f_A(x) = x^2 + 125$ and $f_B(x) = 4x^2$, respectively. The component \mathcal{B}_0^A relinquishes control after reading a and accepts the word b^5 . Dually, the component \mathcal{B}_0^B relinquishes control after reading b and accepts the word a^3 . The components A and B relinquish control after reading a and b , respectively. Thus, a correct Player 2 design must assign initial control to either \mathcal{B}_0^A or \mathcal{B}_0^B , and then assign control two times to A and four times to B , respectively. The component \mathcal{B}_{ALL} is identical to \mathcal{S}_1 . Thus,

Player 1 has a correct design that uses only it, and we refer to this design as \mathcal{D}_{ALL} . Regardless of what design the other players choose, Player 1 pays $1100 + mn + 2m$ for \mathcal{D}_{ALL} . The component \mathcal{B}_0 has a constant latency function 33. It relinquishes control through a unique exit state if #, &, or C_j , for $j \in [m]$ is read. Thus, a correct Player 1 design that does not use \mathcal{B}_{ALL} must assign initial control to \mathcal{B}_0 . When \mathcal{B}_0 relinquishes control after reading #, control must be assigned to mn variable components, after reading C_j , for some $j \in [m]$, control must be assigned to some component $\mathcal{B}_{C_j, \ell_j^a, \ell_j^b}$ as described in the above, and after reading &, control must be assigned to A for six consecutive times and then to B for one time.

We prove that the reduction is correct. Assume that $\varphi \notin \text{NAE}$, thus there is an assignment $\eta : \{x_1, \dots, x_n\} \rightarrow \{T, F\}$ such that for every assignment $\rho : \{y_1, \dots, y_n\} \rightarrow \{T, F\}$, we have that $\langle \eta, \rho \rangle$ is illegal for φ . We claim that the profile $P = \langle \mathcal{D}_{ALL}, \mathcal{D}_B, \mathcal{D}_\eta \rangle$ is a PNE. Clearly Player 2 does not deviate. Since Player 1 uses \mathcal{D}_{ALL} it does not use any of the variable components and all the designs of Player 3 cost the same, thus Player 3 does not deviate as well. Assume Player 1 benefits from deviating to a design \mathcal{D}_ρ . Since Player 2 uses \mathcal{D}_B , Player 1 pays for his uses of \mathcal{B}_0 , A , and B a cost of $33 + 966 + 100 = 1099$. Recall that \mathcal{D}_ρ uses $nm + 2m$ variable components. Since $\text{cost}_1(P) = 1100 + nm + 2m$ and we assume that deviating to \mathcal{D}_ρ is beneficial, the load on every variable component that is used in \mathcal{D}_ρ must be 1. Thus, similarly to the proof of Theorem 13, we can show that $\langle \eta, \rho \rangle$ is legal for φ , which is a contradiction. Next, assume that $\varphi \in \text{NAE}$. We show that there is no PNE in G . Consider a profile in which Player 3 chooses the design \mathcal{D}_η , and let ρ be an assignment such that $\langle \rho, \eta \rangle$ is legal for φ . It is not hard to see that choosing \mathcal{D}_η or \mathcal{D}_{ALL} dominates every other choice of design for Player 1. In Table 2 we show that no matter which design Player 2 chooses, there is no PNE, and we are done.

	\mathcal{D}_ρ	\mathcal{D}_{ALL}
\mathcal{D}_A	$1099 + mn + 2m, 400$	$1100 + mn + 2m, 256$
\mathcal{D}_B	$1138 + mn + 2m, 378$	$1100 + mn + 2m, 288$

■ **Table 2** Players costs. Each entry describes the cost of Player 1 followed by the cost of Player 2.

B Splittable (Non-Atomic) Games

In a splittable game each player can split his task among several strategies. This model suits several applications, in particular planning of preemptive production. Splittable games are well-understood in classical and weighted congestion games [29, 8]. We define the corresponding MCG and show that the positive PNE-existence result, known for weighted congestion games, carry over to games with multisets of resources.

Recall that Σ_i is the strategy space of Player i , where $\Sigma_i = \{m_{i,1}, \dots, m_{i,c_i}\}$ and for each $1 \leq j \leq c_i$, $m_{i,j}$ is a multiset over E . A splitted strategy for a player i is given by $\hat{s}_i = \langle \alpha_{i,1}, \dots, \alpha_{i,c_i} \rangle$ such that $\sum_{j=1}^{c_i} \alpha_j = 1$.

A profile of a game G is a set $P = \langle \hat{s}_1, \hat{s}_2, \dots, \hat{s}_k \rangle$ of strategies selected by the players. For a resource $e \in E$, the load that player i generates on e is $L_{e,i}(P) = \sum_{j=1}^{c_i} \alpha_j \cdot m_{i,j}(e)$. The load on e in P , is $L_e(P) = \sum_{1 \leq i \leq k} L_{e,i}(P)$.

The players' costs are defined as in the unsplitable model, that is, given a profile P , a resource $e \in E$, and $1 \leq i \leq k$, the cost of e for Player i in P is $\text{cost}_{e,i}(P) = L_{e,i}(P) \cdot f_e(L_e(P))$. The cost of Player i in the profile P is then $\text{cost}_i(P) = \sum_{e \in E} \text{cost}_{e,i}(P)$.

A profile P is a PNE if for every player i the cost of i when playing $\langle \alpha_{i,1}, \dots, \alpha_{i,c_i} \rangle$ is not higher than playing $\langle \dots, \alpha_{i,v_1} + \delta, \dots, \alpha_{i,v_2} - \delta, \dots \rangle$ for all v_1, v_2 and δ . In particular, it means that the

marginal cost for all strategies with $\alpha_i > 0$ is the same. The proof of the following theorem is identical to the corresponding proof for classical CGs [29].

► **Theorem 15.** *Every splittable MCG has a PNE.*

For example, consider the instance in Example 1, where $\Sigma_i = \langle \{a^2, b\}, \{b^2, c\}, \{c^2, a\} \rangle$ for $i = 1, 2$. The profile $P = \langle \{\frac{1}{2}, \frac{1}{3}, \frac{1}{6}\}, \{\frac{1}{6}, \frac{1}{3}, \frac{1}{2}\} \rangle$ is a (non-unique) PNE. For each $e \in \{a, b, c\}$, we have $L_a(P) = 2 \cdot \frac{1}{2} + 1 \cdot \frac{1}{6} + 2 \cdot \frac{1}{6} + 1 \cdot \frac{1}{2} = 2$, $L_b(P) = 2$ and $L_c(P) = 2$, thus $cost_1(P) = cost_2(P) = 4 \cdot 3 = 12$. This profile is a PNE, since all strategies have the same marginal cost.