

Dynamic Resource Allocation Games

Guy Avni¹, Thomas A. Henzinger¹, and Orna Kupferman²

¹ IST Austria

² The Hebrew University

Abstract. In *resource allocation games*, selfish players share resources that are needed in order to fulfill their objectives. The cost of using a resource depends on the load on it. In the traditional setting, the players make their choices concurrently and in one-shot. That is, a strategy for a player is a subset of the resources. We introduce and study *dynamic resource allocation games*. In this setting, the game proceeds in phases. In each phase each player chooses one resource. A scheduler dictates the order in which the players proceed in a phase, possibly scheduling several players to proceed concurrently. The game ends when each player has collected a set of resources that fulfills his objective. The cost for each player then depends on this set as well as on the load on the resources in it – we consider both congestion and cost-sharing games. A prime application of dynamic resource allocation games is the setting of networks in which players choose their routes edge by edge, with choices depending on earlier choices of other players. We study the stability of dynamic resource allocation games, where the appropriate notion of stability is that of subgame perfect equilibrium, study the inefficiency incurred due to selfish behavior, and also study problems that are particular to the dynamic setting, like constraints on the order in which resources can be chosen or the problem of finding a scheduler that achieves stability.

1 Introduction

Resource allocation games (RAGs, for short) [23] model settings in which selfish agents share resources that are needed in order to fulfill their objectives. The cost of using a resource depends on the load on it. Formally, a k -player RAG G is given by a set E of resources and a set of possible strategies for each player. Each strategy is a subset of resources, fulfilling some objective of the player. Each resource $e \in E$ is associated with a latency function $\ell_e : \mathbb{N} \rightarrow \mathbb{R}$, where $\ell_e(\gamma)$ is the cost of a single use of e when it has load γ . For example, in *network formation games* (NFGs, for short) [2], a network is modeled by a directed graph, and each player has a source and a target vertex. In the corresponding RAG, the resources are the edges of the graph and the objective of each player is to connect his source and target. Thus, a strategy for a player is a set of edges that form a simple path from the source to the target. When an edge e is used by m players, each of them pays $\ell_e(m)$ for his use.

A key feature of RAGs is that the players choose how to fulfill their objectives *in one shot* and *concurrently*. Indeed, a strategy for a player is a subset of the resources – chosen as a whole, and the players choose their strategies simultaneously. In many settings, however, resource sharing proceeds in a different way. First, in many settings, the choices of the players are made resource by resource as the game evolves. For example, when the network in an NFG models a map of roads and players are drivers choosing routes, it makes sense to allow each driver not to commit to a full route in the beginning of the game but rather to choose one road (edge) at each junction (vertex), gradually composing the full route according to the congestion observed. Second, players may not reach the junctions together. Rather, in each “turn” of the game, only a subset of the players (say, these that have a green light) proceed and chose their next road.

As another example to a rich composition and scheduling of strategies, consider the setting of *synthesis from component libraries* [17], where a designer synthesizes a system from existing components. It is shown in [4,6] that when multiple designers use the same library, a RAG arises. Here too, the choice of components may be made during the design process and may evolve according to choices of other designers.

In this work we introduce and study *dynamic resource allocation games*, which allow the players to choose resources in an iterative and non-concurrent manner. A dynamic RAG is given by a pair $\mathcal{G} = \langle G, \nu \rangle$, where G is a k -player RAG and $\nu : \{1, \dots, k\} \rightarrow \{1, \dots, k\}$ is a *scheduler*. A dynamic RAG proceeds in *phases*. In each phase, each player chooses one resource. A phase is partitioned into at most k *turns*, and the scheduler dictates which players proceed in each turn: Player i moves at turn $\nu(i)$. Note that the scheduler may assign the same turn to several players, in which case they choose a resource simultaneously in a phase. Once all turns have been taken, a phase is concluded and a new phase begins. A *strategy* for a player in a dynamic RAG maps the history of choices made by the players so far (that is, the choices of all players in earlier phases as well as the choices of players that proceed in earlier turns in the current phase) to the player's next choice. A player finishes playing once the resources he has chosen forms a strategy in the underlying RAG. In an outcome of the game, each player selects a set of resources. His cost depends on their load and latency functions as in usual RAGs.

We illustrate the intricacy of the selecting the resources in phases in the following example.

Example 1. Consider the 4-player network formation game that is depicted in Figure 1. The interesting edges have names, e.g., a, b, c, \dots , and their latency function is depicted below the edge. For example, we have $\ell_a(x) = x$ and $\ell_{c_1}(x) = 10x$. The other edges have latency function 0. The source and target of a node of Player i are depicted with a node called s and t , respectively, and with a subscript i . For example, Player 2's source is $s_{1,2}$ and he has two targets t_2^L and t_2^R . The players' strategies are paths from one of their sources to one of their targets.

Consider a dynamic version of the game in which Player i chooses an edge at turn i . At first look, it seems that edge g will never be chosen. However, we show that Player 1's optimal strategy uses it. Player 1 has three options in the first turn, either choose g , a , or b . Assume he chooses a (and dually b). Then, we claim that Player 2 will choose b . Note that Players 3 and 4 move opposite of Player 2 no matter how Player 1 moves, as they prefer avoiding a load of 2 on c_1 and c_2 , which costs 20 each, even at the cost of a load of 3 on f , which costs only 3. Knowing this, Player 2 prefers using b alone over sharing a with Player 1. Since the loads on a and e are 1 and 3, respectively, Player 1's cost is $1 + 3 = 4$.

On the other hand, if Player 1 chooses g in the first phase, he postpones revealing his choice between left and right. If Player 2 proceeds left, then Players 3 and 4 proceed right, and Player 1 proceeds left in the second phase. Now, the load on a and e is 2 and 1, respectively, thus Player 1's cost is $\frac{1}{2} + 2 + 1 = 3\frac{1}{2}$. \square




Fig. 1. A network formation game in which it is beneficial to select a path that is not simple.

The concept of what we refer to as a dynamic game is old and dates back to Von Neumann’s work on *extensive form games* [20]. Most work on RAGs consider the simultaneous setting. However, there have been different takes on adding dynamicity to RAGs. In [19], the authors refine the notion of NE by considering *lookahead* equilibria; a player predicts the reactions of the other players to his deviations, and he deviates only if the outcome is beneficial. The depth of lookahead is bounded and is a parameter to the equilibria. A similar setting was applied to RAGs in [7], where the players are restricted to choose a *best-response* move rather than a deviation that might not be immediately beneficial. Concurrent ongoing games are commonly used in formal methods to model the interaction between different components of a system (c.f., [1]). In such a game, multiple players move a token on a graph. At each node, each player selects a move, and the transition function determines the next position of token, given the vector of moves the players selected. The objectives of the players refer to the generated path and no costs are involved. Closest to our model is the model of [16], and its subsequent works [8,10]. They study RAGs in which players arrive and select strategies one by one, yet in one shot.

Our dynamic games differ from all of these games in two aspects. We allow the players to reveal their choice of resources in parts, thus we allow “breaking” the strategies into parts. Moreover, the choices the players make in all these games are either concurrent or sequential, and we allow a mix between the two. These two concepts are natural and general, and can be applied to other games and settings.

The first question that arises in the context of games, and on which we focus in this work, is the existence of a *stable outcome* of the game. In the context of RAGs, the most prominent stability concept is that of a *Nash equilibrium* (NE, for short) – a profile such that no player can decrease his cost by unilaterally deviating from his current strategy. It is well known that every RAG has an NE [23]. The definition of an NE applies to all games, and can also be applied to our dynamic RAGs. As we demonstrate in Example 2, the dynamic setting calls for a different stability concept, as some NEs need not be achievable by rational players in the dynamic setting. To overcome this limitation of NE, the notion of *subgame perfect equilibrium* (SPE, for short) was introduced in [27], which we define formally in Section 2.

Classifying RAGs, we refer to the type of their latency functions as well as the type of the objectives of the players. *Congestion games* [24] are RAGs in which the latency functions are increasing, whereas in *cost-sharing games* [2], each resource has a cost that is split between the players that use it (in particular, the latency functions are decreasing). In terms of objectives, we consider *singleton* RAGs, in which the objectives of the players are singletons of resources, and *symmetric* RAGs, in which all players have the same objective.

Our most interesting results are in terms of equilibrium existence. It is easy to show, and similar results are well known, that every dynamic RAG with a *sequential scheduler* has an SPE. The proof uses backwards induction on the tree of all possible outcomes of the game (see Theorem 1 for details). One could hope to achieve a similar proof also for schedulers that are not sequential, especially given the fact that every RAG has an NE. Quite surprisingly, however, we show that this is not the case. For congestion games, we show examples of a singleton congestion game and a symmetric congestion game with no SPE. Moreover, the latency function in both cases is linear. On the positive side, we show that singleton and symmetric congestion games are guaranteed to have an SPE for every scheduler. For cost-sharing games, we also show an example with no SPE. In the cost-sharing setting, however, we show that singleton objectives are sufficient to guarantee the existence of an SPE in all schedules. It follows that singleton dynamic congestion games are less stable than singleton dynamic cost-sharing games. This is interesting, as in the on-shot concurrent setting, congestion games are known to be more stable than cost-sharing games in various parameters. One would expect that this “order of stability” would carry over to the dynamic setting, as is the case in other extensions of the traditional setting. For

example, an NE is not guaranteed for *weighted* cost-sharing games [9] as well as very restrictive classes of *multiset* cost-sharing games [5], whereas every linear weighted congestion game [12] and even linear multiset congestion game is guaranteed to have an NE [6].

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. In simultaneous games, the standard measures to quantify the inefficiency incurred due to selfish behavior is the *price of anarchy* (PoA) [15] and *price of stability* (PoS) [2]. In both measures we compare against the *social optimum* (SO, for short), namely the cheapest profile. The PoA is the worst-case inefficiency of an NE (that is, the ratio between the cost of a worst NE and the SO). The PoS is the best-case inefficiency of a Nash equilibrium (that is, the ratio between the cost of a best NE and the social optimum). For the dynamic setting we adjust these two measures to consider SPEs rather than NEs, and we refer to them as DPoA and DPoS. We study the equilibrium inefficiency in the classes of games that have SPEs. We show that the DPoA and DPoS in dynamic singleton cost-sharing games as well as dynamic singleton congestion games coincide with the PoA and PoS in the corresponding simultaneous class. As mentioned above, [16,8,10] study games in which players arrive one after. Since their games are sequential, they always have an SPE. They study the *sequential PoA*, and show that it can either be equal, below, or above the PoA of the corresponding class of RAGs.

We then turn to study computational problems for dynamic RAGs. First, we study the problem of deciding whether a given dynamic RAG has an SPE. We show that the problem is PSPACE-complete for both congestion and cost-sharing games. Our lower bound for cost-sharing games implies that finding an SPE in sequential games is PSPACE-hard. To the best of our knowledge, while this problem was solved in [16] for congestion games, we are the first to solve it for cost-sharing games. We also study the problem of finding a schedule that admits an SPE under given constraints on the order the players move, and show that this problem is also PSPACE-complete. Finally, we consider dynamic games in which there is an order on the resources that the players choose. So, if for two resources e_1 and e_2 , we have $e_1 < e_2$, then a player cannot choose e_1 in a later phase than e_2 . The motivation for an order on resources is natural. For example, returning to network formation games, a driver can only extend the path he chooses as the choices are made during driving. We show that all our results carry over to the ordered case.

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

2 Preliminaries

Resource allocation games For $k \geq 1$, let $[k] = \{1, \dots, k\}$. A *resource-allocation game* (RAG, for short) is a tuple $G = \{[k], E, \{\Sigma_i\}_{i \in [k]}, \{\ell_e\}_{e \in E}\}$, where $[k]$ is a set of k players; E is a set of resources; for $i \in [k]$, the set $\Sigma_i \subseteq 2^E$ is a set of objectives³ for Player i ; and, for $e \in E$, we have that $\ell_e : \mathbb{N} \rightarrow \mathbb{R}$ is a latency function. The game proceeds in one-round in which the players select simultaneously one of their objectives. A *profile* $P = \langle \sigma_1, \dots, \sigma_k \rangle \in \Sigma_1 \times \dots \times \Sigma_k$ is a choice of an objective for each player. For $e \in E$, we denote by $nused(P, e)$ the number of times e is used in P , thus $nused(P, e) = |\{i \in [k] : e \in \sigma_i\}|$. For $i \in [k]$, the *cost* of Player i in P , denoted $cost_i(P)$, is $\sum_{e \in \sigma_i} \ell_e(nused(P, e))$.

Classes of RAGs are characterized by the type of latency functions and types of objectives. In *congestion games* (CGs, for short), the latency functions are increasing. An exceptionally stable class of CGs are ones in which the latency functions are linear (c.f., [12,6]); every resource $e \in E$ has two constants a_e and b_e , and the latency function is $\ell_e(x) = a_e \cdot x + b_e$. In *cost-sharing games*

³ We use “objectives” here rather than “strategies” as the second will later be used for dynamic games.

(SG, for short), each resource $e \in E$ has a *cost* c_e and the players that use the resource share its cost, thus the latency function for e is $\ell_e(x) = \frac{c_e}{x}$. In particular, the latency functions are decreasing. We use DCGs and DSGs to refer to dynamic CGs and dynamic SGs, respectively. In terms of objectives, we study *symmetric* games, where the players' sets of objectives are equal, thus $\Sigma_i = \Sigma_j$ for all $i, j \in [k]$, and *singleton* games, where each $\sigma \in \Sigma_i$ is a singleton, for every $i \in [k]$.

Dynamic resource allocation games A *dynamic* RAG is pair $\mathcal{G} = \langle G, \nu \rangle$, where G is a RAG and $\nu : [k] \rightarrow [k]$ is a *scheduler*. Intuitively, in a dynamic game, rather than revealing their objectives at once, the game proceeds in *phases*: in each phase, each player reveals one resource in his objective. Each phase is partitioned into at most k *turns*. The scheduler dictates the order in which the players proceed in a phase by assigning to each player his turn in the phases. If the scheduler assigns the same turn to several players, they select a resource concurrently. Once all players take their turn, a phase is concluded and a new phase begins. There are two “extreme” schedulers: (1) players get different turns, i.e., ν is a permutation, (2) all players move in one turn, i.e., $\nu \equiv 1$. We refer to games with these schedulers as *sequential* and *concurrent*, respectively. Note that ν might not be an onto function. For simplicity, we assume that, for $j > 1$, if turn j is assigned a player, then so is turn $j - 1$. We use t_ν to denote the last turn according to ν , thus $t_\nu = \max_i \nu(i)$.

Let $E_\perp = E \cup \{\perp\}$, where \perp is a special symbol that represents the fact that a player finished playing. Consider a turn $j \in [k]$. We denote by $\text{before}(j)$ the set of players that play before turn j ; thus $\text{before}(j) = \{i \in [k] : \nu(i) < j\}$. A player has full knowledge of the resources that have been chosen in previous phases and the resources chosen in previous turns in the current phase. A strategy for Player i in \mathcal{G} is a function $f_i : (E_\perp^{[k]})^* \cdot (E_\perp^{\text{before}(\nu(i))}) \rightarrow E_\perp$. A *profile* $P = \langle f_1, \dots, f_k \rangle$ is a choice of a strategy for each player. The *outcome* of the game given a profile P , denoted $\text{out}(P)$, is an infinite sequence of functions π^1, π^2, \dots , where for $i \geq 1$, we have $\pi^i : [k] \rightarrow E_\perp$. We define the sequence inductively as follows. Let $m \geq 1$ and $j \in [k]$. Assume $m - 1$ phases have been played as well as $j - 1$ turns in the m -th phase, thus $\pi^1, \pi^2, \dots, \pi^{m-1}$ are defined as well as $\pi_{j-1}^m : \text{before}(j) \rightarrow E_\perp$. We define π_j^m as follows. Consider a player i with $\nu(i) = j$. The resource Player i chooses in the m -th phase is $f_i(\pi^1, \dots, \pi^{m-1}, \pi_{j-1}^m)$. Finally, we define $\pi^m = \pi_{t_\nu}^m$.

We restrict attention to *legal* strategies for the players, namely ones in which the collection of resources chosen by Player i in all phases is an objective in Σ_i . Also, once Player i chooses \perp , then he has finished playing and all his choices in future phases must also be \perp . Formally, for a profile $P = \langle f_1, \dots, f_k \rangle$ with $\text{out}(P) = \pi^1, \pi^2, \dots$ and $i \in [k]$, let $\text{out}_i(P)$ be $\pi^1(i), \pi^2(i), \dots$. For $j \geq 1$, let $e_j = \pi^j(i)$ be the resource Player i selects in the j -th phase. Thus, $\text{out}_i(P)$ is an infinite sequence over E_\perp . We say that f_i is *legal* if (1) there is an index m such that $e_j \in E$ for all $j < m$ and $e_j = \perp$ for all $j \geq m$, and (2) the set $\{e_1, \dots, e_{m-1}\}$ is an objective in Σ_i . (In particular, a player cannot select a resource multiple times nor a resource that is not a member in his chosen objective). We refer to an outcome in which the players use legal strategies as a *legal outcome* and a prefix of a legal outcome as a *legal history*.

In $\text{out}(P)$, every player selects a set of resources. The cost of a player is calculated similarly to RAGs. That is, his cost for a resource e , assuming the load on it is γ , is $\ell_e(\gamma)$, and his total cost is the sum of costs of the resources he uses. When the outcome of a profile P in a dynamic RAG coincides with the outcome of a profile Q in a RAG G , we say that P and Q are *matching* profiles.

Equilibria concepts A Nash equilibrium⁴ (NE, for short) in a game is a profile in which no player has an incentive to unilaterally deviate from his strategy. Formally, for a profile P , let $P[i \leftarrow f'_i]$ be the profile in which Player i switches to the strategy f'_i and all other players use their strategies in P . Then, a profile P is a NE if for every $i \in [k]$ and every legal strategy f'_i for Player i , we have $\text{cost}_i(P) \leq \text{cost}_i(P[i \leftarrow f'_i])$. It is well known that every RAG is guaranteed to have an NE [23].

The definition of NE applies to all games, in particular to dynamic ones. Every NE Q in a RAG G matches an NE in a dynamic game $\langle G, \nu \rangle$, for some scheduler ν , in which the players ignore the history of the play and follow their objectives in Q . However, such a strategy is not rational. Thus, one could argue that an NE is not necessarily achievable in a dynamic setting. We illustrate this in the following example.

Example 2. Consider a two-player DCG with resources $\{a, b\}$, latency functions $\ell_a(x) = x$ and $\ell_b(x) = 1.5x$, and objectives $\Sigma_1 = \Sigma_2 = \{\{a\}, \{b\}\}$. Consider the sequential scheduling in which Player 1 moves first followed by Player 2. Since the players' objectives are singletons, the dynamic game consists of one phase. Consider the Player 2 strategy f_2 that “promises” to select the resource a no matter what Player 1 selects, thus $f_2(a) = f_2(b) = a$. Let f_1^a and f_1^b be the Player 1 strategies in which he selects a and b , respectively, thus $f_1^a(\epsilon) = a$ and $f_1^b(\epsilon) = b$. Note that these are all of Player 1's possible strategies. The profile $P = \langle f_1^b, f_2 \rangle$ is an NE. Indeed, Player 2 pays 1, which is the least possible payment, so he has no incentive to deviate. Also, by deviating to f_1^a , Player 1's payoff increases from 1.5 to 2, so he has no incentive to deviate either. Note, however, that this strategy of Player 2 is not rational. Indeed, when it is Player 2's turn, he is aware of Player 1's choice. If Player 1 plays f_1^a , then a rational Player 2 is not going to choose a , as this results in a cost of 2, whereas by b , his cost will be 1.5. Thus, an NE profile with f_2 may not be achievable. \square

To overcome this issue, the notion of *subgame perfect equilibrium* (SPE, for short) was introduced. In order to define SPE, we need to define a subgame of a dynamic game. Let $\mathcal{G} = \langle G, \nu \rangle$, where $G = \langle [k], E, \{\Sigma_i\}_{i \in [k]}, \{c_e\}_{e \in E} \rangle$. It is helpful to consider the *outcome tree* $\mathcal{T}_{\mathcal{G}}$ of \mathcal{G} , which is a finite rooted tree that contains all the legal histories of \mathcal{G} . Each internal node in $\mathcal{T}_{\mathcal{G}}$ corresponds to a legal history, its successors correspond to possible extensions of the history, and each leaf corresponds to a legal outcome. Consider a legal history h . We define a dynamic RAG \mathcal{G}_h , which, intuitively, is the game as \mathcal{G} after the history h has been played. More formally, the outcome tree of \mathcal{G}_h is the subtree $\mathcal{T}_{\mathcal{G}}^h$ whose root is the node h . We define the costs in \mathcal{G}_h so that the costs of the players in the leaves of $\mathcal{T}_{\mathcal{G}}^h$ are the same as the corresponding leaves in $\mathcal{T}_{\mathcal{G}}$. Assume that h ends at the m -th turn. A profile P in \mathcal{G} corresponds to a trimming of $\mathcal{T}_{\mathcal{G}}$ in which the internal node h has exactly one child $h \cdot \bar{\sigma}$, where $\bar{\sigma}$ is the set of choices of the players in $\nu^{-1}(m)$ when they play according to their strategies in P . The profile P induces a profile P^h in \mathcal{G}_h , where the trimming of $\mathcal{T}_{\mathcal{G}}^h$ according to P^h coincides with the trimming of \mathcal{G} according to P . We formally define the outcome tree and a subgame in Appendix A.

Definition 1. A profile P is an SPE if for every legal history h , the profile P^h is a NE in \mathcal{G}_h .

Note that the profile $P = \langle f_1^b, f_2 \rangle$ in the example above is an NE but not an SPE. Indeed, for the history $h = a$, the profile P^h is not a NE in \mathcal{G}_h as Player 2 can benefit from unilaterally deviating as described above.

⁴ Throughout this paper, we consider *pure* strategies as is the case in the vast literature on RAGs.

3 Existence of SPE in Dynamic Congestion Games

It is easy to show that every sequential dynamic game has an SPE by unwinding the outcome tree, and similar results have been shown before (c.f., [16]). The proof can be found in Appendix B.

Theorem 1. *Every sequential dynamic game has an SPE.*

One could hope to prove that a general dynamic game \mathcal{G} also has an SPE using a similar unwinding of $\overline{\mathcal{T}}_{\mathcal{G}}$. Possibly using the well-known fact that every CG is guaranteed to have an NE [23]. Unfortunately, and somewhat surprisingly, we show that this is not possible. We show that (very restrictive) DCGs might not have an SPE. For the good news, we identify a maximal fragment that is guaranteed to have an SPE.

Recall that a CG is singleton when the players' objectives consist of singletons of resources, and a CG is symmetric if all the players agree on their objectives. We start with the bad news and show that symmetric DCGs and singleton DCGs need not have an SPE, even with linear latency functions. We then show that the combination of these two restrictions is sufficient for existence of an SPE in a DCG.

Theorem 2. *There are symmetric and singleton linear DCGs with no SPE.*

Proof. We first describe a linear DCG with no SPE, and then alter it to make it symmetric. The proof for singleton linear DCG is given in Appendix D. Consider the following three-player linear CG G with resources $E = \{a, a', b, b', c\}$ and linear latency functions $\ell_a(x) = \ell_b(x) = x$, $\ell_{a'}(x) = \frac{3}{4}x$, $\ell_{b'}(x) = 1\frac{1}{4}$, and $\ell_c(x) = x + \frac{2}{3}$. Let $\Sigma_1 = \Sigma_2 = \{\{a, a'\}, \{b, b'\}, \{c\}\}$ and $\Sigma_3 = \{\{c\}, \{a', b\}\}$. Consider the dynamic game \mathcal{G} in which Players 1 and 2 move concurrently followed by Player 3. Formally, $\mathcal{G} = \langle G, \nu \rangle$, where $\nu(1) = \nu(2) = 1$ and $\nu(3) = 2$.

We claim that there is no SPE in \mathcal{G} . Note that since the players' objectives are disjoint, then once a player reveals the first choice of resource, he reveals the whole objective he chooses, thus we analyze the game as if it takes place in one phase in which the players' reveal their whole objective. The profiles in which Players 1 and 2 choose the same objective are clearly not a SPE as they are not an NE in the game \mathcal{G}_ϵ . As for the other profiles, in Figure 2, we go over half of them, and show that none of them is an SPE. The other half is dual. The root of each tree is labeled by the objectives of Players 1 and 2, and its branches according to Player 3's objectives. In the leaves we state Player 3's payoff. In an SPE, Player 3 performs a best-response according to the objectives he observes as otherwise the subgame is not in an NE. We depict his choice with a bold edge. Beneath each tree we note the payoffs of all the players in the profile, and the directed edges represent the player that can benefit from unilaterally deviating. In Appendix C, we construct a symmetric DCG \mathcal{G}' by altering the game \mathcal{G} above. We do this by adding a fourth player and three new resources so that \mathcal{G}' simulates \mathcal{G} . \square




Fig. 2. Profiles in a DCG with no SPE. Player 3's choices are depicted as bold edges.




Fig. 3. Profiles in a DSG with no SPE.

We now prove that combining the two restrictions does guarantee the existence of SPE. We note that while our negative results hold for linear DCGs, which tend to be stabler than other DCGs, our positive result holds for every increasing latency functions.

Theorem 3. *Every symmetric singleton DCG has an SPE.*

Proof. Consider a symmetric singleton DCG $\mathcal{G} = \langle G, \nu \rangle$. Recall that since G is a singleton game, every outcome of \mathcal{G} consists of one phase. Let P be an NE in G (recall that according to [23] an NE exists in every CG). Since G is symmetric, we can assume that, for $1 \leq j < k$, the players that move in the j -th turn do not pay more than the players that move after them. Formally, for $i, i' \in [k]$, if $\nu(i) < \nu(i')$, then $\text{cost}_i(P) \leq \text{cost}_{i'}(P)$. In particular, the players who move in the first turn pay the least, and the players that move in the last turn pay the most. We construct a profile Q in \mathcal{G} and show that it is an SPE. Intuitively, in Q , the players follow their objectives in P assuming the previous players also follow it. Since the costs are increasing with turns, if Player i deviates, a following Player j will prefer switching resources with Player i and also switching the costs. Thus, the deviation is not beneficial for Player i . In Appendix E, We construct Q formally and prove that it is an SPE. \square

4 Existence of SPE in Dynamic Cost-sharing Games

Cost sharing games tend to be less stable than congestion games in the concurrent setting; for example, very simple fragments of multiset cost-sharing games do not have an NE [5] while linear multiset congestion games are guaranteed to have an NE [6]. In this section we are going to show that, surprisingly, there are classes of games in which an SPE exists only in the cost-sharing setting. Still, SPE is not guaranteed to exist in general DSGs. We start with the bad news.

Theorem 4. *There is a DSG with no SPE.*

Proof. Consider the following four-player SG G with resources $E = \{a, a', a'', b, b', b'', c, c', c''\}$ and costs $c_a = c_b = c_c = 6$, $c_{a'} = c_{b'} = c_{c'} = 4$, and $c_{a''} = c_{b''} = c_{c''} = 3$. Let $\Sigma_1 = \{\{a, a'\}, \{b, b''\}\}$, $\Sigma_2 = \{\{b, b'\}, \{c, c''\}\}$, $\Sigma_3 = \{\{c, c'\}, \{a, a''\}\}$, and $\Sigma_4 = \{\{a, a'\}, \{b, b'\}, \{c, c'\}\}$. Consider the dynamic game \mathcal{G} in which players 1, 2, and 3 move concurrently followed by Player 4. Formally, $\mathcal{G} = \langle G, \nu \rangle$, where $\nu(1) = \nu(2) = \nu(3) = 1$ and $\nu(4) = 2$.

We claim that there is no SPE in \mathcal{G} . Similar to Theorem 2, since the players' objectives are disjoint, we analyze the game as if it takes place in one phase. In Figure 3, we depict some of the profiles and show that none of them are an SPE. As in Theorem 2, the root of each tree is labeled by the objectives of Players 1, 2, and 3, its branches according to Player 4's choices, and in the leaves we state the cost of Player 4 assuming he chooses his best choice given the other players' choices. There are two types of deviations in the cycle: a player pays the full cost of 10 for his more expensive action, deviates to the cheaper one, which he does not share either, for a cost of 9. And, a player who pays 5 deviates to his other action and reduces his cost to 4. Finally, it is not hard to show that every profile not on the cycle of profiles cannot be an SPE. \square

Recall that singleton DCGs are not guaranteed to have an SPE (Theorem 2). On the other hand, we show below that singleton DSGs are guaranteed to have an SPE. In order to find an SPE in such a game, we use a firmer notion of an equilibria in SGs.

A *strong equilibria* (SE, for short) [3] which is stable against deviations of *coalitions* of players rather than deviations of a single player as in NEs. Formally, consider a singleton SG $G = \langle [k], E, \{\Sigma_i\}_{i \in [k]}, \{c_e\}_{e \in E} \rangle$, a profile $P = \langle \sigma_1, \dots, \sigma_k \rangle$, a coalition of players $C \subseteq [k]$,

and a joint move $S \in \bigcup_{i \in C} \Sigma_i^{\{i\}}$ for the members of the coalition. We denote by $P[C \leftarrow S] = \langle \sigma'_1, \dots, \sigma'_k \rangle$ the profile in which the players in C switch to their objective in S , thus $\sigma'_i = S(i)$ for every $i \in C$, and $\sigma'_i = \sigma_i$ for $i \notin C$. We say that S is beneficial for C if it is beneficial for all the members of the coalition, thus for every $i \in C$, we have $\text{cost}_i(P[C \leftarrow S]) < \text{cost}_i(P)$. We say that P is an SE if there is no coalition that has a beneficial move.

We show a connection between strong equilibria and SPEs in singleton SGs. It is shown in [13] that every singleton SG has an SE.

Theorem 5. *Consider a singleton DSG $\mathcal{G} = \langle G, \nu \rangle$. Then, every strong equilibrium in G matches an SPE of \mathcal{G} . In particular, every singleton DSG has an SPE.*

Proof. We describe the intuition of the proof and the details can be found in Appendix F. Consider a singleton DSG $\mathcal{G} = \langle G, \nu \rangle$, and let Q be an SE in G . We describe a profile P in \mathcal{G} that matches Q , and we claim that it is an SPE. Consider a history h that ends in the i -th turn. Assume the players that play in h follow their objective in Q . Then, the players who play next, namely these in $\nu^{-1}(i+1)$, also follow Q . Thus, P matches Q . The definition of the strategies in P for histories that do not follow Q is inductive: assume only the players in $\nu^{-1}(i)$ choose differently than in Q , then the subgame \mathcal{G}_h is a singleton DSG. We find a strong equilibrium in \mathcal{G}_h and let the players in $\nu^{-1}(i+1)$ choose according to it.

We claim that P is an SPE. Assume towards contradiction that there is a Player i who can benefit from a unilateral deviation to a resource e . Such a deviation initiates subsequent deviations from players who choose in later turns than Player i . We consider the outcome of the game, and the players that play differently than in Q . Let $I, D \subseteq [k]$ be the set of players whose cost increases and decreases, respectively, with respect to their cost in Q . We make several observations. Since Q is an SE, we have $I \neq \emptyset$. Moreover, there must be a player $j_I \in I$ who deviates to e . Assume Player j_I chooses e' in Q . Since \mathcal{G} is a SG, there must be a player $j_D \in D$ who chooses e' in Q , but deviates following Player i 's deviation. We continue recursively and identify a sequence of resources $e = e_1, e_2, \dots$ such that for every $j \geq 1$, there are $j_I \in I$ and $j_D \in D$ such that both players deviate into e_j . Moreover, Players j_I and $(j+1)_D$ use the same objective in Q . A contradiction follows from the fact that since there are finitely many resources, the sequence has a loop. \square

Remark 1. One could suspect that existence of strong equilibria in the underlying RAG implies existence of an SPE in the dynamic game. However, [13] shows that singleton CGs are guaranteed to have an SE, while we show in Theorem 2 that singleton DCGs are not guaranteed to have an SPE. In fact, [13] shows that every NE in a singleton CG is also an SE, while it is shown in [25] that this is not the case for singleton SGs.

One could also suspect that Theorem 5 generalizes to richer types of objectives. That is, we can ask whether, for an DSG $\mathcal{G} = \langle G, \nu \rangle$, an SE in G matches an SPE in \mathcal{G} . Theorem 4 shows that this is not the case as in the SG there, the profile $\langle \{b, b''\}, \{b, b'\}, \{a, a''\}, \{b, b'\} \rangle$ is an SE.

Remark 2. Consider a symmetric DSG $\mathcal{G} = \langle G, \nu \rangle$. The social optimum profile O in G is attained when all the players choose the same cheapest objective, namely the objective with the minimal sum of resource costs. It is not hard to see that O is an NE as a deviation results in a more expensive objective with less sharing. Recall that we study SPE in dynamic games as NE might contain strategies that will not be used by rational players. Consider a profile P in \mathcal{G} that matches O (note that there can be many such profiles, and some can consist of strategies that are chosen by rational players). The same arguments stated above imply that P is an NE. Nevertheless, P may not be an SPE, as \mathcal{G} might contain a subgame with no SPE.

5 Equilibrium Inefficiency

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 define the cost of a profile P , denoted $cost(P)$, to be $\sum_{i \in [k]} cost_i(P)$. We denote by OPT the cost of a social-optimal solution; i.e., $OPT = \min_P cost(P)$. Two standard measures that quantify the inefficiency incurred due to self-interested behavior are the *price of anarchy* (PoA) [15,21] and *price of stability* (PoS) [2,26]. The PoA is the worst-case inefficiency of an NE; The PoA of a game G is the ratio between the cost of the most expensive NE and the cost of the social optimum. The PoS measures the best-case inefficiency of an NE, and is defined similarly with the cheapest NE. The PoA of a family of games \mathcal{F} is $\sup_{G \in \mathcal{F}} PoA(G)$, and the definition is similar for PoS.

In dynamic games we consider SPE rather than NE. We adapt the definitions above accordingly, and we refer to the new measures as *dynamic PoA* and *dynamic PoS* (DPoA and DPoS, for short). We study the equilibrium inefficiency in the classes of games that are guaranteed to have an SPE, namely singleton DSGs and symmetric singleton DCGs.

The lower bounds for the PoA and PoS for singleton SG and singleton symmetric CGs follow to the dynamic setting as we can consider the scheduler in which all players choose simultaneously in the first turn. The upper bounds on the DPoS for singleton symmetric DCGs follow from the fact that every NE in singleton symmetric CGs matches an SPE in the corresponding dynamic game. For singleton DSG, recall that an SE in the traditional game matches an SPE in the dynamic game. It is shown in [29] that singleton SGs have an SE whose cost is at most $\log(k) \cdot OPT$, which coincides with the $\log(k)$ lower bound. Finally, the upper bound on the DPoA for singleton DSGs follows from the same argument as traditional games. For congestion games, it follows by applying a recent result by [10] to our setting. The details can be found in Appendix G.

Theorem 6. *The DPoA and DPoS in singleton DSGs and singleton symmetric DCGs coincide with the PoA and PoS in singleton SGs and singleton symmetric CGs, respectively.*

6 Deciding the Existence of SPE

In the previous sections we showed that dynamic RAGs are not guaranteed to have an SPE. A natural decision problem arises, which we refer to as \exists SPE: given a dynamic RAG, decide whether it has an SPE. We show that the problem is PSPACE-complete in DSGs as well as DCGs. We start with the lower bound. The crux of the proof is given in the following lemma. For DCGs, such a construction is described in [16], which uses a construction by [28] in order to simulate the logic of a NAND gate by means of a CG. For SGs we are not aware of a similar known result. We describe the construction in Appendix H, which is inspired by the construction in [16].

Lemma 1. *Given a QBF instance ψ , there is a fully sequential game \mathcal{G}_ψ that is either a DCG or a DSG, and two constants $\gamma, \delta > 0$, such that in every SPE P in \mathcal{G}_ψ , (1) if ψ is true, then $cost_1(P) < \gamma$, and (2) if ψ is false, then $cost_1(P) > \delta$.*

To conclude the lower-bound proof, we combine between the game that is constructed in Lemma 1 and a game that has no SPE as in the examples we show in the previous sections. For the upper bound, consider a dynamic RAG \mathcal{G} , and let $\mathcal{T}_\mathcal{G}$ be the outcome tree of \mathcal{G} . Recall that there is a one-to-one correspondence between leaves in $\mathcal{T}_\mathcal{G}$ and legal outcomes of \mathcal{G} . In order to decide in PSPACE whether \mathcal{G} has an SPE, we guess a leaf l in $\mathcal{T}_\mathcal{G}$ and verify that it is an outcome of an SPE. Thus, we ask if there is an SPE P in \mathcal{G} whose outcome corresponds to l .

Theorem 7. *The \exists SPE problem is PSPACE-complete for dynamic RAGs.*

7 Efficient Stable Scheduling

The underlying assumption in game theory is that the players are selfish yet an authority may construct components of the game, and its challenge is to do so in a way that leads to stability. In some settings, the RAG is fixed and the authority only has the power to schedule the players. We assume that we are given a set S of constraints on schedulers that the authority can impose. A constraint $s \in S$ is either a *sequential constraint*, of the form $i_1 < i_2$, for $i_1, i_2 \in [k]$, stating that Player i_1 moves before Player i_2 , or a *concurrent constraint*, of the form $i_1 = i_2$, stating that the Players i_1 and i_2 move concurrently. Scheduling has a price. The input also contains a value function $c : S \rightarrow \mathbb{Q}$ that assigns to each constraint a cost or a reward. Intuitively, when $c(s) = \gamma \geq 0$, it means that the authority pays γ in order to force s in ν . Then, when $c(s) = \gamma \leq 0$, it means that the authority is rewarded γ if ν respects s . Consider a scheduler ν . Let $R \subseteq S$ be the set of constraints satisfied by ν . Then, $cost_{S,c}(\nu) = \sum_{s \in R} c(s)$.

The *budgeted scheduling problem* (BS problem, for short) gets as input a RAG G , a set of constraints S , a value function $c : S \rightarrow \mathbb{Q}$, and a budget β . The goal is to decide whether there is a scheduler ν with $cost_{S,c}(\nu) \leq \beta$ such that the dynamic game $\langle G, \nu \rangle$ has an SPE. The proof of the following theorem can be found in Appendix J.

Theorem 8. *The BS problem is PSPACE-complete.*

8 Games with Ordered Resources

In many settings it makes sense to restrict the order in which the players select their objectives. For example, recall that a network formation game is played on a network, and each player chooses a path that connects his source and target vertices. When decisions are taken while the path is being generated, it is often the case that a player cannot select the edges on his path in any order. Rather, he must extend his path one edge at a time. The corresponding constraints state that if a player uses an edge $\langle u, v \rangle$, then, unless u is the source vertex, in a previous phase an edge that ends in u must have been chosen.

Generally speaking, we assume a dynamic game is also given by a partial order $<$ on the resources. A legal strategy is one that does not violate the order. Thus, if the sequence of choices for a player in some outcome is e_1, e_2, \dots, e_m , then there are no $1 \leq i < j \leq m$ such that $e_j < e_i$. We restrict the players to choose only legal strategies and we assume there is at least one legal strategy for each player.

Theorem 9. *Our results in terms of SPE existence, equilibrium inefficiency, and computational complexity coincide for ordered dynamic games and dynamic games.*

Proof. Note that ordered dynamic games generalize dynamic games as we allow using the empty partial order. So, all our bad news follow to this setting. In terms of equilibrium existence and inefficiency, our good news are for singleton games. Such games cannot be ordered. As for computational complexity, our upper bounds can easily be adapted to ordered games. \square

References

1. R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.
2. 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.

3. R. Aumann. Acceptable points in games of perfect information. *Contributions to the Theory of Games*, 4:287–324, 1959.
4. G. Avni and O. Kupferman. Synthesis from component libraries with costs. In *Proc. 25th CONCUR*, pp. 156–172, 2014.
5. G. Avni, O. Kupferman, and T. Tamir. Network-formation games with regular objectives. In *Proc. 17th FoSSaCS*, 8412 LNCS, pp. 119–133. Springer, 2014.
6. G. Avni, O. Kupferman, and T. Tamir. Congestion games with multisets of resources and applications in synthesis. In *Proc. 35th FSTTCS*, pp. 365–379, 2015.
7. V. Bilò, A. Fanelli, and L. Moscardelli. On lookahead equilibria in congestion games. In *Proc. 9th WINE*, pp. 54–67, 2013.
8. J. R. Correa, J. de Jong, B. de Keijzer, and M. Uetz. The curse of sequentiality in routing games. In *Proc. 11th WINE*, pp. 258–271, 2015.
9. H. Chen and T. Roughgarden. Network design with weighted players. *Theory Comput. Syst.*, 45(2):302–324, 2009.
10. J. de Jong and M. Uetz. The sequential price of anarchy for atomic congestion games. In *Proc. 10th WINE*, pp. 429–434, 2014.
11. D. Fotakis. Stackelberg strategies for atomic congestion games. In *Proc. 15th ESA*, pp. 299–310, 2007.
12. T. Harks and M. Klimm. On the existence of pure Nash equilibria in weighted congestion games. *Math. Oper. Res.*, 37(3):419–436, 2012.
13. R. Holzman and N. Law-Yone. Strong equilibrium in congestion games. *Games and Economic Behavior*, 21(1-2):85–101, 1997.
14. R. Koch and M. Skutella. Nash equilibria and the price of anarchy for flows over time. *Theory Comput. Syst.*, 49(1):71–97, 2011.
15. E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *CS Review*, 3(2):65–69, 2009.
16. R. P. Leme, V. Syrgkanis, and E. Tardos. The curse of simultaneity. In *Proc. 3rd ITCS*, pp. 60–67, 2012.
17. Y. Lustig and M.Y. Vardi. Synthesis from component libraries. *STTT*, 15(5-6):603–618, 2013.
18. I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111 – 124, 1996.
19. V. S. Mirrokni, N. Thain, and A. Vetta. A theoretical examination of practical game playing: Lookahead search. In *Proc. 5th SAGT*, pp. 251–262, 2012.
20. J. Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320.
21. C. H. Papadimitriou. Algorithms, games, and the internet. In *Proc. 33rd STOC*, pp. 749–753, 2001.
22. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proc. 16th POPL*, pp. 179–190, 1989.
23. R.W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
24. T. Roughgarden and E. Tardos. How bad is selfish routing? *JACM*, 49(2):236–259, 2002.
25. O. Rozenfeld and M. Tennenholtz. Strong and correlated strong equilibria in monotone congestion games. In *Proc. 2nd WINE*, pp. 74–86, 2006.
26. A. S. Schulz and N. E. Stier Moses. On the performance of user equilibria in traffic networks. In *Proc. 14th SODA*, pp. 86–87, 2003.
27. R. Selten. Spieltheoretische behandlung eines oligopolmodells mit nachfragerträgeit. *Zeitschrift für die gesamte Staatswissenschaft*, 121, 1965.
28. A. Skopalik and B. Vöcking. Inapproximability of pure nash equilibria. In *Proc. 40th STOC*, pp. 355–364, 2008.
29. V. Syrgkanis. The complexity of equilibria in cost sharing games. In *Proc. 6th WINE*, pp. 366–377, 2010.

A Full Definitions: Outcome Trees and Subgames

We denote by $\mathcal{T}_{\mathcal{G}}$, the outcome tree of a dynamic RAG $\mathcal{G} = \langle G, \nu \rangle$. The root of $\mathcal{T}_{\mathcal{G}}$ corresponds to the empty prefix ϵ . Recall that t_{ν} is the maximal turn in ν . Let $m \geq 1$, and $j \in [t_{\nu}]$, and consider a node that corresponds to a prefix h of a legal outcome, after $m - 1$ phases have been played as well as $j - 1$ turns in the m -th phase. Thus, $h = \pi^1, \dots, \pi^{m-1}, \pi_{j-1}^m$, where $\pi_r^l : [k] \rightarrow E_{\perp}$ and $\pi_{j-1}^m : \text{before}(j) \rightarrow E_{\perp}$. We say that h is *controlled* by the players in $\nu^{-1}(j)$. Note that if ν is sequential, then each node is controlled by exactly one player. The children of the node h contain all the possible extensions of h with a legal *joint move* by the players in $\nu^{-1}(j)$. Thus, a child h' of h is a prefix of a legal outcome and it is of the form $h' = \pi^1, \dots, \pi^{m-1}, \pi_j^m$, where $\pi_j^m : \text{before}(j + 1) \rightarrow E_{\perp}$, and π_j^m and π_{j-1}^m agree on $\text{before}(j)$. The edge $\langle h, h' \rangle$ in $\mathcal{T}_{\mathcal{G}}$ corresponds to the joint move $\bar{\sigma} : \nu^{-1}(j) \rightarrow E_{\perp}$, where for every $i \in \nu^{-1}(j)$ we have $\bar{\sigma}(i) = \pi_j^m(i)$. We sometimes use $h \cdot \bar{\sigma}$ to refer to h' . Note that if $j = t_{\nu}$, then h' is of the form $\pi^1, \dots, \pi^m, \pi_1^{m+1}$. Finally, h is a leaf if $j = t_{\nu}$ and all players have finished playing in the m -th phase, thus the choices of all the players in the next phases must be \perp . Clearly, $\mathcal{T}_{\mathcal{G}}$ is a finite tree.

Consider a profile $P = \langle f_1, \dots, f_k \rangle$. It is possible to trim $\mathcal{T}_{\mathcal{G}}$ according to P so that each internal node h has exactly one child $h \cdot \bar{\sigma}$, where $\bar{\sigma}$ is the joint objective in which the players who control h follow P . That is, for each Player i who controls h we have $\bar{\sigma}(i) = f_i(h)$. Note that by trimming $\mathcal{T}_{\mathcal{G}}$ according to P , we leave exactly one leaf l that is reachable from the root. Note that both $\text{out}(P)$ as well as the leaf l correspond to the same profile in the underlying RAG G . For every $i \in [k]$, the cost of Player i in l is $\text{cost}_i(P)$.

We proceed to define a subgame. Let $h = \pi^1, \dots, \pi^{m-1}, \pi_{j-1}^m$ be a legal history. Note that $m - 1$ phases have been played as well as $j - 1$ turns in the m -th phase, and the players that should play next are $\nu^{-1}(j)$. The subgame $\mathcal{G}_h = \langle G_h, \nu_h \rangle$ of \mathcal{G} is a dynamic RAG with $G_h = \langle [k], E, \{\Sigma_i^h\}_{i \in [k]}, \{f_e\}_{e \in E} \rangle$, where the sets Σ_i^h of objectives are defined as follows. Consider $i \in [k]$. Let e_1, \dots, e_m be the choices of Player i in h over E_{\perp} . Let $m' \leq m$ be the last index such that $e_{m'} \in E$. Let $\sigma_i^h = \{e_1, \dots, e_{m'}\}$ be the edges in E that Player i collects during h , and let $\text{nused}(h, e)$ be the load generated on e in h by all the players, thus $\text{nused}(h, e) = |\{i \in [k] : e \in \sigma_i^h\}|$. The set of objectives of Player i in G_h is Σ_i^h . Each objective in Σ_i^h corresponds to an objective in Σ_i (that is, Player i 's objectives in G), minus the resources that have been collected in h , thus $\Sigma_i^h = \{\sigma \setminus \sigma_i^h : \sigma \in \Sigma_i\}$. The cost Player i pays in a profile $P = \langle \sigma_1, \dots, \sigma_k \rangle$ takes into account also the use of resources in the history h , thus $\text{cost}_i(P) = \sum_{e \in (\sigma_i(e) \cup \sigma_i^h(e))} \ell_e(\text{nused}(h, e) + \text{nused}(P, e))$. Recall that $j - 1$ turns have been played in the last phase in h . Thus, we “shift” the scheduler ν_h by j . For $l \in [k]$, the players $\nu^{-1}(l)$ who are scheduled to play in the l -th turn by ν , are scheduled to play in the $((l - j) \bmod k) + 1$ turn by ν_h . In particular, for every $i \in \nu^{-1}(j)$, we have $\nu_h(i) = 1$. Finally, consider a strategy f in \mathcal{G} . We define a strategy f^h in the game \mathcal{G}_h by $f^h(x) = f(h \cdot x)$. Given a profile $P = \langle f_1, \dots, f_k \rangle$, the corresponding profile in \mathcal{G}_h is $P^h = \langle f_1^h, \dots, f_k^h \rangle$.

B Proof of Theorem 1

Given a sequential game \mathcal{G} , we construct an SPE by “unwinding” the outcome tree $\mathcal{T}_{\mathcal{G}}$ in a backwards inductive manner. Consider an internal node h in $\mathcal{T}_{\mathcal{G}}$. Since \mathcal{G} is sequential, there is a unique Player i who controls h . If h is a leaf, then \perp is the only choice Player i can make, and we set his strategy to select \perp . Assume h is an internal node, and let e_1, \dots, e_m be the possible resources Player i can select in h . Let h_1, \dots, h_m be the children of the node h in $\mathcal{T}_{\mathcal{G}}$. Assume by induction that there is an SPE in the games $\mathcal{G}_{h_1}, \dots, \mathcal{G}_{h_m}$, and let $\gamma_1, \dots, \gamma_m$ be Player i 's

costs in these SPEs. We set Player i 's choice in h to be a resource e_i that minimizes his payment, thus $\gamma_i = \min\{\gamma_1, \dots, \gamma_m\}$. Clearly, this profile is an NE in the subgame \mathcal{G}_h . We refer to this choice of Player i as a *best response* to the history h that he observes. Note that every choice of Player i that achieves a higher cost is not an NE. In particular, if at every node h there is a unique best response for the player that controls h , then there is a unique SPE in \mathcal{G} .

C Proof of Theorem 2

To conclude the proof, we construct a symmetric CG G' by altering the game G above. We add a fourth player and three new resources d , e , and f , with latency functions $\ell_d(x) = 10x$, $\ell_e(x) = 25x$, and $\ell_f(x) = 30$. The other resources are as in G . The objectives of the players' are symmetric. They consist of $\Sigma_1 \cup \Sigma_2$, where we add d to every strategy in Σ_1 and e to every strategy in Σ_2 . We also add a singleton objective $\{f\}$. Formally, the objectives are $\{\{f\}\} \cup \{\sigma \cup \{d\} : \sigma \in \Sigma_1\} \cup \{\sigma \cup \{e\} : \sigma \in \Sigma_2\}$. Finally, we define a scheduler ν' that is similar to ν only that Player 4 moves last, thus $\nu'(1) = \nu'(2) = 1$, $\nu'(3) = 2$, and $\nu'(4) = 3$.

We claim that $\mathcal{G}' = \langle G', \nu' \rangle$ has no SPE. We claim that in every profile that is a candidate to be an SPE, the choice of Players 1 and 2 in the first phase is d , the choice of Player 3 is e , and the choice of Player 4 is f . This follows from the following three properties: (1) the costs of these three resources is much higher than that of the other resources, so the players' best response in the first phase is to minimize the cost they pay for them, (2) using d at most twice is more beneficial than using e , and using d three times is more costly than using e once, and (3) using e once is more beneficial than using f once, and using f once is more beneficial than using e twice. Once the first phase is played, the analysis is the same as in the game \mathcal{G} , which we proved not to have an SPE.

D An example of a singleton linear DCG with no SPE

Consider the four-player linear singleton CG G with resources $E = \{a, b, c, d\}$ and linear latency functions $\ell_a(x) = 4\frac{1}{2} \cdot x$, $\ell_b(x) = 2\frac{1}{2} \cdot x$, $\ell_c(x) = 3x$, and $\ell_d(x) = 4x$. Let $\Sigma_1 = \{\{a\}, \{c\}\}$, $\Sigma_2 = \{\{b\}, \{d\}\}$, $\Sigma_3 = \{\{b\}, \{c\}\}$, and $\Sigma_4 = \{\{c\}, \{d\}\}$. Consider the dynamic game \mathcal{G} in which Players 1 and 2 move concurrently, then Player 3, and finally Player 4. Formally, $\mathcal{G} = \langle G, \nu \rangle$, where $\nu(1) = \nu(2) = 1$, $\nu(3) = 2$, and $\nu(4) = 3$.

We go over all the profiles in \mathcal{G} and show that none of them is an SPE. The profiles are depicted in Figure 4. Similar to Theorem 2, the root of each tree is labeled by the objective of Players 1 and 2, its branches according to Players 3 and 4's objectives, and in the leaves we state the payoffs of Players 3 and 4 assuming they choose their best choice given the other players' choices.

E Proof of Theorem 3

We construct a profile $Q = \langle f_1, \dots, f_k \rangle$ in \mathcal{G} and show that it is an SPE. Consider a history h . For a resource $e \in E$, we define $nused(h, e)$ and $nused(P, e)$ to be the loads on e in h and P , respectively. We say that h is *consistent with P* if for every $e \in E$, we have $nused(h, e) \leq nused(P, e)$. When P is clear from the context we do not state it implicitly.

We first define the strategies in Q w.r.t. consistent histories. Consider such a history h that ends before the j -th turn, thus the players who control the node h in $\mathcal{T}_{\mathcal{G}}$ are $\nu^{-1}(j)$. We define the strategies in Q as if the players in $\nu^{-1}(j)$ move in a sequential order. Let i_1, i_2, \dots, i_n be




Fig. 4. The profiles of the singleton DCG with no SPE. Bold edges depict Players 3 and 4’s best choices given the other players choices. Directed edges represent the player that can benefit from unilaterally deviating.

an arbitrary order on the players in $\nu^{-1}(j)$. Then, for $1 \leq l \leq n$, we define the strategy f_{i_l} of Player i_l to perform a best response to the objectives of players in $\text{before}(j)$, whose objectives he observes, and as if he also observes the objectives of the players i_1, \dots, i_{l-1} who also move in the j -th turn.

We formally define f_{i_l} . Let h' be the history h concatenated with the objectives of the players i_1, \dots, i_{l-1} . We say that a resource $e \in E$ is *full* in h' if $nused(h', e) = nused(P, e)$. Recall that f_{i_l} is defined w.r.t. h . We define $f_{i_l}(h)$ to be a resource e that is not full after h' and, assuming all resources will eventually be filled up, choosing e will cost the least for Player i_l , thus e minimizes $\{\ell_e(nused(P, e)) : nused(h', e) < nused(P, e)\}$. Note that since players never choose a resource that is full, the history h concatenated with the choices in the j -th turn, is a consistent history.

We have not yet defined the strategies in Q w.r.t. histories that are inconsistent with P . Still, we can show that for every history h that is consistent with P , the profile Q^h is an NE in \mathcal{G}_h . Assume that h ends before the j -th turn. We first show that for every Player $i \in \nu^{-1}(j)$, choosing a resource that is not full in h *dominates* choosing a resource that is already full. That is, no matter how the other players move, it is always more beneficial to choose a resource that is not yet full over one that is full. Then, all that is left in order to prove that Q^h is an NE, is to show that no player can benefit from deviating to a resource that is not full. Such a deviation results in a history that is consistent with P for which we have defined the strategies in Q . Intuitively, such a deviation is not beneficial as we defined Q so that players that move first pay less. By deviating, a player will “switch” costs with a player that moves after him, thus his cost cannot decrease.

We make an observation before proving the claim. Assume the players $\nu^{-1}(j)$ select a joint objective $\bar{\sigma}$ such that $h \cdot \bar{\sigma}$ is a history consistent with P . It is not hard to see that $\text{out}(Q^{h \cdot \bar{\sigma}})$ is also consistent with P . Thus, for each Player $i \in \nu^{-1}(j)$, we have $\text{cost}_i(Q^h) = \text{cost}_{i'}(P)$, for some $i' \in [k]$ (possibly $i' = i$). Note that Q^h is a profile in the game \mathcal{G}_h whereas P is a profile in the game G .

We proceed to prove that for every $i \in \nu^{-1}(j)$, Player i cannot benefit from unilaterally deviating from the profile Q^h in the game \mathcal{G}_h . Assume towards contradiction that Player i can benefit from unilaterally deviating and choosing a resource $e \in E$. Recall that since \mathcal{G} is a singleton game, Player i does not move again. Let $\bar{\sigma}$ be the joint move at h according to Q^h , and let $h' = h \cdot \bar{\sigma}$. We distinguish between two cases. First, assume the resource e is full in the history h' . Thus, $nused(h', e) = nused(P, e)$. Note that Player i ’s deviation forms a history that is not consistent with P and we have not yet defined the strategies in Q w.r.t such histories.

Still, we can show that such a deviation is not beneficial. Note that the load on e , no matter how the other players move is at least $nused(P, e) + 1$. Recall that $cost_i(Q^h) = cost_{i'}(P)$, for some $i' \in [k]$. Since G is a symmetric game, Player i' can choose the objective e , thus $e \in \Sigma_{i'}$. Since P is an NE, we have $cost_{i'}(P) \leq cost_{i'}(P[i' \leftarrow e])$. Note that the load on e in the profile $P[i' \leftarrow e]$ is exactly $nused(P, e) + 1$. Since G is a congestion game, the cost of e increases with the load on it. Thus, the cost Player i pays after deviating is at least $cost_{i'}(P[i' \leftarrow e])$, which is not beneficial, and we reach a contradiction.

In the second case, the resource e to which Player i deviates is not full in the history h' . Let h'' be the history after Player i 's deviation. Then, h'' is a history consistent with P , and Q is defined w.r.t h'' . Let P' be the profile in G that corresponds to the outcome of Q in the subgame $G_{h''}$. Let $l \in [k]$ be the player that selects e in P' . Thus, $cost_i(P') = cost_l(Q)$. Recall that according to Q , at h , Player i should select the cheapest resource that is not full. Since Player i deviates, the resource e is not that cheapest resource. So, Player l moves after Player i , where by “after” we either mean that $\nu(i) < \nu(l)$ or $\nu(i) = \nu(l)$ but i comes before l in the arbitrary order we fix for the players in $\nu(i)$ -th turn. Thus, $cost_l(Q) \geq cost_i(Q)$, and the deviation is not beneficial for Player i .

To conclude the proof, we complete the definition of the strategies in Q . The definition is inductive. Let h' be a history consistent with P , and let $\bar{\sigma}$ be a joint move such that the history $h = h' \cdot \bar{\sigma}$ is inconsistent. Since G is a singleton symmetric game, so is the game G_h . Thus, we find an NE profile P' in G_h and continue as in the above.

F Proof of Theorem 5

Let $P = \langle f_1, \dots, f_k \rangle$ be the profile that is described in the body of the paper whose outcome coincides with the SE Q . We claim that P is an SPE. Consider a history h that ends before the j -th turn, where assume that the players in h follow their choices in Q . For the other histories the proof is similar. Assume towards contradiction that there is $i \in \nu^{-1}(j)$ such that Player i can benefit from unilaterally deviating to a strategy f'_i . We think of the outcome of $P[i \leftarrow f'_i]$ as a profile $Q' = \langle \sigma'_1, \dots, \sigma'_k \rangle$ in G . Let $C \subseteq [k]$ be the players whose objectives in Q' differ from their objectives in Q . Let $I, D \subseteq C$ be the partition of C to players whose payoff in Q' increase and decrease, respectively, with respect to their outcome in Q . Formally, if $i \in I$, then $cost_i(Q) \geq cost_i(Q')$, and if $i \in D$, then $cost_i(Q) < cost_i(Q')$. Since Q is an SE, the coalition C cannot all benefit, thus $I \neq \emptyset$.

Consider a resource σ such that there is a player $j \in D$ with $\sigma'_j = \sigma$. We make three observations. (1) there is $j_1 \in I$ such that $\sigma'_{j_1} = \sigma$. Otherwise, in the game G , the coalition of players in C that play σ in Q' can benefit from deviating from Q , contradicting the fact that it is an SE.

(2) There is $j_2 \in D$ such that Players j_1 and j_2 choose the same objective in Q , thus $\sigma_{j_1} = \sigma_{j_2}$. Assume otherwise. Let Player j be the first player that “leaves” σ_{j_2} , thus Player j chooses σ_{j_2} in Q and not in Q' and this is the first such player to choose. Note that if $j \in I$, then by staying in σ_{j_2} , his cost cannot increase, thus the deviation is not beneficial, and we reach a contradiction.

(3) Note that $\sigma'_{j_1} \neq \sigma'_{j_2}$. Indeed, players j_1 and j_2 pay the same in Q while Player j_1 's payoff in Q' is higher than it is in Q and Player j_2 's payoff is lower.

Recall that we assume that Player i 's objective in Q differs from his objective in Q' and that $i \in D$. We find a sequence of resources e_1, e_2, \dots , and for every $j \geq 1$, we associate with the resource e_j two players $j_i \in I$ and $j_d \in D$ such that $\sigma'_{j_i} = \sigma'_{j_d} = e_j$. First, we define $e_1 = \sigma'_i$ and we associate i with e_1 , thus $1_d = i$. Consider $j \geq 1$. Assume there is a player $j_d \in D$ with $\sigma'_{j_d} = e_j$. By (1), there is a $j_i \in I$ with $\sigma'_{j_i} = e_j$, and we associate j_i with e_j . By (2),

there is $(j+1)_d \in D$ such that Players $(j+1)_d$ and j_i choose the same resource in P_h , thus $\sigma_{j_i} = \sigma_{(j+1)_d}$. We define the next resource in the sequence e_{j+1} to be $\sigma'_{(j+1)_d}$, which by (3) is different from e_j . We associate Player $(j+1)_d \in D$ with e_{j+1} , and continue inductively.

Since there are finitely many resources, there is a loop $e_r, e_{r+1}, \dots, e_{s-1}, e_r$ in the sequence above. For $r \leq j \leq s-1$, note that $\text{cost}_{j_d}(Q') = \text{cost}_{j_i}(Q')$ and $\text{cost}_{j_i}(Q) = \text{cost}_{(j+1)_d}(Q)$. Moreover, $\text{cost}_{j_i}(Q') \geq \text{cost}_{j_i}(Q)$ and $\text{cost}_{j_d}(Q) > \text{cost}_{j_d}(Q')$. Combining the above we have $\text{cost}_{j_d}(Q) > \text{cost}_{j_d}(Q') = \text{cost}_{j_i}(Q') \geq \text{cost}_{j_i}(Q) = \text{cost}_{(j+1)_d}(Q)$, and we reach a contradiction.

G Proof of Theorem 6

The lower bounds for both measures are easy. Since we consider singleton games, an NE in a singleton RAG is an SPE in the corresponding dynamic RAG with the concurrent scheduler, i.e., all the players move simultaneously in the first round.

We continue to the upper bounds, and start with SGs. In SGs, we have $\text{PoS} = \log(k)$ and $\text{PoA} = k$ [2]. The proof for the upper bound for the DPoA is the same as in RAGs: if a player's cost is more than k times his cost in the social optimum in some SPE, then he can deviate to his objective in the social optimum and reduce his cost. For the upper bound on the DPoS, we note that the proof in Theorem ?? shows that a specific profile of an SG is an outcome of an SPE. It is shown in [29] that the cost of this profile is at most $\log(k) \cdot OPT$, thus we have $\text{DPoS} = \log(k)$.

We conclude by studying CGs. In singleton symmetric CGs, we have $\text{PoA} = 4/3$ [11] and we are not aware of a tight bound on the PoS. For the DPoS, Theorem 3 shows that every NE in a symmetric singleton CG corresponds to an SPE. For the DPoA, we use a claim from [10]. They show that with a sequential scheduler, every outcome of an SPE in a symmetric singleton CG is an NE in the underlying simultaneous game. Their proof works also for schedulers that are not sequential.

H Proof of Lemma 1

Note that the players' objectives are disjoint, so we analyze the game as if it takes place in one phase. We claim that a gate game simulates the semantics of a NAND gate. Assume both input players select their 1 objective, which corresponds to the case in which the input of the gate is two 1's, thus the outputs should be 0. For $j \in [r]$, we show that choosing the 0 objective is dominant for Player O_j . Indeed, if Player O_j plays his 0 objective, the cost of the resource g^j is split between three players, so the total cost for Player O_j is $1 + \epsilon$. On the other hand, the cost of the 1 objective is $1 + 1.1\epsilon$. Assume now that one of the input players chooses his 0 objective, thus the outputs should be 1. Then, choosing the 1 objective is dominant for Player O_j as the cost of the 0 objective is at least $1 + 1.5\epsilon$ while the cost of the 1 objective remains $1 + 1.1\epsilon$.

Next, we describe how to connect two gate games. Let \mathcal{G} and \mathcal{G}' be gate games as in the above. We connect the corresponding gates such that the j -th output of the first gate is fed as the first input to the second gate. In the combined game, the players of \mathcal{G} move before the players in \mathcal{G}' . Also, we merge between the output resources of Player O_j in \mathcal{G} with the input resources of Player I'_1 in \mathcal{G}' . More formally, let $\Sigma_{O_j} = \{\{o_1^j\}, \{g^j, o_0^j\}\}$ in \mathcal{G} and $\Sigma_{I'_1} = \{\{i_0^j\}, \{i_1^j, c^1, \dots, c^r\}\}$ in \mathcal{G}' . Then, in the combined game, we have $o_1^j = i_1^1$ and $o_0^j = i_0^1$. The cost of the first resource is 1 and the second is $1 + 1.1\epsilon$.

We claim that it is dominant for Player I'_1 to match Player O_j 's choice of objective. Intuitively, this follows from the fact that the input and output resources cost much more than the

gate resources, so sharing the cost of the first is more beneficial than sharing the second. More formally, assume Player O_j plays his 1 objective. If Player I_1' chooses his 1 objective, the cost of the resource i_1^1 is split between the two players. So, Player I_1' 's cost, no matter what the other players play is a bit over $\frac{1}{2}$. On the other hand, if he chooses his 0 objective, his cost is 1. Choosing the 1 objective in this case is clearly dominant. The case where Player O_j chooses his 0 objective is dual.

We proceed to describe the game \mathcal{G}_ψ (see for example Figure ??). In \mathcal{G}_ψ , there is a gate game that corresponds to every NAND gate in ψ , and the games are connected as in the above. For example, in Figure ??, for $i \in [4]$, let \mathcal{G}_i be the gate game that corresponds to the gate G_i . Consider the gate G_1 . One of its outputs is fed as input to G_4 and the other to G_3 . In \mathcal{G}_1 , the first pair of output resources are $o_1^{1,1}$ and $o_0^{1,1}$, which also serve as input resources in the gate game \mathcal{G}_4 . Thus, the set of objectives of the first input player in \mathcal{G}_4 is $\{\{o_0^{1,1}\}, \{o_1^{1,1}, g_4\}\}$. Similarly, the second pair of output resources are $o_1^{1,2}$ and $o_0^{1,2}$, which also serve as input resources in the game \mathcal{G}_3 .

We are left to describe the inputs and output of the circuit and how they interact. Assume the inputs to ψ are the variables x_1, \dots, x_n . Then, in addition to the players who simulate the NAND gates, the game \mathcal{G}_ψ includes also n *variable players*. As in the above, each variable player $j \in [n]$ has a 0 and 1 objective, which corresponds to an assignment to the variable x_j . Player x_j serves as the input player in every gate game that x_j appears in. So, Player x_j has a 0 objective which is $\{i_0^j\}$ and a 1 objective, which includes an input resource i_1^j as well as gate resources as in the above. For example, in Figure ??, the variable x_2 is fed as input to the gates C_1 and C_2 , so the set of objectives of Player x_2 is $\{\{i_0^2\}, \{i_1^2, g_1^1, g_1^2, c_2\}\}$.

Recall that there is a partition $E_1, \dots, E_m, A_1, \dots, A_m$ of x_1, \dots, x_n . The scheduler in \mathcal{G}_ψ schedules the variable players that correspond to the set of variables E_1 to move first, followed by the players who correspond to the set of variables A_1 , followed by E_2 , etc. The other players in \mathcal{G}_ψ who simulate the NAND gates follow according to the rules above.

Finally, there is a special NAND gate in ψ with only one output, where the output of this gate is the output of the whole circuit. We refer to this gate as the *final gate* and to the corresponding gate game as the *final gate game*. In Figure ??, the final gate is C_4 . Let Player O be the output player in the final gate game. Recall that each variable player x_j has a 1 objective with a resource i_1^j and a 0 objective with an input resource i_0^j . Further recall that in a gate game, the 0 objective of the output players includes a gate resource. We define Player O 's 0 objective to include the gate resource g of the final gate game as well as all the input resources of the universal variable players, thus it is $\{g\} \cup \{i_0^j, i_1^j : x_j \in A_1, \dots, A_m\}$. The 1 objective of Player O includes the input resources of all the existential variables players as well as another gate resource g' with cost 1.1ϵ , which we use to maintain the gate semantics, thus the 0 objective is $\{g'\} \cup \{i_0^j, i_1^j : x_j \in E_1, \dots, E_m\}$. We assume Wlog that the number of existential and universal variables in ψ is the same. So, Player O 's cost for the input resources is the same in both of his objectives no matter what the other players choose. Thus, the NAND semantics of the output gate is maintained.

Let Player x_1 be the first variable player to move in \mathcal{G}_ψ . We claim that if ψ is true, then in every SPE P , we have $\text{cost}_{x_1}(P) < \frac{1}{2} + \epsilon'$, and if ψ is false, then $\text{cost}_{x_1}(P) \geq 1$. Assume that x_1 is an existential variable, and the proof is dual for universal variables. Note that a cost of slightly over $\frac{1}{2}$ for Player x_j is achieved when Player O shares the cost of the input resource Player x_j uses. Thus, Player x_j , as well as all the existential players, have an incentive that the output of the circuit ψ is 1. Indeed, if the output is 0, the Player O shares the input resources with the universal players. Thus, if ψ is true, then the existential players can follow their assignments and guarantee a cost of slightly over $\frac{1}{2}$. On the other hand, if ψ is false, the universal players can guarantee a cost of slightly over $\frac{1}{2}$, making the cost of the existential players slightly over 1.

I Proof of Theorem 7

We start with the lower bound and the case of dynamic CGs. Given a QBF instance ψ , we apply the construction in Lemma 1 to get a sequential CG \mathcal{G}_ψ and $\gamma, \delta > 0$ such that if ψ is true, then $\text{cost}_1(\mathcal{G}_\psi) < \gamma$, and if ψ is false, then $\text{cost}_1(\mathcal{G}_\psi) > \delta$. Let \mathcal{G}_1 be the first 3-player game with no SPE that is described in Theorem 2. We construct a dynamic CG \mathcal{G}' , by merging \mathcal{G}_ψ and \mathcal{G}_1 , such that \mathcal{G}' has an SPE iff ψ is true.

We proceed to construct \mathcal{G}' . We add two players $k + 1$ and $k + 2$ to \mathcal{G}_ψ that take the roles of the first two players in \mathcal{G}_1 . The role of Player 3 in \mathcal{G}_1 is played by Player 1 in \mathcal{G}_ψ , so we add the objectives $\{a', b\}$ and $\{c\}$ to Σ_1 . Players $k + 1$ and $k + 2$ move first concurrently followed by Player 1 and the rest of the players in \mathcal{G}_ψ , which move sequentially. Recall that in every profile that is a candidate to be an SPE in \mathcal{G}_1 , Player 3's cost is at least $1\frac{2}{3}$ and at most $2\frac{2}{3}$. We alter the latency functions in \mathcal{G}_1 so that in these profiles Player 3 pays at least γ and at most δ . Let P be the SPE in \mathcal{G}_ψ . We claim that if $\text{cost}_1(P) < \gamma$, then \mathcal{G}' has an SPE. Indeed, it is not hard to see that the following profile is an SPE: Players $k + 1$ and $k + 2$ choose $\langle \{a, a'\}, \{c\} \rangle$, and the other players play according to P in every subgame. On the other hand, if $\text{cost}_1(P) > \delta$, then, there is no SPE. Indeed, Player 1 prefers the objectives originating from \mathcal{G}_1 over these in \mathcal{G} , and there is no SPE in the topmost subgame as shown in Theorem 2.

The case of dynamic SGs is similar. Let \mathcal{G}_ψ be the output of the construction in Lemma 1 and \mathcal{G}_2 the dynamic SG with no SPE that is described in Theorem 4. We merge \mathcal{G}_ψ with \mathcal{G}_2 by letting the first player in \mathcal{G}_ψ take the role of the fourth player in \mathcal{G}_2 similar to the above. Note that in every candidate profile for an SPE in \mathcal{G}_2 , Player 4's cost is 4. We alter the costs in \mathcal{G}_ψ so that $\gamma = 4 - \epsilon$ and $\delta = 4 + \epsilon$. Let P be an SPE in \mathcal{G}_ψ . We claim that if $\text{cost}_1(P) < \gamma$, then \mathcal{G}' has an SPE. Indeed, it is not hard to see that the following profile is an SPE: the first three players in \mathcal{G}_2 choose $\langle \{b, b''\}, \{b, b'\}, \{a, a''\} \rangle$, and the other players play according to P in every subgame. On the other hand, if $\text{cost}_1(P) > \delta$, then Player 4 always prefers his objective in \mathcal{G}_2 over his objective in P , and by the reasoning in Theorem 4, there is no SPE in the topmost subgame, and we are done.

We continue to study the upper bound. Consider a dynamic RAG \mathcal{G} , and let $\mathcal{T}_\mathcal{G}$ be the outcome tree of \mathcal{G} . Recall that there is a one-to-one correspondence between leaves in $\mathcal{T}_\mathcal{G}$ and legal outcomes of \mathcal{G} . We guess a leaf l in $\mathcal{T}_\mathcal{G}$ and verify that it is an outcome of an SPE. Thus, we ask if there is an SPE P in \mathcal{G} whose outcome corresponds to l .

In order to decide whether such a profile P exists, we traverse the path from l to the root of $\mathcal{T}_\mathcal{G}$. Consider an internal node h and its child on this path $h \cdot \bar{\sigma}$. Intuitively, our guess of l fixes the joint objective in h to be $\bar{\sigma}$. Consider a Player i who controls h and a resource e'_i different from $\bar{\sigma}(i)$. Thus, $\bar{\sigma}[i \leftarrow e'_i]$ is a unilateral deviation of Player i . We ask if it is possible to define the strategies in the other internal nodes such that the deviation is not beneficial for Player i . Formally, we ask if there is an SPE P' in the subgame $\mathcal{G}_{\bar{\sigma}[i \leftarrow e'_i]}$ in which the cost of Player i is at least his cost in l . This is done by calling the algorithm recursively with a slight change to the base case. Recall that in order for P to be an SPE, it must be an NE in *every* internal node in $\mathcal{T}_\mathcal{G}$. Thus, we verify that an SPE exists in every child $h \cdot \bar{\sigma}'$ of h that we have not yet considered. Thus, $\bar{\sigma}$ and $\bar{\sigma}'$ differ by at least two entries. This is again done by a recursive call to the algorithm. Clearly, the algorithm uses polynomial space.

J Proof of Theorem 8

For the upper bound, given an input $\langle G, \langle S, c \rangle, \beta \rangle$, for G with k players, we go over all schedulers $\nu : [k] \rightarrow [k]$. For each scheduler ν , we check whether $\text{cost}_{S,c}(\nu) \leq \beta$ and if so, we use the algorithm in Theorem 7 in order to decide whether the game $\langle G, \nu \rangle$ has an SPE. In case

it does, we accept. Clearly, the algorithm runs in polynomial space and accepts iff the input is legal.

For the lower bound, we show a reduction from \exists SPE. Given an input $\langle G, \nu \rangle$ to \exists SPE we construct an input $\langle G, \langle S, c \rangle, -(k-1)^2 \rangle$ to the BS problem as follows. For every $i_1, i_2 \in [k]$, we add to S a constraint according to the order of i_1 and i_2 in ν . For example, if $\nu(i_1) < \nu(i_2)$, then we add to S the constraint $i_1 < i_2$. All constraints have cost -1 . Clearly, the only scheduler that has cost at most $-(k-1)^2$ is ν , thus $\langle G, \nu \rangle \in \exists$ SPE iff $\langle G, \langle S, c \rangle, -(k-1)^2 \rangle \in$ BS, and we are done.⁵

⁵ Note that often it is possible to force ν with less than $(k-1)^2$ constraints. We have no reason however to minimize the number of constraints in the reduction.