Dynamic network congestion games

Nathalie Bertrand
Univ Rennes, Inria, CNRS, IRISA, France
nathalie.bertrand@inria.fr

Nicolas Markey
Univ Rennes, Inria, CNRS, IRISA, France
nicolas.markey@irisa.fr

Suman Sadhukhan
Univ Rennes, Inria, CNRS, IRISA, France
suman.sadhukhan@inria.fr

Ocan Sankur
Univ Rennes, Inria, CNRS, IRISA, France
ocan.sankur@irisa.fr

Abstract

Congestion games are a classical type of games studied in game theory, in which n players choose a resource, and their individual cost increases with the number of other players choosing the same resource. In network congestion games (NCGs), the resources correspond to simple paths in a graph, *e.g.* representing routing options from a source to a target. In this paper, we introduce a variant of NCGs, referred to as *dynamic NCGs*: in this setting, players take transitions synchronously, they select their next transitions dynamically, and they are charged a cost that depends on the number of players simultaneously using the same transition.

We study, from a complexity perspective, standard concepts of game theory in dynamic NCGs: social optima, Nash equilibria, and subgame perfect equilibria. Our contributions are the following: the existence of a strategy profile with social cost bounded by a constant is in PSPACE and NP-hard. (Pure) Nash equilibria always exist in dynamic NCGs; the existence of a Nash equilibrium with bounded cost can be decided in EXPSPACE, and computing a witnessing strategy profile can be done in doubly-exponential time. The existence of a subgame perfect equilibrium with bounded cost can be decided in 2EXPSPACE, and a witnessing strategy profile can be computed in triply-exponential time.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory; Theory of computation → Formal languages and automata theory; Theory of computation → Solution concepts in game theory; Theory of computation → Verification by model checking

Keywords and phrases Congestion games, Nash equilibria, Subgame perfect equilibria, Complexity

Digital Object Identifier 10.4230/LIPIcs...116

Related Version https://arxiv.org/abs/2009.13632

1 Introduction

Congestion games model selfish resource sharing among several players [20]. A special case is the one of network congestion games, in which players aim at routing traffic through a congested network. Their popularity is certainly due to the fact that they have important practical applications, whether in transportation networks, or in large communication networks [19]. In network congestion games, each player chooses a set of transitions, forming a simple path from a source state to a target state, and the cost of a transition increases with its load, that is, with the number of players using it.

Network congestion games can be classified into atomic and non-atomic variants. Non-atomic semantics is appropriate for large populations of players, thus seen as a continuum. One then considers portions of the population that apply predefined strategies, and there is no such thing as the effect of an individual player on the cost of others. In contrast, in atomic games, the number of players is fixed, and each player may significantly impact the cost other players incur. We only focus on atomic games in this paper.

Network congestion games. Network congestion games, also called atomic selfish routing games in the literature, were first considered by Rosenthal [20]. These games are defined by a directed graph, a number of pairs of source-target vertices, and non-decreasing cost functions for each edge in the graph. For each source-target pair, a player must choose a route from the source to the target vertex. Given their choice of simple paths, the cost incurred by a player depends on the number of other players that choose paths sharing edges with their path, and on the cost functions of these edges. In this setting, a Nash equilibrium maps each player to a path in such a way that no player has an incentive to deviate: they cannot decrease their cost by choosing a different path.

Rosenthal proved that they are potential games, so that Nash equilibria always exist. Monderer and Shapley [17] studied in a more general way potential games, and explained how to iteratively use best-response strategies to converge to an equilibrium. Interestingly, under reasonable assumptions on the cost functions, Bertsekas and Tsitsiklis established that there is a direct correspondence between equilibria in selfish routing and distributed shortest-path routing, which is used in practice for packet routing in computer networks [7]. We refer the interested reader to [21] for an introduction and many basic results on general routing games.

A natural question is whether selfish routing is very different from a routing strategy decided by a centralized authority. In other words, how far can a selfish optimum be from the social optimum, in which players would cooperate. The notion of price of anarchy, first proposed by Koutsoupias and Papadimitriou [14], is the ratio of the worst cost of a Nash equilibrium and the cost of the social optimum. This measures how bad Nash equilibria can be. In the context of network congestion games, the price of anarchy was first studied by Suri *et al.* [22], establishing an upper bound of $\frac{5}{2}$ when all cost functions are affine. A refined upper bound was provided by Awerbuch *et al.* [5]. Bounds on the dual notion of price of stability, which is the ratio of the cost of a best Nash equilibrium and the cost of the social optimum was also studied for routing games [1].

Timing aspects. Several works investigated refinements of this setting. In [11], the authors study network congestion games in which each edge is traversed with a fixed duration independent of its load, while the cost of each edge depends on the load. The model is thus said to have *time-dependent* costs since the load depends on the times at which players traverse a given edge. The authors prove the existence of Nash equilibria by reduction to the setting of [20]. An extension of this setting with timed constraints was studied in [2, 3].

The setting of fixed durations with time-dependent costs is interesting in applications where the players sharing a resource (an edge) see their quality of service decrease, while the time to use the resource is unaffected [3]. This might be the case, for instance, in some telecommunication and multimedia streaming applications. Timing also appears, for instance, in [18, 15] where the load affects travel times and players' objective is to minimize the total travel time. Other works focus on flow models with a timing aspect [13, 8].

Dynamic network congestion games. In classical network congestion games, including those mentioned above, players choose their strategies (*i.e.,* their *simple* paths) in one shot. However, it may be interesting to let agents choose their paths *dynamically*, that is, step by step, by observing other players' previous choices. In this paper, we study network congestion games with time-dependent costs as in [11], but with unit delays, and in a dynamic setting. More precisely, at each step, each of the players simultaneously selects the edge they want to take; each player is then charged a cost that depends on the load of the edge they selected, and traverses that edge in one step. We name these games *dynamic network congestion games* (dynamic NCGs in short); the behaviour of the players in such games is formalized by means of *strategies*, telling the players what to play depending of the current configuration. Notice that, because congestion effect applies to edges used *simultaneously* by several players, taking cycles can be interesting in dynamic NCGs, which makes our setting more complex than most NCG models [4, 11, 20, 21].

Such a dynamic setting was studied in [4] for resource allocation games, which extends [20] with dynamic choices. A more detailed related work appears at the end of this section.

Standard solution concepts. We study classical solution concepts on dynamic network congestion games. A strategy profile (*i.e.*, a function assigning a strategy to each player) is a *Nash Equilibrium* (NE) when each single strategy is an optimal response to the strategies of the other players; in other terms, under such a strategy profile, no player may lower their costs by unilaterally changing their strategies. Notice that NEs need not exist in general, and when they exist, they may not be unique. In the setting of dynamic games, Nash Equilibria are usually enforced using *punishing strategies*, by which any deviating player will be punished by all other players once the deviation has been detected. However, such punishing strategies may also increase the cost incurred to the punishing players, and hence do not form a credible threat; *Subgame-Perfect Equilibria* (SPEs) refine NEs and address this issue by requiring that the strategy profile is an NE along any play.

NEs and SPEs aim at minimizing the individual cost of each player (without caring of the others' costs); in a collaborative setting, the players may instead try to lower the social cost, i.e., the sum of the costs incurred to all the players. Strategy profiles achieving this are called *social optima* (SO). Obviously, the social cost of NEs and SPEs cannot be less than that of the social optimum; the *price of anarchy* measures how bad selfish behaviours may be compared to collaborative ones.

Our contributions. We take a computational-complexity viewpoint to study dynamic network congestion games. We first establish the complexity of computing the social optimum, which we show is in PSPACE and NP-hard. We then prove that best-response strategies can be computed in polynomial time, and that dynamic NCGs are potential games, thereby showing the existence of Nash equilibria in any dynamic NCG; this also shows that some Nash equilibrium can be computed in pseudo-polynomial time. We then give an EXPSPACE (resp. 2EXPSPACE) algorithm to decide the existence of Nash Equilibria (resp. Subgame-Perfect Equilibria) whose costs satisfy given bounds. This allows us to compute best and worst such equilibria, and then the price of anarchy and the price of stability.

Note that some of the high complexities follow from the binary encoding of the number of players, which is the main input parameter. For instance, the exponential-space complexity drops to pseudo-polynomial time for a fixed number of players. This parameter becomes important since we advocate the study of computational problems, such as computing Nash equilibria with a given cost bound. We also believe that computing precise values for price of anarchy and the price of stability is interesting, rather than providing bounds on the set of all instances as in *e.g.* [22].

Omitted proofs can be found in the corresponding arXiv article [6].

Comparison with related work. The works closest to our setting are [11, 4, 2, 3]. As in [11, 3], we establish the existence of Nash equilibria using potential games. Unlike [11], we cannot obtain this result immediately by reducing our games to congestion games [20] since the lengths of the strategies cannot be bounded *a priori*. Moreover, the best-response problem has a polynomial-time solution in our setting while the problem is NP-hard both in [11, 3]. In [11], this is due to the possibility of having arbitrary durations, while the source of complexity in [2, 3] is due to the use of timed automata. Thus, our setting offers a simpler way of expressing timings, and avoids their high complexity for this problem.

Dynamic choices were studied in [4] but with a different cost model. Moreover, network congestion games can only be reduced to such a setting given an a priori bound on the length of the paths. So we cannot directly transfer any of their results to our setting. Dynamic choices were also studied in [11] in the setting of coordination mechanisms which are local policies that allow one to sequentialize traffic on the edges.

2 Preliminaries

2.1 Dynamic network congestion games

Let $\mathcal{F}$ be the family of non-decreasing functions from $\mathbb{N}$ to $\mathbb{N}$ that are piecewise-affine, with finitely many pieces. We assume that each $\mathsf{f} \in \mathcal{F}$ is represented by the endpoints of intervals, and the coefficients, all encoded in binary. An arena for dynamic network congestion games is a weighted graph $\mathcal{A} = \langle V, E, \mathsf{src}, \mathsf{tgt} \rangle$, where V is a finite set of states, $E \colon V \times V \to \mathcal{F}$ is a partial function defining the cost of edges, and src and tgt are a source- and a target state in V. It is assumed throughout this paper that tgt has only a single outgoing transition, which is a self-loop with constant cost function $x \mapsto 0$. We also assume that tgt is reachable from all other states.

A dynamic network congestion game (dynamic NCG for short) is a pair $\mathcal{G} = \langle \mathcal{A}, n \rangle$ where $\mathcal{A}$ is an arena as above and $n \in \mathbb{N}$ is the (binary-encoded) number of players. In a dynamic network congestion game, all players start from src and simultaneously select the edges they want to traverse, with the aim of reaching the target state with minimal individual accumulated cost. Taking an edge $e = (v, \mathsf{f}, v')$ has a cost $\mathsf{f}(l)$, where l is the number of players simultaneously using edge e. The cost function of edge e is denoted by f_e. We let $\kappa = \max_{e \in E} \mathsf{f}_e(n)$, which is the maximal cost that a player may endure along one edge.

Our setting differs from classical network congestion games [21] mainly in two respects:

- first, the game is played in rounds, during which all players take exactly one transition; the number of players using an edge e is measured *dynamically*, at each round;
- second, during the play, players may adapt their choices to what the other players have been doing in the previous rounds.

▶ Remark 1. In this work, we mainly focus on the *symmetric* case, where all players have the same source and target. This is because we take a parametric-verification point of view, with

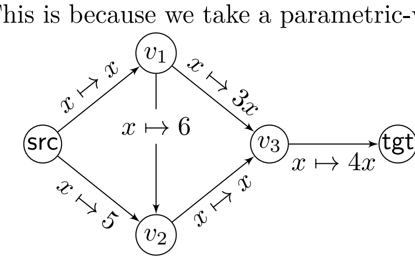

Figure 1 Representation of an arena for a dynamic NCG (loop omitted on tgt)

the (long-term) aim of checking properties of dynamic NCGs for arbitrarily many players. An important consequence of this choice is that the number of players now is encoded in binary, which results in an exponential blow-up in the number of configurations of the game (compared to the asymmetric setting).

Semantics as a concurrent game. For any $k \in \mathbb{N}$, we write $[\![k]\!] = \{i \in \mathbb{N} \mid 1 \leq i \leq k\}$. A *configuration* of a dynamic network congestion game $\langle \mathcal{A}, n\rangle$ is a mapping $c\colon [\![n]\!] \to V$, indicating the position of each player in the arena. We define $c_{\mathsf{src}}\colon i \in [\![n]\!] \mapsto \mathsf{src}$ and $c_{\mathsf{tgt}}\colon i \in [\![n]\!] \mapsto \mathsf{tgt}$ as the initial and target configurations, respectively.

With $\langle \mathcal{A}, n\rangle$, we first associate a multi-weighted graph $\mathcal{M} = \langle C, T\rangle$, where $C = V^{[\![n]\!]}$ is the set of all configurations and $T \subseteq C \times \mathbb{N}^{[\![n]\!]} \times C$ is a set of edges, defined as follows: there is an edge (c, w, c') in T if, and only if, there exists a collection $\mathbf{e} = (e_i)_{i\in[\![n]\!]}$ of edges of E such that for all $i \in [\![n]\!]$, writing $e_i = (v_i, \mathsf{f}_i, v'_i)$ and $u_i = \#\{j \in [\![n]\!] \mid e_j = e_i\}$, we have $c(i) = v_i$, $c'(i) = v'_i$, and $w(i) = \mathsf{f}_i(u_i)$. We denote this edge with $c \xRightarrow{\mathbf{e}} c'$. We may omit to mention $\mathbf{e}$ since it can be obtained from c and c'; similarly, we write $\mathsf{cost}_i(c, c')$ for $w(i)$.

Two edges (c, w, c') and (d, x, d'), in that order, are said to be *consecutive* whenever $d = c'$. Given a configuration c, a *path* from c in a dynamic network congestion game is either the single configuration c (we call this a trivial path) or a non-empty, finite or infinite sequence of consecutive edges $\rho = (t_j)_{1\leq j<|\rho|}$ in $\mathcal{M}$, where t_1 is a transition from c; the size of a path ρ is one for trivial paths, and $|\rho| \in \mathbb{N} \cup \{+\infty\}$ otherwise. We write $\mathsf{Paths}(\langle \mathcal{A}, n\rangle, c)$ and $\mathsf{Paths}^\omega(\langle \mathcal{A}, n\rangle, c)$ for the set of finite and infinite paths from c in $\langle \mathcal{A}, n\rangle$, respectively.

With each path $\rho = (c_j, w_j, c'_j)_j$, and each player $i \in [\![n]\!]$, we associate a *cost*, written $\mathsf{cost}_i(\rho)$, which is zero for trivial paths, $+\infty$ for infinite paths along which $c_j(i) \neq \mathsf{tgt}$ for all j, and $\sum_{j=1}^{|\rho|-1} w_j(i)$ otherwise. We define the *social cost* of ρ, denoted by $\mathsf{soccost}(\rho)$, as $\sum_{i\in[\![n]\!]} \mathsf{cost}_i(\rho)$.

Given a path ρ, an index $1 \leq j < |\rho| + 1$ and a player $i \in [\![n]\!]$, we write $\rho(j)$ for the j-th configuration of ρ, and $\rho(j)(i)$ for the state of Player i in that configuration. For $j \geq 2$, we define $\rho_{\leq j}$ as the prefix of ρ that ends in the j-th configuration; we let $\rho_{\leq 1} = \rho(1)$. Similarly, for $1 \leq j \leq |\rho| - 1$, we let $\rho_{\geq j}$ denote the suffix that starts at the j-th configuration. Finally, if $|\rho|$ is finite, we let $\rho_{\geq|\rho|} = \rho(|\rho|)$.

▶ **Example 2.** Consider the arena $\mathcal{A}$ displayed at Fig. 1 and the dynamic NCG $\langle \mathcal{A}, 2\rangle$ with two players. Assume that Player 1 follows the path $\pi_1\colon \mathsf{src} \to v_1 \to v_3 \to \mathsf{tgt}$ and Player 2 goes via $\pi_2\colon \mathsf{src} \to v_1 \to v_2 \to v_3 \to \mathsf{tgt}$. This gives rise to the following path:

$$\begin{pmatrix} 1 \mapsto \mathsf{src} \\ 2 \mapsto \mathsf{src} \end{pmatrix} \xrightarrow[2\mapsto 2]{1\mapsto 2} \begin{pmatrix} 1 \mapsto v_1 \\ 2 \mapsto v_1 \end{pmatrix} \xrightarrow[2\mapsto 6]{1\mapsto 3} \begin{pmatrix} 1 \mapsto v_3 \\ 2 \mapsto v_2 \end{pmatrix} \xrightarrow[2\mapsto 1]{1\mapsto 4} \begin{pmatrix} 1 \mapsto \mathsf{tgt} \\ 2 \mapsto v_3 \end{pmatrix} \xrightarrow[2\mapsto 4]{1\mapsto 0} \begin{pmatrix} 1 \mapsto \mathsf{tgt} \\ 2 \mapsto \mathsf{tgt} \end{pmatrix}$$

Notice how edge $v_3 \to \mathsf{tgt}$ of $\mathcal{A}$ is used by both players, but not simultaneously, so that the cost of using that edge is 4 for each of them, while it would be 8 in classical NCGs. ◀

We now extend this graph to a concurrent game structure. A *move* for Player $i \in [\![n]\!]$ from configuration c is an edge $e = (v, \mathsf{f}, v') \in E$ such that $v = c(i)$. A *move vector* from c is a sequence $\mathbf{e} = (e_i)_{i\in[\![n]\!]}$ such that for all $i \in [\![n]\!]$, e_i is a move for Player i from c.

A network congestion game $\langle \mathcal{A}, n\rangle$ then gives rise to a concurrent game structure $\mathcal{S} = \langle C, T, M, U\rangle$ where $\langle C, T\rangle$ is the graph defined above, $M\colon C \times [\![n]\!] \to 2^E$ lists the set of possible moves for each player in each configuration, and $U\colon C \times E^{[\![n]\!]} \to T$ is the transition function, such that for every configuration c and every move vector $\mathbf{e} = (e_i)_{i\in[\![n]\!]}$ with $e_i \in M(c, i)$ for all $i \in [\![n]\!]$, $U(c, e) = (c \xRightarrow{\mathbf{e}} c')$.

A *strategy* for Player i in $\mathcal{S}$ from configuration c is a function $\sigma_i \colon \mathsf{Paths}(\langle \mathcal{A}, n\rangle, c) \to E$ that associates, with any finite path ρ from c in $\mathcal{S}$, a move for this player from the last configuration of ρ. A *strategy profile* is a family $\sigma = (\sigma_i)_{i\in[\![n]\!]}$ of strategies, one for each player. We write $\mathfrak{S}$ for the set of strategies, and $\mathfrak{S}^n$ for the set of strategy profiles.

Let c be a configuration, h be a finite path from c and a strategy profile $\sigma = (\sigma_i)_{i\in[\![n]\!]}$ from c. The *residual strategy profile* of σ after h is the strategy profile $\sigma^h = (\sigma^h_i)_{i\in[\![n]\!]}$ from the last configuration of h defined by $\sigma^h_i(h') = \sigma_i(h \cdot h')$, where $h \cdot h'$ is the concatenation of paths h and h'.

The *outcome* of a strategy profile σ from c is the infinite path $\rho = (c_i, w_i, c_{i+1})_{i\geq 1}$, hereafter denoted with $\mathsf{outcome}(\sigma)$, obtained by running the strategy profile; it is formally defined as the only infinite path such that $(c_1, w_1, c_2) = U(c, \sigma(c))$, and such that for any $j \geq 2$, $(c_j, w_j, c_{j+1}) = U(c_j, \sigma(h'))$, where $h' = (c_1, w_1, c_2) \cdots (c_{j-1}, w_{j-1}, c_j)$.

Pick a strategy profile $\sigma = (\sigma_i)_{i\in[\![n]\!]}$, and let $\rho = (t_j)_{j\geq 1}$ be its outcome, writing $t_j = (c_j, (w^i_j)_{i\in[\![n]\!]}, c'_j)$ for all $j \geq 1$. Let $k \in [\![n]\!]$. If $c'_l(k) = \mathsf{tgt}$ for some $l \in \mathbb{N}$, then σ_k is said to be winning for Player k. In that case, we define $\mathsf{cost}_k(\sigma)$ as $\mathsf{cost}_k(\mathsf{outcome}(\sigma))$. If $c'_l(i) = \mathsf{tgt}$ for all $i \in [\![n]\!]$, we define the *social cost* of σ as $\mathsf{soccost}(\sigma) = \mathsf{soccost}(\rho)$.

A strategy σ_i for Player i is said *blind* whenever for any two finite paths ρ and ρ' having same length k, if for any position $0 \leq j < k$ we have $\rho(j)(i) = \rho'(j)(i)$, then $\sigma_i(\rho) = \sigma_i(\rho')$. Intuitively, this means that strategy σ_i follows a path in $\mathcal{A}$, independently of what the other players do. A blind strategy can thus be represented as a path and we write $|\sigma_i|$ for the length of that path (until its first visit to tgt, if any). We write $\mathfrak{B}$ for the set of blind strategies.

▶ **Example 3.** Consider again the arena $\mathcal{A}$ of Fig. 1. The paths π_1 and π_2 from Example 2 are two blind strategies in that dynamic NCG. In a 2-player setting, an example of a non-blind strategy σ consists in first taking the transition $\mathsf{src} \to v_1$, and then either taking $v_1 \to v_3$ if the other player took the same initial transition, or taking $v_1 \to v_2$ otherwise. ◀

Representation as a weighted graph. Another way of representing configurations is to consider their Parikh images. With a configuration $c \in V^{[\![n]\!]}$, we associate an abstract configuration $\bar{c} \in [\![n]\!]^V$ defined as $\bar{c}(v) = \#\{i \in [\![n]\!] \mid c(i) = v\}$.

The *abstract weighted graph* associated with a dynamic network congestion game $\langle \mathcal{A}, n\rangle$ is the weighted graph $\mathcal{P} = \langle A, B\rangle$, where A contains all abstract configurations, and there is an edge (a, w, a') in $B \subseteq A \times \mathbb{N} \times A$ if, and only if, there is a mapping $b \colon E \to [\![n]\!]$ such that $\sum_{e\in E} b(e) = n$ and for all $v \in V$,

$$a(v) = \sum_{e=(v,\mathsf{f},v')} b(e) \qquad w = \sum_{e=(v,\mathsf{f},v')} b(e) \times \mathsf{f}(b(e)) \qquad a'(v) = \sum_{e=(v',\mathsf{f},v)} b(e).$$

Similarly to the representation as multi-weighted graphs, an *abstract path* of a network congestion game is either a single configuration or a non-empty, finite or infinite sequence of consecutive edges in the abstract weighted graph. The *cost* of an abstract path is the sum of the weights of its edges (if any). Then:

▶ **Lemma 4.** *For any $w \in \mathbb{N} \cup \{+\infty\}$, there is an abstract path in $\mathcal{M}$ with social cost w if, and only if, there is an abstract path in $\mathcal{P}$ with cost w.*

2.2 Social optima and equilibria

Consider a dynamic network congestion game $\mathcal{G} = \langle \mathcal{A}, n\rangle$. We recall standard ways of optimizing the strategies of the players, depending on the situation.

In a collaborative situation, all players want to collectively minimize the total cost for having all of them reach the target state of the arena. Formally, a strategy profile $\sigma = (\sigma_i)_{i\in[\![n]\!]}$ realizes the *social optimum* if $\mathsf{soccost}(\sigma) = \inf_{\tau\in\mathfrak{S}^n} \mathsf{soccost}(\tau)$.

In a selfish situation, each player aims at optimizing their response to the others' strategies. Given a strategy profile $\sigma = (\sigma_i)_{i\in[\![n]\!]}$, a player $k \in [\![n]\!]$, and a strategy $\sigma'_k \in \mathfrak{S}$, we denote by $\langle\sigma_{-k}, \sigma'_k\rangle$ the strategy profile $(\tau_i)_{i\in[\![n]\!]}$ such that $\tau_k = \sigma'_k$ and $\tau_i = \sigma_i$ for all $i \in [\![n]\!] \setminus \{k\}$. The strategy σ_k is a *best response* to $(\sigma_i)_{i\in[\![n]\!]\setminus\{k\}}$ if $\mathsf{cost}_k(\sigma) = \inf_{\sigma'_k\in\mathfrak{S}} \mathsf{cost}_k(\langle\sigma_{-k}, \sigma'_k\rangle)$. A strategy profile $\sigma = (\sigma_i)_{i\in[\![n]\!]}$ is a *Nash equilibrium* if for each $k \in [\![n]\!]$, the strategy σ_k is a best response to $(\sigma_i)_{i\in[\![n]\!]\setminus\{k\}}$. In such a case, no player has profitable unilateral deviations, *i.e.*, no player alone can decrease their cost by switching to a different strategy.

Nash equilibria can be defined for subclasses of strategy profiles. In particular, a *blind Nash equilibrium* is a blind strategy profile σ that is a Nash equilibrium *for blind-strategy deviations*: for all $k \in [\![n]\!]$, $\mathsf{cost}_k(\sigma) = \inf_{\sigma'_k\in\mathfrak{B}} \mathsf{cost}_k(\langle\sigma_{-k}, \sigma'_k\rangle)$. *A priori*, a blind Nash equilibrium need not be a Nash equilibrium for general strategies.

In an NE, once a player deviated from their original strategy in the strategy profile, the other players can punish the deviating player, even if this results in increasing their own costs. Indeed, the condition for being an NE only requires that the deviation should not be profitable to the deviating player. Subgame-Perfect Equilibria (SPE) refine NEs and rule out such non-credible threats by requiring that, for any path h in the configuration graph, the residual strategy profile after h is an NE.

▶ **Example 5.** Consider again the dynamic NCG $\langle\mathcal{A}, 2\rangle$, with the arena $\mathcal{A}$ of Fig. 1. Assume that Player 1 plays the blind strategy corresponding to $\pi_3\colon \mathsf{src} \to v_2 \to v_3 \to \mathsf{tgt}$, while Player 2 plays the non-blind strategy σ of Example 3. The cost for Player 1 then is 10, while that of Player 2 is 12.

This strategy profile is an NE: Player 2 could be tempted to play π_1, but they would then synchronize with Player 1 on edge $v_3 \to \mathsf{tgt}$, and get cost 12 again. Similarly, Player 1 could be tempted to play π_1 instead of π_3, but in that case strategy σ would tell Player 2 to follow the same path, and the cost for Player 1 (and 2) would be 16. Notice in particular that this is not an SPE, but that the blind strategy profile $\langle\pi_1, \pi_2\rangle$ (extended to the whole configuration tree in the only possible way) is an SPE in $\langle\mathcal{A}, 2\rangle$. ◀

In Sections 4 and 5, we focus on NEs and SPEs, developing EXPSPACE and 2EXPSPACE-algorithms for deciding the existence of NEs and SPEs respectively of social cost less than or equal to a given bound. Actually, our approach extends to the $\vec{\gamma}$-weighted social cost, where $\vec{\gamma} \in \mathbb{Z}^{[\![n]\!]}$ are coefficients applied to the costs of the respective players when computing the social cost. As a consequence, we can compute best and worst NEs and SPEs, hence also the price of anarchy and price of stability [14]. Before that, in Section 3, we extend classical techniques using blind strategies to compute the social optimum and prove that NEs always exist.

3 Socially-optimal strategy profiles

To compute a socially-optimal strategy profile, it suffices to find a path in the concurrent game structure of the given network congestion game with minimal total cost since one can define a strategy profile that induces any given path. Rather than finding such a path in the concurrent game structure, and in view of Lemma 4, one can look for one in the abstract weighted graph, thereby reducing in complexity. The socially-optimal cost in a dynamic

NCG $\langle \mathcal{A}, n \rangle$ is thus the cost of a shortest path in the associated weighted abstract graph $\mathcal{P}$ from $\bar{c}_{\mathsf{src}}$ to $\bar{c}_{\mathsf{tgt}}$.

Since $\mathcal{P}$ has exponential size, we derive complexity upper bounds for computing a socially-optimal strategy and deciding the associated decision problem. Moreover, adapting [16, Theorem 4.1] which proves NP-hardness in classical NCGs, we provide a reduction from the Partition problem to establish an NP lower-bound.

▶ **Theorem 6.** *A socially-optimal strategy profile can be computed in exponential time. The constrained social-optimum problem is in PSPACE and NP-hard.*

Note, that while $\mathcal{P}$ has size $(n+1)^{|V|}$, it is sufficient to consider paths with a smaller number of transitions when looking for a shortest path:

▶ **Lemma 7.** *There is a shortest path (w.r.t. cost) in $\mathcal{P}$ with size (in terms of its number of transitions) at most $n \cdot |V|$.*

▶ Remark 8. A consequence of Lemma 7 is that deciding the constrained social-optimum problem is in NP for asymmetric games, since in that setting the lists of sources and targets of each player is part of the input, so that n is polynomial in the size of the input. However, our NP-hardness proof only works in the symmetric case.

4 Nash equilibria

In this section, we study the existence of Nash equilibria and give algorithms to compute them under given constraints.

4.1 Existence and computation of (blind) Nash equilibria

To prove that blind Nash equilibria always exist, we establish that dynamic NCGs with blind strategies are potential games [20, 17] which are known to have Nash equilibria.

Consider a dynamic NCG $\langle \mathcal{A}, n \rangle$, a blind strategy profile π, and let N_π denote the maximum length of the paths prescribed by π. We define the following potential function, which is an adaptation of that used in [20]:

$$\psi(\pi) = \sum_{j=1}^{N_\pi} \sum_{e \in E} \sum_{i=1}^{\mathsf{load}_e(\pi,j)} \mathsf{f}_e(i),$$

where $\mathsf{load}_e(\pi, j)$ denotes the number of players that take edge e in the j-step under π, and f_e is the cost function on edge e.

Using the above-defined potential function, one can derive an algorithm to find a Nash equilibrium, by a classical *best-response* iteration. Starting with an arbitrary blind strategy profile, at each step we replace some player's strategy with their best-response, and we continue as long as some player's cost can be decreased. When this procedure terminates, the profile at hand is a blind Nash equilibrium. In dynamic NCGs, best responses exist and can be computed in polynomial time. Indeed, one can construct a game in which all players but Player i follow their fixed strategies given by profile π, using N_π copies of the game in order to distinguish the steps. After the N_π-step, all players in $[\![n]\!] \setminus \{i\}$ have reached their targets. Since it is the only remaining player, the remaining path for Player i should not be longer than $|V|$. Altogether, we obtain the following complexity upper-bound:

▶ **Theorem 9.** *In dynamic NCGs, blind Nash equilibria always exist, and we can compute one in pseudo-polynomial time.*

▶ Remark 10. As an alternative proof to existence of blind NEs, we could have bounded the length of outcomes of blind NEs as follows: all players have a strategy realizing cost at most $|V| \cdot \kappa$, where $\kappa = \max_{e \in E} \mathsf{f}_e(n)$, since the shortest path from src to tgt has length at most $|V|$, and the cost for a player at each step along edge e is at most κ. It follows that no path along which the cost for some player is larger than $|V| \cdot \kappa$ can be the outcome of a blind NE. As a consequence, if a dynamic NCG has a blind NE, then it has one of length at most $|V| \cdot \kappa \cdot |V|^n$ (by removing zero-cycles). Using this bound, we can transform dynamic NCGs into classical congestion games, in which blind NEs always exist [11, 20].

We now show that blind Nash equilibria are in fact Nash equilibria. This is proved using the observation that given a blind strategy profile, the most profitable deviation for any player can be assumed to be a blind strategy.

▶ **Lemma 11.** *In dynamic NCGs, blind Nash equilibria are Nash equilibria.*

Computing some (blind) Nash equilibrium may not be satisfactory for two reasons: one might want to compute the best (or the worst) Nash equilibrium in terms of the social cost; and as Lemma 12 claims, blind Nash equilibria are suboptimal, *i.e.*, a lower social cost can be achieved by Nash equilibria with general strategies. This justifies the study of more complex strategy profiles in the next subsection.

▶ **Lemma 12.** *There exists a dynamic NCG with a Nash equilibrium π such that for all blind Nash equilibria π', we have $\mathsf{cost}(\pi) < \mathsf{cost}(\pi')$.*

The proof is based on the dynamic NCG depicted on Fig. 2, for which we prove there is a Nash equilibrium with total cost 36, while any *blind* Nash equilibrium has higher social cost.

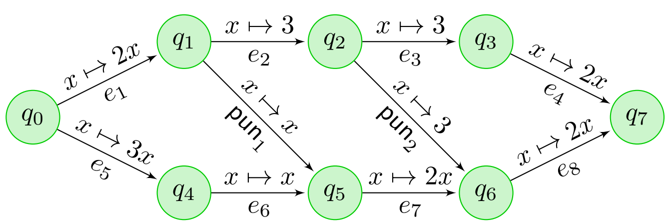

Figure 2 An arena on which blind Nash equilibria are sub-optimal.

4.2 Computation of general Nash equilibria

Characterization of outcomes of Nash Equilibria. Let us consider a dynamic NCG $\langle \mathcal{A}, n \rangle$, and the corresponding game structure $\mathcal{S} = \langle C, T, M, U \rangle$. Given two configurations c, c' with $c \Rightarrow c'$, we let $\mathsf{cost}_i(c, c')$ denote the cost of Player i on this transition from $c(i)$ to $c'(i)$. We define $\mathsf{dev}_i(c, c')$ as the set of all configurations reachable when all players but Player i choose moves prescribed by the given transition $c \Rightarrow c'$:

$$\mathsf{dev}_i(c, c') = \{c'' \in C \mid c \Rightarrow c'' \text{ and } \forall j \in [\![n]\!] \setminus \{i\}.\ c''(j) = c'(j)\}.$$

The *value* of configuration c for Player i is $\mathsf{val}_{i,c} = \sup_{\sigma_{-i} \in \mathfrak{S}^{n-1}} \inf_{\sigma_i \in \mathfrak{S}} \mathsf{cost}_i((\sigma_{-i}, \sigma_i), c)$. Note that the value corresponds to the value of the zero-sum game where Player i plays against the opposing coalition, starting at c. By [12], those values can be computed in polynomial time in the size of the game. Here the game is a 2-player game with state space $|V| \times [\![n-1]\!]^{|V|}$, keeping track of the position of Player i and the abstract position of the coalition. It follows that each $\mathsf{val}_{i,c}$ can be computed in exponential time in the size of the input $\langle \mathcal{A}, n\rangle$. Moreover, memoryless optimal strategies exist (in $\mathcal{S}$), that is, the opposing coalition has a memoryless strategy σ_{-i} to ensure a cost of at least $\mathsf{val}_{i,c}$ from c.

The characterization of Nash equilibria outcomes is given in the following lemma.

▶ **Lemma 13.** *A path ρ in $\langle \mathcal{A}, n\rangle$ is the outcome of a Nash equilibrium if, and only if,*

$$\forall i \in [\![n]\!].\ \forall 1 \leq l < |\rho|.\ \forall c \in \mathit{dev}_i(\rho(l), \rho(l+1)). \quad \mathit{cost}_i(\rho_{\geq l}) \leq \mathit{val}_{i,c} + \mathit{cost}_i(\rho(l), c).$$

The intuition is that if the suffix $\mathsf{cost}_i(\rho_{\geq l})$ of ρ has cost more than $\mathsf{val}_{i,c} + \mathsf{cost}_i(\rho(l), c)$, then Player i has a profitable deviation regardless of the strategy of the opposing coalition, since $\mathsf{val}_{i,c}$ is the maximum cost that the coalition can inflict to Player i at configuration c where the deviation is observed. The lemma shows that the absence of such a suffix means that a Nash equilibrium with given outcome exists, which the proof constructs.

Proof. Consider a Nash equilibrium $\sigma = (\sigma_i)_{i \in [\![n]\!]}$ with outcome ρ. Consider any player i, and any strategy σ'_i for this player. Let ρ' denote the outcome of $\sigma[i \to \sigma'_i]$. Let l denote the index of the last configuration where ρ and ρ' are identical. Since σ is a Nash equilibrium, we have $\mathsf{cost}_i(\rho) \leq \mathsf{cost}_i(\rho')$, that is,

$$\mathsf{cost}_i(\rho_{\geq l}) \leq \mathsf{cost}_i(\rho(l), \rho'(l+1)) + \mathsf{cost}_i(\sigma[i \to \sigma'_i], \rho'_{\leq l+1})$$

where $\mathsf{cost}_i(\sigma[i \to \sigma'_i], \rho'_{\leq l+1})$ is the cost for Player i of the outcome of the residual strategy $(\sigma[i \to \sigma'_i])^{\rho'_{\leq l+1}}$. Since the choice of σ'_i is arbitrary here, we have,

$$\mathsf{cost}_i(\rho_{\geq l}) \leq \mathsf{cost}_i(\rho(l), \rho'(l+1)) + \inf_{\sigma'_i \in \mathfrak{S}} \mathsf{cost}_i(\sigma[i \to \sigma'_i], \rho'_{\leq l+1}).$$

Moreover, we have $\inf_{\sigma'_i \in \mathfrak{S}} \mathsf{cost}_i(\pi[i \to \sigma'_i], \rho'_{\leq l+1}) = \inf_{\sigma'_i \in \mathfrak{S}} \mathsf{cost}_i(\pi[i \to \sigma'_i], \rho'(l+1))$ since memoryless strategies suffice to minimize the cost [12]. We then have

$$\inf_{\sigma'_i \in \mathfrak{S}} \mathsf{cost}_i(\pi[i \to \sigma'_i], \rho'(l+1)) \leq \sup_{\sigma_{-i} \in \mathfrak{S}^{n-1}} \inf_{\sigma_i \in \mathfrak{S}} \mathsf{cost}_i((\sigma_{-i}, \sigma_i), \rho'(l+1)).$$

We obtain the required inequality

$$\begin{aligned}\mathsf{cost}_i(\rho_{\geq l}) &\leq \mathsf{cost}_i(\rho(l), \rho'(l+1)) + \sup_{\sigma_{-i} \in \mathfrak{S}^{n-1}} \inf_{\sigma_i \in \mathfrak{S}} \mathsf{cost}_i((\sigma_{-i}, \sigma_i), \rho'(l+1)) \\ &\leq \mathsf{cost}_i(\rho(l), c) + \mathsf{val}_{i,c}.\end{aligned}$$

Conversely, consider a path ρ that satisfies the condition. We are going to construct a Nash equilibrium having outcome ρ. The idea is that players will follow ρ, and if some player i deviates, then the coalition $-i$ will apply a joint strategy to maximize the cost of Player i, thus achieving at least $\mathsf{val}_{i,c}$, where c is the first configuration where deviation is detected.

Let us define the punishment function $\mathcal{P}_\rho \colon \mathsf{Paths}(\langle \mathcal{A}, n\rangle) \to [\![n]\!] \cup \{\bot\}$ which keeps track of the deviating players and the step where such a player has deviated. For path $h' = h(c, w, c')$, we write

$$\mathcal{P}_\rho(h') = \begin{cases} \bot & \text{if } h' \leq_{\mathsf{pref}} \rho, \\ i & \text{if } h \leq_{\mathsf{pref}} \rho, h(c,w,c') \not\leq_{\mathsf{pref}} \rho, \text{ and } i \in [\![n]\!] \text{ min. s.t. } c'(i) \neq \rho(|h|+1)(i), \\ \mathcal{P}_\rho(h) & \text{otherwise.} \end{cases}$$

Intuitively, $\bot$ means that no players have deviated from ρ in the current path. If $\mathcal{P}_\pi(h) = j$, then Player j was among the first players to deviate from ρ in the path h; so for some l,

$h(l)(j) = \rho(l)(j)$ but $h(l+1)(j) \neq \rho(l+1)(j)$. Notice that if several players deviate at the same step, there are no conditions to be checked, and the strategy can be chosen arbitrarily. For each configuration c and coalition $-i$, let $\sigma_{-i,c}$ be the strategy of coalition $-i$ maximizing the cost of Player i from configuration c; thus achieving at least $\mathsf{val}_{i,c}$. Player j's strategy in this coalition, for $j \neq i$, is denoted $\sigma_{-i,c,j}$. For path $h' = h(c, w, c')$, define

$$\tau_i(h') = \begin{cases} (c'(i), m(i), c''(i)) & \text{if } \mathcal{P}_\rho(h') = \bot, \rho(|h'|+1) = (c', w', c''), \\ & \text{ and } m \in E^n \text{ is such that } T(c', m) = (w', c''), \\ \text{arbitrary} & \text{if } \mathcal{P}_\rho(h') = i, \\ \sigma_{-j,c,i}(h') & \text{if } \mathcal{P}_\rho(h') = j \text{ for some } j \neq i. \end{cases}$$

The first case ensures that the outcome of the profile $(\tau_i)_{i \in [\![n]\!]}$ is ρ. The third case means that Player i follows the coalition strategy $\sigma_{-j,c}$ after Player j has deviated to configuration c. The second case corresponds to the case where Player i has deviated: the precise definition of this part of the strategy is irrelevant.

Let us show that this profile is indeed a Nash equilibrium. Consider any player $j \in [\![n]\!]$ and any strategy τ'_j. Let ρ' denote the outcome of (τ_{-j}, τ'_j), and l the index of the last configuration where ρ and ρ' are identical. We have

$$\begin{aligned} \mathsf{cost}_j((\tau_{-j}, \tau'_j)) &= \mathsf{cost}_j(\rho_{\leq l}) + \mathsf{cost}_j(\rho(l), \rho'(l+1)) + \mathsf{cost}_j((\tau_{-j}, \tau_j), \rho'_{\leq l+1}) \\ &\geq \mathsf{cost}_j(\rho_{\leq l}) + \mathsf{cost}_j(\rho(l), \rho'(l+1)) + \mathsf{val}_{j, \rho'(l+1)(j)} \\ &\geq \mathsf{cost}_j((\tau_i)_{i \in [\![n]\!]}), \end{aligned}$$

where the second line follows from the fact that the coalition switches to a strategies ensuring a cost of at least $\mathsf{val}_{j,\rho'(l)(j)}$ at step l; and the third line is obtained by assumption. This shows that $(\tau_i)_{i \in [\![n]\!]}$ is indeed a Nash equilibrium and concludes the proof. ◀

Algorithm. We define a graph that describes the set of outcomes of Nash equilibria by augmenting the n-weighted configuration graph $\mathcal{M} = \langle C, T \rangle$. For any real vector $\vec{\gamma} = (\gamma_i)_{i \in [\![n]\!]}$, we define the weighted graph $\mathcal{G}_{\langle \mathcal{A}, n \rangle, \vec{\gamma}} = \langle C', T' \rangle$ with $C' = C \times ([\![Y]\!] \cup \{0, \infty\})^n$ where $Y = |V| \cdot \kappa$, and $T' \subseteq C' \times \mathbb{N} \times C'$; remember that all players have a strategy realizing cost at most Y in $\langle \mathcal{A}, n \rangle$. The initial state is $(c_{\mathsf{src}}, \infty^n)$. The set of transitions T' is defined as follows: $((c, b), z, (c', b')) \in T'$ if, and only if, there exists $(c, w, c') \in T$, $z = \vec{\gamma} \cdot w$ (where $\cdot$ is dot product), and for all $i \in [\![n]\!]$,

$$b'_i = \min(b_i - w_i, \min_{c'' \in \mathsf{dev}_i(c(i), c'(i))} \mathsf{cost}_i(c, c'') + \mathsf{val}_{i,c''} - w_i). \qquad (1)$$

Notice that by definition of C', b'_i must be nonnegative for all $i \in [\![n]\!]$, so there are no transitions $((c, b), z, (c', b'))$ if the above expression is negative for some i. Notice also that the size of $\mathcal{G}_{\langle \mathcal{A}, n \rangle, \vec{\gamma}}$ is doubly-exponential in that of the input $\langle \mathcal{A}, n \rangle$, since this is already the case for C, while Y is singly-exponential.

Intuitively, for any path ρ that visits some state (c, b) in this graph, in order for ρ to be compatible with a Nash equilibrium, each player i must have cost no more than b_i in the rest of the path. In fact, the second term of the minimum in (1) is the least cost Player i could guarantee by not following (c, w, c') but going to some other configuration $c'' \in \mathsf{dev}_i(c, c')$, so the bound b_i is used to guarantee that these deviations are not profitable. The definition of b'_i in (1) is the minimum of $b_i - w_i$ and the aforementioned quantity since we check both the previous bound b_i, updated with the current cost w_i (which gives the left term), and the non-profitability of a deviation at the previous state (which is the right term). If this minimum becomes negative, this precisely means that at an earlier point in the current path, there was a strategy for Player i which was more profitable than the current path regardless of the strategies of other players; so the current path cannot be the outcome of a Nash equilibrium. This is why the definition of $\mathcal{G}_{\langle \mathcal{A},n\rangle,\vec{\gamma}}$ restricts the state space to nonnegative values for the b_i.

We prove that computing the cost of a Nash equilibrium minimizing the $\vec{\gamma}$-weighted social cost reduces to computing a shortest path in $\mathcal{G}_{\langle \mathcal{A},n\rangle,\vec{\gamma}}$. In particular, letting $\gamma_i = 1$ for all $i \in [\![n]\!]$, a $\vec{\gamma}$-minimal Nash equilibrium is a best Nash equilibrium (minimizing the social cost), while taking $\gamma_i = -1$ for all $i \in [\![n]\!]$, we get a worst Nash equilibrium (maximizing the social cost).

▶ **Theorem 14.** *For any dynamic NCG $\langle \mathcal{A}, n\rangle$ and vector $\vec{\gamma}$, the cost of the shortest path from $(c_{\mathsf{src}}, \infty^n)$ to some (c_{tgt}, b) in $\mathcal{G}_{\langle \mathcal{A},n\rangle,\vec{\gamma}}$ is the cost of a $\vec{\gamma}$-minimal Nash equilibrium.*

Proof. We show that for each path of $\langle \mathcal{A}, n\rangle$ from c_{src} to c_{tgt}, there is a path in $\mathcal{G}_{\langle \mathcal{A},n\rangle,\vec{\gamma}}$ from $(c_{\mathsf{src}}, \infty^n)$ to some (c_{tgt}, b) with the same cost, and vice versa.

Consider a Nash equilibrium $\pi = (\sigma_j)_{j\in[\![n]\!]}$ with outcome $\rho = (c_j, w_j, c_{j+1})_{1\le j<l}$. We build a sequence $b_1, b_2, \ldots$ such that $\rho' = ((c_j, b_j), \vec{\gamma}\cdot w_j, (c_{j+1}, b_{j+1}))_{1\le j<l}$ is a path of $\mathcal{G}_{\langle \mathcal{A},n\rangle,\vec{\gamma}}$. We set $b_1(j) = \infty$ for all $j \in [\![n]\!]$. For $j \ge 1$, define

$$b_{j+1}(i) = \min\left(b_j(i) - w_j(i), \min_{c''\in \mathsf{dev}_j(c_j(i),c_{j+1}(i))} \mathsf{cost}_i(c_j, c'') + \mathsf{val}_{i,c''} - w_j(i)\right).$$

We are going to show that for all $1 \le j \le l$, $\mathsf{cost}_i(\rho_{\ge j}) \le b_j$, which shows that $b_j \ge 0$, and thus ρ' is a path of $\mathcal{G}_{\langle \mathcal{A},n\rangle,\vec{\gamma}}$.

We show this by induction on j. This is clear for $j = 1$. Assume this holds up to $j \ge 1$. We have, by induction that $\mathsf{cost}_i(\rho_{\ge j}) \le b_j(i)$ for all $i \in [\![n]\!]$. Moreover, since π is a Nash equilibrium, by Lemma 13,

$$\forall i \in [\![n]\!], \mathsf{cost}_i(\rho_{\ge j}) \le \min_{c''\in \mathsf{dev}_i(\rho(j),\rho(j+1))} \mathsf{val}_{i,c''} + \mathsf{cost}_i(\rho(j), c'').$$

Therefore,

$$\begin{aligned}\mathsf{cost}_i(\rho_{\ge j+1}) &= \mathsf{cost}_i(\rho_{\ge j}) - w_j(i)\\ &\le \min(b_j(i) - w_j(i), \min_{c''\in \mathsf{dev}_i(\rho(j),\rho(j+1))} \mathsf{val}_{i,c''} + \mathsf{cost}_i(\rho(j), c'') - w_j(i))\end{aligned}$$

as required, and both paths have the same $\vec{\gamma}$-weighted cost.

Consider now a path $((c_i, b_i), z_i, (c_{i+1}, b_{i+1}))_{1\le i<l}$ in $\mathcal{G}_{\langle \mathcal{A},n\rangle,\vec{\gamma}}$. By the definition of $\mathcal{G}_{\langle \mathcal{A},n\rangle,\vec{\gamma}}$, there exists $w_1, w_2, \ldots$ such that $\rho = (c_j, w_j, c_{j+1})_{1\le j<l}$ is a path of $\langle \mathcal{A}, n\rangle$, and $z_j = \vec{\gamma}\cdot w_j$. So it only remains to show that that ρ is the outcome of a Nash equilibrium. We will show that ρ satisfies the criterion of Lemma 13. We show by backwards induction on $1 \le j \le l$ that for all $i \in [\![n]\!]$,

1. $\mathsf{cost}_i(\rho_{\ge j}) \le b_j(i)$,
2. $\mathsf{cost}_i(\rho_{\ge j}) \le \min_{c''\in \mathsf{dev}_i(\rho(j),c'')} \mathsf{cost}_i(\rho(j), c'') + \mathsf{val}_{i,c''}$.

For $j = l$, we have $\mathsf{cost}_i(\rho_{\ge l}) = 0$ so this is trivial. Assume the property holds down to $j + 1$ for some $1 \le j < l$. By induction hypothesis, we have

$$\mathsf{cost}_i(\rho_{\ge j+1}) \le b_{j+1}(i) = \min\left(b_j(i) - w_j(i), \min_{c''\in \mathsf{dev}_i(\rho(j),c'')} \mathsf{cost}_i(\rho(j), c'') + \mathsf{val}_{i,c''} - w_j(i)\right).$$

Therefore,

$$\mathsf{cost}_i(\rho_{\ge j}) = \mathsf{cost}_i(\rho_{\ge j+1}) + w_j(i) \le \min\left(b_j(i), \min_{c''\in \mathsf{dev}_i(\rho(j),c'')} \mathsf{cost}_i(\rho(j), c'') + \mathsf{val}_{i,c''}\right),$$

as required. By Lemma 13, ρ is the outcome of a Nash equilibrium. ◀

Thanks to Theorem 14, we can compute the costs of the best and worst NEs of $\langle\mathcal{A}, n\rangle$ in exponential space. We can also decide the existence of an NE with constraints on the costs (both social and individual), by non-deterministically guessing an outcome and checking in $\mathcal{G}_{\langle\mathcal{A},n\rangle,\vec{\gamma}}$ that it is indeed an NE. We obtain the following conclusion:

▶ **Corollary 15.** *In dynamic NCGs, the constrained Nash-equilibrium problem is in* EXPSPACE*.*

Proof. As noted earlier, the number of vertices in $\mathcal{G}_{\langle\mathcal{A},n\rangle,\vec{\gamma}}$ is doubly exponential since $|C| = |V|^n$ is doubly exponential. Storing a configuration and computing its successors can be performed in exponential space. One can thus guess a path of size at most the size of the graph and check whether its cost is less than the given bound. This can be done using exponential-space counters, and provides us with an EXPSPACE algorithm. ◀

Note that one can effectively compute a Nash-equilibrium strategy profile satisfying the constraints in doubly-exponential time by finding the shortest path of $\mathcal{G}_{\langle\mathcal{A},n\rangle,\vec{\gamma}}$, and applying the construction of (the proof of) Lemma 13.

▶ Remark 16. The exponential complexity is due to the encoding of the number of players in binary. If we consider asymmetric NCGs, in which the source-target pairs would be given explicitly for all players, the size of $\mathcal{G}_{\langle\mathcal{A},n\rangle,\vec{\gamma}}$ would be singly-exponential, and the constrained Nash-equilibrium problem would be in PSPACE.

5 Subgame-perfect equilibria

In this section, we characterize the outcomes of SPEs and decide the existence of SPEs with constraints on the social cost. We follow the approach of [9], extending it to the setting of concurrent weighted games, which we need to handle dynamic NCGs.

Characterization of outcomes of SPE. Consider a dynamic NCG $\langle\mathcal{A}, n\rangle$, and the associated configuration graph $\mathcal{M} = \langle C, T\rangle$. We partition the set C of configurations into $(X_j)_{0\le j\le n}$ such that a configuration c is in X_j if, and only if, $j = \#\{i \in [\![n]\!] \mid c(i) = \mathsf{tgt}\}$. Since tgt is a sink state in $\mathcal{A}$, if there is a transition from some configuration in X_j to some configuration in X_k, then $k \ge j$. We define $X_{\ge j} = \bigcup_{i\ge j} X_i$, $Z_j = \{(c, w, c') \in T \mid c \in X_j\}$ and $Z_{\ge j} = \{(c, w, c') \in T \mid c \in X_{\ge j}\}$.

Following [9], we inductively define a sequence $(\lambda^{j^*})_{0\le j\le n}$, where each $\lambda^{j^*} = \langle\lambda_i^{j^*}\rangle_{i\in[\![n]\!]}$ is a n-tuple of labeling functions $\lambda_i^{j^*} : Z_{\ge j} \to \mathbb{N} \cup \{-\infty, +\infty\}$. This sequence will be used to characterize outcomes of SPEs through the notion of λ-*consistency*:

▶ **Definition 17.** *Let $j \le n$, and $\lambda = (\lambda_i)_{i\in[\![n]\!]}$ be a family of functions such that $\lambda \colon Z_{\ge j} \to \mathbb{N} \cup \{-\infty, +\infty\}$ Let $c \in X_{\ge j}$. A finite path $\rho = (t_k)_{1\le k<|\rho|}$ from c ending in c_{tgt} is said to be λ-consistent whenever for any $i \in [\![n]\!]$ and any $1 \le k < |\rho|$, it holds $\mathsf{cost}_i(\rho_{\ge k}) \le \lambda_i(t_k)$. We write $\Gamma_\lambda(c)$ for the set of all λ-consistent paths from c.*

We now define λ^{j^*} for all $0 \le j \le n$ in such a way that, for all $c \in X_{\ge j}$, $\Gamma_{\lambda^{j^*}}(c)$ is the set of all outcomes of SPEs in the subgame rooted at c. The case where $j = n$ is simple: we have $X_{\ge n} = \{c_{\mathsf{tgt}}\}$ and $Z_{\ge n} = \{(c_{\mathsf{tgt}}, 0^n, c_{\mathsf{tgt}})\}$; there is a single path, which obviously is the outcome of an SPE since no deviations are possible. For all $i \in [\![n]\!]$, we let $\lambda_i^{n^*}(c_{\mathsf{tgt}}, 0^n, c_{\mathsf{tgt}}) = 0$.

Now, fix $j < n$, assuming that $\lambda^{(j+1)^*}$ has been defined. In order to define λ^{j^*}, we introduce an intermediary sequence $(\mu_i^k)_{k\geq 0, i\in[\![n]\!]}$, with $\mu_i^k\colon Z_{\geq j} \to \mathbb{N}\cup\{-\infty,+\infty\}$, of which $(\lambda_i^{j^*})_{i\in[\![n]\!]}$ will be the limit.

Functions μ_i^k mainly operate on $Z_j = Z_{\geq j}\setminus Z_{\geq j+1}$: for any $\mathbf{e}\in Z_{\geq j+1}$, we let $\mu_i^k(\mathbf{e}) = \lambda_i^{(j+1)^*}(\mathbf{e})$. Now, for $\mathbf{e} = (c,w,c')\in Z_j$, $\mu_i^k(e)$ is defined inductively as follows:

- $\mu_i^0(\mathbf{e}) = 0$ if $c(i) = \mathsf{tgt}$, and $\mu_i^0(\mathbf{e}) = +\infty$ otherwise;
- for $k > 0$, μ^k is defined from μ^{k-1} following three cases: if $c(i) = \mathsf{tgt}$, then $\mu_i^k(\mathbf{e}) = 0$; if $\Gamma_{\mu^{k-1}}(c') = \emptyset$ for some $(c,w',c')\in T$, then $\mu_i^k(\mathbf{e}) = -\infty$; otherwise,

$$\mu_i^k(\mathbf{e}) = \min_{c''\in\mathsf{dev}_i(c,c')}\ \sup_{\rho\in\Gamma_{\mu^{l-1}}(c'')} (\mathsf{cost}_i(c,c'') + \mathsf{cost}_i(\rho))$$

We can then prove that for any $e\in Z_{\geq j}$ and any $k>0$, $\mu_i^k(e) \geq \mu_i^{k-1}(e)$. It follows that the sequence $(\mu^k)_{k\geq 0}$ stabilizes, and we can define λ^{j^*} as its limit. Let $\Gamma^* = \Gamma_{\lambda^{0^*}}$. Then:

▶ **Theorem 18.** *A path ρ in $\mathcal{G} = \langle\mathcal{A}, n\rangle$ is the outcome of an SPE if, and only if, $\rho\in\Gamma^*(c_{\mathsf{src}})$.*

Algorithm. It remains to compute the sequence $(\mu^k)_{k\geq 0}$ (which will include checking non-emptiness of the corresponding Γ-sets), and to bound the stabilization time. To this aim, with any family $\mu = (\mu_i)_{i\in[\![n]\!]}$ of functions as above and any configuration c, we associate an infinite-state *counter graph* $\mathbb{C}[\mu, c] = \langle C', T'\rangle$ to capture all μ-consistent paths from c:

- the set of vertices is $C' = C\times(\mathbb{N}\cup\{+\infty\})^{[\![n]\!]}$;
- T' contains all edges $((d,b),w,(d',b'))$ for which (d,w,d') is an edge of $\mathcal{M}$ and for all $i\in[\![n]\!]$, $b'(i) = 0$ if $d(i) = \mathsf{tgt}$, and $b'_i = \min\{b_i - w_i, \mu_i(d,w,d') - w_i\}$ otherwise (provided that $b'_i \geq 0$ for all i, in order for (d',b') to be an edge of $\mathbb{C}[\mu,c]$).

With the initial configuration c, we associate b^c such that $b^c_i = 0$ if $c(i) = \mathsf{tgt}$ and $b^c_i = +\infty$ otherwise: this configuration imposes no constraint, since no edges has been taken yet. Intuitively, in configuration (d,b), b is used to enforce μ-consistency: each edge taken along a path imposes a constraint on the cost of the players for the rest of the path; this constraint is added to the constraints of the earlier edges, and propagated along the path. We can prove that the number of reachable states from (c,b^c) in $\mathbb{C}[\mu,c]$, which we denote with $|C'|_r$, is bounded by $|C|\cdot(n\cdot|V|^n\cdot\kappa)^{|V|}$.

Computing λ^{j^*} from $\lambda^{(j+1)^*}$ amounts to inductively computing $(\mu_i^{k+1})_{i\in[\![n]\!]}$ from μ^k for edges $e = (c,w,c')\in Z_j$, until stabilization. Since $\mathbb{C}[\mu^k,d]$ can be proved to capture μ^k-consistent paths from d, the computation mainly amounts to checking the existence of paths in such counter graphs, which can be performed in doubly-exponential space. Stabilization can be shown to occur within $|V|(1+n\cdot\kappa\cdot|E|^n)$ steps. In the end:

▶ **Theorem 19.** *The existence of SPEs in a dynamic NCG can be decided in 2*EXPSPACE.

▶ Remark 20. Again, our algorithm is not specific to the symmetric setting of our dynamic NCGs; in an asymmetric context, where the number of players would be given in unary, our algorithm would run in EXPSPACE.

Existence of constrained SPEs. The algorithm above can be extended to compute the cost of the best and worst SPEs, and to include constraints on the costs (both social and individual) of the SPEs we are looking for.

First, for any vector $\vec{\gamma} = (\gamma_i)_{i\in[\![n]\!]}$, we define the $\vec{\gamma}$*-counter graph* $\mathbb{C}[\lambda^*, c_{\mathsf{src}}, \vec{\gamma}]$, which is obtained from $\mathbb{C}[\lambda^*, c_{\mathsf{src}}]$ by replacing the cost vector w on the edges with $\vec{\gamma}\cdot w$.

We can then compute the cost of a $\vec{\gamma}$-minimal SPE by checking existence of a path from $(c_{\mathsf{src}}, b^{c_{\mathsf{src}}})$ to (c_{tgt}, b) in $\mathbb{C}[\lambda^*, c_{\mathsf{src}}, \vec{\gamma}]$, which minimizes the $\vec{\gamma}$-weighted social cost. Again, letting $\gamma_i = 1$ for all $i \in [\![n]\!]$, a $\vec{\gamma}$-minimal SPE is a best SPE, while taking $\gamma_i = -1$ for all $i \in [\![n]\!]$, we get a worst SPE (maximizing social cost).

We can also solve the constrained-SPE-existence problem by non-deterministically guessing an outcome and checking that it is a path in $\mathbb{C}[\lambda^{0^*}, c_{\mathsf{src}}]$ and that it satisfies the constraints. In each case, we can inductively build a strategy profile witnessing the fact that the selected path is the outcome of an SPE.

6 Conclusion and future works

In this paper, we introduced dynamic network congestion games, and studied the complexity of various decision and computation problems concerning social optima, Nash equilibria and subgame perfect equilibria. Our algorithms allow us to compute the price of anarchy and price of stability for those games.

There are couple of areas that are left open in our discussion: possibly the foremost one being the complexity gaps of the decision problems we talked about. As of yet, we do not have interesting lower bounds for constraint NE or constraint SPE problem, so definitely one direction is there for completing the picture. Another aspect of what we do not address in this paper is to obtain *bounds* on PoA/PoSs of our model. Even though we are specifically interested in the measure(s) for a given instance, nonetheless obtaining such bounds could be interesting.

What we are mostly interested in as future work, is to compute how the price of anarchy and the price of stability (and costs of equilibria and social optimum) evolve when the number of players, seen as a parameter, grows.

References

1 Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.

2 Guy Avni, Shibashis Guha, and Orna Kupferman. Timed network games. In Kim Guldstrand Larsen, Hans L. Bodlaender, and Jean-François Raskin, editors, *Proceedings of the 42nd International Symposium on Mathematical Foundations of Computer Science (MFCS'17)*, volume 84 of *Leibniz International Proceedings in Informatics*, pages 37:1–37:16. Leibniz-Zentrum für Informatik, August 2017.

3 Guy Avni, Shibashis Guha, and Orna Kupferman. Timed network games with clocks. In Igor Potapov, Paul G. Spirakis, and James Worrell, editors, *Proceedings of the 43rd International Symposium on Mathematical Foundations of Computer Science (MFCS'18)*, volume 117 of *Leibniz International Proceedings in Informatics*, pages 23:1–23:18. Leibniz-Zentrum für Informatik, August 2018.

4 Guy Avni, Thomas A. Henzinger, and Orna Kupferman. Dynamic resource allocation games. *Theoretical Computer Science*, 807:42–55, February 2020.

5 Baruch Awerbuch, Yossi Azar, and Amir Epstein. Large the price of routing unsplittable flow. In Harold N. Gabow and Ronald Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on the Theory of Computing (STOC'05)*, pages 57–66. ACM Press, May 2005.

6 Nathalie Bertrand, Nicolas Markey, Suman Sadhukhan, and Ocan Sankur. Dynamic network congestion games. *CoRR*, abs/2009.13632, 2020.

7 Dimitri P. Bertsekas and John N. Tsitsiklis. *Parallel and distributed computation: numerical methods.* Prentice Hall, 1989.

8 Umang Bhaskar, Lisa Fleischer, and Elliot Anshelevich. A Stackelberg strategy for routing flow over time. *Games and Economic Behavior*, 92:232–247, July 2015.

9 Thomas Brihaye, Véronique Bruyère, Aline Goeminne, Jean-François Raskin, and Marie Van den Bogaard. The complexity of subgame perfect equilibria in quantitative reachability games. In Wan J. Fokkink and Rob van Glabbeek, editors, *Proceedings of the 30th International Conference on Concurrency Theory (CONCUR'19)*, volume 140 of *Leibniz International Proceedings in Informatics*, pages 13:1–13:16. Leibniz-Zentrum für Informatik, August 2019.

10 Thomas Brihaye, Véronique Bruyère, Noëmie Meunier, and Jean-François Raskin. Weak subgame perfect equilibria and their application to quantitative reachability. In Stephan Kreutzer, editor, *Proceedings of the 24th EACSL Annual Conference on Computer Science Logic (CSL'15)*, volume 41 of *Leibniz International Proceedings in Informatics*, pages 504–518. Leibniz-Zentrum für Informatik, September 2015.

11 Martin Hoefer, Vahab S. Mirrokni, Heiko Röglin, and Shang-Hua Teng. Competitive routing over time. *Theoretical Computer Science*, 412(39):5420–5432, September 2011.

12 Leonid Khachiyan, Endre Boros, Konrad Borys, Khaled Elbassioni, Vladimir Gurvich, Gabor Rudolf, and Jihui Zhao. On short paths interdiction problems: Total and node-wise limited interdiction. *Theory of Computing Systems*, 43(2):204–233, August 2008.

13 Ronald Koch and Martin Skutella. Nash equilibria and the price of anarchy for flows over time. *Theory of Computing Systems*, 49(1):71–97, July 2011.

14 Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, May 2009.

15 Elias Koutsoupias and Katia Papakonstantinopoulou. Contention issues in congestion games. In Artur Czumaj, Kurt Mehlhorn, Andrew Pitts, and Roger Wattenhofer, editors, *Proceedings of the 39th International Colloquium on Automata, Languages and Programming (ICALP'12) – Part II*, volume 7392 of *Lecture Notes in Computer Science*, pages 623–635. Springer-Verlag, July 2012.

16 Carol A. Meyers and Andreas S. Schulz. The complexity of welfare maximization in congestion games. *Networks*, 59(2):252–260, March 2012.

17 Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, May 1996.

18 Michal Penn, Maria Polukarov, and Moshe Tennenholtz. Random order congestion games. *Mathematics of Operations Research*, 34(3):706–725, August 2009.

19 Lili Qiu, Richard Yang, Yin Zhang, and Scott Shenker. On selfish routing in internet-like environments. *IEEE Transactions on Computers*, 14(4):725–738, August 2006.

20 Robert W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, December 1973.

21 Tim Roughgarden. Routing games. In Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, chapter 18, page 461–486. Cambridge University Press, 2007.

22 Subhash Suri, Csaba D. Tóth, and Yunhong Zhou. Selfish load balancing and atomic congestion games. *Algorithmica*, 47(1):79–96, 2007.

A Proofs of Section 3

▶ **Theorem 6.** *A socially-optimal strategy profile can be computed in exponential time. The constrained social-optimum problem is in PSPACE and NP-hard.*

Proof. A socially-optimal strategy can be obtained by computing a shortest path in $\mathcal{P}$ from $\bar{c}_{\mathsf{src}}$ to $\bar{c}_{\mathsf{tgt}}$. The graph $\mathcal{P}$ has $O(n^{|V|})$ states, and shortest paths can be computed in polynomial time. We can thus compute the (abstract) outcome of a socially-optimal strategy profile in exponential time. We easily derive a socially-optimal (blind) strategy profile.

An abstract configuration can be stored in space $O(|V| \cdot \log(n))$, and deciding whether there is an edge of a given weight between two abstract configurations can be checked in PSPACE. Since shortest paths in $\mathcal{P}$ have length at most $n \cdot |V|$, a non-deterministic polynomial-space algorithm can guess a path step-by-step in $\mathcal{P}$, and check that it reaches the target configuration with social-cost at most b.

By Savitch's theorem, this proves PSPACE membership.

To prove the NP-hardness we provide a reduction from the Partition problem which, given a family of integer numbers $(r_i)_{1 \le i \le m}$ that sum up to $2S$, asks whether the family can be split into two subfamilies that both sum up to S. Consider an instance L of Partition, and let $M = 14S + 12m + 1$, and $n = 2S + 2m$. For any $r \in \mathbb{N}$, we define threshold cost function T_r as $\mathsf{T}_r(i) = 1$ if $i \le r$, and $\mathsf{T}_r(i) = M$ otherwise We construct a dynamic NCG $\mathcal{A}_L$ with states $V = \{\mathsf{src}, \mathsf{tgt}, d_1, d_2\} \cup \{s_i, a_{i,1}, a_{i,2}\}_{i \in [\![m]\!]}$. The transitions are defined as follows:

- for each $i \in [\![m]\!]$, there is a transition from src to s_i which cost function T_{r_i+2};
- from each s_i, there is a transition to $a_{i,1}$ and another one to $a_{i,2}$ with the same cost function that assigns cost 2 for one player, and 4 for more;
- for $i \in [\![m]\!]$ and $j \in \{1, 2\}$, there is a transition from $a_{i,j}$ to d_j with constant cost 1, and a transition to tgt with cost function T_1;
- from each d_j, there is a transition to tgt with cost function T_S.

The construction is illustrated in Figure 3.

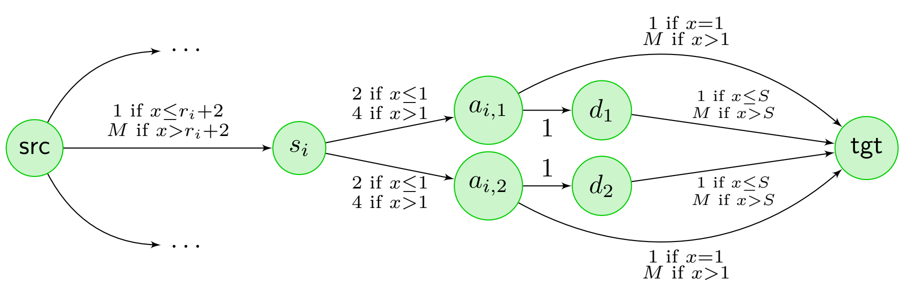

Figure 3 The reduction of Theorem 6. The figure shows the edges involving states $s_i, a_{i,1}, a_{i,2}$. The full graph is obtained by reproducing the shown construction for all $i \in [\![m]\!]$.

We prove that there exists a strategy profile in $\langle \mathcal{A}, n \rangle$ with social cost less than M if, and only if, L is a positive instance of the partition problem. Assume that the instance L has a solution $\langle L_1, L_2 \rangle$. We describe the behaviour of the strategy profile in $\mathcal{A}_L$ achieving social cost less than M.

- Initially, all players are at src. During the first step, for each $i \in [\![m]\!]$, $(r_i + 2)$ players move to s_i. This incurs a total cost of $\sum_{i \in [\![m]\!]} r_i + 2 = \sum_{i \in [\![m]\!]} r_i + 2 \times m = 2S + 2m$;
- Consider $i \in [\![m]\!]$, and let $j \in \{1, 2\}$ be such that $r_i \in L_j$. From state s_i we let $r_i + 1$ players move to $a_{i,j}$, and one player to $a_{i,3-j}$. The $r_i + 1$ players each pay a cost of 4,

and other player a cost of 2. Overall, we get a total cost of $\sum_{i\in[\![m]\!]}((r_i+1)\times 4+1\times 2) = 8S+4m+2m = 8S+6m$;

- From each $a_{i,j}$, one player takes the transition directly to tgt. All other players move to d_j. Overall, $2m$ players directly move to tgt, while all other players, that are $2S$ many of them, move to $\{d_1, d_2\}$. The total cost of this step is therefore $2S+2m\times 2 = 2S+4m$;
- From each d_j, players necessarily move directly to tgt. For each $i \in [\![m]\!]$ and $j \in \{1,2\}$, exactly r_i players arrive to d_j from $a_{i,j}$. Thus there are exactly S players in each d_j, so that the total cost of this step is $2S$.

Summing over all steps, the social cost of this strategy profile is $14S + 12m < M$.

Now for the opposite direction, consider any strategy profile π in $\langle \mathcal{A}, n\rangle$ with social cost less than M. We are going to construct a partition of L. First we divide the path taken by players from src to tgt in three phases: **Phase 1:** from src to some s_i; **Phase 2:** from some s_i to some $a_{i,j}$; **Phase 3:** from $a_{i,j}$ to tgt, either directly or via d_j. We now analyze the cost incurred by players under profile π in all three phases.

In phase 1, each player pays either cost 1 or cost M. By assumption on π, all players must have a cost of 1, and this is only possible if r_i+2 players move to s_i, for each $i \in [\![m]\!]$. The total cost of this phase is thus $2S+2m$.

In phase 3, a player pays a cost of either 2, M or $M+1$. For the latter two cases are not possible by assumption. So phase 3's contribution to the social cost has to be $2\times n = 4S+4m$. Then, the social cost of phase 2 is strictly less than $(14S+12m+1)-(2S+2m)-(4S+4m) = 8S+6m+1$. By phase 1, there are r_i+2 players at each s_i, so minimum contribution to the social cost for these r_i+2 players is $2\times 1+4\times(r_i+2) = 4r_i+6$. Summing overall $i\in[\![m]\!]$, this yields $\sum_{r_i\in L}(4r_i+6) = 8S+6m$. Thus, the social cost of phase 2 is $8S+6m$. But this cost is achieved only when from each s_i, one transition is taken by one player and the other one by r_i+1 players.

Therefore, each $a_{i,j}$ contains either one player or (r_i+1) players under π. Given the cost functions, at most $2S$ players can move to tgt via d_1 or d_2 (otherwise players would get a cost of M). So $2m$ remaining players must take the transition from some $a_{i,j}$ to tgt. But at each $a_{i,j}$, there is a unique player that takes this transition due to the cost function. If $a_{i,j}$ contains r_i+1 players, r_i players take the path via d_j, each paying a cost of 2. As there are in total $2S$ players taking the route via some d_j, and each d_j can contain at most S players because of the cost functions, there must be exactly S many players which arrive at each d_j under $\mathcal{P}$. That is, for each i, j, there are exactly r_i many players coming from some $a_{i,j}$ to d_j, and their total is S. This defines the required partition of L. ◀

▶ **Lemma 7.** *There is a shortest path (w.r.t. cost) in $\mathcal{P}$ with size (in terms of its number of transitions) at most $n\cdot|V|$.*

Proof. We begin with proving that, for any path ρ in the multi-weighted graph $\mathcal{M}$ in which no player reaches the target state during the first n steps, there is a path ρ' with social cost at most the social cost of ρ in which at least one player reaches the target state. This obviously extends to $\mathcal{P}$.

To this aim, we consider, for each state s of $\mathcal{A}$, the shortest path from s to tgt for a single player (i.e., when taking for each edge the value of the cost function when a single player takes that edge). Clearly enough, this is the best a player can hope from state s. Now, we label all state s with the number $\mathsf{opt}(s)$ of edges of the shortest (in terms of the number of edges) of the shortest (in terms of its single-player cost) paths. Then those values are at most $|V|-1$, and $\mathsf{opt}(s) = 0$ if, and only if, $s = \mathsf{tgt}$.

Now, consider a path ρ in $\mathcal{M}$, in which we assume no player reaches tgt during the first $|V|$ steps. We build a table H of size $n \times |V|$, in which cell (i, j) contains $j + \mathsf{opt}(\rho(j)(i))$. Then for any i, $H(i, 1)$ is at most $|V|$, and $H(i, |V|)$ is at least $|V| + 1$, since we assume no players have reached tgt after $|V|$ steps. Hence there is an index $1 \leq j_0 < |V|$ that contains an element smaller than or equal to $|V|$, and being the largest so. Pick i_0 such that $H(i_0, j_0) \leq |V|$, and consider the path ρ' in $\mathcal{M}$ obtained from ρ by making Player i_0 follow an optimal path from $\rho(j_0)(i_0)$. Along this new path ρ', Player i_0 will reach tgt in $H(i_0, j_0)$ steps, which is less than or equal to $|V|$. Moreover, after step j_0, Player i_0 can be sure no other player takes the same edges, since this would mean that for some player i_1, at some step $j_1 > j_0$, we have $H(i_1, j_1) \leq |V|$. Hence the cost for Player i_0 is at least as good along ρ' as it was along ρ; also, the cost of all other players may only have decreased. It follows that the social cost of ρ' is at least as good as that of ρ.

Now, consider a shortest path in $\mathcal{P}$, and a corresponding path ρ in $\mathcal{M}$. Applying the arguments above, we can build a path ρ' in $\mathcal{M}$ with social cost at least as good as that of ρ, and in which at least one player reaches tgt. We can then apply the same arguments recursively to the other players, each in a slide of at most $|V|$ steps. In the end, we get a path with minimal social cost and of length at most $n \times |V|$. ◀

B Proofs of Section 4

▶ **Lemma B.1.** *In every dynamic NCG, there exists a blind Nash equilibrium.*

Proof. Consider a dynamic NCG $\langle \mathcal{A}, n \rangle$ with $\mathcal{A} = \langle V, E, \mathsf{src}, \mathsf{tgt} \rangle$, and a blind strategy profile $\pi = (\sigma_i)_{i \in [\![n]\!]}$. Recall that each blind strategy σ_i prescribes a path independently of the other players moves. Let us show that ψ is indeed a potential function, that is, for each profile π' obtained from π by modifying some player k's strategy, we have

$$\psi(\pi) - \psi(\pi') = \mathsf{cost}_k(\pi) - \mathsf{cost}_k(\pi').$$

Consider such a profile π' obtained by replacing player k's strategy from σ_k to some other blind strategy σ'_k. Let us denote by $e_1 e_2 \ldots$ the path of σ_k, and by $e'_1 e'_2 \ldots$ that of σ'_k. Let E' denote the set of edges of both paths.

$$\begin{aligned}
\psi(\pi) - \psi(\pi') &= \sum_{j=1}^{N_\pi} \sum_{e \in E} \sum_{i=1}^{\mathsf{load}_e(\pi,j)} \mathsf{cost}_e(i) - \sum_{j=1}^{N_{\pi'}} \sum_{e \in E} \sum_{i=1}^{\mathsf{load}_e(\pi',j)} \mathsf{cost}_e(i), \\
&= \sum_{e \in E'} \sum_{j=1}^{N_\pi} \sum_{i=1}^{\mathsf{load}_e(\pi,j)} \mathsf{cost}_e(i) - \sum_{e \in E'} \sum_{j=1}^{N_{\pi'}} \sum_{i=1}^{\mathsf{load}_e(\pi',j)} \mathsf{cost}_e(i), \\
&= \sum_{j=1}^{N_\pi} \mathsf{cost}_{e_j}(\mathsf{load}_{e_j(\pi,j)}) - \sum_{j=1}^{N_{\pi'}} \mathsf{cost}_{e'_j}(\mathsf{load}_{e_j(\pi',j)}) \\
&= \mathsf{cost}_k(\pi) - \mathsf{cost}_k(\pi').
\end{aligned}$$

Here, the second line is the definition; the third line follows from the fact that both terms are equal for edges $e \notin E'$ since there is no change on the load. The fourth line is due to the following observation. Assume that $e_j \neq e'_j$. In this case, the load of e_j is decreased by 1 in π, and that of e'_j is increased by 1 in π'. We have

$$\sum_{i=1}^{\mathsf{load}_e(\pi,j)} \mathsf{cost}_e(i) - \sum_{i=1}^{\mathsf{load}_e(\pi',j)} \mathsf{cost}_e(i) = \mathsf{cost}_e(\mathsf{load}_e(\pi, j)) - \mathsf{cost}_e(\mathsf{load}_e(\pi', j));$$

hence the result. ◀

▶ **Lemma 11.** *In dynamic NCGs, blind Nash equilibria are Nash equilibria.*

Proof. Consider an arbitrary blind strategy profile $\pi = (\sigma_i)_{i \in [\![n]\!]}$. Assume that Player i has a profitable deviation, that is, there exists some (non-blind) strategy σ'_i such that $\mathsf{cost}_i(\pi[i \to \sigma'_i]) < \mathsf{cost}_i(\pi)$. Define the blind strategy σ''_i as the strategy that follows Player-i's path in the outcome of $\pi[i \to \sigma'_i]$. Then $\pi[i \to \sigma'_i]$ and $\pi[i \to \sigma''_i]$ have the same outcomes, hence $\mathsf{cost}_i(\pi[i \to \sigma''_i]) < \mathsf{cost}_i(\pi)$.

In other terms, if there is a profitable deviation strategy for player i, then the new strategy can be chosen to be blind. This shows that a blind Nash equilibrium is a Nash equilibrium. ◀

▶ **Lemma B.2.** *Given a dynamic NCG $\langle \mathcal{A}, n\rangle$, a blind strategy profile π, and $i \in [\![n]\!]$, Player i has a blind strategy that is their best response to π. This strategy has size at most $N_\pi + |V|$ and can be computed in time $O(|V|^2 \cdot N_\pi^2)$.*

Proof. We let $\mathcal{A} = \langle V, E, \mathsf{src}, \mathsf{tgt}\rangle$. We define a weighted graph $\mathcal{G}$ obtained by concatenating $N_\pi + 1$ copies of $\mathcal{A}$, in which the moves of all other players (but Player i) are hard-coded. Formally, $\mathcal{G} = \langle V \times [\![N_\pi + 1]\!], E'\rangle$ where

- for each edge (v, f, v') in E, there is an edge $((v, N_\pi + 1), \mathsf{f}(1), (v', N_\pi + 1))$ in E': this encodes the fact that, after $N_\pi + 1$ steps, all other players have reached the target state, and Player i plays alone;
- for each edge $e = (v, \mathsf{f}, v')$ in E and each index $1 \leq k \leq N_\pi$, there is an edge $((v, k), w, (v', k + 1))$ in E' with $w = \mathsf{f}(load_e(\pi_{-i}, k) + 1)$, where $\mathsf{load}_e(\pi_{-i}, k)$ is the number of players (except Player i) taking edge e at step k when following π_{-i}. This way, w corresponds to the cost incurred to Player i if they were to take edge e at step k.

By construction, any blind strategy σ'_i for Player i in $\langle \mathcal{A}, n\rangle$ corresponds to a path ρ' from $(\mathsf{src}, 1)$ to $(\mathsf{tgt}, N_\pi + 1)$ in $\mathcal{G}$ such that the cost of ρ' is equal to $\mathsf{cost}_i(\sigma_{-i}, \sigma'_i)$: Player i can follow the exact same path as they would in $\mathcal{S}$ by ignoring the second component in the state space of $\mathcal{G}$. Conversely, any path in $\mathcal{G}$ from $(\mathsf{src}, 1)$ to $(\mathsf{tgt}, N_\pi + 1)$ corresponds to a path from src to tgt in $\mathcal{A}$, thus to some blind strategy σ'_i. The cost of the path is equal to $\mathsf{cost}_i(\sigma_{-i}, \sigma'_i)$ by construction.

This shows that the best-response strategy can be computed by a shortest path computation in $\mathcal{G}$. This path has size at most $N_\pi + |V|$: after N_π steps, the path enters configurations of the form $(c, N_\pi + 1)$, so all other players have already reached the target; since the shortest path is acyclic, the bound follows.

This computation can be done with Dijkstra's algorithm, which runs in $O((|V| \cdot N_\pi)^2)$. ◀

▶ **Theorem 9.** *In dynamic NCGs, blind Nash equilibria always exist, and we can compute one in pseudo-polynomial time.*

Proof. We apply the previous lemma as follows. Consider an initial strategy profile π_0 assigning to each player any acyclic path (thus, of length at most $|V|$) from src to tgt. We have $\psi(\pi_0) \leq |V| \cdot \sum_{e \in E} \sum_{i=1}^{n} \mathsf{cost}_e(i) \leq |V| n \max_{e \in E} \mathsf{cost}_e(n)$. This quantity is pseudo-polynomial in the size of $\langle \mathcal{A}, n\rangle$.

Let $\pi_1, \pi_2, \ldots$ denote the strategy profiles generated by this iterative procedure. Let $m_i = N_{\pi_i}$.

By Lemma B.2, we have $m_i \leq m_1 + (i-1)|V|$. We then have

$$\sum_{i=1}^{k} |V|^2 m_i^2 \leq |V|^2 \sum_{i=1}^{k} (m_1 + i|V|)^2$$
$$\leq |V|^2 \sum_{i=1}^{k} (n|V| + i|V|)^2$$
$$\leq |V|^4 \cdot k \cdot (n+k)^2$$

where the second step follows from $m_1 \leq n|V|$. Applying Lemma B.2 again, the running time of the iterative procedure until the k-th step is $O(|V|^4 \cdot k \cdot (n+k)^2)$. In the worst case, we stop after $k = \psi(\pi_0)$ steps, so that the algorithm runs in pseudo-polynomial time. ◀

▶ **Lemma 12.** *There exists a dynamic NCG with a Nash equilibrium π such that for all blind Nash equilibria π', we have $\mathsf{cost}(\pi) < \mathsf{cost}(\pi')$.*

Proof of Lemma 12. We consider the arena $\mathcal{A}$ shown in Fig. 4 with $n = 3$ players, with $\mathsf{src} = q_0, \mathsf{tgt} = q_7$.

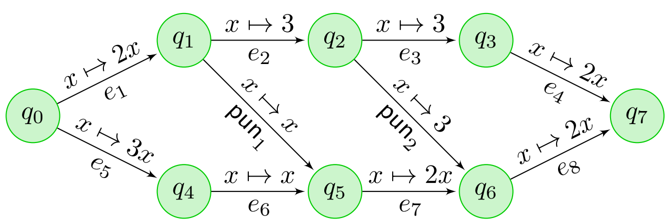

Figure 4 An arena on which blind Nash equilibria are sub-optimal.

The strategy profile $\pi = (\sigma_i)_{i \in [\![3]\!]}$ is defined as follows.

For $i \in \{1,2\}$, strategy σ_i chooses $e_1e_2e_3e_4$ with the following exception: if Player $3-i$ picks e_5 at q_0, then take pun_1 at q_1; if Player $3-i$ picks pun_1 at q_1, then take pun_2 at q_2. Strategy σ_3 follows $e_5e_6e_7e_8$.

We have $\mathsf{cost}_i(\pi) = 14$ for $i \in \{1,2\}$, and $\mathsf{cost}_3(\pi) = 8$, so $\mathsf{cost}(\pi) = 36$.

We first show that π is a Nash equilibrium.

First, Player 3 has no incentive to deviate since taking e_1 at q_0 would alone cost 6, and the rest of the path, either through e_2 or $\mathsf{pun}_1 e_7$, has a cost more than 3 so that the total cost is more than $\mathsf{cost}_3(\pi)$. For Players 1 and 2, there are three states where they can deviate. If Player $i \in \{1,2\}$ chooses e_5 at q_0, then the cost is $6+2+6+6 = 20$ as Player $3-i$ chooses pun_1. Similarly, if they choose pun_1 at q_1 then Player $3-i$ chooses pun_2 at q_2 which yields a cost of $4+1+4+6 = 15 > \mathsf{cost}_i(\pi)$. Last, if Player i chooses pun_2 at q_2, their cost is $4+3+3+4 = 14 = \mathsf{cost}_i(\pi)$.

We now show that the profile π has a lower cost than any blind Nash equilibrium.

There are only four possible blind strategies in our game following one of the four paths $\rho_1 = e_1e_2e_3e_4, \rho_2 = e_5e_6e_7e_8, \rho_3 = e_1\mathsf{pun}_1e_7e_8$ and $\rho_4 = e_1e_2\mathsf{pun}_2e_8$. We will represent these profiles as tuples of paths, *e.g.* (ρ_1, ρ_1, ρ_2). We are going to consider all possible tuples and show that each tuple is either not a Nash equilibrium or has cost more than $\mathsf{cost}(\pi) = 36$.

Observe that when we permute the players in a Nash equilibrium, it remains a Nash equilibrium with the same social cost. So we only need to consider the case where all players use the same strategy (4 possibilities), the case where they all choose distinct strategies ($\binom{4}{3}$), and the case where two of them choose the same strategy and the other one a different one

$(\binom{4}{1} \times \binom{3}{1})$. So there are 20 profiles to check, and the rest of the profiles are permutations of these.

The following is an exhaustive list of the costs of these profiles.

$$\begin{aligned}
&\mathsf{cost}(\rho_1,\rho_1,\rho_1) = 18 \times 3 = 54 && \mathsf{cost}(\rho_1,\rho_1,\rho_4) = 14 \times 2 + 14 = 42\\
&\mathsf{cost}(\rho_2,\rho_2,\rho_2) = 24 \times 3 = 72 && \mathsf{cost}(\rho_2,\rho_2,\rho_1) = 16 \times 2 + 10 = 42\\
&\mathsf{cost}(\rho_3,\rho_3,\rho_3) = 21 \times 3 = 63 && \mathsf{cost}(\rho_2,\rho_2,\rho_3) = 20 \times 2 + 15 = 55\\
&\mathsf{cost}(\rho_4,\rho_4,\rho_4) = 18 \times 3 = 54 && \mathsf{cost}(\rho_2,\rho_2,\rho_4) = 18 \times 2 + 14 = 50\\
&\mathsf{cost}(\rho_1,\rho_2,\rho_3) = 12 + 13 + 12 = 37 && \mathsf{cost}(\rho_3,\rho_3,\rho_1) = 16 \times 2 + 14 = 46\\
&\mathsf{cost}(\rho_1,\rho_3,\rho_4) = 12 + 14 + 15 = 41 && \mathsf{cost}(\rho_3,\rho_3,\rho_2) = 18 \times 2 + 16 = 52\\
&\mathsf{cost}(\rho_1,\rho_2,\rho_4) = 12 + 12 + 13 = 37 && \mathsf{cost}(\rho_3,\rho_3,\rho_4) = 18 \times 2 + 15 = 51\\
&\mathsf{cost}(\rho_2,\rho_3,\rho_4) = 14 + 15 + 16 = 45 && \mathsf{cost}(\rho_4,\rho_4,\rho_1) = 16 \times 2 + 14 = 46\\
&\mathsf{cost}(\rho_1,\rho_1,\rho_2) = 14 \times 2 + 8 = \mathbf{36} && \mathsf{cost}(\rho_4,\rho_4,\rho_2) = 16 \times 2 + 12 = 44\\
&\mathsf{cost}(\rho_1,\rho_1,\rho_3) = 16 \times 2 + 11 = 43 && \mathsf{cost}(\rho_4,\rho_4,\rho_3) = 18 \times 2 + 15 = 51
\end{aligned}$$

All profiles have cost at least 36, and the only one that matches 36 is (ρ_1, ρ_1, ρ_2). However, this is not a Nash equilibrium. In fact, the cost of Player 1 here is 14, but they could profit from deviating to ρ_3 since $\mathsf{cost}_1(\rho_3, \rho_1, \rho_2) = 13$. ◀

C Proofs of Section 5

Equivalence between SPE and very-weak SPE [10]. A strategy σ'_i is said to be *first-shot deviating* from another strategy σ_i if they coincide on all histories except the empty history. A strategy profile $\sigma = (\sigma_i)_{i\in[\![n]\!]}$ is called a *very-weak Nash equilibrium* if for all $i \in [\![n]\!]$, for every first-shot deviating strategy σ'_i, it holds $\mathsf{cost}_i(\langle\sigma_{-i}, \sigma'_i\rangle) \geq \mathsf{cost}_i(\sigma)$. Notice that this need not be an NE [10]. Finally, a strategy profile is a *very-weak subgame-perfect equilibrium* if it is a very-weak NE from every finite history of the game.

We have the following equivalence:

▶ **Theorem C.3.** *In a dynamic NCG, a strategy profile is an SPE if, and only if, it is a very-weak SPE.*

We introduce the following notations to ease the writing of the following proofs: Given a strategy profile σ and a history h, we write $\mathsf{outcome}(\sigma, h)$ for the outcome of the residual strategy σ^h from the last configuration of h. Similarly, we let $\mathsf{cost}_k(\sigma, h) = \mathsf{cost}_k(\mathsf{outcome}(\sigma, h))$.

Proof. It is easily checked that NEs are very-weak NEs, so that SPEs are very-weak SPEs.

For the opposite direction of the equivalence, let us consider a very-weak SPE $\sigma = (\sigma)_{i\in[\![n]\!]}$ of a dynamic NCG $\mathcal{G}$. Suppose that σ is not an SPE; then there exists a history h of $\mathcal{G}$ such that the residual strategy σ^h is not an NE.

This implies that for some player i, there exists a strategy σ'_i such that $\mathsf{cost}_i(\sigma, h) > \mathsf{cost}_i(\langle\sigma_{-i}, \sigma'_i\rangle, h)$. In particular, $\mathsf{cost}_i(\langle\sigma_{-i}, \sigma'_i\rangle, h)$ is finite, and along the outcome of this strategy profile, Player i reaches tgt.

We pick σ'_i as a profitable deviating strategy (after history h) with minimum number of deviations from σ_i. Because Player i reaches tgt along $\mathsf{outcome}(\langle\sigma_{-i}, \sigma'_i\rangle, h)$, there exist profitable strategies with finitely many deviations.

We write h' for the longest finite history along the path $\mathsf{outcome}(\langle\sigma_{-i}, \sigma'_i\rangle, h)$ where σ'_i deviates from σ_i, i.e, $\sigma_i(hh') \neq \sigma'_i(hh') = c$ and $\mathsf{outcome}(\sigma, hh'c) = \mathsf{outcome}(\langle\sigma_{-i}, \sigma'_i\rangle, hh'c)$.

Now if $\mathsf{cost}_i(\sigma, hh') \leq \mathsf{cost}_i(\langle\sigma_{-i}, \sigma'_i\rangle, hh')$ holds, then deviating at history hh' is unnecessary for Player i for having profitable deviating strategy from h, and that contradicts the minimality of the number of deviations of σ'_i from σ_i. Hence we obtain $\mathsf{cost}_i(\sigma, hh') > \mathsf{cost}_i(\langle\sigma_{-i}, \sigma'_i\rangle, hh')$. But that implies we have a first-shot profitable deviating strategy for Player i, σ'_i, from history hh', which contradicts to the assumption that π is a very-weak SPE. ◀

Characterization of the outcomes of SPEs. Consider a dynamic NCG $\langle\mathcal{A}, n\rangle$, and the associated configuration graph $\mathcal{M} = \langle C, T\rangle$. We partition the set C of configurations into $(X_j)_{0\leq j\leq n}$ such that a configuration c is in X_j if, and only if, $j = \#\{i \in [\![n]\!] \mid c(i) = \mathsf{tgt}\}$. Since tgt is a sink state in $\mathcal{A}$, if there is a transition from some configuration in X_j to some configuration in X_k, then $k \geq j$. We define $X_{\geq j} = \bigcup_{i\geq j} X_i$, and $Z_j = \{(c, w, c') \in T \mid c \in X_j\}$ and $Z_{\geq j} = \{(c, w, c') \in T \mid c \in X_{\geq j}\}$.

We inductively define a sequence of families of labeling functions $(\lambda^{j^*})_{0\leq j\leq n}$, where for all $0 \leq j \leq n$, $\lambda^{j^*} = (\lambda^{j^*}_i)_{i\in[\![n]\!]}$, and for all $i \in [\![n]\!]$, $\lambda^{j^*}_i : Z_{\geq j} \to \mathbb{N} \cup \{-\infty, +\infty\}$.

For any family $\lambda = \langle\lambda_i\rangle_{i\in[\![n]\!]}$ with $\lambda^{j^*}_i : Z_{\geq j} \to \mathbb{N} \cup \{-\infty, +\infty\}$ and any $c \in X_{\geq j}$, a path $\rho = (t_j)_{j\geq 1}$ from c visiting c_{tgt} is said to be λ-consistent if for any $i \in [\![n]\!]$ and for any $1 \leq k \leq |\rho|$, it holds $\mathsf{cost}_i(\rho_{\geq k}) \leq \lambda_i(t_k)$. We write $\Gamma_\lambda(c)$ to denote the set of all λ-consistent paths from c.

We now define the sequence $(\lambda^{j^*})_{0\leq j\leq n}$. For $j = n$, we have $X_{\geq n} = \{c_{\mathsf{tgt}}\}$ and $Z_{\geq n} = \{(c_{\mathsf{tgt}}, 0^n, c_{\mathsf{tgt}})\}$; there is a single path, which obviously is the outcome of an SPE since no deviations are possible. For all $i \in [\![n]\!]$, we let $\lambda^{n^*}_i(c_{\mathsf{tgt}}, 0^n, c_{\mathsf{tgt}}) = 0$.

Now, fix $j < n$, assuming that $\lambda^{(j+1)^*}$ has been defined. In order to define λ^{j^*}, we introduced an intermediary sequence $(\mu^k_i)_{k\geq 0, i\in[\![n]\!]}$ with $\mu^k_i : Z_{\geq j} \to \mathbb{N} \cup \{-\infty, +\infty\}$, of which $(\lambda^{j^*}_i)_{i\in[\![n]\!]}$ will be the limit. We call these μ^k's as the λ^{j^*}- *building functions.*

Functions μ^k_i mainly operate on $Z_j = Z_{\geq j} \setminus Z_{\geq j+1}$: for any $\mathbf{e} \in Z_{\geq j+1}$, we let $\mu^k_i(\mathbf{e}) = \lambda^{(j+1)^*}_i(\mathbf{e})$. Now, for $\mathbf{e} = (c, w, c') \in Z_j$, the value $\mu^k_i(\mathbf{e})$ is defined inductively as follows:

- $\mu^0_i(\mathbf{e}) = 0$ if $c(i) = \mathsf{tgt}$, and $\mu^0_i(\mathbf{e}) = +\infty$ otherwise;
- for $k > 0$, assuming μ^{k-1}_i has been defined, we let $\mu^k_i(\mathbf{e}) = 0$ if $c(i) = \mathsf{tgt}$, and

$$\mu^k_i(\mathbf{e}) = \begin{cases} \min\limits_{c''\in\mathsf{dev}_i(c,c')} \sup\limits_{\rho\in\Gamma_{\mu^{l-1}}(c'')} (\mathsf{cost}_i(c, c'') + \mathsf{cost}_i(\rho)) & \text{if } \forall (c, \tilde{w}, \tilde{c}) \in T,\ \Gamma_{\mu^{l-1}}(\tilde{c}) \neq \emptyset \\ -\infty & \text{otherwise} \end{cases}$$

▶ **Lemma C.4.** *For any $\mathbf{e} \in T$, any $i \in [\![n]\!]$, and any $k > 0$, we have $\mu^k_i(\mathbf{e}) \leq \mu^{k-1}_i(\mathbf{e})$.*

Proof. We prove the statement by induction on k, starting with $k = 1$. Write $\mathbf{e} = (c, w, c')$. If $c' = \mathsf{tgt}$, then by definition $\mu^k_i(\mathbf{e}) = 0$ for all k; otherwise, $\mu^0_i(\mathbf{e}) = \infty$; in both cases, the result holds.

Now, let us assume that for some $k > 0$, the following holds: for all $i \in [\![n]\!]$ and all $\mathbf{e} \in X_{\geq j}$, it holds $\mu^l_i(\mathbf{e}) \leq \mu^{l-1}_i(\mathbf{e})$. Then for any $c \in C$, it implies $\Gamma_{\mu^k}(c, w, c') \subseteq \Gamma_{\mu^{k-1}}(c, w, c')$.

Fix an arbitrary $\mathbf{e} = (c, w, c') \in T$. For any $c'' \in \mathsf{dev}_i(c, c')$, by induction hypothesis, we have $\Gamma_{\mu^k}(c'') \subseteq \Gamma_{\mu^{k-1}}(c'')$.

Then:

$$\sup_{\rho\in\Gamma_{\mu^k}(c'')} (\mathsf{cost}_i(c, c'') + \mathsf{cost}_i(\rho)) \leq \sup_{\rho\in\Gamma_{\mu^{k-1}}(c'')} (\mathsf{cost}_i(c, c'') + \mathsf{cost}_i(\rho)).$$

It follows

$$\min_{c''\in \mathsf{dev}_i(c,c')} \sup_{\rho\in\Gamma_{\mu^k}(c'')} (\mathsf{cost}_i(c,c'')+\mathsf{cost}_i(\rho)) \leq \min_{c''\in \mathsf{dev}_i(c,c')} \sup_{\rho\in\Gamma_{\mu^{k-1}}(c'')} (\mathsf{cost}_i(c,c'')+\mathsf{cost}_i(\rho)),$$

which shows $\mu^k(\mathbf{e}) \leq \mu^{k-1}(\mathbf{e})$. ◀

We write Γ^* for $\Gamma_{\lambda^{0*}}$. Then:

▶ **Theorem 18.** *A path ρ in $\mathcal{G} = \langle \mathcal{A}, n\rangle$ is the outcome of an SPE if, and only if, $\rho \in \Gamma^*(c_{\mathsf{src}})$.*

Proof. Let us first consider an SPE σ of $\mathcal{G}$, of which ρ is the outcome. We shall show that for any history h ending in configuration c, $\mathsf{outcome}(\sigma, h) \in \Gamma^*(c)$. In fact, we prove by induction on j that $\mathsf{outcome}(\sigma, h) \in \Gamma_{\lambda^{j*}}(c)$ if $c \in X_j$.

For $j = n$, for any history h ending in X_n, $\mathsf{outcome}(\sigma, h)$ is a path looping on c_{tgt}. We defined $\lambda^{n^*}(\mathbf{e}) = 0$ for $\mathbf{e} \in X_n$, so that the result holds.

Assume the statement is true for some $j + 1 \leq n$: for any history h ending in some configuration $c \in X_{j+1}$, $\mathsf{outcome}(\sigma, h) \in \Gamma^{(j+1)^*}(c)$.

To prove the result at level j, we start another induction to show that for any k, for any history h whose last configuration c is in X_j, the path $\mathsf{outcome}(\sigma, h)$ belongs to $\Gamma_{\mu^k}(c)$.

We begin with the case $k = 0$. Fix a history h ending in $c \in X_j$. Write $\mathsf{outcome}(\sigma, h) = (c_r, w_r, c_{r+1})_{r\geq 1}$. Let r' be the least integer such that $c_{r'}$ in $\mathsf{outcome}(\sigma, h)$ belongs to $X_{>j}$. We need to show that for any $i \in [\![n]\!]$ and any r, $\mathsf{cost}_i((\mathsf{outcome}(\sigma, h))_{\geq r}) \leq \mu_i^0(c_r, w_r, c_{r+1})$. If $r < r'$, then either $\mu^0(c_r, w_r, c_{r+1}) = \infty$, or $\mu^0(c_r, w_r, c_{r+1}) = 0$ and $c_r(i) = \mathsf{tgt}$; in both cases, the result also holds. If $r \geq r'$, we have $\mu_i^0(c_p, w_p, c_{p+1}) = \lambda_i^{(j+1)^*}(c_p, w_p, c_{p+1})$, and by our outer induction hypothesis, $\mathsf{cost}_i((\mathsf{outcome}(\sigma, h)_{\geq r})) \leq \lambda^{(j+1)^*}(c_r, w_r, c_{r+1})$.

We now prove the induction step. We assume that $\mathsf{outcome}(\sigma, h) \in \Gamma_{\mu^{k-1}}(c)$ for any history h ending in $c \in X_j$. Write $\mathsf{outcome}(\sigma, h) = (c_r, w_r, c_{r+1})_{r\geq 1}$. We first observe that $\mu_i^k(c_r, w_r, c_{r+1}) \neq -\infty$ for any $r \geq 1$. This is because for any $(c_r, \tilde{w}, \tilde{c}) \in T$ and for any $c_{\mathsf{src}} \overset{\tilde{h}}{\Rightarrow} \tilde{c}$, by our current induction hypothesis $\Gamma_{\mu^{k-1}}(\tilde{c})$ contains $\mathsf{outcome}(\sigma, \tilde{h})$, hence it is not empty.

Therefore, the only thing left to show now is that, for all $i \in [\![n]\!]$, for all r, we have $\mathsf{cost}_i(\mathsf{outcome}(\sigma, h)_{\geq r}) \leq \mu_i^k(c_r, w_r, c_{r+1})$. This is clear if $c_r(i) = \mathsf{tgt}$. As previously, we let r' be the least integer such that configuration $c_{r'}$ along $\mathsf{outcome}(\sigma, h)$ belongs to $X_{>j}$.

For $r < r'$: assume $\mathsf{cost}_i(\mathsf{outcome}(\sigma, h)_{\geq r}) > \mu_i^k(c_r, w_r, c_{r+1})$; then,

$$\mathsf{cost}_i(\mathsf{outcome}(\sigma, h)_{\geq r}) > \min_{c'\in \mathsf{dev}_i(c_r,c_{r+1})} \sup_{\rho\in\Gamma_{\mu^{k-1}}(c')} \{\mathsf{cost}_i(c_r, c') + \mathsf{cost}_i(\rho)\},$$

which implies the existence of an edge $(c_r(i), \mathsf{f}, v')$ in E such that

$$\mathsf{cost}_i(\mathsf{outcome}(\sigma, h)_{\geq r}) > \mathsf{cost}_i(c_r, c_{r+1}[i \to v']) + \mathsf{cost}_i(\sigma, h \cdot (c_r, w', c_{r+1}[i \to v'])),$$

because by our inner induction $\mathsf{outcome}(\sigma, h \cdot (c_r, w', c_{r+1}[i \to v'])) \in \Gamma_{\mu^{k-1}}(c_r, w', c_{r+1}[i \to v'])$. This shows a profitable first-shot deviation for Player i, contradicting to our hypothesis that σ is an SPE.

Otherwise for $r \geq r'$, the condition follows directly from the outer induction hypothesis, because μ^k coincides with λ^{j^*} for edges in those regions.

We now prove the converse implication.

Picking $\rho \in \Gamma^*(c_{\mathsf{src}})$, we build a very-weak SPE $\sigma = (\sigma_i)_{i\in[\![n]\!]}$ (which is an SPE by Theorem C.3), step-by-step; for notational convenience, instead of defining $\sigma_i(h)$ for any history h, we directly define what $\mathsf{outcome}(\sigma, h)$ would be (if not already defined).

As a first step, we let $\mathsf{outcome}(\sigma, c_{\mathsf{src}}) = \rho$. Now let us construct rest of the σ. We pick an arbitrary non-initial history $h' = h \cdot (c, w, c')$ for which $\mathsf{outcome}(\sigma, h')$ is not defined; we assume that $\mathsf{outcome}(\sigma, h)$ is already defined as $(c_j, w_j, c_{j+1})_{j\geq 1}$ such that $c' \neq c_2$.

If there exists $i_1 \neq i_2 \in [\![n]\!]$ such that $c'(i_j) \neq c_2(i_j)$ for both $j = 1$ and $j = 2$, then pick any $\tilde{\rho} \in \Gamma^*(c')$ (which we know is not empty) and define $\mathsf{outcome}(\sigma, h \cdot (c, w, c')) = \tilde{\rho}$.

Otherwise c_2 differs from c' only at a single player's position, say Player i's. In that case we define

$$\mathsf{outcome}(\sigma, h \cdot (c, w, c')) = \arg\max_{\rho\in\Gamma^*(c')}\{\mathsf{cost}_i(c_1, c') + \mathsf{cost}_i(\rho)\} \tag{2}$$

Note that, here we can take arg max because $\lambda_i^*(\mathbf{e})$ is finite (which we shall establish in Corollary C.9) for any $\mathbf{e} \in \mathcal{E}$, so there are finitely many plays in $\Gamma^*(c)$. This ends our definition of the strategy profile σ, which we now prove is an SPE.

Consider an arbitrary history h and denote $\mathsf{outcome}(\sigma, h)$ by $(c_j, w_j, c_{j+1})_{1\leq j}$. Pick a configuration c' such that $(c_1, w', c') \in T$, $c'(j) = c_2(j)$ for all $j \neq i$, but $c'(i) \neq c_2(i)$. That is, c' is a configuration obtained by Player i's deviation from σ_i at the end of h.

We show that such a single player deviation is not beneficial if, after that deviation, all players continue to play following σ:

$$\begin{aligned}
\mathsf{cost}_i(c_1, c') + \mathsf{cost}_i(\sigma, h \cdot (c_1, w', c')) &= \sup_{\rho\in\Gamma^*(c')}\{\mathsf{cost}_i(\rho) + \mathsf{cost}_i(c_1, c')\} && \text{(by (2))}\\
&\geq \min_{\tilde{c}\in\mathsf{dev}_i(c_1,c_2)} \sup_{\rho\in\Gamma^*(\tilde{c})}\{cost_i(\rho) + cost_i(c, \tilde{c})\}\\
&= \lambda_i^*((c_1, w_1, c_2))\\
&\geq \mathsf{cost}_i(\sigma, h)
\end{aligned}$$

So no such configuration c' would be beneficial for Player i. Hence from any history h, σ is very-weak NE. Therefore, we can conclude that σ is an SPE. ◀

Counter Graph. Having this correspondence between outcomes of SPE and Γ^*-paths, our job to decide existence of SPE now reduces to checking non-emptiness of $\Gamma^*(c_{\mathsf{src}})$. To this aim, with any such family $\mu = \langle\mu_i\rangle_{i\in[\![n]\!]}$ and configuration c, we associate an infinite-state *counter graph* $\mathbb{C}[\mu, c] = \langle C', T'\rangle$ to capture μ-consistent paths from c:

- the set of vertices is $C' = C \times (\mathbb{N} \cup \{+\infty\})^{[\![n]\!]}$
- T' contains all edges $((d, b), w, (d', b'))$ such that (d, w, d') is an edge of $\mathcal{M}$ and for all $i \in [\![n]\!]$, either $b'(i) = 0$ if $d(i) = \mathsf{tgt}$, or $b'_i = \min\{b_i - w_i, \mu_i(d, w, d') - w_i\}$ otherwise (provided that $b'_i \geq 0$ for all i, in order for (d', b') to be an edge of $\mathbb{C}[\mu, c]$).

Intuitively, in configuration (d, b), b is used to enforce μ-consistency: each edge taken along a path imposes a constraint on the cost of the players for the rest of the path; this constraint is added to the constraints of the earlier edges, and propagated along the path.

With the initial configuration c, we associate b^c such that $b^c_i = 0$ if $c(i) = \mathsf{tgt}$ and $b^c_i = +\infty$ otherwise: this configuration imposes no constraints since no edges have been taken yet.

Notice that $\mathbb{C}[\mu, c]$ is infinite, but as we show below, only finitely many states are reachable from the initial configuration. We write $|C'|_r$ for the number of reachable states in C'.

We extend region decomposition of $\mathcal{M} = \langle C, T\rangle$ to any counter graph $\mathbb{C}[\mu, c] = \langle C', T'\rangle$ in the natural way: $(c', b') \in X'_j$ if $c' \in X_j$, and an edge $((c', b'), w', (c'', b'')) \in Z'_{\geq j}$ if $(c', w', c'') \in Z_{\geq j}$.

We call a path π from (c, b^c) to (c_{tgt}, b) (for some b) as a *valid* path in $\mathbb{C}[\mu, c]$. For any path $\pi \in \mathbb{C}[\mu, c]$, (where $c \in X_j$), we write $\pi = \pi[j] \cdot \pi[j+1] \cdots \pi[n]$ where $\pi[l]$ denotes the (possibly empty) section of path in Z'_l. We also use the notation $\max b_l(\mu, c)$ to denote the maximum finite counter value that appears in the vertices reachable from (c, b^c) and belonging to X'_j: more precisely,

$$\max b_l(\mu, c) = \max\{m \in \mathbb{N} \mid \exists (c', b') \in X_l \text{ s.t. } (c, b^c) \to^* (c', b') \text{ and } \exists i \in [\![n]\!].\ b'(i) = m\}.$$

We extend this notation to $\max b_{\geq l}(\mu, c)$, which denotes $\max_{l \geq j} \max b_l(\mu, c)$. When all the counter values are in $\{0, \infty\}$, we take $\max b_l(\mu, c) = 0$.

▶ **Lemma C.5.** *There exists a path $\pi = ((c_j, b_j), w_j(c_{j+1}, b_{j+1}))_j$ from (c, b^c) to a vertex (c_{tgt}, b) (for some b) in the counter graph $\mathbb{C}[\lambda, c] = \langle C', T' \rangle$ of $\mathcal{G}$ if, and only if, there is a λ-consistent path $\rho = (c_j, w_j, c_{j+1})_j$ from c in $\mathcal{G}$.*

Proof. Assume that such a path π exists. Along π, for each player i, let $k(i)$ be the least index such that $c_{k(i)}(i) = \mathsf{tgt}$. Then it is enough to show that for $1 \leq k < k(i)$, we have $\mathsf{cost}_i(\rho_{\geq k}) \leq \lambda_i(c_k, w_k, c_{k+1})$. And indeed we have

$$\begin{aligned}
0 \leq b_{k(i)}(i) &\leq b_{k(i)-1}(i) - \mathsf{cost}_i(c_{k(i)-1}, c_{k(i)}) \\
&\leq b_{k(i)-2}(i) - \mathsf{cost}_i(c_{k(i)-2}, c_{k(i)-1}) - \mathsf{cost}_i(c_{k(i)-1}, c_{k(i)}) \\
&\ \vdots \\
&\leq b_{k+1}(i) - \sum_{j=k+1}^{k(i)-1} \mathsf{cost}_i(c_j, c_{j+1})
\end{aligned}$$

so that $\mathsf{cost}_i(\rho_{\geq k}) \leq b_{k+1}(i) \leq \lambda_i(c_k, w_k, c_{k+1})$,

Conversely, if there is a λ-consistent path $\rho = (c_j, w_j, c_{j+1})_{1 \leq j < |\rho|}$ from c, we define $\pi = ((c_j, b_j), w_j, (c_{j+1}, b_{j+1}))_{1 \leq j < |\rho|}$ with:

$$b_1(i) = \begin{cases} 0 & \text{if } c_1(i) = \mathsf{tgt} \\ +\infty & \text{otherwise} \end{cases}$$

$$b_j(i) = \min\{b_{j-1}(i) - \mathsf{cost}_i(c_{j-1}, c_j), \lambda_i(c_{j-1}, w_{j-1}, c_j) - \mathsf{cost}_i(c_{j-1}, c_j)\} \quad \text{for } 1 < j < |\rho|$$

For π to be a valid path in $\mathbb{C}[\lambda, c]$, we have to prove that $b_j(i) \geq 0$ for all $1 < j \leq |\rho|$ and all $i \in [\![n]\!]$. For $j = 1$, it is evident.

For $j > 1$, the second element of the minimum defining $b_j(i)$ is non-negative, because $\mathsf{cost}_i(c_{j-1}, c_j) \leq \mathsf{cost}_i(\rho_{\geq j-1}) \leq \lambda_i(c_{j-1}, w_{j-1}, c_j)$. Suppose the first element $b_{j-1}(i) - \mathsf{cost}_i(c_{j-1}, c_j)$ is negative. Using definition of $b_{j-1}(i)$, we can say the above inequality can be true only if either $b_{j-2}(i) - \sum_{k=j-2}^{j-1} \mathsf{cost}_i(c_k, c_{k+1}) < 0$, or $\lambda_i(c_{j-2}, w_{j-2}, c_{j-1}) - \sum_{k=j-2}^{j-1} \mathsf{cost}_i(c_k, c_{k+1}) < 0$. But again the latter cannot be true because it directly contradicts λ-consistency of ρ. Hence, our supposition can only be true if the former constraint holds. We repeat the process with $b_{j-2}(i)$, coming down to the condition $b_2(i) - \sum_{k=2}^{j-1} \mathsf{cost}_i(c_k, w_k, c_{k+1}) < 0$ to make our supposition $b_{j-1}(i) - \mathsf{cost}_i(c_{j-1}, c_j) < 0$ true. But $b_2(i) = \lambda_i(c_1, w_1, c_2) - \mathsf{cost}_i(c_1, c_2)$, so that the inequality above entails $\lambda_i(c_1, w_1, c_2) < \sum_{k=1}^{j-1} \mathsf{cost}_i(c_k, c_{k+1})$. This contradicts λ-consistency condition at the beginning of ρ. It follows that $b_j(i) \geq 0$ for all $1 \leq j \leq |\rho|$ and all $i \in [\![n]\!]$. ◀

▶ **Lemma C.6.** *$\sup_{\rho\in\Gamma_{\mu^k}(c)} \mathsf{cost}_i(\rho) = +\infty$ if, and only if, there exists a valid path π in $\mathbb{C}[\mu^k, c]$ with the following conditions:*

- *π is of the form $h\cdot\beta\cdot h'$, where β is a cycle in $\mathbb{C}[\mu^k, c]$;*
- *Player i's (constant) counter value $b(i)$ is positive throughout β.*

Proof. $\sup_{\rho\in\Gamma_{\mu^k}(c)} \mathsf{cost}_i(\rho) = +\infty$ implies there exists a sequence of paths $(\rho_m)_{m\geq 1}$ in $\Gamma_{\mu^k}(c)$ such that $\mathsf{cost}_i(\rho) \to \infty$ as m grows. By Lemma C.5, for each of those ρ_m, there exists corresponding π_m from (c, b^c) to (c_{tgt}, b) in $\mathbb{C}[\mu^k, c]$. As a player's cost in a single edge is bounded, the length of π_m has to grow unboundedly. First of all, that is only possible if there is a cycle β in $\mathbb{C}[\mu^k, c]$ - we have the first condition now. Now Player i's counter value cannot be 0 in β because then the cycle doesn't contribute to make $\mathsf{cost}_i(\rho) \to 0$. Thus the second condition also holds true.

Conversely, for $\pi = h\cdot\beta\cdot h'$ in $\mathcal{C}[\mu^k, c]$, where β is a cycle and Player i's cost > 0 in β, we construct a sequence of paths $\pi^m = h\cdot\beta^{m+1}\cdot h'$ for $m \geq 0$. By Lemma C.5, corresponding to this sequence of paths, there exists a sequence of paths $(\rho_m)_{m\geq 0}$ in $\Gamma_{\mu^{k-1}}(c'')$, and for those paths, $\mathsf{cost}_i(\rho_m) \to \infty$ as m grows. Hence, $\mu_i^k(c, w, c') = +\infty$ if these condition holds. ◀

▶ **Lemma C.7.** *For any edge $\mathbf{e} = (c, w, c') \in Z_j$, the $(|V|+1)$-st family $\mu^{|V|}$ of λ^{j^*}-building functions satisfies $\mu_i^{|V|}(\mathbf{e}) \leq |V| \times \kappa$, where $\kappa = max_{e\in E} \mathsf{f}_e(n)$.*

Proof. We prove a slightly stronger statement. Let m be the length of the shortest path from $c(i)$ to tgt in $\mathcal{A}$. We show (by induction) that $\mu_i^m(\mathbf{e}) \leq m \times \kappa$. As $m \leq |V|$ and by Lemma C.4, this entails our current lemma.

If $m = 1$, then there is an edge e from $c[i]$ to tgt in $\mathcal{A}$. Let us consider an edge $\mathbf{e}' = (c, w', c'')$ such that $c''[j] = c'[j]$ for all $j \neq i$, and $c''[i] = \mathsf{tgt}$. Now if $\Gamma_{\mu^{m-1}}(\tilde{c}) = \emptyset$ for any $(c, \tilde{w}, \tilde{c}) \in T$, then anyway $\mu_i^m(\mathbf{e}) = -\infty$, and the result holds. Otherwise, we have $\Gamma_{\mu^{m-1}}(c'') \neq \emptyset$, and for any $\rho \in \Gamma_{\mu^{m-1}}(c'')$, we have $\mathsf{cost}_i(\rho) = 0$. Therefore, $\mu_i^m(\mathbf{e}) = \mathsf{cost}_i(c, c'')$, which is bounded by κ.

Now assume that the induction hypothesis holds up to step $m-1$, and consider a configuration c such that the length of a shortest path from $c(i)$ to tgt in $\mathcal{A}$ is m. Fix an edge $(c, w, c') \in T$. Write $(c(i), \mathsf{f}, v') \in E$ for the first edge of a shortest path from $c(i)$ to tgt. We consider a configuration c'' such that $c''(j) = c'(j)$ for all $j \neq i$, and $c''(i) = v'$. By construction, there is a path from $c''(i)$ to tgt of length $\leq m-1$. By induction hypothesis, for any edge $(c'', \tilde{w}, \tilde{c}) \in T$, we have $\mu_i^{m-1}((c'', \tilde{w}, \tilde{c})) \leq (m-1) \times \kappa$. This implies for any path $\rho = (t_j)_{j\geq 1} \in \Gamma_{\mu^{m-1}}(c'')$, we have $\mathsf{cost}_i(\rho) \leq \mu_i^{m-1}(t_1) \leq (m-1) \times \kappa$. Therefore, $\mu_i^m(c, w, c') \leq \mathsf{cost}_i(c, c'') + (m-1) \times \kappa \leq m \times \kappa$. ◀

Lemmas C.4 and C.7 provide a bound on the number of steps until any sequence of λ^{j^*}-computing functions stabilize:

▶ **Corollary C.8.** *Any sequence $(\mu^k)_{k\geq 0}$ of λ^{j^*}-computing functions stabilizes after at most $|V|(1 + n\cdot\kappa\cdot|E|^n)$ steps.*

From Lemma C.7, we also get that the sequence $(\mu^k)_{k\geq 0}$ built for computing λ^{j^*} cannot stabilize unless for all $i \in [\![n]\!]$, for all $\mathbf{e} \in Z_{\geq j}$, we have $\mu_i^k(\mathbf{e}) \leq |V| \times \kappa$. By Lemma C.4:

▶ **Corollary C.9.** *For any $0 \leq i, j \leq n$, and any $\mathbf{e} \in Z_{\geq j}$, we have $\lambda_i^{j*}(\mathbf{e}) \leq |V| \times \kappa$.*

At this point, we have bounded the finite values that λ^{j^*}'s can take. But in the transitioning from $\lambda^{(j+1)^*}$ to λ^{j^*}, μ_i^k can take larger values when $k < |V|$. In the sequel, we bound the values that any family $\mu^k = \langle\mu_i^k\rangle_{i\in[\![n]\!]}$ can return.

To this aim, we begin with working on the supremum of the cost for Player i of the paths in $\Gamma_{\mu^k}(c)$.

When we consider a $\max b_l(\mu^k, c)$ with $c \in X_j$ of $\mathcal{M}$, it is implicit that $l \geq j$. For $l > j$, we have

$$\max b_l(\mu^k, c) \leq \max b_l(\mu^0, c) \leq \max_{\substack{\mathbf{e}\in Z_l \\ i \in [\![n]\!]}} \lambda_i^{l^*}(\mathbf{e}) \leq |V| \times \kappa$$

because $\mu_i^k(\mathbf{e}) = \mu_i^0(\mathbf{e}) = \lambda_i^{l^*}(\mathbf{e})$ for those $\mathbf{e} \in Z_{\geq l}$ and for all $i \in [\![n]\!]$. So it remains to bound $\max b_j(\mu^k, c)$; but for that too, when $k \geq |V|$, we have $\max b_j(\mu^k, c) \leq |V| \times \kappa$. Therefore, we only need to provide a bound for $\max b_j(\mu^k, c)$ when $k < |V|$.

▶ **Lemma C.10.** *For an edge $\mathbf{e} = (c, w, c') \in Z_j$ and a λ^{j^*}-building function $\mu^k = \langle \mu_i^k \rangle_{i \in [\![n]\!]}$, if $\mu_i^k(\mathbf{e})$ is non-zero finite then it holds*

$$\mu_i^k(\mathbf{e}) \leq (n|C| + 2|V|) \times \sum_{l=1}^{k} (n|C|)^{l-1} \cdot \kappa^l$$

Moreover, the above bound also applies to $\max b_j(\mu^k, c'')$ for any $c'' \in X_j$:

$$\max b_j(\mu^k, c) \leq (n|C| + 2|V|) \times \sum_{l=1}^{k} (n|C|)^{l-1} \cdot \kappa^l.$$

Proof. We can claim that the finite maximum counter value appeared in X'_j of any counter graph $\mathbb{C}[\mu^k, c'']$ (where $c \in X_j$) is bounded by the finite maximum $\mu_i^k(\mathbf{e})$ value appeared the same region (maximum over i and $\mathbf{e} \in Z_j$), i.e,

$$\max b_j(\mu^k, c'') \leq \max\{\mu_i^k(c, w, c) \in \mathbb{N} \mid ((c, b), w, (c', b')) \in Z'_j \text{ of } \mathbb{C}[\mu^k, c''],\ i \in [\![n]\!]\}.$$

This is justified because for the initial vertex (c, b^c) of X'_j, $b^c \in \{0, \infty\}$. Hence, a counter value becomes non-zero finite, when some $\mu_i^k(\mathbf{e})$ for $\mathbf{e} \in Z_j$ becomes non-zero finite. But by definition of counter graph, $((c, b), w, (c', b')) \in Z'_j$ implies $(c, w, c') \in Z_j$, hence we can obtain,

$$\max b_j(\mu^k, c'') \leq \max\{\mu_i^k(c, w, c') \in \mathbb{N} \mid (c, w, c') \in Z_j, i \in [\![n]\!]\}$$

By induction on k, we prove that for any $c'' \in X_j$ and any $i \in [\![n]\!]$,

$$\max b_j(\mu^k, c'') \leq \max\{\mu_i^k(\mathbf{e}) \in \mathbb{N} \mid \mathbf{e} \in Z_j\} \leq (n|C| + 2|V|) \times \sum_{l=1}^{k} (n|C|)^{l-1} \cdot \kappa^l.$$

As $\mu_i^0(\mathbf{e}) \in \{0, \infty\}$ for all $\in Z_j$, the given bounds hold for $k = 0$.

We fix an arbitrary $\mathbf{e} = (c, w, c') \in Z_j$ and a player $i \in [\![n]\!]$ such that $\mu_i^k(\mathbf{e}) \in \mathbb{N} \setminus \{0\}$. That $\mu_i^k(\mathbf{e})$ is a non-zero finite value means that $\sup_{\rho \in \Gamma_{\mu^{k-1}}(\tilde{c})} \mathsf{cost}_i(\rho) \in \mathbb{N}$ for some $\tilde{c} \in \mathsf{dev}_i(c, c')$. If $\tilde{c}(i) = \mathsf{tgt}$ then $\sup_{\rho \in \Gamma_{\mu^{k-1}}(\tilde{c})} \mathsf{cost}_i(\rho) = 0$, and that makes $\mu_i^k(\mathbf{e}) \leq \mathsf{cost}_i(c, \tilde{c}) \leq \kappa$, satisfying the given bound.

Otherwise, $\tilde{c}(i) \neq \mathsf{tgt}$, and depending whether $\tilde{c}$ belongs to X_j or $X_{>j}$, we analyze two cases:

- $\tilde{c} \in X_l$ for some $l > j$: Then

$$\begin{aligned}\mu_i^k(\mathbf{e}) &\leq \sup_{\rho \in \Gamma_{\mu^{k-1}}(\tilde{c})} \{\mathsf{cost}_i(c, \tilde{c}) + \mathsf{cost}_i(\rho)\} \\ &\leq \mathsf{cost}_i(c, \tilde{c}) + \max_{\substack{\mathbf{e}'\in Z_l \\ i \in [\![n]\!]}} \mu_i^{k-1}(\mathbf{e}') \leq (1 + |V|) \times \kappa.\end{aligned}$$

- $\tilde{c} \in X_j$: Now consider $\mathbb{C}[\mu^{k-1}, \tilde{c}]$, and its initial vertex $(\tilde{c}, b^{\tilde{c}})$. By Lemma C.5, for any path $\rho \in \Gamma_{\mu^{k-1}}(\tilde{c})$, we have a corresponding valid path π_ρ in $\mathbb{C}[\mu^{k-1}, \tilde{c}]$. We also consider the region decomposition $\pi_\rho = \pi_\rho[j] \dots \pi_\rho[n]$. From the fact that $\sup_{\rho \in \Gamma_{\mu^{k-1}}(\tilde{c})} < +\infty$, we argue that in π_ρ, from $(\tilde{c}, b^{\tilde{c}})$ within each $|C|$ step either one counter value strictly decreases, or Player i reaches the tgt. Otherwise, there would have been a cycle in $\mathbb{C}[\mu^{k-1}, \tilde{c}]$ resulting the supremum $+\infty$ (thanks to lemma C.6). Recall by design, counter values of π_ρ in X'_j lie in $[\![\max b_j(\mu^{k-1}, \tilde{c})]\!] \cup \{0, +\infty\}$, and when a counter value decreases along an edge, it decreases at least by 1. Therefore, within $n \times |C| \times (\max b_j(\mu^{k-1}, \tilde{c}) + 1)$ steps from $(\tilde{c}, b^{\tilde{c}})$ at least one of the counter value becomes 0. When a player-l counter value becomes 0, π_ρ must reaches the next region (making the corresponding $c(l) = \mathsf{tgt}$), otherwise π_ρ is not a valid path. Moreover, from the next region, $\mathsf{cost}_i(\rho[j+1] \dots \rho[n])$ is bounded by $\mu_i^{k-1}(\mathbf{e}')$, where $\mathbf{e}'$ denotes the first edge of ρ which belongs to $Z_{>j}$. Therefore,

$$\mathsf{cost}_i(\rho) \le (n \times |C| \times (\max b_j(\mu^{k-1}, \tilde{c}) + 1)) \times \kappa + \max_{l>j} \max b_l(\mu^{k-1}, \tilde{c})$$
$$\le (n \times |C| + |V|) \times \kappa + (n \times |C|) \times \max b_j(\mu^{k-1}, \tilde{c}) \times \kappa.$$

If Player i reaches tgt within X'_j, $\mathsf{cost}_i(\rho)$ would be much smaller. In the above, we have used the bound from corollary C.9 as $\mu^{k-1}(\mathbf{e}) = \lambda^{l^*}(\mathbf{e})$ for $\mathbf{e} \in X_l$ for $l > j$. Now as $\tilde{c} \in X_j$, we can use induction hypothesis for giving bound to $\max b_j(\mu^{k-1}, \tilde{c})$. As the above shown bound works for any $\rho \in \Gamma_{\mu^{k-1}}(\tilde{c})$, it works for the supremum too, hence we have

$$\mu_i^k(c, w, c') \le \sup_{\rho \in \Gamma_{\mu^{k-1}}(\tilde{c})} \mathsf{cost}_i(c', \tilde{c}) + \mathsf{cost}_i(\rho)$$
$$\le |V| \times \kappa + (n \times |C| + |V|) \times \kappa + (n \times |C|) \times \max b_j(\mu^{k-1}, \tilde{c}) \times \kappa$$
$$\le (n|C| + 2|V|) \times \kappa + (n|C| \times \kappa) \times \{(n|C| + 2|V|) \times \sum_{l=1}^{k-1} (n|C|)^{l-1} \cdot \kappa^l$$
$$= (n|C| + 2|V|) \times \sum_{l=1}^{k} (n|C|)^{l-1} \cdot \kappa^l$$

◀

This ends having bound for any $\max b_l(\mu, c)$ for any λ^{j^*}-building function μ, $c \in X_j$, and $l \ge j$. Hence, we can conclude that the counter graph $\mathbb{C}[\mu, c] = \langle C', T' \rangle$ can be made finite, by taking $C' = \{(d, b) \in C \times ([0; Y] \cup \{+\infty\})^{[\![n]\!]} \mid c \Rightarrow^* d\}$, with $Y = n^{|V|} \cdot |V|^{n.|V|} \cdot \kappa^{|V|}$, which is doubly-exponential in the encoding of n. Note that, this makes $|C'|_r$ at most double-exponential too - which will be one of the key arguments in the complexity analysis of the final algorithm.

Algorithm. Now we describe in details how we check non-emptiness of $\Gamma_{\lambda^*}(c_{\mathsf{src}})$. Starting from λ^{n^*}, we inductively compute $\lambda^{j^*} = \langle \lambda_i^{j^*} \rangle_{i \in [\![n]\!]}$. When computing λ^{j^*}, we use an intermediary induction (following the definition): we initialize with the definition of $\mu^0 = \langle \mu_i^0 \rangle_{i \in [\![n]\!]}$. Then, assuming we have already computed $\{\mu_i^{k-1}(\mathbf{e}) \mid \mathbf{e} \in T\}$, a general step to compute $\mu_i^k(c, w, c')$ is as follows:

- First we need to check that $\mu_i^k(c, w, c')$ is $-\infty$. For that, we check, one by one, whether for each of $(c, w', c'') \in T$, $\Gamma^{\mu^{k-1}}(c'')$ is not $\emptyset$. In order to do that,we guess a valid path (but do not store) in $\mathbb{C}[\mu^{k-1}, c'']$. If there is such a path in $\mathbb{C}[\mu^{k-1}, c'']$, there will be one, length of which is bounded by $|C|$, which is doubly-exponential in the size of input. Hence we keep a counter which we increase by 1 every time we correctly guess a new edge along the path, and the counter stops either when we reach a (c_{tgt}, b), or at the latest when it crosses $|C|$ marks - this counter requires exponential space to encode.

- if the previous check failed, we know that $\mu_i^k(c, w, c')$ would be in $\mathbb{N} \cup \{+\infty\}$. We now check if $\mu_i^k(c, w, c') = +\infty$. For that we need to verify the conditions stated in Lemma C.6, i.e, for each $c'' \in \mathsf{dev}_i(c, c')$, we guess a valid path π of the form $h \cdot \beta \cdot h'$ in $\mathbb{C}[\mu^{k-1}, c'']$, where β is a cycle and Player i's counter value is > 0 inside β. We first guess the first vertex of β, say (c_1, b_1) (with $b_1(i) > 0$), from which the cycle starts, then guess a cycle on (c_1, b_1), keeping only the current edge in memory. Now if there is a cycle on (c_1, b_1), there has to be a cycle within $|C|$ length because within a cycle no counter value can change. Then we guess a path from $(c'', b^{c''})$ to (c_1, b_1), and another path from (c_1, b_1) to (c_{tgt}, b) (for some b). The length of the part up to (c_1, b_1) is bounded by $|C'|_r$, whilst we can always get a path of length at most $|C|$ (if a path exists) from (c_1, b_1) to (c, b). This discrepancy between two bounds is mainly because for the latter part of the path, we do not exactly fix the final vertex for the latter part (b can be any tuple of n nonnegative values), we just want the first component to be c_{tgt}, while for the former part it gets fixed to (c_1, b_1) when we guessed earlier. If we can guess such a path π for each of $c'' \in \mathsf{dev}_i(c, c')$, we return $\mu_i^k(c, w, c') = +\infty$.
- otherwise, we have at least 1 and at most $|V|$ many configurations $c'' \in \mathsf{dev}_i(c, c')$ such that $\sup_{\rho \in \Gamma_{\mu^{k-1}}(c'')} \mathsf{cost}_i(\rho) \in \mathbb{N}$. We call this set of configurations as $\mathsf{dev}_i(c, c')|_{\mathsf{finsup}}$ To compute the finite value of $\mu_i^k(c, w, c')$, we compute $M_j = \sup_{\rho \in \Gamma_{\mu^{k-1}}(c''_j)} \mathsf{cost}_i(\rho)$ for each of $c''_j \in \mathsf{dev}_i(c, c')|_{\mathsf{finsup}}$, and then take the minimum. In order to compute M_j, we first tentatively set $M_j = \min_{e \in E} \mathsf{f}_e(1)$, and proceed iteratively as follows: we guess a valid path π of length at most $|C'|_r$ of $\mathbb{C}[\mu^{k-1}, c''_j]$ such that $\mathsf{cost}_i(\rho) \geq M_j$, where ρ is the corresponding path of π in $\mathcal{G}$. If such a path exists, we increase M_j by 1, and repeat. At some point, we get M_j such that there doesn't exist a valid path π of length at most $|C'|_r$ with Player i's cost larger than or equal to $M_j + 1$, but there exists a valid path with Player i's cost larger than or equal to M_j. We store that M_j. When all values M_j have been computed, we return $M = \min_{j \in [\![|\mathsf{dev}_i(c,c')|_{\mathsf{finsup}}|]\!]} \{\mathsf{cost}_i(c, c''_j) + M_j\}$. We use at most doubly-exponential space in the procedure of guessing a path, and we reuse that space for guessing next path in the above algorithm.

We keep another binary counter throughout transitioning from μ^{k-1} to μ^k to flag whether the fixpoint has been reached. Once we reach $\lambda^* = \langle \lambda_i^* \rangle$, we finally check whether $\Gamma_{\lambda^*}(c_{\mathsf{src}})$ is empty by guessing a path in $\mathbb{C}[\lambda^*, c_{\mathsf{src}}]$ from $(c_{\mathsf{src}}, b^{c_{\mathsf{src}}})$ to (c_{tgt}, b).

In conclusion, we use doubly-exponential space: (1) to store $\{\mu_i^{k-1}(\mathbf{e}) \mid i \in [\![n]\!], \mathbf{e} \in T\}$ which is double-exponential in the encoding of the number of players, and (2) to encode a counter which keeps checking whether the length of our guessed paths does not exceed $|C'|_r$.

▶ **Theorem 19.** *The existence of SPEs in a dynamic NCG can be decided in 2EXPSPACE.*