

Non-Blind Strategies in Timed Network Congestion Games*

Aline Goeminne, Nicolas Markey, and Ocan Sankur
firstname.lastname@irisa.fr

Univ Rennes, Inria, CNRS, IRISA – Rennes, France

Abstract. Network congestion games are a convenient model for reasoning about routing problems in a network: agents have to move from a source to a target vertex while avoiding congestion, measured as a cost depending on the number of players using the same link. Network congestion games have been extensively studied over the last 40 years, while their extension with timing constraints were considered more recently.

Most of the results on network congestion games consider *blind* strategies: they are static, and do not adapt to the strategies selected by the other players. We extend the recent results of [Bertrand *et al.*, Dynamic network congestion games. FSTTCS'20] to timed network congestion games, in which the availability of the edges depend on (discrete) time. We prove that computing Nash equilibria satisfying some constraint on the total cost (and in particular, computing the best and worst Nash equilibria), and computing the social optimum, can be achieved in exponential space. The social optimum can be computed in polynomial space if all players have the same source and target.

1 Introduction

Network congestion games allow one to model situations in which agents compete for resources such as routes or bandwidth [Ros73], e.g. in communication networks [AKH09, QYZS06]. In these games, the objective of each agent is to go from a source to a target vertex, traversing a number of edges that represent resources. The cost incurred by the player for the use of each resource is a function of the *load*, that is, the number of agents using the same resource. One of the fundamental notions studied in these games is that of *Nash equilibria* which is used to model stable situations. A strategy profile is a Nash equilibrium if none of the players can reduce their cost by unilaterally changing their strategy.

It is well-known that these games can be *inefficient* in the sense that there are Nash equilibria whose social cost (*i.e.*, the sum of the agents' costs) is bounded away from the optimum that can be achieved by arbitrary profiles (that may not be Nash equilibria). Research has been focused on proving bounds on this inefficiency, formalized by the *price of anarchy (PoA)* [KP09]. A tight bound of $\frac{5}{2}$

* This work was partially funded by ANR project Ticktac (ANR-18-CE40-0015).

on the price of anarchy was given in [AAE05, CK05]. The *price of stability (PoS)* is dual to PoA: it is the ratio between the cost of the best Nash equilibrium and the social cost, and was introduced in [ADK⁺04]. Bounds on PoA and PoS have been studied for restricted classes of graphs or types of cost functions [NRTV07].

Timed network games. Extensions of these games with real-time constraints have been considered. In the setting of [HMRT11], each edge is traversed with a fixed duration, independent of its load, while the cost is still a function of the load at the traversal time. The model thus has *time-dependent* costs since the load depends on the times at which players traverse a given edge. The existence of Nash equilibria is proven by reduction to [Ros73]. The work in [AGK17] propose another real-time extension, in which transitions are instantaneous, and can only be taken during some intervals of time. Time elapses in vertices, which incurs a cost that is proportional to the load and the delay. In those works, time is continuous; *boundary strategies* are strategies that take transitions at dates that are boundaries of the constraining intervals. It is shown that *boundary Nash equilibria* always exist, but need not capture optimal Nash equilibria. The prices of anarchy and stability are shown to be bounded by $5/2$ and $1 + \sqrt{3}/3$, respectively, and computing *some* Nash equilibrium is PLS-complete. This study was further extended to richer timing constraints (involving clocks) in [AGK18].

Non-blind strategies. In all the works mentioned above, the considered strategies are *blind*: each player chooses a path to follow at the beginning of the game, and cannot adapt their behaviors to their observations during the game. *Non-blind* strategies which allow players to choose their next moves depending on the whole history were studied in [BMSS20]. The advantage of non-blind players is that these allow one to obtain Nash equilibria that have a lower social cost than with blind profiles. Thus, in general, the price of anarchy is lower with non-blind strategies. In [BMSS20], the existence of Nash equilibria was established for these strategies, and an algorithm was given to decide the existence of Nash equilibria satisfying given constraints on the costs of the players.

Our contributions. We pursue the study of the real-time extensions of network congestion games by considering non-blind strategies. We consider timed network congestion games similar to [AGK17] (albeit with a discrete-time semantics), in which the edges in the network are guarded by a time interval which defines the time points at which the edges can be traversed. Moreover, each vertex is endowed with a cost function depending on the load, and the players incur a cost for each time unit spent on the vertex. We consider both the *symmetric* case in which all players source and target vertex pairs are identical, and the *asymmetric* case where these pairs vary. We formally define the semantics of our setting with non-blind strategies, show how to compute the social optimum in PSPACE in the symmetric case and in EXPSPACE in the asymmetric case. Moreover, we give an algorithm to decide the existence of Nash equilibria satisfying a given set of cost constraints in EXPSPACE, for both the symmetric and asymmetric cases. We can then compute the prices of anarchy and of stability in EXPSPACE.

Related Works. The existence of Nash equilibria in all atomic congestion games is proven using potential games. The notion of potential was generalized by Monderer and Shapley [MS96] who showed how to iteratively use best-response computation to obtain a Nash equilibrium. We refer the interested reader to [Rou07] for an introduction and main results on the subject.

Timing constraints are also considered in [PPT09, KP12] where travel times also depend on the load. Other works focus on flow models with a timing aspect [KS11, BFA15].

2 Preliminaries

2.1 Timed network congestion game.

Timed network. A *timed network* \mathcal{A} is a tuple $(V, E, \text{wgt}, \text{guard})$ where V is a set of vertices, $E \subseteq V \times V$ is a set of edges, $\text{wgt}: V \rightarrow (\mathbb{N}_{>0} \rightarrow \mathbb{N}_{>0})$ associates with each vertex a non-decreasing weight function, and $\text{guard}: E \rightarrow \mathcal{P}(\mathbb{N})$ associates with each edge the dates at which that edge is available. We require in the sequel that for all $e \in E$, $\text{guard}(e)$ is a finite union of disjoint intervals with bounds in $\mathbb{N} \cup \{+\infty\}$. A guard $\text{guard}(e)$ is said to be satisfied at date d if $d \in \text{guard}(e)$.

A *trajectory* in a timed network $\mathcal{A} = (V, E, \text{wgt}, \text{guard})$ is a (finite or infinite) sequence $(v_j, d_j)_{0 \leq j \leq l}$ such that for all $0 \leq j < l$, $(v_j, v_{j+1}) \in E$ and $d_{j+1} \in \text{guard}(v_j, v_{j+1})$. We write $\text{Traj}_{\mathcal{A}}(v, d, v')$ for the set of trajectories with $v_0 = v$, $d_0 = d$, and $v_l = v'$, possibly omitting to mention \mathcal{A} when no ambiguity arises.

Timed network game. A *timed network congestion game* (a.k.a. *timed network game*, or TNG for short) is a model to represent and reason about the congestion caused by the simultaneous use of some resources by several users.

Definition 1. A timed network game \mathcal{N} is a tuple $(n, \mathcal{A}, (\text{src}_i, \text{tgt}_i)_{1 \leq i \leq n})$ where

- n is the number of players (in binary). We write $\llbracket n \rrbracket$ for the set of players;
- $\mathcal{A} = (V, E, \text{wgt}, \text{guard})$ is a timed network;
- for each $1 \leq i \leq n$, the pair $(\text{src}_i, \text{tgt}_i) \in V^2$ is the objective of Player i .

Symmetric TNGs are TNGs in which all players have the same objective (src, tgt) , i.e., $\text{src}_i = \text{src}$ and $\text{tgt}_i = \text{tgt}$ for all $i \in \llbracket n \rrbracket$.

Remark 1. Several encodings can be used for the description of $(\text{src}_i, \text{tgt}_i)_{1 \leq i \leq n}$, which may impact the size of the input: for symmetric TNGs, the players' objectives would naturally be given as a single pair (src, tgt) . For the asymmetric case, the players' objectives could be given explicitly as a list of (src, tgt) -pairs (the size is then linear in n and logarithmic in $|V|$), or as a function $V^2 \rightarrow \mathbb{N}$ (size quadratic in $|V|$ and logarithmic in n). Usually, the number of players is large (and can be seen as a parameter), and the size of the input is at least linear in $|V|$, so that the latter representation is preferred. \triangleleft

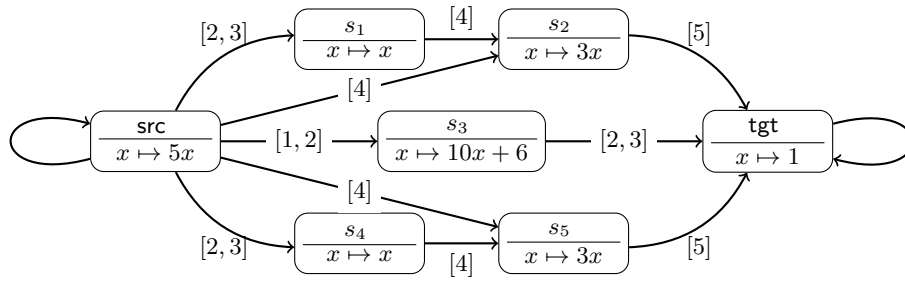


Fig. 1. Example of a timed network game

A *configuration* of a TNG \mathcal{N} is a mapping $c: \llbracket n \rrbracket \rightarrow V$ that maps each player to some vertex. We write Conf for the set of all configurations of \mathcal{N} . There are two particular configurations: the *initial configuration* init , in which all players are in their source vertices, *i.e.*, $\text{init}: i \mapsto \text{src}_i$; and the final configuration final in which all players are in their target vertices, *i.e.*, $\text{final}: i \mapsto \text{tgt}_i$. A timed configuration of \mathcal{N} is a pair (c, d) where c is a configuration and $d \in \mathbb{N}$ is the current time. We write TConf for the set of all timed configurations of \mathcal{N} , and st for the starting timed configuration $(\text{init}, 0)$.

Remark 2. We assume in the sequel that for each $0 \leq i \leq n$, there is a trajectory from $(\text{src}_i, 0)$ to the target vertex tgt_i ; this can be checked in polynomial time by exploring \mathcal{N} . We also require that for each vertex $v \in V$ and each date d , there exist a vertex v' and a time $d' > d$ such that $(v, v') \in E$ and $d' \in \text{guard}(v, v')$. \triangleleft

Example 1. Figure 1 represents a timed network \mathcal{A} (that we will use throughout this section to illustrate our formalism). The guards label the corresponding edges with two conventions: guards of the form $[d, d]$ for $d \in \mathbb{N}$ are denoted by $[d]$ and edges without label have $[0, +\infty)$ as a guard, *e.g.*, $\text{guard}(\text{src}, s_4) = [2, 3]$ and $\text{guard}(\text{src}, \text{src}) = [0, +\infty)$.

The weight function wgt is defined inside the corresponding vertex, *e.g.*, $\text{wgt}(\text{src}): x \mapsto 5x$ and $\text{wgt}(\text{tgt}): x \mapsto 1$. As explained in the sequel, this function indicates the cost of spending one time unit in each vertex, depending on the number of players. For instance, if two players spend one time unit in src , they both pay 10. Notice that to comply with Remark 2, extra edges with guard $[6, +\infty)$ are added from each vertex s_k ($1 \leq k \leq 5$) to an additional sink vertex (omitted in the figure for clarity). \triangleleft

Concurrent game associated with a timed network game. Consider a TNG $\mathcal{N} = (n, \mathcal{A}, (\text{src}_i, \text{tgt}_i)_{1 \leq i \leq n})$, with $\mathcal{A} = (V, E, \text{wgt}, \text{guard})$. The semantics runs intuitively as follows: a timed configuration represents the positions of the players in the network at a given date. At each round, each player selects the transition they want to take in \mathcal{A} , together with the date at which they want to take it. All players selecting the earliest date will follow the edges they selected, giving rise to a new timed configuration, from where the next round will take place.

We express this semantics as an n -player infinite-state weighted concurrent game $\mathcal{G} = (\text{States}, \text{Act}, (\text{Allow}_i)_{1 \leq i \leq n}, \text{Upd}, (\text{cost}_i)_{1 \leq i \leq n})$ [AHK02]. The set **States** of states of \mathcal{G} is the set **TConf** of timed configurations. The set **Act** of actions is the set $\{(v, d) \mid v \in V, d \in \mathbb{N}_{\geq 0}\}$. Functions Allow_i return the set of allowed actions for Player i from each configuration: for a timed configuration (c, d) , we have $(v, d') \in \text{Allow}_i(c, d)$ whenever $d' > d$ and $(c(i), v) \in E$ and $d' \in \text{guard}(c(i), v)$. Notice that by Remark 2, $\text{Allow}_i(c, d)$ is never empty. An *action vector* in \mathcal{G} is a vector $\vec{a} = (a_i)_{1 \leq i \leq n}$ of actions, one for each player. We write **Mov** for the set of action vectors. An action vector \vec{a} is valid from state (c, d) if for all $1 \leq i \leq n$, action $a_i \in \text{Allow}_i(c, d)$. We write $\text{Valid}(c, d)$ for the set of valid action vectors from state (c, d) .

From a state (c, d) , a valid action vector $\vec{a} = (v_i, d_i)_{1 \leq i \leq n}$ leads to a new state (c', d') where $d' = \min\{d_i \mid 1 \leq i \leq n\}$ and for all $1 \leq i \leq n$, $c'(i) = v_i$ if $d_i = d'$, and $c'(i) = c(i)$ otherwise. This is encoded in the update function **Upd** by letting $\text{Upd}((c, d), \vec{a}) = (c', d')$. This models the fact that the players proposing the shortest delay apply their action once this shortest delay has elapsed. We write $\text{Select}(\vec{a}) = \{1 \leq i \leq n \mid \forall 1 \leq j \leq n. d_i \leq d_j\}$.

We let $\text{Trans} = \{(s, \vec{a}, s') \in \text{States} \times \text{Mov} \times \text{States} \mid \vec{a} \in \text{Valid}(s), \text{Upd}(s, \vec{a}) = s'\}$ be the set of *transitions* of \mathcal{G} . We write G for the graph structure $(\text{States}, \text{Trans})$.

The delays elapsed in the vertices of \mathcal{N} come with a cost for each player, computed as follows: given a configuration c , we write $\text{load}_c: V \rightarrow \mathbb{N}$ for the load of each vertex of the timed network, i.e., the number of players sitting in that vertex: $\text{load}_c(v) = \#\{1 \leq i \leq n \mid c(i) = v\}$. For each $1 \leq i \leq n$, the cost for Player i of taking a transition $t = ((c, d), \vec{a}, (c', d'))$, moving from timed configuration (c, d) to timed configuration (c', d') , then is $\text{cost}_i(t) = (d' - d) \cdot \text{wgt}(c(i))(\text{load}_c(c(i)))$: it is proportional to the time elapsed and to the load of the vertex where Player i is waiting. Notice that it does not depend on \vec{a} (nor on c').

Remark 3. With our definition, players remain in the game even after they have reached their targets. This may be undesirable in some situations. One way of avoiding this is to add another limitation on the availability of transitions to players: transitions would be allowed only to some players, and in particular, from their target states, players would only be allowed to go to a sink state. All our results can be adapted to this setting. \triangleleft

Example 2. We illustrate those different notions on a two-player symmetric timed network game $\mathcal{N} = (2, \mathcal{A}, (\text{src}, \text{tgt}))$, where \mathcal{A} is the timed network of Example 1.

Let $(c, d) = ((\text{src}, s_1), 2)$ be the timed configuration in **TConf**, with $c(1) = \text{src}$, $c(2) = s_1$ and $d = 2$. From this timed configuration, the set of allowed actions for Player 1 is $\text{Allow}_1(c, d) = \{(s_1, 3), (s_4, 3), (s_2, 4), (s_5, 4), (\text{src}, d) \mid d \geq 3\}$. If we consider the valid action vector $\vec{a} = ((s_1, 3), (s_2, 4))$, then $\text{Upd}((c, d), \vec{a}) = (c', d')$ with $c' = (s_1, s_1)$ and $d' = 3$. Indeed even if both players choose their actions simultaneously, we have $\text{Select}(\vec{a}) = \{1\}$, and Player 1 is the only player who moves at time 3. Moreover, writing $t = ((c, d), \vec{a}, (c', d'))$, we have $\text{cost}_1(t) = 5$, because $\text{load}_{(c, d)}(\text{src}) = 1$. Finally, let us consider another example of cost computation: let $(c, d) = ((s_3, s_3), 1)$ and $(c', d') = ((\text{tgt}, \text{tgt}), 3)$, and $\vec{a} = ((\text{tgt}, 3), (\text{tgt}, 3))$; we have that $\text{cost}_1((c, d), \vec{a}, (c', d')) = 2 \cdot (10 \cdot 2 + 6) = 52$. \triangleleft

Plays and histories. A play $\rho = (s_k, \vec{a}_k, s'_k)_{k \in \mathbb{N}} \in \text{Trans}^\omega$ in \mathcal{G} (also called play in \mathcal{N}) is an infinite sequence of transitions such that for all $k \in \mathbb{N}$, $s'_k = s_{k+1}$. We denote by Plays the set of plays of \mathcal{G} . We denote by $\text{Plays}(s)$ the set of plays that start in state $s \in \text{States}$, i.e. $\text{Plays}(s) = \{(s_k, \vec{a}_k, s'_k)_{k \geq 0} \in \text{Plays} \mid s_0 = s\}$.

A history is a finite prefix of a play. We denote the set of histories by Hist , and the subset of histories starting in a given state s by $\text{Hist}(s)$.

Given a play $\rho = (s_k, \vec{a}_k, s'_k)_{k \geq 0} \in \text{Plays}$ and an integer $j \in \mathbb{N}$, we write ρ_j for the timed configuration s_j , $\rho_{\geq j}$ for the suffix $(s_{j+k}, \vec{a}_{j+k}, s'_{j+k})_{k \geq 0} \in \text{Plays}(s_j)$, and $\rho_{< j}$ for the history $(s_k, \vec{a}_k, s'_k)_{k < j} \in \text{Hist}(s_0)$. For a history $h = (s_k, \vec{a}_k, s'_k)_{0 \leq k < j}$ in $\text{Hist}(s_0)$, we write $\text{last}(h) = s'_{j-1}$ when $j > 0$ (we may also write it as s_j when no ambiguity arises), and $\text{last}(h) = s_0$ otherwise.

Cost functions. For each player $i \in \llbracket n \rrbracket$, we define a cost function $\text{cost}_i: \text{Plays} \rightarrow \mathbb{N} \cup \{+\infty\}$ such that for all $\rho = (s_k, \vec{a}_k, s'_k)_{k \in \mathbb{N}} \in \text{Plays}$,

$$\text{cost}_i(\rho) = \begin{cases} +\infty & \text{if } s_k(i) \neq \text{tgt}_i \text{ for all } k \in \mathbb{N} \\ \sum_{k=0}^{\ell-1} \text{cost}_i(s_k, \vec{a}_k, s'_k) & \text{if } \ell \text{ is the least index such that } s_\ell(i) = \text{tgt}_i \end{cases}$$

This function is extended to histories in the natural way. For all $\rho \in \text{Plays} \cup \text{Hist}$, the vector $\text{cost}(\rho) = (\text{cost}_i(\rho))_{1 \leq i \leq n}$ is the *cost profile* of ρ .

Strategies. Given a concurrent game \mathcal{G} , a state s in \mathcal{G} , and $1 \leq i \leq n$, a *strategy* for Player i from s is a function $\sigma_i: \text{Hist}(s) \rightarrow \text{Act}$ such that $\sigma_i(h) \in \text{Allow}_i(\text{last}(h))$ for all $h \in \text{Hist}(s)$. We denote by $\Sigma_i(s)$ the set of strategies of Player i from s (we may omit to mention s when it is the initial timed configuration st). Given a state s , a subset $I \subseteq \llbracket n \rrbracket$ of players, and strategies $(\sigma_i)_{i \in I}$ from s for those players, a play $\rho = (s_k, \vec{a}_k, s'_k)_{k \in \mathbb{N}}$ from s is consistent with $(\sigma_i)_{i \in I}$ if for all $k \in \mathbb{N}$ and all $i \in I$, it holds $\vec{a}_{k,i} = \sigma_i(\rho_{< k})$.

A *strategy profile* $\sigma = (\sigma_i)_{1 \leq i \leq n}$ from s is a tuple of strategies from s , one for each player. We write $\Sigma(s)$ for the set of strategy profiles from s . Given a strategy profile σ from s , there is a unique play ρ from s that is *consistent* with that strategy profile. This play is denoted by $\langle \sigma \rangle_s$ and is called the *outcome* of the strategy profile σ from s .

Given a strategy profile σ from s , a player $i \in \llbracket n \rrbracket$ and a strategy σ'_i from s for Player i , we write (σ_{-i}, σ'_i) for the strategy profile $(\tau_j)_{1 \leq j \leq n}$ for which $\tau_j = \sigma_j$ when $j \neq i$ and $\tau_i = \sigma'_i$. Given a player $1 \leq i \leq n$ and a strategy σ_i from s for Player i , for all $h \in \text{Hist}(s)$, we denote by $\sigma_{i \uparrow h}$ the strategy of Player i from $\text{last}(h)$ such that $\sigma_{i \uparrow h}: h' \in \text{Hist}(\text{last}(h)) \mapsto \sigma_i(hh')$. This is extended to strategy profiles in the natural way.

Blind strategies. An important class of strategies is the class of *blind strategies*. Intuitively, a blind strategy follows a single trajectory, without looking at how the other players are moving. In order to define blind strategies, we first introduce a notion of *projection* of plays on the actions of each player.

For any play $\rho = (s_k, \vec{a}_k, s'_k)_{k \in \mathbb{N}}$ from st and any $1 \leq i \leq n$, writing $\vec{a}_k = (v_{k,i}, d_{k,i})_{1 \leq i \leq n}$ for all $k \in \mathbb{N}$, we define $\mu_i: \mathbb{N} \mapsto \mathbb{N} \cup \{+\infty\}$ inductively as

$\mu_i(0) = -1$ and, for all j for which $\mu_i(j) < +\infty$,

$$\mu_i(j+1) = \inf\{k > \mu_i(j) \mid d_{k,i} = \min\{d_{k,l} \mid 1 \leq l \leq n\}\}.$$

In other terms, μ_i returns the indices where Player i proposed a minimal delay, and thus could move to the vertex they proposed. The trajectory of Player i along ρ , denoted with $\text{traj}_i(\rho)$, is then defined as the trajectory $(v_{\mu_i(j),i}, d_{\mu_i(j),i})_{j \geq 0}$, with the convention that $(v_{-1,i}, d_{-1,i}) = (\text{src}_i, 0)$. Notice that this trajectory could be finite. Functions traj_i are extended to histories in the natural way.

Definition 2. A strategy σ_i for Player i is blind if, for any two histories h and h' such that $\text{traj}_i(h) = \text{traj}_i(h')$, it holds $\sigma_i(h) = \sigma_i(h')$.

As the next lemma suggests, playing a blind strategy amounts to following a fixed trajectory in the timed network, independently of the actions of the other players.

Lemma 3. Let $i \in \llbracket n \rrbracket$ be a player and σ_i be a strategy of Player i from some state s_0 . If σ_i is blind, then for all strategies σ_{-i} and σ'_{-i} from s_0 , we have

$$\text{traj}_i(\langle \sigma_{-i}, \sigma_i \rangle_{s_0}) = \text{traj}_i(\langle \sigma'_{-i}, \sigma_i \rangle_{s_0}).$$

In view of this lemma, any blind strategy σ_i for Player i from some state s_0 can be associated with a trajectory, denoted with $\langle \sigma_i \rangle_{s_0}$, and defined as $\langle \sigma_i \rangle_{s_0} = \text{traj}_i(\langle \sigma_{-i}, \sigma_i \rangle_{s_0})$ (for any $\sigma_{-i} \in \Sigma_{-i}(s_0)$.) Conversely, for any trajectory π from (v_0, d_0) , any player $i \in \llbracket n \rrbracket$, and any timed configuration $s_0 = (c, d_0)$ with $c(i) = v_0$, there exists a blind strategy σ_i for Player i from s_0 whose associated trajectory is π .

A play $\rho = (s_k, \vec{a}_k, s'_k)_{k \in \mathbb{N}}$ from the initial timed configuration st is said *winning* for Player i if there exists $k \in \mathbb{N}$ such that $s_k = (c_k, d_k)$ with $c_k(i) = \text{tgt}_i$. A strategy σ_i for Player i from st is winning if any play consistent with that strategy is winning for Player i . For each $1 \leq i \leq n$, as we assumed $\text{Traj}(\text{src}_i, 0, \text{tgt}_i) \neq \emptyset$, we get that there exists a winning (blind) strategy σ_i for Player i .

Example 3. We consider the two-player symmetric TNG described in Example 2.

The infinite sequence of transitions $\rho: ((\text{src}, \text{src}), 0) \xrightarrow{\begin{bmatrix} (s_1, 2) \\ (s_2, 4) \end{bmatrix}} ((s_1, \text{src}), 2) \xrightarrow{\begin{bmatrix} (s_2, 4) \\ (s_2, 4) \end{bmatrix}}$

$((s_2, s_2), 4) \left(\xrightarrow{\begin{bmatrix} (\text{tgt}, 5+k) \\ (\text{tgt}, 5+k) \end{bmatrix}} ((\text{tgt}, \text{tgt}), 5+k) \right)_{k \geq 0}$ is a play in \mathcal{N} such that $\text{cost}_1(\rho) = 2 \cdot (5 \cdot 2) + 2 \cdot (1 \cdot 1) + 1 \cdot (3 \cdot 2) = 28$ and $\text{cost}_2(\rho) = 2 \cdot (5 \cdot 2) + 2 \cdot (5 \cdot 1) + 1 \cdot (3 \cdot 2) = 36$.

Let us now consider two trajectories $\pi_1: (\text{src}, 0)(s_1, 2)(s_2, 4)(\text{tgt}, 5+k)_{k \geq 0}$ and $\pi_2: (\text{src}, 0)(s_2, 4)(\text{tgt}, 5+k)_{k \geq 0}$ together with σ_1 , a blind strategy of Player 1, and σ_2 , a blind strategy of Player 2, such that $\langle \sigma_1 \rangle_{\text{st}} = \pi_1$ and $\langle \sigma_2 \rangle_{\text{st}} = \pi_2$. The outcome of the strategy profile (σ_1, σ_2) from st is ρ , i.e., $\langle \sigma_1, \sigma_2 \rangle_{\text{st}} = \rho$. \triangleleft

Social optima and Nash equilibria. Let \mathcal{G} be a concurrent game. The *social welfare* of a play ρ in \mathcal{G} is the sum of the costs of the players along ρ : $\text{SW}(\rho) = \sum_{1 \leq i \leq n} \text{cost}_i(\rho)$. The social welfare of a strategy profile is the social welfare of its outcome.

A strategy profile σ from the starting configuration st is a social optimum (SO) whenever $\text{SW}(\langle\sigma\rangle_{\text{st}}) = \inf_{\sigma' \in \Sigma(\text{st})} \text{SW}(\langle\sigma'\rangle_{\text{st}})$. When no ambiguity arises, social optimum may also refer to the social welfare of socially-optimal strategy profiles. Notice that since we consider discrete time, this value is an integer, and a strategy profile realizing the social optimum always exists.

A strategy profile σ from st is a *Nash equilibrium* (NE) if for all $1 \leq i \leq n$, for all strategies $\sigma'_i \in \Sigma_i(\text{st})$ of Player i , $\text{cost}_i(\langle\sigma\rangle_{\text{st}}) \leq \text{cost}_i(\langle\sigma_{-i}, \sigma'_i\rangle_{\text{st}})$: for all $1 \leq i \leq n$, the strategy σ_i of Player i is (one of) the best strategies against the strategies σ_{-i} of the other players.

Example 4. We show an example of a Nash equilibrium from st in the two-player symmetric TNG described in Example 2. The strategy of Player 1, denoted by σ_1 , is given by a blind strategy associated with the trajectory $\pi_1 = (\text{src}, 0)(s_3, 1)(\text{tgt}, 2 + k)_{k \geq 0}$. The strategy of Player 2, denoted by σ_2 , is a more involved. The first action of Player 2 from the timed configuration st is $(\text{src}, 1)$, then at time 1 they observe if Player 1 has complied with their strategy, *i.e.*, if Player 1 is in s_3 . If so, Player 2 moves to s_4 at time 2, to s_5 at time 4, and ends in tgt at time 5; otherwise, they wait in src in order to observe the exact deviation of Player 1, and punish them adequately, by moving from src to s_5 at time 4 if Player 1 is in s_4 or to s_2 otherwise, and ending in tgt at time 5. The outcome $\langle\sigma_1, \sigma_2\rangle_{\text{st}}$ of this strategy profile from st is:

$$\rho: \text{st} \xrightarrow{\begin{bmatrix} (s_3, 1) \\ (\text{src}, 1) \end{bmatrix}} ((s_3, \text{src}), 1) \xrightarrow{\begin{bmatrix} (\text{tgt}, 2) \\ (s_4, 2) \end{bmatrix}} ((\text{tgt}, s_4), 2) \xrightarrow{\begin{bmatrix} (\text{tgt}, 3) \\ (s_5, 4) \end{bmatrix}} ((\text{tgt}, s_4), 3) \\ \xrightarrow{\begin{bmatrix} (\text{tgt}, 4) \\ (s_5, 4) \end{bmatrix}} ((\text{tgt}, s_5), 4) \left(\xrightarrow{\begin{bmatrix} (\text{tgt}, 5+k) \\ (\text{tgt}, 5+k) \end{bmatrix}} ((\text{tgt}, \text{tgt}), 5+k) \right)_{k \geq 0}.$$

We have that $\text{Cost}_1(\rho) = 26$ and $\text{cost}_2(\rho) = 20$. In particular, the social welfare of ρ is $\text{SW}(\rho) = 46$.

We can prove that (σ_1, σ_2) is a Nash equilibrium. We only explain here the interest of the threat of punishment of Player 2 by considering the deviating strategy σ'_1 of Player 1 such that he moves from src to s_1 at time 2, to s_2 at time 4 and ends in tgt at time 5. The outcome $\langle\sigma'_1, \sigma_2\rangle_{\text{st}}$ of this new strategy profile from st is

$$\rho': \text{st} \xrightarrow{\begin{bmatrix} (s_1, 2) \\ (\text{src}, 1) \end{bmatrix}} ((\text{src}, \text{src}), 1) \xrightarrow{\begin{bmatrix} (s_1, 2) \\ (\text{src}, 2) \end{bmatrix}} ((s_1, \text{src}), 2) \xrightarrow{\begin{bmatrix} (s_2, 4) \\ (s_2, 4) \end{bmatrix}} \\ ((s_2, s_2), 4) \left(\xrightarrow{\begin{bmatrix} (\text{tgt}, 5+k) \\ (\text{tgt}, 5+k) \end{bmatrix}} ((\text{tgt}, \text{tgt}), 5+k) \right)_{k \geq 0}.$$

The new cost of Player 1 is $\text{cost}_1(\rho') = 28$ proving that σ'_1 is not a provitable deviation for Player 1 w.r.t. (σ_1, σ_2) .

Remark that if Player 2 does not apply this punishment and moves to s_4 at time 2 whatever the behavior of Player 1, playing the deviating strategy σ'_1 would be a profitable deviation for Player 1. \triangleleft

2.2 Studied problems

Given a timed network game \mathcal{N} , the price of anarchy of \mathcal{N} , denoted by $\text{PoA}_{\mathcal{N}}$, is the ratio between the worst social welfare of a Nash equilibrium and the social optimum. Similarly, the *price of stability* of \mathcal{N} , denoted by $\text{PoS}_{\mathcal{N}}$, is the ratio between the best social welfare of a Nash equilibrium and the social optimum. Those values measure the impact of playing selfishly. In this paper, we address the following three problems:

Problem 1 (Constrained social welfare). Given a timed network game \mathcal{N} and a threshold $x \in \mathbb{N}$, is the social optimum in \mathcal{N} less than or equal to x ?

Problem 2 (Constrained existence of a Nash equilibrium). Given a timed network game \mathcal{N} with n players and a family of linear constraints $(\phi_q)_q$ over n variables, does there exist a Nash equilibrium σ from st such that all constraints $\phi_q(\text{cost}(\langle \sigma \rangle_{\text{st}}))$ are satisfied?

Problem 3 (Constrained price of anarchy and stability). Given a timed network game \mathcal{N} and a threshold $x \in \mathbb{Q}$, is the price of anarchy (resp. of stability) in \mathcal{N} less than or equal to x ?

Theorem 4. *The constrained-social-welfare problem is in PSPACE in the symmetric case, and in EXPSPACE in the asymmetric setting. The constrained existence of a Nash equilibrium and constrained price of anarchy and stability are in EXPSPACE in both the symmetric and asymmetric cases.*

3 Existence and computation of Nash equilibria

We first notice that, similarly to the untimed case [BMSS20], there are more Nash equilibria in non-blind strategies than when restricting to blind strategies:

Proposition 5. *There exists a timed network game \mathcal{N} that admits a Nash equilibrium σ from st , and whose all blind Nash equilibria τ from st are such that $\text{SW}(\langle \sigma \rangle_{\text{st}}) < \text{SW}(\langle \tau \rangle_{\text{st}})$.*

In this section, we transform the infinite concurrent game associated with a timed network congestion game into a finite one, preserving the set of costs of Nash equilibria. We use this finite game to solve the constrained-Nash-equilibrium problem, and then explain how to compute witnessing Nash equilibria.

3.1 Transformation into an equivalent finite concurrent game

Let $\mathcal{N} = (n, \mathcal{A}, (\text{src}_i, \text{tgt}_i)_{1 \leq i \leq n})$ be a timed network game, and M be the largest integer appearing in the guards of \mathcal{A} . In our assumption of Remark 2 that, for all $1 \leq i \leq n$, there must exist a trajectory $(v_j, d_j)_{j \in \mathbb{N}}$ from $(\text{src}_i, 0)$ to (tgt_i, δ_i) for some date δ_i , we can additionally require that $\delta_i \leq M + |V|$. Indeed, consider

such a winning trajectory with minimal δ_i : then either $\delta_i \leq M$, or for some j_0 , $d_{j_0-1} < M$ and $d_{j_0} \geq M$. Since M is the maximal constant appearing in the guard of \mathcal{A} , we can then modify the dates after d_{j_0} so that $d_{j_0} = M$ and $d_{j_0+k} = M + k$, while still satisfying the guards.

Let $\kappa = \max_{v \in V} \text{wgt}(v)(n)$ be the maximum cost per time unit that may occur in the game. With the arguments above, all players can always ensure a cost of at most $K = \kappa \cdot (M + |V|)$. Since each time unit costs at least one cost unit, $\kappa \cdot (M + |V|)$ is also a bound on the maximum time within which any player must have reached their target vertex in any Nash equilibrium: if this were not the case, then that player would have a profitable deviation. It follows:

Lemma 6. *Let $\rho = (s_k, \vec{a}_k, s'_k)_{k \geq 0}$ be a play in \mathcal{G} . For all $1 \leq i \leq n$, if $\text{cost}_i(\rho)$ is finite, then there exists a position $k^* \leq \text{cost}_i(\rho)$ such that $s_{k^*} = (c_{k^*}, d_{k^*})$ with $c_{k^*}(i) = \text{tgt}_i$ and $d_{k^*} \leq \text{cost}_i(\rho)$.*

This applies to the particular case where ρ is the outcome of a Nash equilibrium, since in this case, for all $1 \leq i \leq n$, $\text{cost}_i(\rho)$ is finite, and bounded by K .

Definition of the finite concurrent game. From the lemma above, when looking for Nash equilibrium, we can unfold the concurrent game \mathcal{G} into a tree (actually, a directed acyclic graph) and prune all subtrees at depth K . Formally:

Definition 7. *Let \mathcal{N} be a timed network game, and $K = \kappa \cdot (M + |V|)$. With this timed network game, we associate the (finite) concurrent game $\mathcal{G}^F = (\text{States}^F, \text{Act}^F, (\text{Allow}_i^F)_{1 \leq i \leq n}, \text{Upd}^F, (\text{cost}_i^F)_{1 \leq i \leq n})$ defined as follows:*

- $\text{States}^F = \{(c, d) \in \text{TConf} \mid 0 \leq d \leq K + 1\}$ is the set of time-bounded timed configurations;
- $\text{Act}^F = \{(v, d) \in \text{Act} \mid d \leq K + 1\} \cup \{(\perp, K + 1)\}$ is the finite set of actions;
- for any $1 \leq i \leq n$ and any state $(c, d) \in \text{States}^F$ with $d < K$, we have action $(v, d') \in \text{Allow}_i^F(c, d)$ if, and only if, $d < d' \leq K + 1$ and $(c(i), v) \in E$ and $d' \in \text{guard}(c(i), v)$. For states (c, d) with $d \geq K$, we have $\text{Allow}_i^F(c, d) = \{(\perp, K + 1)\}$. We write Mov^F for the set of action vectors. An action vector $\vec{a} = (a_i)_{1 \leq i \leq n}$ is in $\text{Valid}^F(c, d)$ if each a_i is in $\text{Allow}_i^F(c, d)$;
- for a state $s = (c, d)$ and a valid action vector $\vec{a} = (v_i, d'_i)_{1 \leq i \leq n}$, writing $m = \min\{d'_i \mid 1 \leq i \leq n\}$, we have $\text{Upd}^F(s, \vec{a}) = (c, K + 1)$ if $m = K + 1$, and $\text{Upd}^F(s, \vec{a}) = \text{Upd}(s, \vec{a})$ otherwise;
- in the same way as for (infinite) concurrent games, we let $\text{Trans}^F = \{(s, \vec{a}, s') \in \text{States}^F \times \text{Mov}^F \times \text{States}^F \mid \vec{a} \in \text{Valid}^F(s) \text{ and } s' = \text{Upd}^F(s, \vec{a})\}$ be the set of transitions of \mathcal{G}^F . Cost functions cost_i^F are defined on Trans^F in the same way as for \mathcal{G} : we set $\text{cost}_i^F((c, d), \vec{a}, (c', d'))$ as $(d' - d) \cdot \text{wgt}(c(i))(\text{load}_c(c(i)))$ if $d' \leq K$, and as 0 otherwise.

Notice that States^F has size doubly-exponential in the size of the input, since the number of players is given in binary. We write G^F for the graph $(\text{States}^F, \text{Trans}^F)$. Plays, histories, costs and strategies are defined as for infinite concurrent games.

Remark 4. By construction of \mathcal{G}^F , any play $\rho = (s_k, \vec{a}_k, s'_k)_{k \geq 0}$ in \mathcal{G}^F ends in a self-loop on some configuration $(c, K + 1)$, where the only valid action vector \vec{a}

is such that $a_i = (\perp, K + 1)$ for all $1 \leq i \leq n$. The prefix of ρ before entering this loop corresponds to a prefix of a play ρ' in \mathcal{G} , with the difference that some actions (with $d > K$) in \mathcal{G} may not be available in \mathcal{G}^F , but can be modified by setting $d = K + 1$. As a consequence, Lemma 6 also holds for plays in \mathcal{G}^F .

Also, since the set of actions is finite, and that there is a single available action in the terminal self-loop, there is a finite number of strategies in \mathcal{G}^F . \triangleleft

We now establish a correspondence between the Nash equilibria of \mathcal{G}^F and \mathcal{G} :

Theorem 8. *Let \mathcal{N} be a timed network game and let \mathcal{G} and \mathcal{G}^F be its associated infinite and finite concurrent games. Let $x \in (\mathbb{N} \cup \{+\infty\})^n$. Then there exists a Nash equilibrium σ in \mathcal{G} with $\text{cost}(\langle \sigma \rangle_{st}) = x$ if, and only if, there exists a Nash equilibrium τ in \mathcal{G}^F with $\text{cost}^F(\langle \tau \rangle_{st}) = x$*

We prove this result by establishing simulation relations between \mathcal{G} and \mathcal{G}^F , satisfying special properties for Proposition 9 below to apply. We let $\sim_{\text{BT}} \subseteq \text{States} \times \text{States}^F$ such that, for any $s_1 = (c_1, d_1) \in \text{States}$ and for any $s_2 = (c_2, d_2) \in \text{States}^F$, $s_1 \sim_{\text{BT}} s_2$ if, and only if, either (i) $d_1 = d_2 < K$ and $c_1 = c_2$, or (ii) $d_1 > K$ and $d_2 > K$.

We then use this relation (and its inverse) in the following generic proposition in order to prove the correspondence between the Nash equilibria in \mathcal{G} and \mathcal{G}^F :

Proposition 9. *Let $\mathcal{G} = (\text{States}, \text{Act}, (\text{Allow}_i)_{1 \leq i \leq n}, \text{Upd}, (\text{cost}_i)_{1 \leq i \leq n})$ and $\mathcal{G}' = (\text{States}', \text{Act}', (\text{Allow}'_i)_{1 \leq i \leq n}, \text{Upd}', (\text{cost}'_i)_{1 \leq i \leq n})$ be two n -player concurrent games, $s_0 \in \text{States}$ and $s'_0 \in \text{States}'$, and $\triangleleft \subseteq \text{States} \times \text{States}'$ be a relation such that*

1. $s_0 \triangleleft s'_0$;
2. there exists $\lambda \in \mathbb{N}$ such that for any NE σ in \mathcal{G} from s_0 , for any $1 \leq i \leq n$, it holds $\text{cost}_i(\langle \sigma \rangle_{s_0}) \leq \lambda$;
3. for all plays $\rho \in \text{Plays}_{\mathcal{G}}$ and $\rho' \in \text{Plays}_{\mathcal{G}'}$ such that $\rho \triangleleft \rho'$ (i.e., $\rho_j \triangleleft \rho'_j$ for all $j \in \mathbb{N}$), all $1 \leq i \leq n$, if $\text{cost}_i(\rho) \leq \lambda$ or $\text{cost}'_i(\rho') \leq \lambda$, then $\text{cost}_i(\rho) = \text{cost}'_i(\rho')$;
4. For all $s \in \text{States}$, for all $\vec{a} \in \text{Valid}(s)$, for all $s' \in \text{States}'$, if $s \triangleleft s'$, then there exists $\vec{a}' \in \text{Valid}'(s')$ such that:
 - (a) $\text{Upd}(s, \vec{a}) \triangleleft \text{Upd}'(s', \vec{a}')$;
 - (b) for all $1 \leq i \leq n$, for all $b'_i \in \text{Allow}'_i(s')$, there exists $b_i \in \text{Allow}_i(s)$ such that

$$\text{Upd}(s, (a_{-i}, b_i)) \triangleleft \text{Upd}'(s', (a'_{-i}, b'_i)).$$

Then for any Nash equilibrium σ in \mathcal{G} from s_0 , there exists a Nash equilibrium σ' in \mathcal{G}' from s'_0 such that $\text{cost}(\langle \sigma \rangle_{s_0}) = \text{cost}'(\langle \sigma' \rangle_{s'_0})$.

3.2 Existence of Nash equilibria

Theorem 10. *Let \mathcal{N} be a timed network game and let \mathcal{G} be its associated concurrent game. There exists a Nash equilibrium σ from st in \mathcal{G} .*

We prove this result using a *potential function* Ψ [MS96]; a potential function is a function that assigns a non-negative real value to each strategy profile, and decreases when the profile is *improved* (in a sense that we explain below). Since those improvements are performed among a finite set of strategy profiles, this entails convergence of the sequence of improvements.

An *improvement* consists in changing the strategy of one of the players by a better strategy for that player (if any), in the sense that their individual cost decreases. Because the game admits a potential function over finitely many strategies, this *best-response dynamics* must converge in finitely many steps, and the limit is a strategy profile where no single player can improve their strategy, *i.e.*, a Nash equilibrium.

We prove that timed network games admit a potential function when considering winning blind strategies; hence there exists Nash equilibria for that set of strategies. We then prove that a Nash equilibrium in this restricted setting remains a Nash equilibrium w.r.t. the set of all strategies.

For all $1 \leq i \leq n$, there are a finite number of strategies of Player i in \mathcal{G}^F . It follows that the set of winning blind strategies of Player i is also finite; we let

$$\Sigma_i^{WB} = \{\sigma_i \mid \sigma_i \text{ is a winning blind strategy of Player } i\}.$$

Let $\Sigma^{WB} = \Sigma_1^{WB} \times \dots \times \Sigma_n^{WB}$; Σ^{WB} is a finite set of strategy profiles. We consider a restriction of \mathcal{G}^F in which for all $1 \leq i \leq n$, Player i is only allowed to play a strategy in Σ_i^{WB} . A Σ^{WB} -Nash equilibrium then is a strategy profile in Σ^{WB} such that for all $1 \leq i \leq n$, Player i has no profitable deviation in Σ_i^{WB} . The next proposition implies that \mathcal{G}^F always admits Σ^{WB} -Nash equilibria [MS96].

Proposition 11. *Let \mathcal{N} be a timed network game and let \mathcal{G}^F be its associated finite concurrent game and Σ^{WB} be the set of winning blind strategy profiles. The game \mathcal{G}^F restricted to strategy profiles in Σ^{WB} has a potential function.*

Moreover each Σ^{WB} -Nash equilibrium corresponds to a Nash equilibrium with the same cost profile (Proposition 12). Notice that the converse result fails to hold, as proved in Proposition 5.

Proposition 12. *Let \mathcal{N} be a timed network game. Let \mathcal{G}^F be its associated finite concurrent game, and Σ^{WB} the set of winning blind strategy profiles. If there exists a Σ^{WB} -Nash equilibrium σ in \mathcal{G}^F from st , then there exists a Nash equilibrium τ from st in \mathcal{G}^F with the same costs for all players.*

3.3 Computation of Nash equilibria

Characterization of outcomes of Nash equilibria. In this section, we develop an algorithm for deciding (and computing) the existence of a Nash equilibrium satisfying a given constraint on the costs of the outcome. We do this by computing the maximal punishment that a coalition can inflict to a deviating player. This value can be computed using classical techniques in a two-player zero-sum concurrent game: from a timed configuration $s = (c, d)$, this value is defined as

$$\underline{\text{Val}}_i(s) = \sup_{\sigma_{-i} \in \Sigma_{-i}(s)} \inf_{\sigma_i \in \Sigma_i(s)} \text{cost}_i(\langle \sigma_{-i}, \sigma_i \rangle_s).$$

For all histories $h = (s_k, \vec{a}_k, s'_k)_{0 \leq k < \ell}$ in $\text{Hist}_{\mathcal{G}^F}(s)$, we let $\text{Visit}(h)$ be the set of players who visit their target vertex along h . Formally, for the empty history from $s = (c, d)$, we let $\text{Visit}(h) = \{1 \leq i \leq n \mid c(i) = \text{tgt}_i\}$. If $h = (s_k, \vec{a}_k, s'_k)_{0 \leq k < \ell}$ is non-empty, writing $s_k = (c_k, d_k)$ for all $0 \leq k \leq \ell$, we have $\text{Visit}(h) = \{1 \leq i \leq n \mid \exists 1 \leq k \leq \ell. c_k(i) = \text{tgt}_i\}$.

The following theorem is a characterization of outcomes of Nash equilibria. Similar characterizations were proven in [BMSS20] for (untimed) network congestion games, and in [KLŠT12, AAB18] for generic concurrent games:

Theorem 13. *Let \mathcal{N} be a timed network congestion game and \mathcal{G}^F be its associated finite concurrent game. A play $\rho = (s_k, \vec{a}_k, s'_k)_{k \in \mathbb{N}} \in \text{Plays}_{\mathcal{G}^F}(st)$ is the outcome of a Nash equilibrium from st in \mathcal{G}^F if, and only if,*

$$\forall 1 \leq i \leq n. \forall k \in \mathbb{N}. \forall b_i \in \text{Allow}_i^F(s_k). i \notin \text{Visit}(\rho_{<k}) \implies \text{cost}_i^F(\rho_{\geq k}) \leq \underline{\text{Val}}_i(s') + \text{cost}_i^F(s_k, (\vec{a}_{k,-i}, b_i), s') \quad (1)$$

where $s' = \text{Upd}^F(s_k, (\vec{a}_{k,-i}, b_i))$.

The values $\underline{\text{Val}}_i(s)$ can be computed by transforming the finite game \mathcal{G}^F into a two-player game, since the deviating player competes against the coalition of all the other players. In this transformation, we do not need to keep track of the position of all the other players individually, since their aim now only is to maximize the cost for the deviating player. This two-player game thus has size exponential, and each $\underline{\text{Val}}_i(s)$ can be computed in exponential time. Then:

Proposition 14. *Problem 2 can be decided in EXPSpace, both for the symmetric and for the asymmetric cases.*

Proof. Let \mathcal{N} be a timed network game and $(\phi_q)_q$ be a set of linear constraints. We prove that we can decide in exponential space the existence of a Nash equilibrium σ in \mathcal{G}^F such that $\text{cost}^F(\langle \sigma \rangle_{st})$ satisfies all linear constraints $(\phi_q)_q$. By Theorem 8, this entails the same result in \mathcal{G} .

As already argued, each play ρ in \mathcal{G}^F ends up in a loop after at most K steps. Our algorithm will guess such a play until the loop; the play can be stored in exponential space. The algorithm will then check that it is a valid play, that Eq. 1 holds at each step (which requires to compute $\underline{\text{Val}}$), compute the costs paid by the players until they reach their targets, and check that the cost constraints $(\phi_q)_q$ are satisfied.

This algorithm uses pseudo-polynomial space, as it has to store a polynomial quantity of data for each player. Notice that the algorithm would not store all $\underline{\text{Val}}$ values as the number of values would be doubly-exponential; instead, it will re-compute those values on-demand. \square

Remark 5. Our definition of plays and histories include the action vectors at each step, which implies that strategies observe the actions of all players and

can base their decisions on that information. In our setting, however, it is possible to identify one of the deviating players based only on the configurations; in particular, if a single player deviates, they can be identified and punished. As a consequence, our results still hold in the setting where strategies may only depend on the sequence of timed configurations. \triangleleft

4 Social optimality and prices of anarchy and stability

In this section, given a timed network game \mathcal{N} , we study the social optimum $\text{SO}_{\mathcal{N}}$. In order to obtain the best social welfare, the players aim at minimizing the sum of their costs whatever their selfish interest. Thus we want to find a play from st in the concurrent game \mathcal{G} such that the sum of the costs of all the players is as small as possible. We also consider the price of stability $\text{PoS}_{\mathcal{N}}$ (resp. of anarchy $\text{PoA}_{\mathcal{N}}$) to know how far is the social optimum from the best (resp. worst) social welfare of a Nash equilibrium from st in \mathcal{N} .

Before studying the social optimum and the prices of anarchy and stability, we explain how to solve the constrained-social-welfare problem both in asymmetric and symmetric timed network games. First notice that the following lemma is a consequence of Lemma 6:

Lemma 15. *For all plays $\rho \in \text{Plays}_{\mathcal{G}}(\text{st})$ such that $\text{SW}(\rho)$ is finite, for all $i \in \llbracket n \rrbracket$, there exists $k_i \leq \text{SW}(\rho)$ such that $\rho_{k_i} = (c_{k_i}, d_{k_i})$ with (i) $c_{k_i}(i) = \text{tgt}_i$ and (ii) $d_{k_i} \leq \text{SW}(\rho)$.*

4.1 Constrained-social-welfare problem: asymmetric case

First of all, let us assume that the players' objectives are asymmetric. In this setting, Problem 1 can be solved by non-deterministically guessing a (finite) play in \mathcal{G} step-by-step: Lemma 15 gives a polynomial bound on the size of the configurations to be guessed; keeping track of the set of players who reached their targets requires exponential space. By Savitch's theorem, we get:

Proposition 16. *The constrained-social-welfare problem (Problem 1) can be decided in EXPSpace if the players' objectives are asymmetric.*

4.2 Constrained-social-welfare problem: symmetric case

For symmetric TNGs, the objectives of the players are identical: there exist src and tgt in V such that $\text{src}_i = \text{src}$ and $\text{tgt}_i = \text{tgt}$ for all $1 \leq i \leq n$.

We could of course reuse the algorithm we developed for the asymmetric case, resulting in an EXPSpace algorithm. Nevertheless, we can refine this approach by considering a weighted graph in which we only take into account *abstract timed configurations*. An abstract timed configuration \tilde{c} is a tuple $(P_A, P_W, d) \in [0, n]^V \times [0, n]^V \times \mathbb{N}$ where (i) P_A maps each vertex to the number of *active players* (players who have not visited the target vertex yet) in that vertex; (ii) P_W maps each vertex to the number of *winning players* (players who have already visited their target set) in that vertex and (iii) d is the current time.

Abstract timed configurations store enough information to compute the social welfare of a play in symmetric TNGs and give rise to a weighted graph $\mathcal{W} = (A, B, \tilde{w})$ (see App. C for a formal definition). The set A is the set of abstract timed configurations, $B \subseteq A \times A$ is the set of edges such that there exists an edge $(\tilde{c}_1, \tilde{c}_2) \in B$ between two abstract timed configurations \tilde{c}_1 and \tilde{c}_2 if there exists a valid action for each player regarding their position given by \tilde{c}_1 such that updating P_A, P_B and the current time w.r.t. this action vector leads to the abstract configuration \tilde{c}_2 . The weight function $\tilde{w} : B \rightarrow \mathbb{N}$ represents the sum of the costs of active players for an edge. Notice that the winning players are taken into account to compute the cost of an active player since their presence in a vertex influences the load in that vertex.

A path p in \mathcal{W} is a finite sequence of abstract timed configurations consistent with the graph structure (A, B) , starting from the initial vertex $(\tilde{\text{st}}; \{0\}^V; 0)$ (assuming $\text{src} \neq \text{tgt}$) with $\tilde{\text{st}} : V \rightarrow [0, n] : v \mapsto \#\{1 \leq i \leq n \mid \text{init}(i) = v\}$. The cost of a path p in \mathcal{W} is either the sum of the weights $\tilde{w}(a_1, a_2)$ along the path until visiting a final vertex (where $P_A(v) = 0$ for all $v \in V$), or $+\infty$ if no such vertices appear along p . Clearly enough, the abstract weighted graph encodes the trajectories of \mathcal{N} in the following sense:

Lemma 17. *Let \mathcal{N} be a timed network game and \mathcal{W} be its associated abstract weighted graph. For all $c \in \mathbb{N}$, there exists a play $\rho \in \text{Plays}_{\mathcal{G}}(\tilde{\text{st}})$ such that $\text{SW}(\rho) = c$ if, and only if, there exists a path p in \mathcal{W} with cost c .*

The constrained-social-welfare problem for symmetric objectives can then be solved by non-deterministically guessing the successive vertices of a path p in \mathcal{W} , step-by-step. The constraint c gives a bound on the length of the path, so that the algorithm runs in polynomial space.

Proposition 18. *The constrained-social-welfare problem (Problem 1) can be decided in PSPACE if the players' objectives are symmetric.*

4.3 Social Optimum and Prices of Anarchy and Stability

Optimum Social. We now explain how we can compute the exact social optimum: noticing that the social optimum can be bounded by $n \cdot K$, it can be computed by performing a binary search, iteratively applying the algorithm above. Computing the social optimum can thus be performed in polynomial space in the symmetric case, and exponential space in the asymmetric case.

Prices of Anarchy and Stability. The constrained-price-of-anarchy (resp. stability) problem can now be solved using our algorithms for solving the constrained-social-welfare and constrained-Nash-equilibrium problems: thanks to the pseudo-polynomial bound $n \cdot K$ on the social welfare of the social optimum and on the social welfare of any Nash equilibrium, and the fact that those values are integers, we can perform binary searches for the exact social welfare of the social optimum and for the worst (resp. best) social welfare of a Nash equilibrium. This only requires a polynomial number of iterations, so that the whole algorithm runs in exponential space.

References

- [AAB18] Shaull Almagor, Rajeev Alur, and Suguman Bansal. Equilibria in quantitative concurrent games, 2018.
- [AAE05] Baruch Awerbuch, Yossi Azar, and Amir Epstein. The price of routing unsplittable flow. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, STOC’05, page 57–66, New York, NY, USA, 2005. Association for Computing Machinery.
- [ADK⁺04] Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. In *Proceedings of the 45th Annual Symposium on Foundations of Computer Science (FOCS’04)*, pages 295–304. IEEE Comp. Soc. Press, October 2004.
- [AGK17] 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.
- [AGK18] 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.
- [AHK02] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, September 2002.
- [AKH09] Eitan Altman, Anurag Kumar, and Yezekael Hayel. A potential game approach for uplink resource allocation in a multichannel wireless access network. In *Proceedings of the Fourth International ICST Conference on Performance Evaluation Methodologies and Tools*, pages 1–9, 2009.
- [BFA15] Umang Bhaskar, Lisa Fleischer, and Elliot Anshelevich. A Stackelberg strategy for routing flow over time. *Games and Economic Behavior*, 92:232–247, July 2015.
- [BMSS20] Nathalie Bertrand, Nicolas Markey, Suman Sadhukhan, and Ocan Sankur. Dynamic network congestion games. In Nitin Saxena and Sunil Simon, editors, *FSTTCS 2020*, volume 182 of *LIPICs*, pages 40:1–40:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [CK05] George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, STOC’05, page 67–73, New York, NY, USA, 2005. Association for Computing Machinery.
- [HMRT11] Martin Hoefer, Vahab S. Mirrokni, Heiko Röglin, and Shang-Hua Teng. Competitive routing over time. *Theoretical Computer Science*, 412(39):5420–5432, September 2011.
- [KLŠT12] Miroslav Klimoš, Kim G. Larsen, Filip Štefaňák, and Jeppe Thaarup. Nash equilibria in concurrent priced games. In Adrian-Horia Dediu and Carlos Martín-Vide, editors, *Language and Automata Theory and Applications*, pages 363–376, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

- [KP09] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009.
- [KP12] Elias Koutsoupias and Katia Papakonstantinou. 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.
- [KS11] 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.
- [LMO06] François Laroussinie, Nicolas Markey, and Ghassan Oreiby. Model-checking timed ATL for durational concurrent game structures. In Eugene Asarin and Patricia Bouyer, editors, *Proceedings of the 4th International Conferences on Formal Modelling and Analysis of Timed Systems (FORMATS'06)*, volume 4202 of *Lecture Notes in Computer Science*, pages 245–259. Springer-Verlag, September 2006.
- [MS96] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, May 1996.
- [NRTV07] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
- [PPT09] Michal Penn, Maria Polukarov, and Moshe Tennenholtz. Random order congestion games. *Mathematics of Operations Research*, 34(3):706–725, August 2009.
- [QYZS06] 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.
- [Ros73] Robert W. Rosenthal. The network equilibrium problem in integers. *Networks*, 3:53–59, 1973.
- [Rou07] Tim Roughgarden. Routing games. In Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, chapter 18, pages 461–486. Cambridge University Press, 2007.

A Proofs of Section 2

Lemma 3. *Let $i \in \llbracket n \rrbracket$ be a player and σ_i be a strategy of Player i from some state s_0 . If σ_i is blind, then for all strategies σ_{-i} and σ'_{-i} from s_0 , we have*

$$\text{traj}_i(\langle \sigma_{-i}, \sigma_i \rangle_{s_0}) = \text{traj}_i(\langle \sigma'_{-i}, \sigma_i \rangle_{s_0}).$$

Proof. For a contradiction, assume that this were not the case, and consider two strategy profiles (σ_{-i}, σ_i) and (σ'_{-i}, σ_i) , and the first position j where $\text{traj}_i(\langle \sigma_{-i}, \sigma_i \rangle_{s_0})$ and $\text{traj}_i(\langle \sigma'_{-i}, \sigma_i \rangle_{s_0})$ differ. This position corresponds to two (possibly different) positions $k = \mu_i(j)$ and $k' = \mu'_i(j)$ along $\langle \sigma_{-i}, \sigma_i \rangle_{s_0}$ and $\langle \sigma'_{-i}, \sigma_i \rangle_{s_0}$.

The prefixes h and h' of $\langle \sigma_{-i}, \sigma_i \rangle_{s_0}$ and $\langle \sigma'_{-i}, \sigma_i \rangle_{s_0}$ up to positions $k - 1$ and $k' - 1$ are projected on the same trajectory for Player i , so that σ_i returns the same action for both histories. Moreover, by construction of k and k' , the delay proposed in this action must be shorter than (or equal to) the delays proposed by the strategies in σ_{-i} and σ'_{-i} after histories h and h' ; then the same edge is applied for Player i from position $k - 1$ of $\langle \sigma_{-i}, \sigma_i \rangle_{s_0}$ and position $k' - 1$ of $\langle \sigma'_{-i}, \sigma_i \rangle_{s_0}$, contradicting the fact that $\text{traj}_i(\langle \sigma_{-i}, \sigma_i \rangle_{s_0})$ and $\text{traj}_i(\langle \sigma'_{-i}, \sigma_i \rangle_{s_0})$ differ at position j . \square

B Proofs of Section 3

B.1 Details of the proof of Proposition 5

Proposition 5. *There exists a timed network game \mathcal{N} that admits a Nash equilibrium σ from st , and whose all blind Nash equilibria τ from st are such that $\text{SW}(\langle \sigma \rangle_{st}) < \text{SW}(\langle \tau \rangle_{st})$.*

Proof. We consider the timed network game $\mathcal{N} = (3, \mathcal{A}, (\text{src}, \text{tgt}))$ such that its timed network is given in Figure 2. First of all, we focus on strategy profiles composed only by blind strategies.




Fig. 2. Timed network game with a lower social welfare of an NE with non-blind strategies

There are four different trajectories of blind strategies:

1. $\pi_1 = (\text{src}, 0)(s_1, 1)(s_2, 2)(s_3, 3)(s_4, 4)(s_5, 5)(s_6, 6)(\text{tgt}, 7 + k)_{k \geq 0}$
2. $\pi_2 = (\text{src}, 0)(s_7, 1)(s_8, 2)(s_9, 3)(s_{10}, 4)(s_{11}, 5)(s_{12}, 6)(\text{tgt}, 7 + k)_{k \geq 0}$
3. $\pi_3 = (\text{src}, 0)(s_1, 1)(s_{14}, 2)(s_9, 3)(s_{10}, 4)(s_{11}, 5)(s_{12}, 6)(\text{tgt}, 7 + k)_{k \geq 0}$
4. $\pi_4 = (\text{src}, 0)(s_1, 1)(s_2, 2)(s_3, 3)(s_{15}, 4)(s_{11}, 5)(s_{12}, 6)(\text{tgt}, 7 + k)_{k \geq 0}$

By abuse of notation, in the rest of this example we identify a trajectory π with a blind strategy σ such that $\langle \sigma \rangle_{\text{st}} = \pi$.

We consider the social welfare of the following strategy profiles:

- | | |
|--|--|
| <ul style="list-style-type: none"> – $(\pi_1, \pi_1, \pi_1) \rightsquigarrow 19 \cdot 3 = 57$ – $(\pi_2, \pi_2, \pi_2) \rightsquigarrow 25 \cdot 3 = 75$ – $(\pi_3, \pi_3, \pi_3) \rightsquigarrow 22 \cdot 3 = 66$ – $(\pi_4, \pi_4, \pi_4) \rightsquigarrow 20 \cdot 3 = 60$ | <ul style="list-style-type: none"> – $(\pi_3, \pi_3, \pi_1) \rightsquigarrow 17 \cdot 2 + 15 = 49$ – $(\pi_3, \pi_3, \pi_2) \rightsquigarrow 19 \cdot 2 + 17 = 55$ – $(\pi_3, \pi_3, \pi_4) \rightsquigarrow 19 \cdot 2 + 20 = 58$ |
| <ul style="list-style-type: none"> – $(\pi_1, \pi_1, \pi_2) \rightsquigarrow 15 \cdot 2 + 9 = \mathbf{39}$ – $(\pi_1, \pi_1, \pi_3) \rightsquigarrow 17 \cdot 2 + 12 = 46$ – $(\pi_1, \pi_1, \pi_4) \rightsquigarrow 17 \cdot 2 + 16 = 50$ | <ul style="list-style-type: none"> – $(\pi_4, \pi_4, \pi_1) \rightsquigarrow 18 \cdot 2 + 15 = 51$ – $(\pi_4, \pi_4, \pi_2) \rightsquigarrow 18 \cdot 2 + 13 = 49$ – $(\pi_4, \pi_4, \pi_3) \rightsquigarrow 20 \cdot 2 + 16 = 56$ |
| <ul style="list-style-type: none"> – $(\pi_2, \pi_2, \pi_1) \rightsquigarrow 17 \cdot 2 + 11 = 45$ – $(\pi_2, \pi_2, \pi_3) \rightsquigarrow 21 \cdot 2 + 16 = 58$ – $(\pi_2, \pi_2, \pi_4) \rightsquigarrow 19 \cdot 2 + 16 = 54$ | <ul style="list-style-type: none"> – $(\pi_1, \pi_2, \pi_3) \rightsquigarrow 13 + 13 + 14 = 40$ – $(\pi_1, \pi_2, \pi_4) \rightsquigarrow 13 + 11 + 16 = 40$ – $(\pi_1, \pi_3, \pi_4) \rightsquigarrow 15 + 14 + 18 = 47$ – $(\pi_2, \pi_3, \pi_4) \rightsquigarrow 15 + 16 + 18 = 49$ |

We prove that the social welfare of all Nash equilibria with blind strategies is greater than 39. All social welfare of those blind strategy profiles are greater than 39 except (π_1, π_1, π_2) but this is not a Nash equilibrium since with this strategy profile Player 2 has a cost of 15 and he has an incentive to deviate from π_1 to π_3 and obtains a cost of 14.

We now prove that there exists a Nash equilibrium from st with social welfare equal to 39 and such that one player uses a non-blind strategy. We consider the strategy profile $\sigma = (\sigma_1, \sigma_2, \sigma_3)$ where, roughly speaking, Player 1 and Player 2 follow the trajectory π_1 and Player 3 follows the trajectory π_2 . If Player 3- i , for $i \in \{1, 2\}$, deviates in src then Player i follows trajectory π_3 . If Player 3- i , for $i \in \{1, 2\}$, deviates in s_1 then Player i follows trajectory π_4 .

The outcome of σ from st is:

$$\begin{array}{c}
\begin{array}{ccccccc}
\left[\begin{array}{c} (s_1, 1) \\ (s_1, 1) \\ (s_7, 1) \end{array} \right] & & \left[\begin{array}{c} (s_2, 2) \\ (s_2, 2) \\ (s_8, 2) \end{array} \right] & & \left[\begin{array}{c} (s_3, 3) \\ (s_3, 3) \\ (s_9, 3) \end{array} \right] & & \left[\begin{array}{c} (s_4, 4) \\ (s_4, 4) \\ (s_{10}, 4) \end{array} \right] \\
\text{st} \xrightarrow{\quad} & ((s_1, s_1, s_7), 1) & \xrightarrow{\quad} & ((s_2, s_2, s_8), 2) & \xrightarrow{\quad} & ((s_3, s_3, s_9), 3) & \xrightarrow{\quad}
\end{array} \\
\begin{array}{ccc}
\left[\begin{array}{c} (s_5, 5) \\ (s_5, 5) \\ (s_{11}, 5) \end{array} \right] & & \left[\begin{array}{c} (s_6, 6) \\ (s_6, 6) \\ (s_{12}, 6) \end{array} \right] \\
\xrightarrow{\quad} & ((s_5, s_5, s_{11}), 5) & \xrightarrow{\quad} & ((s_6, s_6, s_{12}), 6)
\end{array} \\
\left(\begin{array}{c} \left[\begin{array}{c} (\text{tgt}, 7+k) \\ (\text{tgt}, 7+k) \\ (\text{tgt}, 7+k) \end{array} \right] \\ \xrightarrow{\quad} ((\text{tgt}, \text{tgt}, \text{tgt}), 7+k) \end{array} \right)_{k \geq 0} .
\end{array}$$

with cost (15, 15, 9) thus the social welfare is equal to 39.

We prove that $(\sigma_1, \sigma_2, \sigma_3)$ is a Nash equilibrium.

For Player i with $i \in \{1, 2\}$:

- if he deviates in v_0 then the cost of Player i is 21;
- if he deviates in v_1 then the cost of Player i is equal to 16;
- if he deviates in v_3 then the cost of Player i is equal to 16.

For Player 3:

- if he deviates and follows trajectory π_1 , then the cost of Player 3 is 19;
- if he deviates and follows trajectory π_3 , then the cost of Player 3 is 12;
- If he deviates and follows trajectory π_4 , then the cost of Player 3 is 16.

That proves that there is no profitable deviation for any player. \square

B.2 Proofs of Section 3.1

Theorem 8. *Let \mathcal{N} be a timed network game and let \mathcal{G} and \mathcal{G}^F be its associated infinite and finite concurrent games. Let $x \in (\mathbb{N} \cup \{+\infty\})^n$. Then there exists a Nash equilibrium σ in \mathcal{G} with $\text{cost}(\langle \sigma \rangle_{\text{st}}) = x$ if, and only if, there exists a Nash equilibrium τ in \mathcal{G}^F with $\text{cost}^F(\langle \tau \rangle_{\text{st}}) = x$*

Simulations between the graph structures. In order to prove Theorem 8, we have to define how to simulate the actions of the players from \mathcal{G} to \mathcal{G}^F and vice-versa.

We first define an equivalence relation $\sim_{\text{BT}} \subseteq \text{States} \times \text{States}$ such that, for any $s_1 = (c_1, d_1) \in \text{States}$ and for any $s_2 = (c_2, d_2) \in \text{States}$, $s_1 \sim_{\text{BT}} s_2$ if, and only if, either (i) $d_1 = d_2 < K$ and $c_1 = c_2$, or (ii) $d_1 > K$ and $d_2 > K$.

Notice that since $\text{States}^F \subseteq \text{States}$, the relation \sim_{BT} is also well-defined on $\text{States} \times \text{States}^F$.

Proposition 19 (Action vector simulation from \mathcal{G} to \mathcal{G}^F). For all $s_1 \in \text{States}$ and $s'_1 \in \text{States}^F$ such that $s_1 \sim_{\text{BT}} s'_1$, for all $\vec{a} \in \text{Valid}(s_1)$, there exists $\vec{a}' \in \text{Valid}^F(s'_1)$ such that

1. $\text{Upd}(s_1, \vec{a}) \sim_{\text{BT}} \text{Upd}^F(s'_1, \vec{a}')$;
2. for all $i \in \llbracket n \rrbracket$, for all $b'_i \in \text{Allow}_i^F(s'_1)$, there exists $b_i \in \text{Allow}_i(s_1)$ such that

$$\text{Upd}(s_1, (a_{-i}, b_i)) \sim_{\text{BT}} \text{Upd}^F(s'_1, (a'_{-i}, b'_i)).$$

Proof. Let $s_1 = (c_1, d_1) \in \text{States}$, let $\vec{a} = (v_i, t_i)_{1 \leq i \leq n} \in \text{Valid}(s_1)$, and let $s'_1 = (c'_1, d'_1) \in \text{States}^F$ such that $s_1 \sim_{\text{BT}} s'_1$. Let $s_2 = (c_2, d_2) = \text{Upd}(s_1, \vec{a})$.

We have to define $\vec{a}' = (v'_i, t'_i)_{1 \leq i \leq n} \in \text{Valid}^F(s'_1)$ such that, letting $s'_2 = \text{Upd}^F(s'_1, \vec{a}')$, it holds $s_2 \sim_{\text{BT}} s'_2$. This depends on d_1 :

1. **if $d_1 < K$:** since $s_1 \sim_{\text{BT}} s'_1$, we have $s_1 = s'_1$. We let $a'_i = (v_i, \min\{t_i, K+1\})$ for all $1 \leq i \leq n$. Notice that $\vec{a}' \in \text{Valid}^F(s'_1)$. Indeed, for all $1 \leq i \leq n$, if $\min\{t_i, K+1\} = t_i$, since $(v_i, t_i) \in \text{Allow}_i(s_1)$, then also $(v_i, t_i) \in \text{Allow}_i^F(s'_1)$ (as $s_1 = s'_1$); otherwise, if $\min\{t_i, K+1\} = K+1$, then $t_i \geq K+1 \geq M$, and since $(v_i, t_i) \in \text{Allow}_i(s_1)$, then also $(v_i, K+1) \in \text{Allow}_i^F(s'_1)$ (as $s_1 = s'_1$, and M is the largest integer constant appearing in the guards). We let $s'_2 = (c'_2, d'_2) = \text{Upd}^F(s'_1, \vec{a}')$, $t^* = \min\{t_i \mid 1 \leq i \leq n\}$, and $t'^* = \min\{t'_i \mid 1 \leq i \leq n\}$. Remember that $\text{Select}(\vec{a}) = \{1 \leq i \leq n \mid t_i = t^*\}$.
 - (a) If $t^* \geq K+1$: we then have $d_2 = t^* \geq K+1$, hence $t'^* = K+1$ and $d'_2 = K+1$. It follows that $s_2 \sim_{\text{BT}} s'_2$.
 - (b) otherwise, $t^* = t'^*$, and it follows that $\text{Select}(\vec{a}) = \text{Select}(\vec{a}')$. Thus for all $i \in \text{Select}(\vec{a})$, we have $a_i = a'_i$. It follows $\text{Upd}(s_1, \vec{a}) = \text{Upd}^F(s'_1, \vec{a}')$, and $s_2 \sim_{\text{BT}} s'_2$.

That proves the first assertion. We now prove the second assertion. Take $1 \leq i \leq n$, and $b'_i \in \text{Allow}_i^F(s'_1)$. Since $d'_1 < K$, we have $b'_i = (v'_i, T'_i)$ for some $T'_i \leq K+1$ and some $v'_i \in V$. We let $b_i = b'_i$. Since $s_1 = s'_1$, we have $b_i \in \text{Allow}_i(s_1)$.

- (a) If $T'_i \leq K$, then $\min(\{t'_j \mid j \neq i\} \cup \{T'_i\}) = \min(\{t_j \mid j \neq i\} \cup \{T'_i\})$ since $t'_j = t_j$ for all $t'_j \leq K$. Thus $\text{Select}(a_{-i}, b_i) = \text{Select}(a'_{-i}, b'_i)$ and since $v'_j = v_j$ for all $j \neq i$, we get $\text{Upd}(s_1, (a_{-i}, b_i)) = \text{Upd}^F(s'_1, (a'_{-i}, b'_i))$.
 - (b) otherwise, $T'_i = K+1$; in that case, if there exists $j \neq i$ such that $t'_j \leq K$, then $t_j = t'_j$, and $i \notin \text{Select}(a'_{-i}, b'_i)$ and $i \notin \text{Select}(a_{-i}, b_i)$. It follows that $\text{Upd}(s_1, \vec{a}) = \text{Upd}(s_1, (a_{-i}, b_i))$, and $\text{Upd}^F(s'_1, \vec{a}') = \text{Upd}^F(s'_1, (a'_{-i}, b'_i))$. Since $\text{Upd}(s_1, \vec{a}) \sim_{\text{BT}} \text{Upd}^F(s'_1, \vec{a}')$, we get $\text{Upd}(s_1, (a_{-i}, b_i)) = \text{Upd}^F(s'_1, (a'_{-i}, b'_i))$. Now, if $t'_j = K+1$ for all $j \neq i$, then also $t_j \geq K+1$ for all $j \neq i$. It follows that, writing $(c_3, d_3) = \text{Upd}(s_1, (a_{-i}, b_i))$ and $(c'_3, d'_3) = \text{Upd}^F(s'_1, (a'_{-i}, b'_i))$, we have $d_3 \geq K+1$ and $d'_3 = K+1$. Thus, $(c_3, d_3) \sim_{\text{BT}} (c'_3, d'_3)$.
2. **if $d_1 \geq K$:** we have that $d'_1 \geq K$ and $d_2 > K$. It follows that the only possible move from (c'_1, d'_1) is \vec{a}' with $a'_i = (\perp, K+1)$ for all $1 \leq i \leq n$. If $s'_2 = (c'_2, d'_2) = \text{Upd}^F(s'_1, \vec{a}')$, then $d'_2 = K+1$, and thus $s_2 \sim_{\text{BT}} s'_2$. That proves the first assertion. We now prove the second assertion. Take $0 \leq i \leq n$, and $b'_i \in \text{Allow}_i^F(s'_1)$. We let $b_i = (v_i, K+1)$ for some $v_i \in V$ such

that $(v_i, K + 1) \in \text{Allow}_i(s_1)$. Notice that by Remark 2, there always exists $v_i \in V$ such that $K + 1 \models \text{guard}(c_1(i), v_i)$. Let $(c_3, d_3) = \text{Upd}(s_1, (a_{-i}, b_i))$ and $(c'_3, d'_3) = \text{Upd}^F(s'_1, (a'_{-i}, b'_i))$. Since d_1 and d'_1 are larger than or equal to K , then $d_3 \geq K + 1$ and $d'_3 = K + 1$. It follows that $(c_3, d_3) \sim_{\text{BT}} (c'_3, d'_3)$. \square

We now prove the converse simulation property:

Proposition 20 (Action vector simulation from \mathcal{G}^F to \mathcal{G}). *For all $s'_1 \in \text{States}^F$, $\vec{a}' \in \text{Mov}^F(s'_1)$, and $s_1 \in \text{States}$ such that $s_1 \sim_{\text{BT}} s'_1$, there exists $\vec{a} \in \text{Mov}(s_1)$ such that*

1. $\text{Upd}(s_1, \vec{a}) \sim_{\text{BT}} \text{Upd}^F(s'_1, \vec{a}')$.
2. for all $1 \leq i \leq n$, for all $b_i \in \text{Mov}_i(s_1)$, there exists $b'_i \in \text{Mov}_i^F(s'_1)$ such that

$$\text{Upd}(s_1, (a_{-i}, b_i)) \sim_{\text{BT}} \text{Upd}^F(s'_1, (a'_{-i}, b'_i)).$$

Proof. Let $s'_1 = (c'_1, d'_1) \in \text{States}^F$, let $\vec{a}' = (v'_i, t'_i)_{1 \leq i \leq n} \in \text{Valid}^F(s'_1)$ (possibly with $v'_i = \perp$), and let $s_1 = (c_1, d_1) \in \text{States}$ such that $s_1 \sim_{\text{BT}} s'_1$. Let $s'_2 = (c'_2, d'_2) = \text{Upd}^F(s'_1, \vec{a}')$.

We have to define $\vec{a} = (v_i, t_i)_{1 \leq i \leq n} \in \text{Valid}(s_1)$ such that, letting $s_2 = \text{Upd}(s_1, \vec{a})$, then $s_2 \sim_{\text{BT}} s'_2$. This depends on d'_1 :

1. **if $d'_1 < K$:** we have $s_1 = s'_1$. We let $\vec{a} = \vec{a}'$. Notice that since $d'_1 < K$, for all $1 \leq i \leq n$, we have $v'_i \neq \perp$. Moreover, because $\vec{a}' \in \text{Valid}^F(s'_1)$ and $s_1 = s'_1$, we have $\vec{a} \in \text{Valid}(s_1)$. In this way, letting $s_2 = \text{Upd}(s_1, \vec{a})$, we have $s_2 = s'_2$, and in particular $s_2 \sim_{\text{BT}} s'_2$, which proves the first assertion. It remains to prove the second assertion.
Pick $1 \leq i \leq n$ and $b_i = (v_i, T_i) \in \text{Allow}_i(s_1)$. We define $b'_i = (v_i, \min\{T_i, K + 1\})$. Once again, as $d_1 = d'_1 < K$ and $s_1 = s'_1$, we have $b'_i \in \text{Valid}^F(s'_1)$. Let $\vec{ab} = (a_{-i}, b_i)$ and $\vec{ab}' = (a'_{-i}, b'_i)$.
(a) If $T_i \leq K$, then $\min\{T_i, K + 1\} = T_i$. Thus $\min\{t'_j \mid j \neq i\} \cup \{T_i\} = \min\{t_j \mid j \neq i\} \cup \{T_i\}$, since for all $j \neq i$, $t'_j = t_j$. It follows that $\text{Select}(\vec{ab}') = \text{Select}(\vec{ab})$, and for all $j \in \text{Select}(\vec{ab}')$, the next vertices chosen in \vec{ab}'_j and in \vec{ab}_j are the same. It follows that $\text{Upd}(s_1, \vec{ab}) = \text{Upd}^F(s'_1, \vec{ab}')$.
(b) If $T_i \geq K + 1$: if there exists $j \neq i$ such that $t'_j \leq K$, then $t_j = t'_j \leq K$ and $i \notin \text{Select}(\vec{ab})$ nor $i \notin \text{Select}(\vec{ab}')$. It follows that $\text{Upd}(s'_1, \vec{a}') = \text{Upd}(s'_1, \vec{ab}')$ and $\text{Upd}^F(s_1, \vec{a}) = \text{Upd}^F(s_1, \vec{ab})$. On the other hand, if $t'_j \geq K + 1$ for all $j \neq i$, then letting $(c_3, d_3) = \text{Upd}(s_1, \vec{ab})$ and $(c'_3, d'_3) = \text{Upd}^F(s'_1, \vec{ab}')$, we have $(c_3, d_3) \sim_{\text{BT}} (c'_3, d'_3)$, since $d_3 \geq K + 1$ and $d'_3 \geq K + 1$.
2. **if $d'_1 \geq K$:** the only allowed move \vec{a}' is such that $a'_i = (\perp, K + 1)$ for all $1 \leq i \leq n$. Moreover, $d_1 \geq K$. We then choose \vec{a} such that $a_i = (v_i, d'_1 + 1)$ for some $v_i \in V$ such that $(v_i, d_1 + 1) \in \text{Allow}_i(s_1)$. Letting $s_2 = (c_2, d_2) = \text{Upd}(s_1, \vec{a})$, we have $d_2 = d_1 + 1 > K$. Since $d'_2 = K + 1$, we get that $s_2 \sim_{\text{BT}} s'_2$, which proves the first assertion.

We now prove the second assertion. Take $1 \leq i \leq n$, and $b_i = (v_i, T_i) \in \text{Allow}_i(s_1)$. Since $d'_1 \geq K$, we have $T_i \geq K + 1$. We fix $b'_i = (\perp, K + 1)$ (there are no other possible choices). Letting $(c_3, d_3) = \text{Upd}(s_1, (a_{-i}, b_i))$ and $(c'_3, d'_3) = \text{Upd}^F(s'_1, (a'_{-i}, b'_i))$, we have $(c_3, d_3) \sim_{\text{BT}} (c'_3, d'_3)$ since $d_3 \geq K + 1$ and $d'_3 \geq K + 1$. \square

We now establish a generic proposition showing a correspondence between Nash equilibria in two game structures:

Proposition 9. *Let $\mathcal{G} = (\text{States}, \text{Act}, (\text{Allow}_i)_{1 \leq i \leq n}, \text{Upd}, (\text{cost}_i)_{1 \leq i \leq n})$ and $\mathcal{G}' = (\text{States}', \text{Act}', (\text{Allow}'_i)_{1 \leq i \leq n}, \text{Upd}', (\text{cost}'_i)_{1 \leq i \leq n})$ be two n -player concurrent games, $s_0 \in \text{States}$ and $s'_0 \in \text{States}'$, and $\triangleleft \subseteq \text{States} \times \text{States}'$ be a relation such that*

1. $s_0 \triangleleft s'_0$;
2. *there exists $\lambda \in \mathbb{N}$ such that for any NE σ in \mathcal{G} from s_0 , for any $1 \leq i \leq n$, it holds $\text{cost}_i(\langle \sigma \rangle_{s_0}) \leq \lambda$;*
3. *for all plays $\rho \in \text{Plays}_{\mathcal{G}}$ and $\rho' \in \text{Plays}_{\mathcal{G}'}$ such that $\rho \triangleleft \rho'$ (i.e., $\rho_j \triangleleft \rho'_j$ for all $j \in \mathbb{N}$), all $1 \leq i \leq n$, if $\text{cost}_i(\rho) \leq \lambda$ or $\text{cost}'_i(\rho') \leq \lambda$, then $\text{cost}_i(\rho) = \text{cost}'_i(\rho')$;*
4. *For all $s \in \text{States}$, for all $\vec{a} \in \text{Valid}(s)$, for all $s' \in \text{States}'$, if $s \triangleleft s'$, then there exists $\vec{a}' \in \text{Valid}'(s')$ such that:*
 - (a) $\text{Upd}(s, \vec{a}) \triangleleft \text{Upd}'(s', \vec{a}')$;
 - (b) *for all $1 \leq i \leq n$, for all $b'_i \in \text{Allow}'_i(s')$, there exists $b_i \in \text{Allow}_i(s)$ such that*

$$\text{Upd}(s, (a_{-i}, b_i)) \triangleleft \text{Upd}'(s', (a'_{-i}, b'_i)).$$

Then for any Nash equilibrium σ in \mathcal{G} from s_0 , there exists a Nash equilibrium σ' in \mathcal{G}' from s'_0 such that $\text{cost}(\langle \sigma \rangle_{s_0}) = \text{cost}'(\langle \sigma' \rangle_{s'_0})$.

Proof. Let σ be a Nash equilibrium from s_0 in \mathcal{G} , and $\rho = \langle \sigma \rangle_{s_0}$. By condition (2), for all $1 \leq i \leq n$, we have $\text{cost}_i(\rho) \leq \lambda$ for some $\lambda \in \mathbb{N}$. By (1) and (4a), there exists $\rho' \in \text{Plays}_{\mathcal{G}'}(s'_0)$ such that $\rho \triangleleft \rho'$. By (3), for all $1 \leq i \leq n$, $\text{cost}_i(\rho) = \text{cost}'_i(\rho')$. We will construct a strategy profile σ' in (\mathcal{G}') that is a Nash equilibrium from s'_0 and such that $\langle \sigma' \rangle_{s'_0} = \rho'$.

For any history $h = ((s_k, \vec{a}_k, s'_k)_{0 \leq k < \ell})$ in $\text{Hist}_{\mathcal{G}}(s_0)$, for any $1 \leq i \leq n$, and for any strategy profile σ , we define $\text{Dev}(h, i, \sigma)$ to be **true** if, and only if, for all $0 \leq k < \ell$ and all $j \neq i$, $a_{k,j} = \sigma_j(\rho_{<k})$. By convention, we let $\text{Dev}(s_0, i, \sigma)$ be **true**. We define Dev in the same way on histories of \mathcal{G}' . For all $h \in \text{Hist}_{\mathcal{G}}$ and all strategy profiles σ , we let

$$\mathcal{D}(h, \sigma) = \{1 \leq i \leq n \mid \text{Dev}(h, i, \sigma) = \text{true}\}.$$

Notice that $\mathcal{D}(h, \sigma)$ can be either $\llbracket n \rrbracket$, in case no players deviated, or a singleton $\{i\}$, in case only Player i deviated, or empty, in case at least two players deviated.

We will construct σ' step-by-step as follows: (i) we define σ' such that its outcome is ρ' and (ii) we extend σ' on histories, by induction on their length. During the construction of σ' , we define a partial function $\mathcal{R}: \text{Hist}_{\mathcal{G}'}(s'_0) \rightarrow$

$\text{Hist}_{\mathcal{G}}(s_0)$, which associates, with each history $h' \in \text{Hist}_{\mathcal{G}'}(s'_0)$ such that $\mathcal{D}(h'c', \tau)$ is non-empty, a representative history in \mathcal{G} . At the end of the procedure we want that \mathcal{R} satisfies the following properties: for all h' of length k in $\text{Hist}_{\mathcal{G}'}(s'_0)$ such that $\mathcal{D}(h', \sigma')$ is non-empty:

- P1. for all $\ell \leq k$, $\mathcal{R}(h'_{<\ell}) = \mathcal{R}(h')_{<\ell}$;
P2. $\mathcal{D}(h', \sigma') = \mathcal{D}(\mathcal{R}(h'), \sigma)$ and $\mathcal{R}(h') \triangleleft h'$.

We first let $\sigma'_i(\rho'_{<k}) = a'_{k,i}$ and $\mathcal{R}(\rho'_{<k}) = \rho_{<k}$ for all $1 \leq i \leq n$ and all $k \in \mathbb{N}$, so that $\langle \sigma' \rangle_{s'_0} = \rho'$. Thus by construction, for all $k \in \mathbb{N}$, $\mathcal{D}(\rho'_{<k}, \sigma') = \llbracket n \rrbracket$ and $\rho'_{<k}$ satisfies Properties (P1) and (P2).

We now extend the definition of σ' step-by-step. At step k :

1. we define σ' on all histories $h' \in \text{Hist}_{\mathcal{G}'}(s'_0)$ of length k that are not prefixes of ρ' , and \mathcal{R} on the resulting outcomes, of length $k + 1$;
2. we prove that Properties (P1) and (P2) are satisfied for all histories h' of length $k + 1$ for which $\mathcal{D}(h', \sigma')$ is non-empty.

At step 0: The only history in $\text{Hist}_{\mathcal{G}'}(s'_0)$ of length 0 is s'_0 , which is a prefix of ρ' . By hypothesis, we have that $s_0 \triangleleft s'_0$. By construction, $\mathcal{R}(s'_0) = s_0$.

We define \mathcal{R} on all histories $h' = (u_0, \vec{a}', u'_0)$ of length 1 that are not prefixes of ρ' and for which $\mathcal{D}(h', \sigma')$ is non-empty (hence $\mathcal{D}(h', \sigma')$ is a singleton $\{i\}$ for some $1 \leq i \leq n$, since there must have been a deviation for h' not to be a prefix of ρ').

By Condition (4b), there exists $b_i \in \text{Allow}_i(s_0)$ such that $s_1 = \text{Upd}(s_0, (\sigma_{-i}(s_0), b_i))$ is such that $v_1 \triangleleft v'_1$. We let $\mathcal{R}(h') = (s_0, (\sigma_{-i}(s_0), b_i), s_1)$.

Property (P1) clearly holds, and since $\mathcal{D}(h', \sigma') = \{i\} = \mathcal{D}(\mathcal{R}(h'), \sigma)$ and $\mathcal{R}(h') \triangleleft h'$, Property (P2) is also satisfied.

At step k : Let us assume that step $k - 1$ has been completed, and that Properties (P1) and (P2) hold for histories of length k .

We first define σ' on histories of length k that are not prefixes of ρ' . For such a history $h' = (u'_j, \vec{a}'_j, v'_j)_{0 \leq j < k}$, $\mathcal{D}(h', \sigma')$ is either empty (if at least two players deviated from σ'), or it is a singleton $\{i\}$ (in case only Player i deviated).

In the first case ($\mathcal{D}(h', \sigma') = \emptyset$), we let $\sigma'(h') = \vec{b}'$ for some $\vec{b}' \in \text{Valid}'(\text{last}(h'))$. In the second case ($\mathcal{D}(h', \sigma') = \{i\}$), we let $h = (u_j, \vec{a}_j, v_j)_{0 \leq j < k} = \mathcal{R}(h')$. By induction hypothesis, we have $h \triangleleft h'$ and $\mathcal{D}(h, \sigma) = \{i\}$. Take $\vec{b} = \sigma(h)$: by (4a), there exists $\vec{b}' \in \text{Valid}(u'_k)$ such that, writing $v_k = \text{Upd}(u_k, \vec{b})$ and $v'_k = \text{Upd}(u'_k, \vec{b}')$, $v_k \triangleleft v'_k$ (remember that we write u_k for v_{k-1}). We then let $\sigma'(h') = \vec{b}'$, and $\mathcal{R}(h' \cdot (u'_k, \vec{b}', v'_k)) = h \cdot (u_k, \vec{b}, v_k)$.

We now prove that Properties (P1) and (P2) are satisfied for all histories h' of length $k + 1$ for which $\mathcal{D}(h', \sigma')$ is non-empty. Let $h' = (u'_j, \vec{a}'_j, v'_j)_{0 \leq j < k+1}$ be a history of length $k + 1$ in $\text{Hist}_{\mathcal{G}'}(s'_0)$.

If $\mathcal{D}(h', \sigma') = \llbracket n \rrbracket$, then h' is a prefix of ρ' , and we already proved that Properties (P1) and (P2) are satisfied.

We now focus on the case where $\mathcal{D}(h', \sigma') = \{i\}$ for some $1 \leq i \leq n$:

- if $\mathcal{R}(h')$ was already defined, then the last move was played by all players according to σ' . Thus $\mathcal{D}(h'_{<k}, \sigma') = \{i\}$ and for all $j \in \llbracket n \rrbracket$, $a'_k = \sigma'(h'_{<k})$. By definition of $\mathcal{R}(h')$, Properties (P1) and (P2) are satisfied.
- if $\mathcal{R}(h')$ is not yet defined, it means that $a'_k(i) \neq \sigma'_i(h'_{<k})$ and $a'_k(j) = \sigma'_j(h'_{<k})$ for all $j \neq i$. By induction hypothesis, $\mathcal{D}(h'_{<k}, \sigma') = \mathcal{D}(\mathcal{R}(h'_{<k}), \sigma)$ and $\mathcal{R}(h'_{<k}) \triangleleft h'_{<k}$.
By (4b), there exists $b_i \in \text{Allow}_i(u_k)$ such that $v_k = \text{Upd}(u_k, (\sigma_{-i}(h), b_i) \triangleleft v'_k$. We fix $\mathcal{R}(h') = \mathcal{R}(h'_{<k}) \cdot (u_k, (\sigma_{-i}(h), b_i), v_k)$. By the two properties above, Properties (P1) and (P2) are satisfied for h' .

This concludes the construction of σ' such that $\langle \sigma' \rangle_{s'_0} = \rho'$. It remains to prove that σ' is a Nash equilibrium in \mathcal{G}' from s'_0 . Towards a contradiction, we assume that there exists $1 \leq i \leq n$ and a strategy τ'_i of Player i in \mathcal{G}' from s'_0 such that $\text{cost}_i(\langle \sigma'_{-i}, \tau'_i \rangle_{s'_0}) < \text{cost}_i(\langle \sigma' \rangle_{s'_0})$. In particular, $\text{cost}_i(\langle \sigma'_{-i}, \tau'_i \rangle_{s'_0}) < \lambda$.

Let $\bar{\rho}' = \langle \sigma'_{-i}, \tau'_i \rangle_{s'_0}$. We will build $\bar{\rho} \in \text{Plays}_{\mathcal{G}}(s_0)$ such that $\bar{\rho} \triangleleft \bar{\rho}'$ and for all $k \in \mathbb{N}$, $\mathcal{D}(\bar{\rho}_{<k}, \sigma) = \mathcal{D}(\bar{\rho}'_{<k}, \sigma')$. This entails that there exists a strategy τ_i of Player i in \mathcal{G} from s_0 such that $\langle \sigma_{-i}, \tau_i \rangle_{v_0} = \bar{\rho}$. It follows that

$$\text{cost}_i(\bar{\rho}) = \text{cost}_i(\bar{\rho}') < \text{cost}_i(\rho') = \text{cost}_i(\rho).$$

Thus τ_i is a profitable deviation for Player i w.r.t. σ , which contradicts the fact that σ was chosen to be a Nash equilibrium.

It remains to show to construct $\bar{\rho}$ from $\bar{\rho}'$. We proceed by induction.

- we let $\bar{\rho}_0 = s_0$; by hypothesis $s_0 \triangleleft s'_0$. Moreover, $\mathcal{D}(\bar{\rho}_{<0}, \sigma) = \llbracket n \rrbracket$ and $\mathcal{D}(\bar{\rho}'_{<0}, \sigma') = \llbracket n \rrbracket$;
- We assume that for all $0 \leq \ell \leq k$, we have $\bar{\rho}_{<\ell} \triangleleft \bar{\rho}'_{<\ell}$, $\mathcal{D}(\bar{\rho}_{<\ell}, \sigma) = \mathcal{D}(\bar{\rho}'_{<\ell}, \sigma')$ and $\bar{\rho}_{<\ell} = \mathcal{R}(\bar{\rho}'_{<\ell})$.

We show how to define the last step of $\bar{\rho}_{<k+1}$. Notice that for all prefix $h' < \bar{\rho}'$, $\mathcal{R}(h')$ is well-defined, since $\mathcal{D}(h', \sigma')$ is either $\llbracket n \rrbracket$ or $\{i\}$. Let $h = (u_j, \vec{a}_j, v_j)_{0 \leq j < k+1} = \mathcal{R}(\bar{\rho}'_{<k+1})$. By construction, $\mathcal{D}(h, \sigma) = \mathcal{D}(\bar{\rho}'_{<k+1}, \sigma')$ and $h \triangleleft \bar{\rho}'_{<k+1}$. Thus we choose $\bar{\rho}_{<k+1} = \bar{\rho}_{<k} \cdot (u_k, \vec{a}_k, v_k)$.

In this way, we obtain that for all $k \in \mathbb{N}$, $\bar{\rho}'_{<k} \triangleleft \bar{\rho}_{<k}$ and $\mathcal{D}(\bar{\rho}_{<k}, \sigma) = \mathcal{D}(\bar{\rho}'_{<k}, \sigma')$. \square

Proof (Proof of Theorem 8). Theorem 8 can now be proven by applying Proposition 9 twice. We first apply Proposition 9 to \mathcal{G} and \mathcal{G}^F , with $\triangleleft = \sim_{\text{BT}}$.

We prove that \mathcal{G} , \mathcal{G}^F and \sim_{BT} satisfy the four hypotheses of Proposition 9, which will allow us to conclude that any Nash equilibrium in \mathcal{G} from $s_0 = \text{st}$ has a corresponding Nash equilibrium from $s'_0 = \text{st}$ in \mathcal{G}^F , with the same cost profile:

1. by definition of \sim_{BT} , we have $s_0 \sim_{\text{BT}} s'_0$;
2. we choose $\lambda = K$, and the assertion holds by Lemma 6.
3. let $\rho = (u_k, \vec{a}_k, v_k)_{k \in \mathbb{N}} \in \text{Plays}_{\mathcal{G}}(s_0)$ and $\rho' = (u'_k, \vec{a}'_k, v'_k)_{k \in \mathbb{N}} \in \text{Plays}_{\mathcal{G}^F}(s'_0)$ such that $\rho \sim_{\text{BT}} \rho'$. Let $1 \leq i \leq n$.

- If $\text{cost}_i(\rho) \leq K$, by Lemma 6 there exists $k_i^* \leq K$ such that $u_{k_i^*} = (\text{tgt}_i, d_i)$ with $d_i \leq K$. It follows that for all $\ell \leq k_i^*$, $\rho_{<\ell} = \rho'_{<\ell}$, and $\text{cost}_i(\rho) = \text{cost}_i(\rho')$.
 - If $\text{cost}_i(\rho') \leq K$, applying Lemma 6 to plays in \mathcal{G}^F , the same reasoning applies.
4. by Proposition 19.

Conversely, we apply Proposition 9 to \mathcal{G}^F and \mathcal{G} , with $\triangleleft = \sim_{\text{BT}}^{-1}$. We have to prove that they satisfy the four conditions of Proposition 9:

1. we have $s_0 \sim_{\text{BT}} s'_0$;
2. we choose $\lambda = K$ and the assertion holds by Lemma 6.
3. Same proof as the other implication.
4. By Proposition 20.

This concludes the proof of Theorem 8. □

B.3 Proofs of Section 3.2

Lemma 21. *For all σ_i and τ_i in $\Sigma_i(s_0)$, for all $\sigma_{-i} \in \Sigma_{-i}(s_0)$ such that for all $j \in \llbracket n \rrbracket \setminus \{i\}$, σ_j is a blind strategy, we have: if $\text{traj}_i(\langle \sigma_{-i}, \sigma_i \rangle_{s_0}) = \text{traj}_i(\langle \sigma_{-i}, \tau_i \rangle_{s_0})$, then $\text{cost}_i(\langle \sigma_{-i}, \sigma_i \rangle_{s_0}) = \text{cost}_i(\langle \sigma_{-i}, \tau_i \rangle_{s_0})$.*

Proof. The trajectories are the same for all players in both outcomes, hence the sequences of timed configurations visited along both outcomes are the same. Now, the cost of a transition only depends on the timed configurations it comes from and goes to, and not on the exact action vector that is being played. The result follows. □

Proposition 12. *Let \mathcal{N} be a timed network game. Let \mathcal{G}^F be its associated finite concurrent game, and Σ^{WB} the set of winning blind strategy profiles. If there exists a Σ^{WB} -Nash equilibrium σ in \mathcal{G}^F from st , then there exists a Nash equilibrium τ from st in \mathcal{G}^F with the same costs for all players.*

Proof. Let us assume that there exists a Σ^{WB} -Nash equilibrium σ in \mathcal{G}^F from st . Notice that for all $1 \leq i \leq n$, strategy σ_i is a winning blind strategy for Player i from st . We prove that σ is also a Nash equilibrium from st in \mathcal{G}^F .

Ad absurdum, we assume that there exist $1 \leq i \leq n$ and a strategy τ_i of Player i in \mathcal{G}^F from st such that

$$\text{cost}_i^F(\langle \sigma_{-i}, \tau_i \rangle_{\text{st}}) < \text{cost}_i^F(\langle \sigma \rangle_{\text{st}}). \quad (2)$$

Consider the trajectory $p = \text{traj}_i(\langle \sigma_{-i}, \tau_i \rangle_{\text{st}})$. Writing $p = (v_k, d_k)_{k \in \mathbb{N}}$, since $\text{cost}_i(\langle \sigma_{-i}, \tau_i \rangle_{\text{st}})$ is finite, there exists $k^* \in \mathbb{N}$ such that $v_{k^*} = \text{tgt}_i$. Thus there exists a winning blind strategy $\sigma'_i \in \Sigma_i^{WB}$ corresponding to this trajectory.

By hypothesis σ_j is a blind strategy for all $j \in \llbracket n \rrbracket \setminus \{i\}$. By Lemma 21, it follows:

$$\text{cost}_i^F(\langle \sigma_{-i}, \sigma'_i \rangle_{\text{st}}) = \text{cost}_i^F(\langle \sigma_{-i}, \tau_i \rangle_{\text{st}}).$$

Thus $\sigma'_i \in \Sigma_i^{WB}$ is a profitable blind deviation for Player i w.r.t. σ , contradicting the fact that σ is a Σ^{WB} -Nash equilibrium. □

Proposition 11. *Let \mathcal{N} be a timed network game and let \mathcal{G}^F be its associated finite concurrent game and Σ^{WB} be the set of winning blind strategy profiles. The game \mathcal{G}^F restricted to strategy profiles in Σ^{WB} has a potential function.*

Proof. Let $R = \{(v, d) \mid v \in V, 0 \leq d \leq K\}$. For each strategy profile $\sigma \in \Sigma^{WB}$ and each pair $(v, d) \in R$, writing $\langle \sigma \rangle_{\text{st}} = ((c_k, d_k), \vec{a}_k, (c'_k, d'_k))_{k \in \mathbb{N}}$ and k^* for the largest index for which $d_k \leq d$, we define $\text{load}_\sigma(v, d) = \#\{1 \leq i \leq n \mid c_{k^*}(i) = v\}$. In other terms, $\text{load}_\sigma(v, d)$ is the number of players standing in vertex v at time d along the outcome of σ from st .

For any strategy profile σ in Σ^{WB} and any subset $J \subseteq R$, we define

$$\Psi_J(\sigma) = \sum_{(v,d) \in J} \sum_{k=1}^{\text{load}_\sigma(v,d)} \text{wgt}(v)(k).$$

We let $\Psi = \Psi_R$, and prove that it is a potential function: assume that some player i deviates from their strategy σ_i in σ and follows some other strategy $\sigma'_i \in \Sigma_i^{WB}$ instead; write $\sigma' = (\sigma_{-i}, \sigma'_i)$. We prove that:

$$\Psi(\sigma') - \Psi(\sigma) = \text{cost}_i^F(\langle \sigma' \rangle_{\text{st}}) - \text{cost}_i^F(\langle \sigma \rangle_{\text{st}}).$$

Write P for the set of plays in \mathcal{G}^F along which all players visit their target vertices. For such a play $\rho = ((c_k, d_k), \vec{a}_k, (c'_k, d'_k))_{k \in \mathbb{N}}$ in P , and for each $1 \leq i \leq n$, we define the set of positions occupied by Player i along ρ (before reaching their target vertex) as follows: we first let k_i be the least index k such that $c_k(i) = \text{tgt}_i$; then

$$\text{Pos}_i(\rho) = \{(v, d) \mid \exists 0 \leq j < k_i. d_j \leq d < d_{j+1} \text{ and } v = v_j\}.$$

Now, let $V = \text{Pos}_i(\langle \sigma \rangle_{\text{st}})$ and $V' = \text{Pos}_i(\langle \sigma' \rangle_{\text{st}})$. Notice that since σ and σ' are in Σ^{WB} , both V and V' are well-defined. Then let

$$B = V \cap V' \quad O = V \setminus B \quad N = V' \setminus B \quad X = R \setminus (B \cup O \cup N).$$

Notice that X , B , O and N are disjoint. Moreover, since all players but Player i follow the same blind strategy in σ and in σ' , we have

$$\begin{aligned} \text{load}_\sigma(v, d) &= \text{load}'_\sigma(v, d) && \text{for all } (v, d) \in B \cup X \\ \text{load}_\sigma(v, d) &= \text{load}'_\sigma(v, d) + 1 && \text{for all } (v, d) \in O \\ \text{load}_\sigma(v, d) &= \text{load}'_\sigma(v, d) - 1 && \text{for all } (v, d) \in N. \end{aligned}$$

We can then write:

$$\begin{aligned} \text{cost}_i^F(\langle \sigma' \rangle_{\text{st}}) - \text{cost}_i^F(\langle \sigma \rangle_{\text{st}}) &= \sum_{(v,d) \in V'} \text{wgt}(v)(\text{load}_{\sigma'}(v, d)) - \sum_{(v,d) \in V} \text{wgt}(v)(\text{load}_\sigma(v, d)) \\ &= \sum_{(v,d) \in N} \text{wgt}(v)(\text{load}_{\sigma'}(v, d)) - \sum_{(v,d) \in O} \text{wgt}(v)(\text{load}_\sigma(v, d)). \end{aligned}$$

Indeed, the terms $\sum_{(v,d) \in B} \text{wgt}(v)(\text{load}_\sigma(v,d))$ and $\sum_{(v,d) \in B} \text{wgt}(v)(\text{load}_{\sigma'}(v,d))$ cancel out by the first equality above.

Similarly,

$$\begin{aligned} \Psi(\sigma') - \Psi(\sigma) &= \Psi_X(\sigma') + \Psi_B(\sigma') + \Psi_O(\sigma') + \Psi_N(\sigma') - (\Psi_X(\sigma) + \Psi_B(\sigma) + \Psi_O(\sigma) + \Psi_N(\sigma)) \\ &= \sum_{(v,d) \in N} \text{wgt}(v)(\text{load}_{\sigma'}(v,d)) - \sum_{(v,d) \in O} \text{wgt}(v)(\text{load}_\sigma(v,d)). \end{aligned}$$

The equality proving that Ψ is a potential function follows. \square

We can now prove Theorem 10:

Proof. Let \mathcal{N} be a timed network game and let \mathcal{G}^F be its finite associated concurrent game. Take a winning blind strategy profile $\sigma \in \Sigma^{WB}$ (which must exist thanks to our hypotheses). Then $\Psi(\sigma)$ is finite. Then, as long as is possible, replace the winning blind strategy of some player i with one that achieves a (strictly) better cost for that player. From Proposition 11, the value of Ψ (strictly) decreases at each step. Since Ψ takes non-negative integer values, this process must terminate. Upon convergence, the strategy profile we obtain is a winning blind strategy profile in which no player can improve their cost by a unilateral deviation, hence it is a Σ^{WB} -Nash equilibrium. By Proposition 12, it is a Nash equilibrium in \mathcal{G}^F , and by Theorem 8, it is also a Nash equilibrium in \mathcal{G} . \square

B.4 Proofs of Section 3.3

Theorem 13. *Let \mathcal{N} be a timed network congestion game and \mathcal{G}^F be its associated finite concurrent game. A play $\rho = (s_k, \vec{a}_k, s'_k)_{k \in \mathbb{N}} \in \text{Plays}_{\mathcal{G}^F}(\text{st})$ is the outcome of a Nash equilibrium from st in \mathcal{G}^F if, and only if,*

$$\begin{aligned} \forall 1 \leq i \leq n. \forall k \in \mathbb{N}. \forall b_i \in \text{Allow}_i^F(s_k). i \notin \text{Visit}(\rho_{<k}) \implies \\ \text{cost}_i^F(\rho_{\geq k}) \leq \text{Val}_i(s') + \text{cost}_i^F(s_k, (\vec{a}_{k,-i}, b_i), s') \end{aligned} \quad (1)$$

where $s' = \text{Upd}^F(s_k, (\vec{a}_{k,-i}, b_i))$.

Proof. We begin with the first direction, showing that the outcome of a Nash equilibrium satisfies Property (1).

Let σ be a Nash equilibrium in \mathcal{G}^F , and $\rho = \langle \sigma \rangle_{\text{st}}$. Let $1 \leq i \leq n$ and $k \in \mathbb{N}$ such that $i \notin \text{Visit}(\rho_{<k})$, and $b_i \in \text{Allow}_i^F(s_k)$. Let $s' = \text{Upd}(s_k, (\vec{a}_{k,-i}, b_i))$ be the vertex reached after the deviation of Player i in s_k .

We define a set $D_i \subseteq \Sigma_i(\text{st})$ of all strategies for Player i that follow σ_i along the first $k-1$ steps of ρ , and play b_i at s_k . When $k=0$, D_i contains all strategies that play b_i from st .

Let $h = \rho_{<k}$ be the prefix of ρ until s_k , and $h' = h \cdot (s_k, (\vec{a}_{k,-i}, b_i), s')$. Since σ is a Nash equilibrium from st in \mathcal{G}^F , we have that for all $\tau_i \in D_i$:

$$\text{cost}_i^F(h \cdot \langle \sigma \upharpoonright h \rangle_{s_k}) = \text{cost}_i^F(\langle \sigma \rangle_{\text{st}}) \leq \text{cost}_i^F(\langle \sigma_{-i} \rangle_{\text{st}}, \tau_i) = \text{cost}_i^F(h \cdot \langle \sigma_{-i} \upharpoonright h, \tau_i \upharpoonright h \rangle_{\text{st}})$$

where $\sigma_{\uparrow h}$ is the residual strategy of σ after history h . Since $i \notin \text{Visit}(h)$, we get $\text{cost}_i(\langle \sigma_{\uparrow h} \rangle_{s_k}) \leq \text{cost}_i(\langle \sigma_{-i \uparrow h}, \tau_i \uparrow h \rangle_{\text{st}})$. It follows that, for all $\tau_i \in D_i$:

$$\begin{aligned} \text{cost}_i^F(\langle \sigma_{\uparrow h} \rangle_{s_k}) &\leq \text{cost}_i^F(\langle \sigma_{-i \uparrow h}, \tau_i \uparrow h \rangle_{s_k}) \\ &= \text{cost}_i^F((s_k, (\vec{a}_{k,-i}, b_i), s') \cdot \langle \sigma_{-i \uparrow h'}, \tau_i \uparrow h' \rangle_{s'}) \\ &= \text{cost}_i^F(s_k, (\vec{a}_{k,-i}, b_i), s') + \text{cost}_i^F(\langle \sigma_{-i \uparrow h'}, \tau_i \uparrow h' \rangle_{s'}). \end{aligned}$$

Since this must hold for any strategy τ_i in D_i , we get:

$$\begin{aligned} \text{cost}_i^F(\langle \sigma_{\uparrow h} \rangle_{s_k}) &\leq \text{cost}_i^F(s_k, (\vec{a}_{k,-i}, b_i), s') + \inf_{\tau_i \in D_i} \text{cost}_i^F(\langle \sigma_{-i \uparrow h'}, \tau_i \uparrow h' \rangle_{s'}) \\ &= \text{cost}_i^F(s_k, (\vec{a}_{k,-i}, b_i), s') + \inf_{\tau_i \in \Sigma_i} \text{cost}_i^F(\langle \sigma_{-i \uparrow h'}, \tau_i \uparrow h' \rangle_{s'}) \quad (3) \\ &\leq \text{cost}_i^F(s_k, (\vec{a}_{k,-i}, b_i), s') + \sup_{\nu_{-i} \in \Sigma_{-i}} \inf_{\tau_i \in \Sigma_i} \text{cost}_i^F(\langle \tau_{-i \uparrow h'}, \nu_i \uparrow h' \rangle_{s'}) \\ &= \text{cost}_i^F(s_k, (\vec{a}_{k,-i}, b_i), s') + \underline{\text{Val}}_i(s'). \end{aligned}$$

Equality (3) holds because only the residual part of τ_i after h' is used, and there are no constraints on that part in D_i . This concludes the proof of this implication.

We now turn to the converse implication. Take $\rho = (s_k, \vec{a}_k, s'_k)_{k \in \mathbb{N}} \in \text{Plays}_{\mathcal{G}^F}(\text{st})$ satisfying Property (1). We build a strategy profile σ from st in \mathcal{G}^F that is a Nash equilibrium and such that $\langle \sigma \rangle_{\text{st}} = \rho$.

The main idea is that all players follow ρ , until one player deviates. If we assume that Player i is the first player who deviates, then after this deviation the strategy of the coalition $-i$ of the other players will follow the punishing strategy in the corresponding two-player zero-sum game.

For all $1 \leq i \leq n$, for all $j \neq i$, and for all state s in \mathcal{G}^F , we denote by $\sigma_{j,i}^s$ the strategy of Player j obtained from the optimal strategy of coalition $-i$ from s in the two-player zero-sum game in which Player i aims at minimizing their cost. In such two-player zero-sum games with weights in \mathbb{N} , such optimal strategies exist, and for any state s , we have

$$\inf_{\mu_i \in \Sigma_i(s)} \text{cost}_i^F(\langle (\sigma_{j,i}^s)_{j \neq i}, \mu_i \rangle_s) \geq \underline{\text{Val}}_i(s). \quad (4)$$

First, we define $P_\rho: \text{Hist}_{\mathcal{G}^F}(\text{st}) \rightarrow (\text{Hist}_{\mathcal{G}^F}(\text{st}) \times \llbracket n \rrbracket) \cup \{\perp\}$ to keep track of the first deviation (w.r.t. ρ) along a history h . Formally, for $h = (u_k, \vec{b}_k, u'_k)_{0 \leq k < \ell}$, we let

$$P_\rho(h) = \begin{cases} P_\rho(h_{< \ell-1}) & \text{if } h_{< \ell-1} \text{ is not a prefix of } \rho \\ (h, i) & \text{if } h_{< \ell-1} \text{ is a prefix of } \rho \text{ but } h \text{ is not, and} \\ & i \text{ is the least index for which } b_{\ell,i} \neq a_{\ell,i} \\ \perp & \text{otherwise} \end{cases}$$

Hence $P_\rho(h) = (h', i)$ indicates that there has been a deviation w.r.t. ρ along h , and that the first deviation occurred after prefix h' , and that Player i is one of the players who deviated at that point. If there has been no deviation (hence h is a prefix of ρ), then $P_\rho(h) = \perp$.

We can now define σ : for all $h \in \text{Hist}_{\mathcal{G}^F}(\text{st})$ of length ℓ and for all $i \in \llbracket n \rrbracket$, we let

$$\sigma_i(h) = \begin{cases} a_{\ell,i} & \text{if } h \text{ is a prefix of } \rho \\ \text{any allowed move} & \text{if } P_\rho(h) = (h', i) \text{ for some } h' \\ \sigma_{i,j}^s(h \setminus h') & \text{if } P_\rho(h) = (h', j) \text{ for some } h' \text{ and some } j \neq i, \\ & \text{with } s = \text{last}(h') \text{ and } h \setminus h' \text{ is the suffix of } h \text{ after } h' \end{cases}$$

Let us prove that σ is a Nash equilibrium in \mathcal{G}^F . Let $1 \leq i \leq n$, let $\tau_i \in \Sigma_i(\text{st})$. Let $\rho = (s_k, \vec{a}_k, s'_k)_{k \in \mathbb{N}} = \langle \sigma \rangle_{\text{st}}$ and $\rho' = (u_k, \vec{b}_k, u'_k)_{k \in \mathbb{N}} = \langle \sigma_{-i}, \tau_i \rangle_{\text{st}}$. We have to prove that $\text{cost}_i^F(\rho) \leq \text{cost}_i^F(\rho')$.

If both outcomes are identical, the result is trivial. If not, let $k \in \mathbb{N}$ be the largest index such that $\rho_{<k} = \rho'_{<k}$: this means that the first deviation of Player i occurs in state s_k along ρ .

If $i \in \text{Visit}(\rho_{<k})$, then $\text{cost}_i^F(\rho) = \text{cost}_i^F(\rho')$, and we are done. Otherwise:

$$\begin{aligned} \text{cost}_i^F(\rho') &= \text{cost}_i^F(h' \cdot \langle \sigma_{-i \uparrow h'}, \tau_i \uparrow h' \rangle_{\rho'_{k+1}}) && \text{where } h' = \rho'_{<k+1} \\ &= \text{cost}_i^F(h' \cdot \langle (\sigma_{j,i \uparrow h'}^{u_{k+1}})_{j \neq i}, \tau_i \uparrow h' \rangle_{u_{k+1}}) && \text{by definition of } \sigma \\ &= \text{cost}_i^F(\rho_{<k}) + \text{cost}_i^F(u_k, \vec{b}_k, u'_k) + \text{cost}_i^F(\langle (\sigma_{j,i \uparrow h'}^{u_{k+1}})_{j \neq i}, \tau_i \uparrow h' \rangle_{u_{k+1}}) \\ & && \text{because } i \notin \text{Visit}(\rho_{<k}) \\ &\geq \text{cost}_i^F(\rho_{<k}) + \text{cost}_i^F(u_k, \vec{b}_k, u'_k) + \inf_{\nu_i \in \Sigma_i(u_{k+1})} \text{cost}_i^F(\langle (\sigma_{j,i}^{u_{k+1}})_{j \neq i}, \nu_i \rangle_{u_{k+1}}) \\ &\geq \text{cost}_i^F(\rho_{<k}) + \text{cost}_i^F(u_k, \vec{b}_k, u'_k) + \underline{\text{Val}}_i(u_{k+1}) && \text{by 4} \\ &\geq \text{cost}_i^F(\rho_{<k}) + \text{cost}_i^F(\rho_{\geq k}) && \text{by 1} \\ &= \text{cost}_i^F(\langle \sigma \rangle_{\text{st}}) && \square \end{aligned}$$

Remark 6. Our definition of plays and histories include the action vectors at each step, which implies that strategies observe the actions of all players and can base their decisions on those informations. In particular, they can easily detect changes in strategies. This is used in the definition of P_ρ , and thus of σ_i , in the proof of Theorem 13.

We could easily adapt our results to handle the case where strategies are not allowed to depend on action vectors. Indeed, in our setting of network congestion games, the relevant single-player deviations can be detected by observing only the sequence of configurations. With *relevant*, we mean those deviations that modify the outcome (other deviations would never be profitable).

Indeed, assume that a single player deviates from their strategy from some configuration (c, d) , so that the play goes to (c_2, d_2) instead of (c_1, d_1) . In case $d_2 < d_1$, then d_2 must be the date proposed by the deviating player, and they will be the only player applying their action; they are thus easily identified. In case $d_2 > d_1$, it must be the case that d_1 was proposed by the deviating player alone, and they were the only player applying their action between c and c_1 . Finally,

in case $d_2 = d_1$, by hypothesis $c_1 \neq c_2$, so if only one player, say Player i , deviated, then i will be the only index for which $c_1(i) \neq c_2(i)$, which again permits to identify the deviating player.

In the end, in case strategies are not allowed to depend on move vectors, the proof of Theorem 13 can be adapted, and our result still holds. \triangleleft

Proposition 22. *For all $1 \leq i \leq n$ and for all $s \in \text{States}^F$, the value $\text{Val}_i(s)$ can be computed in exponential time.*

Proof. From the finite concurrent game \mathcal{G}^F , we build a two-player zero-sum concurrent game \mathcal{Z} in which the two players are the player i , which has the same possible actions as in \mathcal{G}^F , and the coalition of other players $-i$, which chooses the actions of all the other players. The states of \mathcal{Z} are $V \times [0, n]^V \times [0, K + 1]$: such a state keeps track of the current position of Player i , an abstract representation of the current position of all players of the coalition $-i$ (*i.e.*, the number of players of the coalition in each vertex) and the current date. The objective of Player i is to reach tgt_i (starting from src_i) while minimizing their cost, and the goal of coalition player $-i$ is opposite, *i.e.*, to maximize the cost of Player i .

Computing Val_i in such a game can be performed in time polynomial in the size of the game [LMO06], hence in time exponential in $|V|$ and pseudo-polynomial in K . \square

C Proofs of Section 4

Proposition 16. *The constrained-social-welfare problem (Problem 1) can be decided in EXPSPACE if the players' objectives are asymmetric.*

Proof. Let \mathcal{N} be a timed network game and let $x \in \mathbb{N}$ be a threshold. We want to decide if there exists a play ρ in \mathcal{N} beginning in st such that $\text{SW}(\rho) \leq x$.

In view of Lemma 15, if there exists such a play, there exists a history h such that for all timed configurations (c, d) along h , $d \leq x$ and $\llbracket n \rrbracket = \text{Visit}(h)$.

During the procedure we use a counter S that stores the current sum of the players' costs. We also need to keep in memory the set of players who have already visited their target vertex; we write I for this set, and $s = (c, d)$ for the current timed configuration.

The algorithm begins with $S = 0$, $I = \{1 \leq i \leq n \mid \text{src}_i = \text{tgt}_i\}$, and $s = (\text{init}, 0)$. Then, step-by-step:

1. the algorithm guesses the next timed configuration $s' = (c, d)$ with $d \leq x$ (if such a successor exists);
2. the counter S is updated for all players $i \notin I$;
3. I is augmented with $\{1 \leq i \leq n \mid c(i) = \text{tgt}_i\}$;
4. the algorithm forgets s , and continues from s' , unless $I = \llbracket n \rrbracket$.

This procedure may stop for two reasons: (*i*) if at some point $I = \llbracket n \rrbracket$: then we check if $S \leq x$. If it is the case, then we have found a history h such that

$\text{SW}(h) \leq x$. This history can be extended to a play ρ with $\text{SW}(\rho) \leq x$; (ii) if from some timed configuration $s = (c, d)$, no successor (c', d') with $d' \leq x$ exists. That means that we have failed to find a history h such that all players visit their target vertex within time x , and in particular due to Lemma 15, there does not exist a play ρ such that $\text{SW}(\rho) \leq x$.

The counter S is incremented by at most $n \cdot (x \cdot \max_{v \in V} \text{wgt}(v)(n))$ and there are at most x steps. Thus the value of S can always be encoded in polynomial space. On the other hand, the algorithm also has to store two consecutive timed configurations and the action vector between them; if we assume that the objectives of the players are given as a function $V^2 \rightarrow \mathbb{N}$ (with binary encoding), this takes exponential space. Similarly, keeping track of which players have visited their objectives takes exponential space. \square

Remark 7. Notice that if the objectives are given explicitly as a list of (src, tgt)-pairs, one for each player, then the input would be exponentially larger, and our algorithm would then be in polynomial space in the size of the input. \triangleleft

Proposition 18. *The constrained-social-welfare problem (Problem 1) can be decided in PSPACE if the players' objectives are symmetric.*

The proof of this result relies on the following notion of *abstract weighted graph*, which stores the number of active and winning players in each vertex:

Definition 23 (Abstract weighted graph). *The abstract weighted graph $\mathcal{W} = (A, B, \bar{w})$ of a symmetric TNG \mathcal{N} is defined as follows:*

- the set of vertices is $A = [0, n]^V \times [0, n]^V \times \mathbb{N}$. The first part $[0, n]^V$ gives the number of active players in each vertex (i.e., those players that have not visited the target vertex yet); the second component gives the number of winning players in each vertex (those who have already visited the target vertex); the last integer is the current time;
- the set $B \subseteq A \times A$ is the set of edges of the graph: given two vertices $a_1 = (P_A^1: V \rightarrow [0, n], P_W^1: V \rightarrow [0, n], d_1)$ and $a_2 = (P_A^2: V \rightarrow [0, n], P_W^2: V \rightarrow [0, n], d_2)$, there is an edge (a_1, a_2) in B whenever
 - $d_1 < d_2$;
 - there exist functions $b_A: E \rightarrow [0, n]$ and $b_W: E \rightarrow [0, n]$, representing the number of active and winning players taking each edge of \mathcal{N} , such that
 - * $1 \leq \sum_{e \in E} (b_A(e) + b_W(e)) \leq n$ (only the player(s) proposing the shortest delay will move);
 - * for all $e \in E$, if $b_A(e) + b_W(e) > 0$, then $d_2 \models \text{guard}(e)$;
 - * for all $v \in V$, $\sum_{e=(v,v')} b_A(e) \leq P_A^1(v)$ and $\sum_{e=(v,v')} b_W(e) \leq P_W^1(v)$;
 - * $P_A^2(\text{tgt}) = 0$ and for all $v \in V \setminus \{\text{tgt}\}$, $P_A^2(v) = P_A^1(v) - \sum_{e=(v,v')} b_A(e) + \sum_{e=(v',v)} b_A(e)$;
 - * similarly, for P_W^2 : $P_W^2(\text{tgt}) = P_W^1(\text{tgt}) - \sum_{e=(\text{tgt},v')} b_W(e) + \sum_{e=(v',\text{tgt})} (b_W(e) + b_A(e))$, and for all $v \in V \setminus \{\text{tgt}\}$, $P_W^2(v) = P_W^1(v) - \sum_{e=(v,v')} b_W(e) + \sum_{e=(v',v)} b_W(e)$.

– the weight function \tilde{w} is defined for each edge $(a_1, a_2) \in B$ as

$$\tilde{w}(a_1, a_2) = \sum_{v \in V | P_A^1(v) \neq 0} \text{wgt}(v)(P_A^1(v) + P_W^1(v)) \cdot (d_2 - d_1).$$

Proof (of Proposition 18). The non-deterministic algorithm guesses the successive vertices of a path p in the abstract graph \mathcal{W} step-by-step. The constraint on the social welfare also gives a bound on the length of the path to be guessed, so that this algorithm runs in polynomial space. \square