

Memory Requirements in Non-Zero-Sum Games

Yoav Feinstein 

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

Orna Kupferman 

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

Abstract

The interaction between a system and the components modeling its environment is traditionally modeled by a *multi-player game* played on a finite graph. In *zero-sum* games, the players have conflicting objectives, and it is clear that increasing the *memory* of the environment players can only make it harder for the system to win. In *non-zero-sum* games, the objectives of the players may overlap. There, typical questions concern the stability of the game and the equilibria the players may reach. In particular, in *rational synthesis* (RS), the goal is to find an equilibrium that satisfies the objective of the system.

We study how the memory of the environment players may affect the existence of an RS solution. As we show, the picture is diverse, even when the objectives of all players are memoryless. On the one hand, when stability amounts to a *Nash equilibrium* (NE), then increasing the memory of the environment may only help the system to suggest an RS solution. On the other hand, when the notion of stability involves deviations by *coalitions* of environment players, for example in a *strong Nash equilibrium* (SNE), then increasing their memory may sometimes enable and sometimes prevent the existence of an RS solution. We study memory bounds for the players, showing that the memory required may be polynomial in an NE-RS solution and exponential in an SNE-RS solution. We also solve the SNE-RS problem, show that it is PSPACE-complete, and relate the differences between NE and SNE with the differences between *cooperative* and *non-cooperative* RS.

2012 ACM Subject Classification Theory of computation → Formal languages and automata theory; Theory of computation → Semantics and reasoning

Keywords and phrases Non-Zero-Sum Games, Synthesis, Memory

Digital Object Identifier 10.4230/LIPIcs.CSL.2026.34

Related Version *Full Version*: <https://www.cs.huji.ac.il/~ornak/publications/csl26.pdf>

1 Introduction

Synthesis is the automated construction of systems from their specifications [28]. Modern systems often consist of interacting components. The interaction is modeled by a *multi-player game* played on a finite graph. In the *turn-based* setting, the vertices of the game graph are partitioned among the players. A token is placed on an initial vertex, and in each turn, the player that owns the vertex with the token moves it to a successor vertex. Each player has a *strategy* that directs her how to move the token when it reaches vertices she owns. A *profile* is a vector of strategies, one for each player. The outcome of a profile is a *play* – an infinite path in the game graph, obtained when the players follow their strategies. The goal of each player is to direct the game into a play that satisfies her objective. Each objective α defines a subset of V^ω [25], where V is the set of vertices of the game graph. For example, in games with *Büchi* objectives, α is a subset of V , and a play satisfies α if it visits vertices in α infinitely often.

In *zero-sum* games, the players compete with each other on the satisfaction of contradicting objectives. Zero-sum games with ω -regular objectives are *determined*: in every game, exactly one player has a winning strategy – one that achieves her objective against all strategies of the other players [25]. Deciding a zero-sum game amounts to finding this player. In contrast,



© Yoav Feinstein and Orna Kupferman;

licensed under Creative Commons License CC-BY 4.0

34th EACSL Annual Conference on Computer Science Logic (CSL 2026).

Editors: Stefano Guerrini and Barbara König; Article No. 34; pp. 34:1–34:23

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

in *non-zero-sum* games, the objectives of the players may overlap [12, 31]. There, typical questions concern the stability of the game and the equilibria the players may reach [32]. The most common notion of stability is *Nash equilibrium* (NE) [26]. A profile of strategies is an NE if no (single) player can benefit from unilaterally changing her strategy.

Two-player zero-sum games model the interaction between a system and its environment. The system aims to satisfy its specification in all environments. Accordingly, the environment is assumed to be hostile, as if its objective is to violate the specification [4]. Often, however, the components composing the environment have objectives of their own, and they act to achieve their objectives. For example, clients interacting with a server typically have objectives other than to fail the server. Accordingly, the setting induces a non-zero-sum game, where the objectives of the players are induced from their specifications. In *rational synthesis*, we let the system exploit the rationality of the environment. In particular, in *cooperative* rational synthesis (CRS) [18], the desired output is an NE profile whose outcome satisfies the objective of the system. Thus, in CRS, we assume that we can suggest strategies to the environment players, and if they have no incentive to deviate from these strategies, they follow them. Rational synthesis has been extensively studied in various settings and variants [33, 13, 1, 22, 9, 10]. Recall that strategies for the players direct them how to move the token in vertices they own. A strategy may depend on the history of the game so far. Thus, in different visits of the token to the same vertex v , a strategy may direct the owner of v to move the token to different successors. Extensive research has concerned the *memory requirements* for strategies in zero-sum games with ω -regular objectives [30, 14, 5, 11, 8]. For example, it is well known that a winning strategy for a Büchi objective can be *memoryless*, thus it may depend only on the current vertex. On the other hand, a winning strategy for a conjunction of k Büchi objectives may require memory k [14]. Researchers have also studied games in which the memory of the players is bounded [29, 15, 17, 20, 23]. Clearly, increasing the memory of the system or reducing the memory of the environment in a zero-sum game can only help the system to win a game.

For non-zero-sum games, the situation is less clear. First, as we show in Section 3, even when its objective is memoryless, the system may need memory for its strategy in a CRS solution. Moreover, unlike the situation in zero-sum games, an increase of memory to the environment players may be helpful for the system. That is, even when the objectives are memoryless, the only possible CRS solutions may require the environment players to have memory. Intuitively, increasing the memory of the players in non-zero-sum games enables them to satisfy multiple objectives, which may be essential for achieving both stability and the satisfaction of the system's objective [31]. In fact, we show that when the objectives of the environment players are memoryless, then adding memory to the environment players may only help the system to win.

Our basic observations above raise several interesting questions about memory bounds in non-zero-sum games. We first prove that a CRS solution in a k -player non-zero-sum game with *sink objectives* (that is, ones that can be specified with all reachability or ω -regular objectives) may require each of the players to have memory $O(k)$, matching the known upper bound [24, 31]. Further questions concern richer settings, detailed below.

The notion of an NE corresponds to deviations of single players. In some applications, a *coalition* of players may deviate together. For example, protocols for voting, mechanisms for exchange of messages, allocation and construction of shared resources – all should take into account the possibility of players that deviate together. The different applications induce different notions of stability. The first such definition is of *strong Nash equilibrium* (SNE). A profile is an SNE if no subset of players can deviate in a way that benefits all its members [3].

Then, for $b_1, b_2 \geq 0$, a profile is a (b_1, b_2) -robust equilibrium if no coalition of size b_1 can deviate in a way that benefits at least one of its members without harming the other members, and no coalition of size b_2 can deviate in a way that harms the other players [6]. Finally, a profile is a *strong secure equilibria* (SSE) if every deviation of a coalition of players that harms some player, also harms a player in the coalition [7].

We study the CRS problem when stability is defined with respect to deviations by a coalition of players. For example, in the *SNE-CRS problem*, we seek a profile of strategies that satisfies the objective of the system and is an SNE. The contributions in [6, 7] include a study of the complexity of the CRS problem when the solution concepts are robust-equilibrium and SSE. PSPACE upper bounds for the problems involve a reduction to a two-player game, termed *the deviator game* [6], which we easily extend to SSE-CRS. For the lower-bound, our contribution is more interesting and involves relating non-cooperation in rational synthesis with deviations of coalitions in SNEs: Recall that in cooperative rational synthesis, the desired output is an NE that satisfies the system objective. In *non-cooperative* rational synthesis (NRS) [21], the desired output is a strategy for the system player such that the objective of the system is satisfied in the outcome of all NE profiles that include this strategy. Thus, in NRS, the environment players are rational, but we cannot suggest them a strategy. As shown in [1], the cooperative and non-cooperative approaches are related to the two stability-inefficiency measures of *price of stability* [2] and *price of anarchy* [19, 27]. We relate NE-NRS (that is, NRS when stability amounts to being an NE) with SNE-CRS, showing that the challenge of coping with deviations of a coalition is similar to the challenge of coping with non-cooperation. Intuitively, in both cases, all the environment players may deviate simultaneously, as long as these deviations are beneficial for them. The relation implies that the PSPACE-hardness of the NE-NRS problem [13] applies also to the SNE-CRS problem.

Back to the study of memory requirements, we examine how the transition to solution concepts that involve deviations by a coalition of players affects these requirements. Since players in a coalition care for the satisfaction of the objectives of all the players in the coalition, their strategies have to satisfy multiple objectives. Since the latter typically requires memory, the study of the memory requirements in non-zero-sum games with deviations by coalitions is more interesting and involved. We start with some observations about the effect of increasing the memory to the environment players and show that, unlike the case of NE-CRS, here an increase to the memory of the environment players may prevent the existence of an SNE-CRS solution even when the objectives are memoryless.

On the other hand, for some games, the existence of an SNE-CRS solution requires the environment players to have memory, and we examine the memory requirements for them. For the upper bound, it is not hard to extend the analysis in [6] and describe an exponential upper bound for the required memory. Our main technical contribution is a matching lower bound. We show that even in k -player games with sink objectives, an SNE-CRS solution may require $O(k)$ players to have memory $2^{\Theta(k)}$. Moreover, our bounds apply also to the solution concepts of $(k, 0)$ -robust-equilibrium and SSE, completing the picture to all known solution concepts with deviations of a coalition of players.

2 Preliminaries

Games. For $k \geq 1$, let $[k] = \{1, \dots, k\}$. A k -player (turn-based) game graph is a tuple $G = \langle \{V_i\}_{i \in [k]}, v_0, E \rangle$, where $V_1, \dots, V_k$ are disjoint sets of vertices. For every $i \in [k]$, the vertices in V_i are owned by Player i , and we let $V = \bigcup_{i \in [k]} V_i$. Then, $v_0 \in V$ is an initial vertex, and $E \subseteq V \times V$ is a total edge relation, thus for every $v \in V$, there is $u \in V$ such that $\langle v, u \rangle \in E$. For $v \in V$, we denote by $\text{owner}(v)$ the player $i \in [k]$ such that $v \in V_i$. The size of G , denoted $|G|$, is $|E|$, namely the number of edges in it.

A game is a tuple $\mathcal{G} = \langle G, \{\alpha_i\}_{i \in [k]} \rangle$, where G is a k -player game graph, and α_i , for $i \in [k]$, is a winning condition (a.k.a. *objective*) for Player i . In the beginning of a play in the game, a token is placed on v_0 . Then, in each turn, the player that owns the vertex that hosts the token chooses a successor vertex and moves the token to it. Together, the players generate a *play* $\rho = v_0, v_1, v_2 \dots \in V^\omega$ in $\mathcal{G}$, namely an infinite path that starts in v_0 and respects E : for all $i \geq 0$, we have that $\langle v_i, v_{i+1} \rangle \in E$.

Each winning condition α_i defines a subset of V^ω . The objective of Player i is to cause the interaction to generate a play that satisfies α_i . We describe some types of winning conditions below. For a play $\rho = v_0, v_1 \dots$, we denote by $\text{reach}(\rho)$ the set of vertices visited at least once along ρ , and by $\text{inf}(\rho)$ the set of vertices visited infinitely often along ρ . That is, $\text{reach}(\rho) = \{v \in V : \text{there exists some } i \geq 0 \text{ such that } v_i = v\}$ and $\text{inf}(\rho) = \{v \in V : \text{there are infinitely many } i \geq 0 \text{ such that } v_i = v\}$. A *reachability* objective is given by a set of vertices $\alpha \subseteq V$, and it requires some vertex in α to be visited at least once; thus a play $\rho \in V^\omega$ satisfies α iff $\text{reach}(\rho) \cap \alpha \neq \emptyset$. A *Büchi* objective is given by a set of vertices $\alpha \subseteq V$, and it requires some vertex in α to be visited infinitely often; thus ρ satisfies α iff $\text{inf}(\rho) \cap \alpha \neq \emptyset$. The objectives dual to reachability and Büchi are *avoid* (also known as *safety*) and *co-Büchi*, respectively. Formally, a play ρ satisfies an avoid objective $\alpha \subseteq V$ iff $\text{reach}(\rho) \cap \alpha = \emptyset$, and satisfies a co-Büchi objective $\alpha \subseteq V$ iff $\text{inf}(\rho) \cap \alpha = \emptyset$. A *generalized* Büchi objective is a set $\alpha = \{\alpha_1, \dots, \alpha_m\}$ of Büchi objectives. A play $\rho \in V^\omega$ satisfies α if it satisfies all the objectives in α ; thus if for all $j \in [m]$, we have that $\text{inf}(\rho) \cap \alpha_j \neq \emptyset$. Generalized reachability objectives are defined similarly, requiring all underlying reachability objectives to be satisfied.

Strategies, profiles, and equilibria. For $i \in [k]$, a strategy for Player i is a function $f_i : V^* \cdot V_i \rightarrow V$ that maps prefixes of plays that end in a vertex owned by Player i to possible extensions in a way that respects E . That is, for every history $h \in V^*$ and $v \in V_i$, we have that $\langle v, f_i(h \cdot v) \rangle \in E$. Intuitively, a strategy for Player i directs her how to move the token, and the direction may depend on the history of the play so far.

A *profile* is a tuple $\pi = \langle f_1, \dots, f_k \rangle$ of strategies, one for each player. The *outcome* of a profile $\pi = \langle f_1, \dots, f_k \rangle$ is the play obtained when the players follow their strategies. Formally, $\text{outcome}(\pi) = v_0, v_1, v_2, \dots \in V^\omega$ is such that for all $j \geq 0$, we have that $v_{j+1} = f_{\text{owner}(v_j)}(v_0 \dots v_j)$. Consider a game $\mathcal{G}$ and a profile π . The set of winners in $\mathcal{G}$ when the players follow π , denoted $\text{Win}(\mathcal{G}, \pi)$, is the set of players whose objectives are satisfied in $\text{outcome}(\pi)$. Formally, $i \in \text{Win}(\pi, \mathcal{G})$ iff $\text{outcome}(\pi)$ satisfies α_i . The set of losers in π , denoted $\text{Lose}(\mathcal{G}, \pi)$, is then $[k] \setminus \text{Win}(\pi)$, namely the set of players whose objectives are not satisfied in $\text{outcome}(\pi)$. When $\mathcal{G}$ is known from the context we write $\text{Win}(\pi)$ and $\text{Lose}(\pi)$ respectively.

For a subset $S \subseteq [k]$ of players, an *S-profile* is a set of strategies, one for each player in S . We say that a profile π extends an *S-profile* π' if the players in S use in π their strategies in π' . For a profile $\pi = \langle f_1, \dots, f_k \rangle$, a non-empty subset $C \subseteq [k]$, and a *C-profile* $\pi'_C = \bigcup_{i \in C} \{f'_i\}$, we denote by $\pi[C \leftarrow \pi'_C]$ the profile in which the players in C follow their strategies in π'_C and the players in $[k] \setminus C$ follow their strategies in π . Formally, $\pi[C \leftarrow \pi'_C] = \langle g_1, \dots, g_k \rangle$, where for every $i \in [k]$, we have that $g_i = f'_i$, if $i \in C$, and $g_i = f_i$, otherwise. When $C = \{i\}$ is a singleton, for some $i \in [k]$, we simplify the notation and use $\pi[i \leftarrow f'_i]$ rather than $\pi[\{i\} \leftarrow \{f'_i\}]$.

A profile $\pi = \langle f_1, \dots, f_k \rangle$ is a *Nash Equilibrium* (NE, for short) [26] if no single player has an incentive to deviate from π . Formally, π is an NE if for every $i \in [k]$, if $i \in \text{Lose}(\pi)$, then for every strategy f'_i for Player i , we have that $i \in \text{Lose}(\pi[i \leftarrow f'_i])$. The notion of an NE

assumes deviation by single players. In some applications, a *coalition* of players may deviate together. The different applications induce different definitions of stability. We consider three definitions. First, a profile π is a *Strong Nash Equilibrium* (SNE, for short) [3] if no coalition of players can jointly deviate in a way that strictly benefits all its members. Formally, π is an SNE if for every non-empty subset $C \subseteq \text{Lose}(\pi)$, and every C -profile π'_C , there exists $j \in C$ such that $j \in \text{Lose}(\pi[C \leftarrow \pi'_C])$. Note that an NE is a special case of an SNE in which only deviations of coalitions of size 1 are possible. Then, for $b_1, b_2 \geq 0$, a profile is a (b_1, b_2) -robust equilibrium if no coalition of size b_1 can deviate in a way that benefits at least one of its members without harming at least one of the other members, and no coalition of size b_2 can deviate in a way that harms at least one of the other players [6].¹ Formally, π is (b_1, b_2) -robust if π is b_1 -resilient: for every subset $C \subseteq [k]$ of size at most b_1 , and every C -profile π'_C , if $\text{Lose}(\pi) \cap \text{Win}(\pi[C \leftarrow \pi'_C]) \neq \emptyset$, then $\text{Win}(\pi) \cap \text{Lose}(\pi[C \leftarrow \pi'_C]) \neq \emptyset$, and π is b_2 -immune: for every subset $C \subseteq [k]$ of size at most b_2 , and every C -profile π'_C , we have that $\text{Win}(\pi) \cap \text{Lose}(\pi[C \leftarrow \pi'_C]) \cap ([k] \setminus C) \neq \emptyset$. Finally, a profile is a *strong secure equilibria* (SSE) if every deviation of a coalition of players that harms some player, also harms a player in the coalition [7]. Formally, π is an SSE if for every subset $C \subseteq [k]$, and every C -profile π'_C , if $\text{Win}(\pi) \cap \text{Lose}(\pi[C \leftarrow \pi'_C]) \cap ([k] \setminus C) \neq \emptyset$, then $\text{Win}(\pi) \cap \text{Lose}(\pi[C \leftarrow \pi'_C]) \cap C \neq \emptyset$.

For a subset $W \subseteq [k]$ of players, we say that π is a W -NE if π is an NE with $W = \text{Win}(\pi)$, and similarly for the other solution concepts.

Rational Synthesis. We consider a setting in which the players model a controllable system and its rational environment. Technically, we assume that Player 1 models the system (a system may be composed from several components, but since the system is controllable, we can merge them to a single player), and the other players model the components of the environment. Let $\text{Env} = \{2, \dots, k\}$. The basic problem we consider is the existence and finding stable profiles that satisfy the objective of the system.

We refine the notions of NE and SNE to take into account our ability to control the system. For a profile π , we say that π is a *1-fixed NE* if no player in Env can benefit from unilaterally changing her strategy. Likewise, we say that π is a *1-fixed SNE* if no coalition of players in Env can jointly deviate in a way that strictly benefits all its members.

Consider a k -player game $\mathcal{G}$. The problem of *NE (SNE) cooperative rational synthesis*, denoted NE-CRS (resp., SNE-CRS) is to return a 1-fixed NE (resp., SNE) in $\mathcal{G}$ in which Player 1 wins. As in traditional synthesis, one can also define the corresponding decision problems, of rational realizability, where we only need to decide whether the desired strategies exist. In order to avoid additional notations, we sometimes refer to NE-CRS and SNE-CRS also as decision problems.

Finite-Memory Strategies. A strategy $f_i : V^* \rightarrow V$ is *finite-memory* if it is possible to replace the unbounded histories in V^* by finitely many memories. Formally, a memory structure for a game $\mathcal{G} = \langle G, \{\alpha_i\}_{i \in [k]} \rangle$ with $G = \langle \{V_i\}_{i \in [k]}, v_0, E \rangle$ is $\mathcal{M} = \langle M, \mu_0, \delta \rangle$, consisting of a finite set M of memory states, an initial memory state $\mu_0 \in M$, and an update function $\delta : M \times E \rightarrow M$. A memory structure is similar to an automaton with alphabet E , which is executed in parallel to the game: it starts in μ_0 and reads the edges traversed by the token. Then, a strategy for Player i that relies on $\mathcal{M}$ replaces the dependency on the

¹ The setting in [6] considers weighted objectives, where rather than winning or losing, each profile induces a payoff for each player, which enables also a quantitative definition of “harm” and “benefit”. Here we consider ω -regular Boolean objectives, inducing a Boolean interpretations for “harm” and “benefit”.

history of the play by dependency on the current memory state of $\mathcal{M}$. Thus, the strategy is given by a function $f_i : M \times V_i \rightarrow V$, such that for all $\mu \in M$ and $v \in V_i$, we have that $\langle v, f_i(\mu, v) \rangle \in E$. When the current memory state is μ and the token is in vertex $v \in V_i$, Player i moves the token to $f_i(\mu, v)$ and $\mathcal{M}$ moves to state $\delta(\mu, \langle v, f_i(\mu, v) \rangle)$. A strategy f_i is *memoryless* if it only depends on the current vertex. That is, if for every two histories $h, h' \in V^*$ and vertex $v \in V_i$, we have that $f_i(h \cdot v) = f_i(h' \cdot v)$. Note that a memoryless strategy can be viewed as a function $f_i : V_i \rightarrow V$, and corresponds to the case $|M| = 1$. An objective type γ is called *memoryless* if in every zero-sum game with an objective of type γ , if Player 1 wins, then she has a memoryless winning strategy. Reachability, avoid, Büchi, and co-Büchi objectives are all memoryless [30].

For $l \geq 1$, we say that a strategy $f_i : V^* \cdot V_i \rightarrow V$ *uses memory l* if a memory structure that generates f_i needs l states. Given a profile $\pi = \langle f_1, \dots, f_k \rangle$, we say that Player i uses memory l (in π) if f_i uses memory l .

3 Is Memory for the Environment Helpful?

In zero-sum games, increasing the memory of the environment may only decrease the ability of the system to satisfy the specification. Formally, for every zero-sum game $\mathcal{G}$ and for every bound $m \geq 1$, if Player 1 wins against Player 2 that uses memory m , then for every $m' \leq m$, Player 1 also wins against Player 2 that uses memory m' , and possibly there is $m'' > m$ such that Player 1 loses against Player 2 that uses memory m'' .

In this section we show that the picture in non-zero-sum games is different and more involved. First, the system may need memory in order to have a CRS solution in a game with memoryless objectives for all players. In addition, memory for the environment is required for the existence of a CRS solution in some cases yet prevents the existence of a CRS solution in other cases. Intuitively, memory for the environment enables the system to suggest to the environment richer strategies, but also enables the environment to have richer deviations.


We first describe cases in which increasing the memory of the system and the environment is required for a CRS solution, even in games with memoryless objectives for all players. Our examples are with two-player games, and thus apply to both NE-CRS and SNE-CRS.

► **Theorem 1.** *There are two-player Büchi (or reachability) games $\mathcal{G}_1$ and $\mathcal{G}_2$ such that the following hold.*

1. *There is a CRS solution for $\mathcal{G}_1$ in which Player 1 uses a memory of size 2 and there is no CRS solution for $\mathcal{G}_1$ in which Player 1 is memoryless.*
2. *There is a CRS solution for $\mathcal{G}_2$ in which Player 2 uses a memory of size 2 and there is no CRS solution for $\mathcal{G}_2$ in which Player 2 is memoryless.*

Proof. We describe $\mathcal{G}_1$ and $\mathcal{G}_2$ with Büchi objectives. The same games and considerations apply when the games have reachability objectives. Consider the Büchi game $\mathcal{G}_1$ in Figure 1 (left). Let $\alpha_1 = \{v_3\}$ and $\alpha_2 = \{v_1, v_4\}$. Drawing two-player games we use circles and boxes to describe the vertices in V_1 and V_2 , respectively.

Consider the strategy f_1 for Player 1 that, in v_2 , alternates between v_3 and v_4 . That is $f_1(v_0, v_2, (v_3, v_2, v_4, v_2)^*) = v_3$ and $f_1(v_0, v_2, (v_3, v_2, v_4, v_2)^*, v_3, v_2) = v_4$. Note that in order to implement the alternation, the strategy f_1 requires a memory of size 2. Consider the strategy f_2 for Player 2 that, in v_0 , takes the token down to v_2 . Note that the profile $\langle f_1, f_2 \rangle$ is a CRS solution. Indeed, its outcome is $v_0, (v_2, v_3, v_2, v_4)^\omega$, which visits both v_3 and v_4 infinitely often. Thus, both objectives are satisfied, and the profile is a CRS solution.



■ **Figure 1** The games $\mathcal{G}_1$ and $\mathcal{G}_2$.

On the other hand, consider a memoryless strategy for Player 1. If, in v_0 , Player 2 moves the token to v_1 , then the play reaches and stays forever in v_1 no matter what the strategy of Player 1 is, and α_1 is not satisfied. If, in v_0 , Player 2 moves the token to v_2 , then either, in v_2 , Player 1 always moves the token to v_4 , in which case the play never visits v_3 , and so α_1 is not satisfied, or Player 1 always moves the token to v_3 , in which case α_2 is not satisfied, causing Player 2 to deviate in v_0 . Thus, there is no CRS solution in which Player 1 uses a memoryless strategy.

Consider now the Büchi game $\mathcal{G}_2$ in Figure 1 (right). Let $\alpha_1 = \{v_3\}$ and $\alpha_2 = \{v_1, v_4\}$. Note that all the vertices with more than one successor belong to Player 2, and so there is a single strategy f_1 for Player 1 in the game. Consider the strategy f_2 for Player 2 that, in v_0 , takes the token down to v_2 , and, in v_2 , alternates between v_3 and v_4 . That is $f_2(v_0, v_2, (v_3, v_2, v_4, v_2)^*) = v_3$ and $f_2(v_0, v_2, (v_3, v_2, v_4, v_2)^*, v_3, v_2) = v_4$. Note that in order to implement the alternation, the strategy f_2 requires a memory of size 2. It is easy to see that the profile $\langle f_1, f_2 \rangle$ is a CRS solution. Indeed, its outcome is $v_0, (v_2, v_3, v_2, v_4)^\omega$, which satisfies both objectives.

On the other hand, consider a memoryless strategy for Player 2. If, in v_0 , Player 2 moves the token to v_1 , then the play reaches and stays forever in v_1 and α_1 is not satisfied. If, in v_0 , Player 2 moves the token to v_2 , then either, in v_2 , Player 2 always moves the token to v_4 , in which case α_1 is not satisfied, or Player 2 always moves the token to v_3 , in which case α_2 is not satisfied, causing Player 2 to deviate in either v_0 or v_2 . Thus, there is no CRS solution in which Player 2 uses a memoryless strategy. ◀

The example of $\mathcal{G}_2$ in Theorem 1 shows that increasing the memory of the environment may help the system to achieve a CRS solution. We continue and examine whether this is always the case. We first need some notations. Consider a non-zero-sum game $\mathcal{G} = \langle G, \{\alpha_i\}_{i \in [k]} \rangle$. For $m \geq 1$, we say that a profile $\pi = \langle f_1, f_2, \dots, f_k \rangle$ is an *m-bounded 1-fixed NE* if for every $2 \leq i \leq k$, the strategy f_i uses memory at most m , and if $i \in \text{Lose}(\pi)$, then for every strategy f'_i for Player i that uses memory at most m , we have that $i \in \text{Lose}(\pi[i \leftarrow f'_i])$. Thus, environment players are restricted to strategies that use memory at most m , in both π and their deviations. Likewise, π is an *m-bounded 1-fixed SNE* if all the strategies of the environment players in π use memory at most m and no coalition of environment players can jointly deviate to strategies that use memory at most m in a way that strictly benefits all its members. Then, we say that π is an *m-bounded NE-CRS (SNE-CRS) solution* if π is an *m-bounded 1-fixed NE (SNE, respectively)* that satisfies α_1 . Note that the usual CRS problem coincides with the case $m = \infty$.

We first show that, unsurprisingly, when the objectives of the environment players require memory, in particular when they consist of a conjunction of objectives, then a system may have a CRS solution only thanks to bounds on the memory of the environment players.

► **Theorem 2.** *There is a two-player generalized-Büchi (or generalized reachability) game $\mathcal{G}_3$ such that there is no CRS solution for $\mathcal{G}_3$, yet there is a 1-bounded CRS solution for $\mathcal{G}_3$.*


Proof. We prove the theorem for the Büchi case. The same game and considerations apply for generalized-reachability objectives.² Consider the generalized Büchi game $\mathcal{G}_3$ played on the graph of $\mathcal{G}_2$ (Figure 1, right), now with $\alpha_1 = \{\{v_1\}\}$ and $\alpha_2 = \{\{v_3\}, \{v_4\}\}$. Note that Player 1 has a Büchi objective.

Recall that all the vertices with more than one successor belong to Player 2, and so there is a single strategy f_1 for Player 1 in the game. Consider the strategy f_2 for Player 2 that, in v_0 , takes the token to v_1 , where it loops forever. Clearly, the induced play satisfies α_1 . When Player 2 is restricted to memoryless strategies, it cannot satisfy her objectives, making this profile a 1-bounded CRS solution. On the other hand, when Player 2 is not memoryless, she can take the token down to v_2 and then alternate between v_3 and v_4 . Hence, when Player 2 has memory 2 or more, she would deviate from every profile that leads to v_0 , and so there is no CRS solution in $\mathcal{G}_3$. ◀

The example of $\mathcal{G}_3$ in the proof of Theorem 2 heavily relies on Player 2 having an objective that requires memory. We now show that for SNE-CRSs, when a deviation of a player may need to be beneficial to several players, a system may have an SNE-CRS solution only thanks to bounds on the memory of the environment players, even when all players have memoryless objectives.

► **Theorem 3.** *There is a 3-player Büchi game $\mathcal{G}_4$ such that there is no SNE-CRS solution for $\mathcal{G}_4$, yet there is a 1-bounded SNE-CRS solution for $\mathcal{G}_4$.*

Proof. We prove the theorem for the Büchi case. The same game and considerations apply for reachability objectives. Consider the 3-players Büchi game $\mathcal{G}_4$ in Figure 2. We use diamonds to denote vertices controlled by Player 3. Let $\alpha_1 = \{v_1\}$, $\alpha_2 = \{v_3\}$, and $\alpha_3 = \{v_4\}$.



■ **Figure 2** The game $\mathcal{G}_4$.

Note that all the vertices with more than one successor belong to Player 2 or Player 3, and so there is a single strategy f_1 for Player 1 in the game. We first describe a 1-bounded SNE-CRS solution for $\mathcal{G}_4$. Let f_2 be the memoryless strategy for Player 2 that, in v_0 , takes the token to v_1 and, in v_2 , take the token to v_3 . Let f_3 be the memoryless strategy for Player 3 that loops in v_5 . Clearly, $\text{outcome}(\pi) = v_0, v_1^\omega$, and so $\text{Win}(\pi) = \{1\}$. We prove that

² The application to reachability is less straightforward here. In particular, for Büchi, one could give up the vertex v_2 and let v_0 have three successors. For reachability, we need the decision of Player 2 about not visiting v_1 in her first transition to be un-recoverable.

Player 2 and Player 3 cannot deviate to memoryless strategies in a way that causes their objectives to be satisfied. Clearly, a deviation by Player 3 alone cannot affect the outcome of the game. Also, a deviation by Player 2 alone may not cause the outcome to reach v_3 or v_4 . Consider now a joint deviation of Player 2 and Player 3. When Player 2 uses a memoryless strategy, it cannot cause the outcome of the profile to visit both v_3 and v_4 infinitely often. Thus, the deviation is not beneficial to either Player 2 and Player 3. Hence, π is a 1-bounded SNE-CRS solution.

On the other hand, when Player 2 has memory $m \geq 2$, then for every profile whose outcome reaches v_1 , Player 2 and Player 3 would deviate to an outcome that satisfies their both objectives, and so no SNE-CRS solution exists for $\mathcal{G}_4$. $\blacktriangleleft$

We now complete the picture and show that when the objective of the environment players are memoryless and deviations are allowed only for single players, then adding memory to the environment may only help the system. Thus, for NE-CRS with memoryless objectives, we cannot have examples as those in Theorems 2 and 3.

► **Theorem 4.** *For every Büchi (or reachability) game $\mathcal{G}$, if there is a 1-bounded NE-CRS solution in $\mathcal{G}$, then there is also an NE-CRS solution in $\mathcal{G}$.*

Proof. Consider a Büchi (or reachability) game $\mathcal{G} = \langle G, \{\alpha_i\}_{i \in [k]} \rangle$. Let $\pi = \{f_1, f_2, \dots, f_k\}$ be a 1-bounded NE-CRS solution. We prove that there exists a strategy g_1 for Player 1 such that the profile $\pi' = \pi[1 \leftarrow g_1]$ is an NE-CRS solution.³

For every $i \in \text{Env}$, let $G^{\pi, i}$ be the graph obtained from G by removing edges that leave vertices in V_j that do not agree with the memoryless strategy f_j , for all $j \in [k] \setminus \{1, i\}$. Consider the zero-sum two-player game $\mathcal{G}^{\pi, i} = \langle G^{\pi, i}, \alpha_i \rangle$ between Player i and Player 1. Let $W_i \subseteq V$ be the winning region of Player i in $\mathcal{G}^{\pi, i}$, and let f'_i be a memoryless strategy for Player i for the vertices in W_i . That is, $v \in W_i$ iff for every strategy f'_1 for Player 1, the outcome in $\mathcal{G}^{\pi, i}$ of f'_i and f'_1 from v satisfies α_i . Likewise, let f_1^i be a memoryless strategy for Player 1 in $\mathcal{G}^{\pi, i}$ that is winning in all vertices not in W_i . That is, $v \notin W_i$ iff for every strategy g_i for Player i , the outcome in $\mathcal{G}^{\pi, i}$ of g_i and f_1^i from v does not satisfy α_i . Note that since α_i is a Büchi objectives, memoryless strategies f'_i and f_1^i exist.

Let $\rho = v_0, v_1, v_2, \dots = \text{outcome}(\pi)$. We first argue that for all $i \in \text{Lose}(\pi)$ and vertices $v \in \text{reach}(\rho) \cap V_i$, we have that $v \notin W_i$. To see this, assume by way of contradiction that there exists $i \in \text{Lose}(\pi)$ and a vertex $v \in \text{reach}(\rho) \cap V_i$ such that $v \in W_i$. Let $v = v_j$ be the first such vertex. That is, $v_0, \dots, v_{j-1}$ are all not in W_i and $v_j \in W_i$. Consider the strategy g_i for Player i that agrees with f_i in the vertices $v_0, \dots, v_{j-1}$ and agrees with f'_i in all other vertices. Consider the profile $\pi'_i = \pi[i \leftarrow g_i]$. Note that g_i is memoryless, $\text{outcome}(\pi'_i)$ has a prefix $v_0, \dots, v_j$, and, as f'_i is winning in $\mathcal{G}^{\pi, i}$, the outcome continues in a way that satisfies α_i . Thus, $i \in \text{Win}(\pi[i \leftarrow g_i])$, contradicting the fact that π is a 1-bounded 1-fixed NE.

We can now define the strategy g_1 , as follows. As long as the generated play follows ρ , then g_1 agrees with f_1 . If for some $i \in \text{Env}$ and vertex $v \in \text{reach}(\rho) \cap V_i$, Player i deviates and moves the token to a successor of v that is different from $f_i(v)$, then g_1 follows the strategy f_1^i for the rest of the game. Note that g_1 need not be memoryless (even when f_1 is memoryless).

³ In Appendix A.1, we show that the transition to g_1 is essential, thus π need not be an NE-CRS solution. In fact, the example is stronger, showing that a profile π may be a 1-bounded NE-CRS solution and still no profile in which the system follows its strategy in π is an NE-CRS solution. We also show that the strategy of Player 1 in a 1-bounded NE-CRS solution may require memory of size 2.

We argue that the profile $\pi' = \pi[1 \leftarrow g_1]$ is an NE-CRS solution. First, since g_1 agrees with f_1 as long as the environment players follow their strategies in π , then $\text{outcome}(\pi') = \text{outcome}(\pi)$, and so $\text{Win}(\pi') = \text{Win}(\pi)$. Thus, $1 \in \text{Win}(\pi')$. It is left to show that π' is a 1-fixed NE. Consider a player $i \in \text{Lose}(\pi)$ and a strategy f'_i for Player i . Let $h \cdot v \in V^* \cdot V_i$ be the longest prefix of $\text{outcome}(\pi'[i \leftarrow f'_i])$ that agrees with ρ . Thus, $f'_i(h \cdot v) \neq f_i(v)$, and so, by its definition, the strategy g_1 starts to follow the strategy f_1^i after the history $h \cdot v$. Since $v \in \text{reach}(\rho) \cap V_i$, then $v \notin W_i$. Therefore, the strategy f_1^i is winning in $\mathcal{G}^{\pi, i}$ when the game starts in v . Hence, $\text{outcome}(\pi'[i \leftarrow f'_i])$ does not satisfy α_i , and so f'_i is not a beneficial deviation for Player i . Thus, π' is an NE-CRS solution, and we are done. $\blacktriangleleft$

4 Memory Requirements for NE-CRS

In this section we consider the memory requirements for NE-CRS, namely when the solution concept is an NE. Consider a Büchi game $\mathcal{G} = \langle G, \{\alpha_i\}_{i \in [k]} \rangle$. In [31], Ummels shows⁴ that for a desired set W of winners, if there exists a W -NE in $\mathcal{G}$, then there exists a W -NE in which the memory of all players is of size $O(k)$. Since an NE-CRS solution corresponds to a 1-fixed W -NE with $1 \in W$, and once $1 \in W$, then a profile is a 1-fixed W -NE iff it is a W -NE, the result provides an upper bound also to our problem. Below we prove a matching lower bound. The proof is not too complicated, and mainly serves as a warm-up to the study of SNE-CRS.

The lower bound holds already the class of *sink games*, defined below. Consider a graph G . A vertex in G is a *sink* if it has only one outgoing edge, which is a self-loop. Then, a k -player *sink game* $\mathcal{G} = \langle G, \{\alpha_i\}_{i \in [k]} \rangle$ is a game in which the only cycles in G are sinks, and for every $i \in [k]$, the objective $\alpha_i \subseteq V$ is a set of sinks. Note that since once a play reaches a sink it stays there forever, the objective in sink games can be described using reachability, avoid, Büchi, or co-Büchi objectives.


► **Theorem 5.** *For every $k > 2$, we can construct a k -player sink game $\mathcal{G}_k$ such that $\mathcal{G}_k$ has an NE-CRS solution, and every CRS solution for $\mathcal{G}_k$ requires all the environment players to have memory $k - 2$.*

Proof. We define $\mathcal{G}_k$ as follows (see an illustration for the case $k = 4$ in Figure 3).

A play in $\mathcal{G}_k$ starts at the initial vertex d_2 . For each $i \in \text{Env}$, Player i controls the vertex d_i and decides whether to move the token to d_{i+1} or to a vertex c_j for some $j \in \text{Env} \setminus \{i\}$. For each $j \in \text{Env}$, Player j also controls the vertex c_j . From c_j , Player j chooses a successor s_i for some $i \in \text{Env}$. The vertex s_i is a sink that satisfies the objectives of all players in $\text{Env} \setminus \{i\}$. The vertex d_{k+1} is a sink that satisfies the objective of Player 1.

Note that the only play in which Player 1 achieves her objective is $d_2, d_3, \dots, d_k, (d_{k+1})^\omega$, in which all players in Env , when controlling their d_i vertices, choose to move the token towards d_{k+1} . If for some $i \in \text{Env}$, Player i instead chooses to move the token to a vertex c_j for some $j \in \text{Env} \setminus \{i\}$, then Player j can move the token from c_j to s_i , making the deviation non-beneficial for Player i . Thus, an NE-CRS solution for $\mathcal{G}_k$ consists of strategies that direct the token to d_{k+1} and “punish” players that do not follow such a direction. Consider $j \in \text{Env}$. In order to punish Player i by moving the token from c_j to s_i , Player j has to remember the vertex from which the token has reached c_j . Since there are $k - 2$ such possible vertices, $\mathcal{G}_k$

⁴ The study in [31] considers Streett objectives, and includes also the parameter of the number of pairs in the objectives. It also combines the memory of the strategy with the size of the state space. The presentation here includes a straightforward adjustment to our setting.



■ **Figure 3** The game G_4 . Each vertex is labeled by its owner (top). The colors correspond to owners.

has a NE-CRS solution where the strategy of each player uses $k - 2$ memory. Also, if for some $i \in [k]$, the strategy of Player i has a memory smaller than $k - 2$, then there exists a player $j \in \text{Env} \setminus \{i\}$ that can move from d_j to c_i without being punished, which implies that the profile is not an NE-CRS solution.

In Appendix A.2, we describe $\mathcal{G}_k$ and prove the bounds formally. ◀

5 On the SNE-CRS Problem

In this section we analyze the complexity of the SNE-CRS problem and the memory required for the players in an SNE-CRS solution. For both problems, the upper bounds follow easily from the study of robust equilibrium and SSE. Our main contributions are the lower bounds. For the complexity of the SNE-CRS problem, we relate the collaboration of players in SNE with the concept of non-cooperation in rational synthesis. For the results on the memory, lower bounds are open also for the other solution concepts, and we show that our contribution applies for them too.

5.1 Solving SNE-CRS

In this section we prove that the SNE-CRS problem is PSPACE-complete. The upper bound is similar to the one presented for other types of equilibria with respect to deviations by a coalition of players and is based on adjusting the objectives in the *deviator game* of Brenguier [6] to the solution concept of SNE. The adjustment is quite straightforward, and we describe it in the full version. The PSPACE algorithm holds for every objective that can be translated in polynomial space to an *Emerson-Lei* objective⁵ [16]. This clearly includes (but is not limited to) Büchi and co-Büchi objectives.

Our lower bound involves an interesting relation between the collaboration of players in an SNE and the concept of non-cooperation in rational synthesis. We start with a definition of the latter. Recall that in CRS, we assume that the environment players are collaborative, in the sense they would follow a suggested equilibrium. In *non-cooperative rational synthesis*

⁵ An Emerson-Lei objective is given by a Boolean assertion θ over subsets of V . A play $\rho \subseteq V^\omega$ induces an assignments $f_\rho : 2^V \rightarrow \{F, T\}$, where for every set $S \in 2^V$, we have that $f_\rho(S) = T$ iff $\inf(\rho) \cap S \neq \emptyset$. Then, ρ satisfies θ iff f_ρ satisfies θ . For example, the Emerson-Lei objective $\alpha_1 \wedge \alpha_2 \wedge \dots \wedge \alpha_k$ is equivalent to the generalized Büchi objective $\{\alpha_1, \alpha_2, \dots, \alpha_k\}$, and the objective $\alpha_1 \wedge \neg \alpha_2$ is satisfied by plays that visit vertices in α_1 infinitely often and visit vertices in α_2 only finitely often.

(NRS), we cannot suggest a strategy to the environment players and only know they would reach an equilibrium [21]. Accordingly, for NE-NRS, the goal is to return a strategy f_1 for Player 1 such that Player 1 wins in every 1-fixed NE $\langle f_1, f_2, \dots, f_k \rangle$. In other words, Player 1 follows f_1 , and no matter how the environment players behave, then as long as they are rational, and so resulting profile is a 1-fixed NE, the objective of Player 1 is satisfied.

Note that in NE-NRS, the system has to cope only with deviations of single players, yet the players are non-cooperative. On the other hand, in SNE-CRS, the system has to cope with deviations of coalitions of players, yet the players are cooperative, and would follow a suggested 1-fixed SNE. In this section we relate NE-NRS with SNE-CRS, showing that the challenge of coping with deviations of coalitions is similar to the challenge of coping with non-cooperation. Intuitively, in both cases, all the environment players may deviate simultaneously, as long as these deviations are beneficial for them.

We formalize the connection by describing a class $\mathcal{C}$ of games such that for every game $\mathcal{G} = \langle G, \{\alpha_i\}_{i \in [k]} \rangle$ in the class $\mathcal{C}$, there is a game $\mathcal{G}' = \langle G, \{\alpha'_i\}_{i \in [k]} \rangle$ such that there is an NE-NRS in $\mathcal{G}$ iff there is an SNE-CRS in $\mathcal{G}'$. Note that $\mathcal{G}$ and $\mathcal{G}'$ are defined with respect to the same game graph, and only the objectives are different. Moreover, the class $\mathcal{C}$ is strong, in the sense that the PSPACE-hardness of NE-NRS applies already to a game in $\mathcal{C}$ [13]. Accordingly, the results in this section, beyond the interesting connection between non-cooperation and coalitions, also imply PSPACE-hardness for the problem of SNE-CRS.

In the rest of this section we define the class $\mathcal{C}$ and describe the reduction from $\mathcal{G}$ to $\mathcal{G}'$. Recall that in a sink game, the only cycles are sinks, and objectives are set of sinks. The class $\mathcal{C}$ consists of *special sink games*, defined below, which restricts sink games further. Nevertheless, the class of special sink games captures the PSPACE-hardness of NE-NRS for all types of objectives, and we are going to relate NE-NRS solutions in special sink games to SNE-CRS solutions in sink games.

► **Definition 6.** A k -player sink game $\mathcal{G} = \langle G, \{\alpha_i\}_{i \in [k]} \rangle$, with $G = \langle \{V_i\}_{i \in [k]}, v_0, E \rangle$, is *special* if the following holds:

1. $\alpha_1 = \bigcap_{i \in \{3, \dots, k\}} \alpha_i$.
2. $\alpha_2 = V$.
3. For every $i \in \{3, \dots, k\}$ and vertex $v \in V_i$, there is a vertex $u \in \alpha_1$ such that $\langle v, u \rangle \in E$.

Note that the third condition implies that α_1 is not empty and that whenever a token reaches a vertex of Player i , for $i \in \{3, \dots, k\}$, she can move the token to α_1 and cause all players to win.

► **Theorem 7.** Consider a k -player special sink game $\mathcal{G} = \langle G, \{\alpha_1, \alpha_2, \dots, \alpha_k\} \rangle$. The game $\mathcal{G}$ has an NE-NRS solution iff the game $\mathcal{G}' = \langle G, \{\alpha_1, \alpha_2 \setminus \alpha_1, \alpha_3 \setminus \alpha_1, \dots, \alpha_k \setminus \alpha_1\} \rangle$ has an SNE-CRS solution.

Proof. Let $G = \langle \{V_i\}_{i \in [k]}, v_0, E \rangle$, and let $\text{Env} = \{2, \dots, k\}$ be the set of environment players.

Given a play $p \in V^\omega$ in G , let $\text{relevant}(p) \subseteq \text{Env}$ denote the set of environment players that own at least one vertex visited along p . For a profile π , we use $\text{relevant}(\pi)$ to denote $\text{relevant}(\text{outcome}(\pi))$. Note that when we consider deviations in a 1-NE π , only deviations of players in $\text{relevant}(\pi)$ are of interest. Indeed, for every $i \in \text{Env} \setminus \text{relevant}(\pi)$, if Player i changes her strategy, the outcome of the profile is not changed. We say that a play $p \in V^\omega$ is a *trap for Player 1* in $\mathcal{G}$ if it reaches a sink $s \notin \alpha_1$ such that for every $i \in \text{relevant}(p)$, we have that $s \in \alpha_i$. Given a strategy f_i for Player i , and a play $p = v_0, v_1, v_2 \dots$ in G , we say that p *agrees* with f_i (and that f_i agrees with p) if for every $j \geq 0$, if $v_j \in V_i$, then $f_i(v_0, v_1, \dots, v_j) = v_{j+1}$.

In Appendix A.3, we prove that the following three claims are equal. The theorem follows from the equivalence between (C_1) and (C_2) .

- (C_1) The game $\mathcal{G}$ has an NE-NRS solution.
- (C_2) The game $\mathcal{G}'$ has an SNE-CRS solution.
- (C_3) There exists a strategy f_1 for Player 1 such that there does not exist a trap in $\mathcal{G}$ that agrees with f_1 .

Intuitively, both NE-NRS and SNE-CRS solutions consider stable profiles in which Player 1 uses a strategy f_1 as described in (C_3) . In NE-NRS, stability amounts to an NE, and the universal quantification on all stable profiles is explicit in the definition of NRS. In SNE-CRS, we do start with a single profile, but deviations of sets of players capture the many profiles that have to be considered in NRS. Then, as specified in (C_3) , in the setting of special sink games, the relevant deviations correspond to traps, and the two notions coincide. ◀

Recall that Büchi and co-Büchi objectives are special cases of Emerson-Lei objectives, for which the SNE-CRS problem can be solved in PSPACE. Also, sink objectives are special case of both Büchi and co-Büchi objectives. Since the NE-NRS problem is PSPACE-hard already for special sink games [13], Theorem 7 enables us to conclude with the following.

► **Theorem 8.** *The SNE-CRS problem for Büchi and co-Büchi games is PSPACE-complete.*

5.2 Memory Requirements for SNE-CRS

In this section we consider the memory requirements for SNE-CRS. The upper bound is similar to the one known for resilient-CRS and is based on an analysis of the memory required for the players in the corresponding deviator game. Essentially (see details in the full version), the players have to maintain in their memory the set of players who have deviated, as well as memory required for the satisfaction of the objectives of the winning players.

► **Theorem 9.** *Consider a k -player Büchi game $\mathcal{G}$. If there exists a SNE-CRS solution in $\mathcal{G}$, then there also exists an SNE-CRS solution in which the strategy of each player uses memory at most $2^{O(k)}$.*


Our main contribution is a lower bound, which was not studied before, and applies also to other solution concepts.

► **Theorem 10.** *For every $k \geq 4$, and $m \leq k - 2$, there is a sink game $\mathcal{G}_{k,m} = \langle G_k, \{\alpha_i\}_{i \in [k]} \rangle$ such that $\mathcal{G}_{k,m}$ has an SNE-CRS solution, and every SNE-CRS solution requires $O(k)$ environment players to have memory $2^{\Theta(k)}$.*

Proof. We define $\mathcal{G}_{k,m}$ as follows (see $G_{6,2}$ in Figure 4).

Let $\mathcal{S}_m$ denote the set of all subsets of $\{3, \dots, k\}$ of size m . That is, $\mathcal{S}_m = \{A \subseteq \{3, \dots, k\} : |A| = m\}$. For every set of players $A \in \mathcal{S}_m$ and $l \in [m]$, let $A[l]$ be the l -th element in A when the elements are ordered by the usual $\leq$ order on $\{3, \dots, k\}$.

The game begins at the initial vertex q_0 , which is owned by Player 2. From q_0 , Player 2 chooses a subset $A \in \mathcal{S}_m$ and moves the token to the vertex $\langle A, 1 \rangle$. For each subset $A \in \mathcal{S}_m$ and $l \in [m]$, Player $A[l]$ controls the vertex $\langle A, l \rangle$. If $l \neq m$, then the successors of $\langle A, l \rangle$ are the vertices $\perp$ and $\langle A, l + 1 \rangle$. If $l = m$, then the successors of $\langle A, l \rangle$ are the vertices $\perp$ and c_j , for $j \in \{3, \dots, k\}$. For each $i \in \{3, \dots, k\}$, Player i controls the vertex c_i . From c_i , Player i chooses a subset $A \in \mathcal{S}_m$ and moves the token to the vertex $\langle A, i, P \rangle$, where P is a symbol, indicating the game is in the “punishment” layer. For every $A \in \mathcal{S}_m$ and $i \in \{3, \dots, k\}$, the vertex $\langle A, i, P \rangle$ is owned by Player 2, who chooses an index $j \in A$ and moves the token to the vertex $s_{i,j}$. For every $i, j \in \{3, \dots, k\}$, the vertex $s_{i,j}$ is a sink that satisfies the objectives of all players in $[k] \setminus \{1, i, j\}$. The vertex $\perp$ is a sink that satisfies the objective of Player 1.



■ **Figure 4** The sink game $G_{6,2}$. Each vertex is labeled by its owner (top). The colors correspond to owners. Some edges are omitted for clarity. In particular, all the vertices in $\mathcal{S}_2 \times \{1, 2\}$ have an edge to the vertex $\perp$.

The idea behind $\mathcal{G}_{k,m}$ is as follows. Consider a profile $\pi = \langle f_1, \dots, f_k \rangle$. Note that if Player 1 wins in π , then its outcome reaches the vertex $\perp$, in which case all the other players lose in π . Also, Player 2 wins in π iff its the outcome of π reaches a sink of the form $s_{i,j}$, in which case Player 1 loses. Thus, the objectives of Player 1 and Player 2 complement each other. Assume that Player 1 wins in π , thus its outcome reaches $\perp$. In order for π to be an SNE, every deviation of players in $\{2, \dots, k\}$ should result in a profile in which at least one of the players that deviate does not satisfy its objective. Thus, when the coalition of deviators is $C \subseteq \{2, \dots, k\}$, then there should be $i \in C$ such that the outcome of the new profile either continues to reach $\perp$ or, in case $i \neq 2$, reaches a sink $s_{i,j}$ or $s_{j,i}$ for some $j \in \{3, \dots, k\}$. Before we explain why this implies the existence and an SNE-CRS solution, and why strategies for players $3, \dots, k$ in any SNE-CRS solution require exponential memory, let us define $\mathcal{G}_{k,m}$ formally.

We define $G_{k,m} = \langle \langle \{V_i\}_{i \in [k]}, q_0, E \rangle, \{\alpha_i\}_{i \in [k]} \rangle$, as follows.

1. The set of vertices and its partition to owners is as follows.
 - $V_1 = \{\perp\} \cup \{s_{i,j} : i, j \in \{3, \dots, k\}\}$.
 - $V_2 = \{q_0\} \cup \{\langle A, i, p \rangle : A \in \mathcal{S}_m \text{ and } i \in \{3, \dots, k\}\}$.
 - For every $i \in \{3, \dots, k\}$, we define $V_i = \{c_i\} \cup \{\langle A, l \rangle : A \in \mathcal{S}_m, l \in [m], \text{ and } i = A[l]\}$.
2. The set E contains edges of the following types:
 - For every $A \in \mathcal{S}_m$, there is an edge from q_0 to $\langle A, 1 \rangle$.
 - For every $A \in \mathcal{S}_m$ and $l \in [m-1]$, there is an edge from $\langle A, l \rangle$ to $\perp$ and to $\langle A, l+1 \rangle$.
 - For every $A \in \mathcal{S}_m$, there is an edge from $\langle A, m \rangle$ to $\perp$ and to c_j , for all $j \in \{3, \dots, k\}$.
 - For every $A \in \mathcal{S}_m$ and $i \in \{3, \dots, k\}$, there is an edge from c_i to $\langle A, i, p \rangle$.
 - For every $A \in \mathcal{S}_m$, $i \in \{3, \dots, k\}$, and $j \in A$, there is an edge from $\langle A, i, p \rangle$ to $s_{i,j}$.
 - The vertices $\perp$ and $s_{i,j}$, for all $i, j \in \{3, \dots, k\}$, are sinks.
3. The objectives of the players are defined as follows.
 - $\alpha_1 = \{\perp\}$.
 - For every $i \in \{2, \dots, k\}$, we have $\alpha_i = \{s_{l,j} : l, j \in \{3, \dots, k\} \setminus \{i\}\}$.

We first describe an SNE-CRS solution $\pi = \langle f_1, \dots, f_k \rangle$ in $\mathcal{G}_{k,m}$. Since Player 1 only owns sinks, her strategy is straightforward. Moreover, the notions of SNE and 1-fixed SNE coincide in $\mathcal{G}_{k,m}$. For Player 2, the strategy f_2 is an arbitrary memoryless strategy. As for $i \in \{3, \dots, k\}$, the strategy f_i directs Player i to move the token to $\perp$ from vertices in $\mathcal{S}_m \times [m]$ that she owns and to move the token to $\langle A, i, P \rangle$ from c_i . It is easy to see that $\text{outcome}(\pi) = q_0, f_2(q_0), (\perp)^\omega$, and so $\text{Win}(\pi) = \{1\}$. We prove that π is a 1-fixed SNE. Consider a coalition $C \subseteq \{2, \dots, k\}$ and a deviation profile $\{f'_i\}_{i \in C}$. Since Player 2 is losing in π , and win in every outcome that does not reach $\perp$, a deviation for Player 2 that does not reach $\perp$ is always beneficial, and so, we can assume that $2 \in C$. Let $\pi' = \pi[C \leftarrow \{f'_i\}_{i \in C}]$. Let $A \in \mathcal{S}_m$ be the set of players that Player 2 chooses from q_0 in π' ; thus $f'_2(q_0) = \langle A, 1 \rangle$. If there is a player $i \in A$ that in π' moves the token from the vertex $\langle A, l \rangle$ with $A[l] = i$ to the sink $\perp$, then the outcome of π' reaches $\perp$. Thus, $\text{Win}(\pi') = \{1\}$, and the deviation of the players in C is not beneficial. If all players in A do not move the token to $\perp$ in their strategies in π' (in particular, this means that $A \subseteq C$), then let c_i be the vertex to which Player $A[m]$ moves the token from $\langle A, m \rangle$ in π' . Since all the sinks $s_{i,l}$ that are reachable from c_i are losing for Player i , the deviation is not beneficial for Player i , implying that $i \notin C$. Hence, Player i follows f_i , which directs her to move the token from c_i to $\langle A, i, P \rangle$. Since every successor of $\langle A, i, P \rangle$ is losing for some player in A , the deviation is not beneficial to all the players in C , and so π is an SNE in which Player 1 wins.

We continue and prove that every SNE-CRS solution in $\mathcal{G}_{n,k}$ requires each of the players $3, \dots, k$ to have memory $\binom{k-3}{m}$. Assume by contradiction that there exists an SNE profile $\pi = \langle f_1, \dots, f_k \rangle$ in which Player 1 wins and there exists $i \in \{3, \dots, k\}$ such that the strategy f_i has a memory structure $\mathcal{M}_i$ with fewer than $\binom{k-3}{m}$ states. Since there are fewer than $\binom{k-3}{m}$ states in $\mathcal{M}_i$, and there are $\binom{k-3}{m}$ subsets of $\{3, \dots, k\} \setminus \{i\}$ of size m , there exists a set $A \in \mathcal{S}_m$ such that $i \notin A$ and f_i never (that is, no matter what the history along which c_i has been reached) directs Player i to move the token from c_i to $\langle A, i, P \rangle$.

Consider the coalition $C = A \cup \{2\}$, and the deviation profile $f'_C = \{\{f'_j\}_{j \in C}\}$, where for every $j \in A$, the strategy f'_j agrees with f_j , except that from $\langle A, l \rangle$, with $A[l] = j$, the strategy f'_j directs Player j to move the token to $\langle A, l+1 \rangle$, in case $l < m$, and to c_i , in case $l = m$. Finally, the strategy f'_2 agrees with f_2 , except that from q_0 , the strategy f'_2 directs Player 2 to move the token to $\langle A, 1 \rangle$, and for every $A' \in \mathcal{S}_m$ such that $A' \neq A$, the strategy f'_2 directs Player 2 to move the token from $\langle A', i, P \rangle$ to $s_{i, \min A' \setminus A}$. Note that since both A and A' are of size m , the set $A' \setminus A$ is not empty, and thus $\min A' \setminus A$ exists and is in $\{3, \dots, k\}$. Let $\pi' = \pi[C \leftarrow f'_C]$. By the definition of the strategies in π' , there exists $A' \in \mathcal{S}_m$ such that $A' \neq A$ and $\text{outcome}(\pi') = q_0, \langle A, 1 \rangle, \langle A, 2 \rangle, \dots, \langle A, m \rangle, c_i, \langle A', i, P \rangle, (s_{i, \min A' \setminus A})^\omega$. Thus, $\text{Win}(\pi') = [k] \setminus \{1, i, \min A' \setminus A\}$. Since $1, i$, and $\min A' \setminus A$ are all not in A , and thus also not in C , it follows that $C \subseteq \text{Win}(\pi')$. Hence, the deviation to f'_C is beneficial to all the players in C , contradicting the assumption that π is an SNE. $\blacktriangleleft$

6 Memory Requirements for Additional Solution Concepts

Recall that different applications have initiated the study of different solution concepts for settings in which a coalition of players may deviate together. While upper bounds on the concepts of robust equilibria and SSE serve as a basis to our upper bounds here, no lower bounds are known on the memory requirements for CRS solutions with respect to these concepts. In this section we show that the construction in Theorem 10 can be modified to show a lower bound on the memory required to the environment players in CRS solution when the solution concepts are resilient (and hence, robust) equilibria and SSE.

For $k \geq 4$ and $m \in [k - 2]$, consider the k -player sink game $\mathcal{G}'_{k,m}$ obtained from $\mathcal{G}_{k,m}$ by changing the objectives of the players so that now the sink $\perp$ is winning for all players in $[k] \setminus \{2\}$ (rather than for Player 1 only in $\mathcal{G}_{k,m}$). Thus, as in $\mathcal{G}_{k,m}$, the objectives of Player 1 and Player 2 still complement each other, and now, Players $3, \dots, k$ win also in profiles that reach the vertex $\perp$. Also, as in $\mathcal{G}_{k,m}$, all the vertices owned by Player 1 have only one successor, and the notions of 1-fixed equilibrium and (usual) equilibrium coincide.

We start with resilient equilibria, and show that $\mathcal{G}'_{k,m}$ has a resilient-CRS solution, and that every resilient-CRS solution requires $O(k)$ environment players to have memory $2^{\Theta(k)}$. Recall that a profile is a resilient-equilibrium if for every subset $C \subseteq \text{Env}$, every deviation of the players in C that benefits a player in C also harms a player in C . The proof is similar to the proof of Theorem 10. In particular, the profile $\pi = \langle f_1, \dots, f_k \rangle$ described in the proof is a resilient-CRS solution in $\mathcal{G}'_{k,m}$. To see this, recall that $\text{outcome}(\pi)$ reaches $\perp$, and is thus winning in $\mathcal{G}'_{k,m}$ for all the players in $[k] \setminus \{2\}$. Thus, only Player 2 may benefit from a deviation. When Player 2 chooses a set $A \in \mathcal{S}_m$, all the players in A have to deviate on order for $\perp$ to be avoided, and so, by the definition of resilient equilibrium, all of them have to win when the game reach the punishment layer. This, however, is impossible in deviations from π , as for all $i \in \{3, \dots, k\}$, the only way for Player i to make sure she does not lose is to remember the set A and proceed as in π , from c_i to the vertex $\langle A, i, p \rangle$. Moreover, using considerations similar to these in the proof of Theorem 10 (see details in Appendix A.4), if there is $i \in \{3, \dots, k\}$ such that Player i has no memory to keep track of the set A chosen by Player 2, then there is no resilient-CRS. Essentially, in such a case Player 2 can convince a coalition A of players to join the deviation by letting the outcome of the new profile reach c_i , where a player not in A would be punished.

For the solution concept of SSE, we show that $\mathcal{G}'_{k,m}$ has an SSE-CRS solution, and that every SSE-CRS solution requires $O(k)$ environment players to have memory $2^{\Theta(k)}$. Recall that a profile is an SSE if for every coalition $C \subseteq \text{Env}$, every deviation that harms a player not in C also harms a player in C . In $\mathcal{G}'_{k,m}$, a deviation that does not reach the vertex $\perp$ causes Player 1 to lose. Thus, a profile is an SSE iff for every coalition $C \subseteq \text{Env}$ and every deviation that causes the game to reach the punishment layer, at least one player in the coalition loses. Accordingly, the profile π described in the proof of Theorem 10 is an SSE-CRS solution in $\mathcal{G}'_{k,m}$ and in every SSE-CRS solution, every player in $\{3, \dots, k\}$ has to remember the set A chosen by Player 2 from q_0 . Indeed (see details in Appendix A.5), otherwise, a coalition of players in A can deviate by letting the outcome of the new profile reach the vertex c_i , where a player not in A would be punished.

References

- 1 S. Almagor, O. Kupferman, and G. Perelli. Synthesis of controllable Nash equilibria in quantitative objective game. In *Proc. 27th Int. Joint Conf. on Artificial Intelligence*, pages 35–41, 2018.
- 2 E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proc. 45th IEEE Symp. on Foundations of Computer Science*, pages 295–304. IEEE Computer Society, 2004. doi:10.1109/FOCS.2004.68.
- 3 R. Aumann. Acceptable Points in General Cooperative n -Person Games. In *Contributions to the Theory of Games*, volume 4, 1959.
- 4 R. Bloem, K. Chatterjee, and B. Jobstmann. Graph games and reactive synthesis. In *Handbook of Model Checking.*, pages 921–962. Springer, 2018. doi:10.1007/978-3-319-10575-8_27.

- 5 P. Bouyer, N. Fijalkow, M. Randour, and P. Vandenheve. How to play optimally for regular objectives? In *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*, volume 261 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 118:1–118:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.ICALP.2023.118.
- 6 R. Brenguier. Robust equilibria in mean-payoff games. In *Proc. 19th Int. Conf. on Foundations of Software Science and Computation Structures*, volume 9634 of *Lecture Notes in Computer Science*, pages 217–233. Springer, 2016. doi:10.1007/978-3-662-49630-5_13.
- 7 L. Brice, J-F. Raskin, M. Sassolas, G. Scerri, and M. Bogaard. Pessimism of the will, optimism of the intellect: Fair protocols with malicious but rational agents. In *IEEE 38th Computer Security Foundations Symposium*, pages 33–47, 2025.
- 8 T. Brihaye, A. Goeminne, J.C.A. Main, and M. Randour. Reachability games and friends: A journey through the lens of memory and complexity. In *43rd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2023)*, volume 284 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 1:1–1:26. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.FSTTCS.2023.1.
- 9 V. Bruyère, C. Grandmont, and J-F. Raskin. As soon as possible but rationally. In *35th International Conference on Concurrency Theory (CONCUR 2024)*, volume 311 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 14:1–14:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.CONCUR.2024.14.
- 10 V. Bruyère, J-F. Raskin, A. Reynouard, and M. Bogaard. The non-cooperative rational synthesis problem for spes and omega-regular objectives. In *36th International Conference on Concurrency Theory (CONCUR 2025)*, volume 348 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 12:1–12:23. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.CONCUR.2025.12.
- 11 A. Casares. On the minimisation of transition-based rabin automata and the chromatic memory requirements of muller conditions. In *30th EACSL Annual Conference on Computer Science Logic (CSL 2022)*, volume 216 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 12:1–12:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.CSL.2022.12.
- 12 K. Chatterjee, R. Majumdar, and M. Jurdzinski. On Nash equilibria in stochastic games. In *Proc. 13th Annual Conf. of the European Association for Computer Science Logic*, volume 3210 of *Lecture Notes in Computer Science*, pages 26–40. Springer, 2004. doi:10.1007/978-3-540-30124-0_6.
- 13 R. Condurache, E. Filiot, R. Gentilini, and J-F. Raskin. The complexity of rational synthesis. In *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, volume 55 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 121:1–121:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.ICALP.2016.121.
- 14 S. Dziembowski, M. Jurdzinski, and I. Walukiewicz. How much memory is needed to win infinite games. In *Proc. 12th ACM/IEEE Symp. on Logic in Computer Science*, pages 99–110, 1997.
- 15 R. Ehlers. Symbolic bounded synthesis. In *Proc. 22nd Int. Conf. on Computer Aided Verification*, volume 6174 of *Lecture Notes in Computer Science*, pages 365–379. Springer, 2010. doi:10.1007/978-3-642-14295-6_33.
- 16 E.A. Emerson and C.-L. Lei. Modalities for model checking: Branching time logic strikes back. *Science of Computer Programming*, 8:275–306, 1987. doi:10.1016/0167-6423(87)90036-0.
- 17 E. Filiot, N. Jin, and J-F. Raskin. An antichain algorithm for LTL realizability. In *Proc. 21st Int. Conf. on Computer Aided Verification*, volume 5643, pages 263–277, 2009. doi:10.1007/978-3-642-02658-4_22.

- 18 D. Fisman, O. Kupferman, and Y. Lustig. Rational synthesis. In *Proc. 16th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems*, volume 6015 of *Lecture Notes in Computer Science*, pages 190–204. Springer, 2010. doi:10.1007/978-3-642-12002-2_16.
- 19 E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009. doi:10.1016/J.COSREV.2009.04.003.
- 20 O. Kupferman, Y. Lustig, M.Y. Vardi, and M. Yannakakis. Temporal synthesis for bounded systems and environments. In *Proc. 28th Symp. on Theoretical Aspects of Computer Science*, pages 615–626, 2011.
- 21 O. Kupferman, G. Perelli, and M.Y. Vardi. Synthesis with rational environments. *Annals of Mathematics and Artificial Intelligence*, 78(1):3–20, 2016. doi:10.1007/S10472-016-9508-8.
- 22 O. Kupferman and N. Shenwald. Games with trading of control. In *34th International Conference on Concurrency Theory (CONCUR 2023)*, volume 279 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 19:1–19:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.CONCUR.2023.19.
- 23 O. Kupferman and N. Shenwald. Positional-Player Games. In *50th International Symposium on Mathematical Foundations of Computer Science (MFCS 2025)*, volume 345 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 64:1–64:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.MFCS.2025.64.
- 24 J. C. A. Main. Arena-independent memory bounds for nash equilibria in reachability games. In *41st International Symposium on Theoretical Aspects of Computer Science (STACS 2024)*, volume 289 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 50:1–50:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.STACS.2024.50.
- 25 D.A. Martin. Borel determinacy. *Annals of Mathematics*, 65:363–371, 1975.
- 26 J.F. Nash. Equilibrium points in n -person games. In *Proceedings of the National Academy of Sciences of the United States of America*, 1950.
- 27 C. H. Papadimitriou. Algorithms, games, and the internet. In *Proc. 33rd ACM Symp. on Theory of Computing*, pages 749–753, 2001.
- 28 A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proc. 16th ACM Symp. on Principles of Programming Languages*, pages 179–190, 1989.
- 29 S. Schewe and B. Finkbeiner. Bounded synthesis. In *5th Int. Symp. on Automated Technology for Verification and Analysis*, volume 4762 of *Lecture Notes in Computer Science*, pages 474–488. Springer, 2007. doi:10.1007/978-3-540-75596-8_33.
- 30 W. Thomas. On the synthesis of strategies in infinite games. In *Proc. 12th Symp. on Theoretical Aspects of Computer Science*, volume 900 of *Lecture Notes in Computer Science*, pages 1–13. Springer, 1995. doi:10.1007/3-540-59042-0_57.
- 31 M. Ummels. The complexity of Nash equilibria in infinite multiplayer games. In *Proc. 11th Int. Conf. on Foundations of Software Science and Computation Structures*, pages 20–34, 2008.
- 32 J. von Neumann and O. Morgenstern. *Theory of games and economic behavior*. Princeton University Press, 1953.
- 33 M. Wooldridge, J. Gutierrez, P. Harrenstein, E. Marchioni, G. Perelli, and A. Toumi. Rational verification: From model checking to equilibrium checking. In *Proc. of 30th Conf. on Artificial Intelligence*, pages 4184–4190, 2016.

A Missing Proofs and Examples


A.1 On strategies for Player 1 in a 1-bounded NE-CRS solution

We describe two interesting examples. The first is a game $\mathcal{G}_5$ and a CRS solution $\langle f_1, f_2 \rangle$ such that there is no 1-bounded CRS solution $\langle f_1, f'_2 \rangle$ for a memoryless f_2 . The second is a game $\mathcal{G}_6$ such that every 1-bounded CRS solution $\langle f_1, f_2 \rangle$ requires f_1 to have memory of size at least 2.

Consider the Büchi game $\mathcal{G}_5$ in Figure 5 (left). Let $\alpha_1 = \{v_3, v_4\}$ and $\alpha_2 = \{v_4\}$. Consider the following strategy of f_1 of Player 1:

- $f_1(v_0, v_2, (v_3, v_2)^*) = v_3$.
- $f_1(v_0, v_1, v_0, v_2, (v_4, v_2)^*) = v_4$.

Thus, if the token reaches v_2 without visiting v_1 before, Player 1 always direct it to v_3 , where only Player 1 wins, and if the token visits v_1 for one time before reaching v_2 , then Player 1 always direct the token to v_4 , where both players win. Let f_2 be a strategy for Player 2 that proceeds to v_1 once and then to v_2 . It is easy to see that while $\langle f_1, f_2 \rangle$ is a CRS solution, there is no 1-bounded CRS solution $\langle f_1, f'_2 \rangle$ for a memoryless f_2 .



■ **Figure 5** The games $\mathcal{G}_5$ (left) and $\mathcal{G}_6$ (right).

We continue and show that the strategy of Player 1 in a 1-bounded CRS solution may require memory of size 2. Consider the Büchi game $\mathcal{G}_6$ in Figure 5 (right). Let $\alpha_1 = \{v_3\}$ and $\alpha_2 = \{v_4\}$.

Consider a profile $\pi = \langle f_1, f_2 \rangle$ for a memoryless strategy f_1 for Player 1. We claim that π is not a 1-bounded CRS solution. Indeed, if $f_1(v_0) = v_1$, then α_1 is not satisfied, and if $f_1(v_0) = v_2$, then Player 2 can deviate to a memoryless strategy $f'_2(v_2) = v_4$, where only α_2 is satisfied.

Nevertheless, the game $\mathcal{G}_6$ does have a 1-bounded CRS solution $\pi = \langle f_1, f_2 \rangle$, for f_1 that is not memoryless. To see this, consider a strategy f_1 of Player 1 that behaves as follows:

- $f_1((v_0, v_2, v_3)^*, v_0) = v_2$.
- $f_1(v_0, v_2, v_4, v_0) = v_1$.

Thus, when the game starts and as long as Player 2 moves the token from v_2 to v_3 , Player 1 moves the token from v_0 down to v_2 . If Player 2 moves the token from v_2 to v_4 , then on the next visit of the token in v_0 , Player 1 moves it to v_1 , where no player wins. Note that when Player 2 uses a memoryless strategy, then this “next” visit must be the second one, and so the definition of f_1 above covers all possible histories.

It is not hard to see that $\pi = \langle f_1, f_2 \rangle$, for the memoryless strategy f_2 with $f_2(v_2) = v_3$ is a 1-bounded CRS solution

A.2 Missing Details in the proof of Theorem 5

The game graph $G_k = \langle \{V_i\}_{i \in [k]}, d_2, E \rangle$ and the objectives $\{\alpha_i\}_{i \in [k]}$ are defined as follows:

1. The vertex sets of the players are defined as follows:
 - Player 1 controls the set $V_1 = \{d_{k+1}, c_1\} \cup \{s_2, \dots, s_k\}$
 - For every $i \in \{2, \dots, k\}$, Player i controls the set $V_i = \{c_i, d_i\}$.
2. The set E contains edges of the following types:
 - For every $i \in \{2, \dots, k\}$, edges from d_i to d_{i+1} .
 - For every $i \in \{2, \dots, k\}$, there is an edge from d_i to c_j where $j \in [k] \setminus \{i\}$.
 - For every $i \in [k]$, and $j \in \{2, \dots, k\}$, there is an edge from c_i to s_j .
 - The vertices in $\{d_{k+1}\} \cup \{s_2, \dots, s_k\}$ are sink vertices.

3. The objectives are defined as follows:

- The objective of Player 1 is $\alpha_1 = \{d_{k+1}\}$.
- For every $i \in \{2, \dots, k\}$, the objective of Player i is $\alpha_i = \{s_j : j \in [k] \setminus \{i\}\}$.

We begin by describing an NE profile $\pi = \langle f_1, \dots, f_k \rangle$ in which Player 1 wins. For every $i \in \{2, \dots, k\}$, Player i decides to always move from d_i towards d_{k+1} . If there exists $i \in [k] \setminus \{1\}$ such that Player i deviated and moved the token from d_i to c_j for some $j \in [k] \setminus \{i\}$, then Player j , move the token to s_i . The outcome of the profile π is $d_2, d_3, \dots, d_k, d_{k+1}^\omega$, thus, $\text{Win}(\pi) = \{1\}$. If for some $l \in \{2, \dots, k\}$, Player l deviates and move the token from d_l to c_j for some $j \in [k] \setminus \{l\}$, then Player j moves the token to s_l and Player l loses. Meaning that deviation is not beneficial and that π is an NE. Overall π is an NE in which Player 1 wins. Note that the strategy for each players in π can be implemented with a finite memory structure of size $k - 2$. Indeed. each player only needs to remember if a deviation has happened, and if so, which player (besides themselves and Player 1) has deviated.

We now show that in every NE in which Player 1 wins, the strategy of every player uses memory at least $k - 2$. Assume by contradiction that there exists an NE profile $\pi = \langle f_1, \dots, f_k \rangle$ where Player 1 wins and there exists $i \in [k]$ such that the strategy f_i has a memory structure $\mathcal{M}_i$ with less than $k - 2$ states. Then, by the pigeon hole principle, since there are less than $k - 2$ states in $\mathcal{M}_i$ but c_i has $k - 1$ successors, there exists a successors to c_i , which we mark with s_j , such that f_i never chooses s_j as a successor from c_i . Since $j > 2$, we have that $j \in \text{Lose}(\pi)$ and may deviate. Let f'_j be a strategy for Player j that agrees with f_j , except that from d_j , the strategy f'_j always move the token to c_i . Then, since f_i never chooses s_j as a successor from c_i , the profile $\pi' = \pi[i \leftarrow f'_i]$ visits a sink s_m where $m \in [k] \setminus \{1, j\}$. Thus, Player j wins in π' and so π is not stable.

A.3 Missing Details in the proof of Theorem 7

We prove that the following three claims are equal:

- (C₁) The game $\mathcal{G}$ has an NE-NRS solution.
- (C₂) The game $\mathcal{G}'$ has an SNE-CRS solution.
- (C₃) There exists a strategy f_1 for Player 1 such that there does not exist a trap in $\mathcal{G}$ that agrees with f_1 .

We first prove that (C₁) iff (C₃). In fact we prove a stronger claim, namely that for every strategy f_1 for Player 1, we have that f_1 is an NE-NRS solution in $\mathcal{G}$ iff there does not exist a trap in $\mathcal{G}$ that agrees with f_1 .

Consider a strategy f_1 for Player 1. Assume first that f_1 is an NE-NRS solution in $\mathcal{G}$, and assume by way of contradiction that there exists a trap p that agrees with f_1 . Thus, p reaches a sink $s \notin \alpha_1$, and for every $i \in \text{relevant}(p)$, we have that $s \in \alpha_i$. Consider a profile $\pi = \langle f_1, \dots, f_k \rangle$, such that for every $i \in \text{Env}$, the strategy f_i agrees with p . Since $s \notin \alpha_1$, Player 1 loses in π . Since for every $i \in \text{relevant}(p)$, we have that $s \in \alpha_i$, then $\text{relevant}(\pi) \subseteq \text{Win}(\mathcal{G}, \pi)$. Thus, no player in $\text{relevant}(\pi)$ has an incentive to deviate in $\mathcal{G}$. Hence, π is a 1-fixed NE in $\mathcal{G}$ in which Player 1 loses, contradicting the fact f_1 is an NE-NRS solution for $\mathcal{G}$.

For the second direction, assume that there does not exist a trap in $\mathcal{G}$ that agrees with f_1 , and assume by way of contradiction that f_1 is not an NE-NRS solution for $\mathcal{G}$. That is, there is a 1-NE profile $\pi = \langle f_1, \dots, f_k \rangle$ such that $1 \notin \text{Win}(\mathcal{G}, \pi)$. Consider the play $p = \text{outcome}(\pi)$. Since $1 \notin \text{Win}(\mathcal{G}, \pi)$, the play p reaches a sink $s \notin \alpha_1$. Also, since p is not a trap, there exists $i \in \text{relevant}(p)$ such that $s \notin \alpha_i$, which implies that $i \notin \text{Win}(\mathcal{G}, \pi)$. By the third condition on

the structure of special sink games, Player i can deviate and move the token from the first vertex she owns in p to a vertex in α_1 . Since $\alpha_i \subseteq \alpha_1$, such a deviation causes Player i to win in $\mathcal{G}$, contradicting the fact π is a 1-fixed NE.

We continue and prove that (C_2) iff (C_3) . Here too, we prove a stronger claim, namely that for every strategy f_1 for Player 1, we have that there is an SNE-CRS solution $\langle f_1, f_2, \dots, f_k \rangle$ for $\mathcal{G}'$ iff there does not exist a trap in $\mathcal{G}$ that agrees with f_1 .

Assume first that f_1 is such that there is an SNE-CRS solution $\pi = \langle f_1, f_2, \dots, f_k \rangle$ for $\mathcal{G}'$. Recall that for every $i \in \text{Env}$, we have that $\alpha'_i = \alpha_i \setminus \alpha_1$. Thus, as $1 \in \text{Win}(\mathcal{G}', \pi)$, it must be that $\text{Env} \subseteq \text{Lose}(\mathcal{G}', \pi)$.

Assume by way of contradiction there exists a trap p that agrees with f_1 . Let $C = \text{relevant}(p)$. Since $C \subseteq \text{Env}$, then $C \subseteq \text{Lose}(\mathcal{G}', \pi)$. For every $i \in C$, let g_i be a strategy for Player i that agrees with p . Let $g_C = \{g_i : i \in C\}$, and consider the profile $\pi' = \pi[C \leftarrow g_C]$. Since f_1 agrees with p and for all $i \in \text{relevant}(p)$, the strategy g_i agrees with p , we have that $\text{outcome}(\pi') = p$. Since p is a trap, we have that $C \subseteq \text{Win}(\mathcal{G}', \pi')$. Thus, the deviation strictly benefits all the players in C and so π is not a 1-fixed SNE.

Assume now that f_1 is such that there does not exist a trap p that agrees with f_1 . Consider the profile $\pi = \langle f_1, f_2, \dots, f_k \rangle$, where f_2 is some memoryless strategy, and for every $i \in \{3, \dots, k\}$ the strategy f_i is a memoryless strategy that moves the token from every vertex to α_1 . Note that since $\mathcal{G}'$ is a special sink game, such strategies $f_3, \dots, f_k$ exist. We prove that π is an SNE-CRS solution for $\mathcal{G}'$.

Let $p = \text{outcome}(\pi)$. First, we show that p reaches α_1 , which implies that $1 \in \text{Win}(\mathcal{G}', \pi)$. Clearly, if $\text{relevant}(p) \cap \{3, \dots, k\} \neq \emptyset$, then, by the definition of the strategies $f_3, \dots, f_k$, the play p reaches α_1 . Otherwise, $\text{relevant}(p) \subseteq \{1, 2\}$. Then, if p does not reach α_1 , then, by the definition of $\alpha'_2 = V \setminus \alpha_1$, we have that p reaches α'_2 , implying that p is a trap that agrees with f_1 and contradicting the fact that no such trap exists.

Second, we show that π is a 1-fixed SNE in $\mathcal{G}'$. Assume by way of contradiction that π is not a 1-fixed SNE in $\mathcal{G}'$. Then, there exists a coalition $C \subseteq \text{Lose}(\mathcal{G}', \pi)$ and a strategy profile g_C such that for the profile $\pi' = \pi[C \leftarrow g_C]$, we have that $C \subseteq \text{Win}(\mathcal{G}', \pi')$. By the definition of $\alpha'_2, \dots, \alpha'_k$, the latter implies that $\text{outcome}(\pi')$ reaches a sink $s \notin \alpha_1$. Let $p' = \text{outcome}(\pi')$. By the definition of the strategies f_i , it must be that $\text{relevant}(p') \cap \{3, \dots, k\} \subseteq C$. Indeed, otherwise p' would have reached α_1 . Thus, as $C \subseteq \text{Win}(\mathcal{G}', \pi')$, it follows that for every $i \in \text{relevant}(p') \cap \{3, \dots, k\}$, we have that $s \in \alpha_i$. In addition, as $\alpha_2 = V$, we also have that $s \in \alpha_2$. Thus, the play p' reaches a sink $s \notin \alpha_1$, and for every $i \in \text{relevant}(\pi')$, we have that $s \in \alpha_i$. Thus, p' is a trap that agrees with f_1 , contradicting the assumption that no such trap exists.

A.4 Resilient-CRS memory lower bound

We first show that $\pi = \langle f_1, \dots, f_k \rangle$ is a resilient-CRS. Consider a coalition $C \subseteq \{2, \dots, k\}$ and a deviation profile $\{f'_i\}_{i \in C}$. Let $\pi' = \pi[C \leftarrow \{f'_i\}_{i \in C}]$. Let $A \in \mathcal{S}_m$ be the set of players that Player 2 choose from q_0 in π' ; thus $f'_2(q_0) = \langle A, 1 \rangle$. If there is a player $i \in A$ that in π' moves the token from the vertex $\langle A, l \rangle$ with $A[l] = i$ to the sink $\perp$, then the outcome of π' reaches $\perp$. Thus, $\text{Win}(\pi') = [k] \setminus \{2\}$, and the deviation of the players in C is not strictly beneficial to a player in the coalition. If all players in A do not move the token to $\perp$ in their strategies in π' (in particular, this means that $A \subseteq C$), then let c_i be the vertex to which Player $A[m]$ moves the token from $\langle A, m \rangle$ in π' . Since all the sinks $s_{i,l}$ that are reachable from c_i are losing for Player i , the deviation is harmful for Player i , implying that $i \notin C$. Hence, Player j follows f_i , which directs her to move the token from c_i to $\langle A, i, p \rangle$. Since every successor of $\langle A, i, p \rangle$ is losing for some player in A , the deviation harms a player in C , and so π is resilient in which Player 1 wins.

We continue and prove that every resilient-CRS solution in $\mathcal{G}'_{k,m}$ requires each of the players $3, \dots, k$ to have memory $\binom{k-3}{m}$. Assume by contradiction that there exists a resilient profile $\pi = \langle f_1, \dots, f_k \rangle$ in which Player 1 wins and there exists $i \in \{3, \dots, k\}$ such that the strategy f_i has a memory structure $\mathcal{M}_i$ with fewer than $\binom{k-3}{m}$ states. Since there are fewer than $\binom{k-3}{m}$ states in $\mathcal{M}_i$, and there are $\binom{k-3}{m}$ subsets of $\{3, \dots, k\} \setminus \{i\}$ of size m , there exists a set $A \in \mathcal{S}_m$ such that $i \notin A$ and f_i never (that is, no matter what the history along which c_i has been reached) directs Player i to move the token from c_i to $\langle A, i, P \rangle$.

Consider the coalition $C = A \cup \{2\}$, and the deviation profile $f'_C = \{\{f'_j\}_{j \in C}\}$, where for every $j \in A$, the strategy f'_j agrees with f_j , except that from $\langle A, l \rangle$, with $A[l] = j$, the strategy f'_j directs Player j to move the token to $\langle A, l+1 \rangle$, in case $l < m$, and to c_i , in case $l = m$. Finally, the strategy f'_2 agrees with f_2 , except that from q_0 , the strategy f'_2 directs Player 2 to move the token to $\langle A, 1 \rangle$, and for every $A' \in \mathcal{S}_m$ such that $A' \neq A$, the strategy f'_2 directs Player 2 to move the token from $\langle A', i, P \rangle$ to $s_{i, \min A' \setminus A}$. Note that since both A and A' are of size m , the set $A' \setminus A$ is not empty, and thus $\min A' \setminus A$ exists and is in $\{3, \dots, k\}$. Let $\pi' = \pi[C \leftarrow f'_C]$. By the definition of the strategies in π' , there exists $A' \in \mathcal{S}_m$ such that $A' \neq A$ and $\text{outcome}(\pi') = q_0, \langle A, 1 \rangle, \langle A, 2 \rangle, \dots, \langle A, m \rangle, c_i, \langle A', i, P \rangle, (s_{i, \min A' \setminus A})^\omega$. Thus, $\text{Win}(\pi') = [k] \setminus \{1, i, \min A' \setminus A\}$. Since 1, i , and $\min A' \setminus A$ are all not in A , and thus also not in C , it follows that $C \subseteq \text{Win}(\pi')$.

Hence, the deviation to f'_C is strictly beneficial for Player 2, and does not harm any member of A , contradicting the assumption that π is resilient.

A.5 SSE-CRS memory lower bound

We first show that $\pi = \langle f_1, \dots, f_k \rangle$ is an SSE-CRS. Consider a coalition $C \subseteq \{2, \dots, k\}$ and a deviation profile $\{f'_i\}_{i \in C}$. Let $\pi' = \pi[C \leftarrow \{f'_i\}_{i \in C}]$. Let $A \in \mathcal{S}_m$ be the set of players that Player 2 choose from q_0 in π' ; thus $f'_2(q_0) = \langle A, 1 \rangle$. If there is a player $i \in A$ that in π' moves the token from the vertex $\langle A, l \rangle$ with $A[l] = i$ to the sink $\perp$, then the outcome of π' reaches $\perp$. Thus, $\text{Win}(\pi') = [k] \setminus \{2\}$, and the deviation of the players in C does not harm a player outside C . If all players in A do not move the token to $\perp$ in their strategies in π' (in particular, this means that $A \subseteq C$), then let c_i be the vertex to which Player $A[m]$ moves the token from $\langle A, m \rangle$ in π' . Since all the sinks $s_{i,l}$ that are reachable from c_i are losing for Player i , the deviation harms Player i , implying that $i \notin C$. Hence, Player j follows f_i , which directs her to move the token from c_i to $\langle A, i, P \rangle$. Since every successor of $\langle A, i, P \rangle$ is losing for some player in A , the deviation harms a player in C , and so π is SSE in which Player 1 wins.

We continue and prove that every SSE-CRS solution in $\mathcal{G}_{n,k}$ requires each of the players $3, \dots, k$ to have memory $\binom{k-3}{m}$. Assume by contradiction that there exists a SSE profile $\pi = \langle f_1, \dots, f_k \rangle$ in which Player 1 wins and there exists $i \in \{3, \dots, k\}$ such that the strategy f_i has a memory structure $\mathcal{M}_i$ with fewer than $\binom{k-3}{m}$ states. Since there are fewer than $\binom{k-3}{m}$ states in $\mathcal{M}_i$, and there are $\binom{k-3}{m}$ subsets of $\{3, \dots, k\} \setminus \{i\}$ of size m , there exists a set $A \in \mathcal{S}_m$ such that $i \notin A$ and f_i never (that is, no matter what the history along which c_i has been reached) directs Player i to move the token from c_i to $\langle A, i, P \rangle$.

Consider the coalition $C = A \cup \{2\}$, and the deviation profile $f'_C = \{\{f'_j\}_{j \in C}\}$, where for every $j \in A$, the strategy f'_j agrees with f_j , except that from $\langle A, l \rangle$, with $A[l] = j$, the strategy f'_j directs Player j to move the token to $\langle A, l+1 \rangle$, in case $l < m$, and to c_i , in case $l = m$. Finally, the strategy f'_2 agrees with f_2 , except that from q_0 , the strategy f'_2 directs Player 2 to move the token to $\langle A, 1 \rangle$, and for every $A' \in \mathcal{S}_m$ such that $A' \neq A$, the strategy f'_2 directs Player 2 to move the token from $\langle A', i, P \rangle$ to $s_{i, \min A' \setminus A}$. Note that since both A and A' are of size m , the set $A' \setminus A$ is not empty, and thus $\min A' \setminus A$ exists and is in

$\{3, \dots, k\}$. Let $\pi' = \pi[C \leftarrow f'_C]$. By the definition of the strategies in π' , there exists $A' \in \mathcal{S}_m$ such that $A' \neq A$ and $\text{outcome}(\pi') = q_0, \langle A, 1 \rangle, \langle A, 2 \rangle, \dots, \langle A, m \rangle, c_i, \langle A', i, P \rangle, (s_{i, \min A' \setminus A})^\omega$. Thus, $\text{Win}(\pi') = [k] \setminus \{1, i, \min A' \setminus A\}$. Since 1, i , and $\min A' \setminus A$ are all not in A , and thus also not in C , it follows that $C \subseteq \text{Win}(\pi')$.

Hence, the deviation to f'_C does not harm any member of C while harming Player 1, contradicting the assumption that π is SSE.