Admissible Strategies in Infinite Games over Graphs* **

Marco Faella

Università di Napoli "Federico II", Italy

Abstract. We consider games played on finite graphs, whose objective is to obtain a trace belonging to a given set of accepting traces. We focus on the states from which Player 1 cannot force a win. We compare several criteria for establishing what is the preferable behavior of Player 1 from those states, eventually settling on the notion of *admissible* strategy.
As the main result, we provide a characterization of the goals admitting positional admissible strategies. In addition, we derive a simple algorithm for computing such strategies for various common goals, and we prove the equivalence between the existence of positional winning strategies and the existence of positional subgame perfect strategies.

1 Introduction

Games played on finite graphs have been widely investigated in Computer Science, with applications including controller synthesis [PR89,ALW89,dAFMR05], protocol verification [KR01,BBF07], logic and automata theory [EJ91,Zie98], and compositional verification [dAH01].

These games consist of a finite graph, whose set of states is partitioned into Player-1 and Player-2 states, and a *goal*, which is a set of infinite sequences of states. The game consists in the two players taking turns at picking a successor state, eventually giving rise to an infinite path in the game graph. Player 1 wins the game if she manages to obtain an infinite path belonging to the goal, otherwise Player 2 wins. A (deterministic) *strategy* for a player is a function that, given the current history of the game (a finite sequence of states), chooses the next state. A state s is said to be *winning* if there exists a strategy that guarantees victory to Player 1 regardless of the moves of the adversary, if the game starts in s. A state that is not winning is called *losing*.

The main algorithmic concern of the classical theory of these games is determining the set of winning states. In this paper, we shift the focus to *losing* states, since we believe that many applications would benefit from a theory of best-effort strategies which allowed Player 1 to play in a rational way even from losing states.

* This work was supported by the MIUR PRIN Project 2007-9E5KM8.

** Extended from *Proc. of MFCS 2009: 34th International Symposium on Mathematical Foundations of Computer Science,* Lectures Notes in Computer Science. ©Springer-Verlag, 2009. The appendix does not appear in the published version.

For instance, many game models correspond to real-world problems which are not really competitive: the game is just a tool which enables to distinguish internal from external non-determinism. In practice, the behavior of the adversary may turn out to be random, or even cooperative. A strategy of Player 1 which does not "give up", but rather tries its best at winning, may in fact end up winning, even starting from states that are theoretically losing.

In other cases, the game is an over-approximation of reality, giving to Player 2 a wider set of capabilities (i.e., moves in the game) than what most adversaries actually have in practice. Again, a best-effort strategy for Player 1 can thus often lead to victory, even against an adversary which is strictly competitive.

In this paper, we compare several alternative definitions of best-effort strategies, eventually settling on the notion of *admissible* strategy. As a guideline for our investigation, we take the application domain of automated verification and synthesis of open systems. Such a domain is characterized by the fact that, once good strategies for a game have been found, they are intended to be actually implemented in hardware or software.

Best-effort strategies. The classical definition of what a "good" strategy is states that a strategy is *winning* if it guarantees victory whenever the game is started in a winning state [Tho95]. This definition does not put any burden on a strategy if the game starts from a losing state. In other words, if the game starts from a losing state, all strategies are considered equivalent.

A first refinement of the classical definition is a slight modification of the game-theoretic notion of *subgame-perfect equilibrium* [OR94]. Cast in our framework, this notion states that a strategy is good if it enforces victory whenever the game history is such that victory can be enforced. We call such strategies *strongly winning*, to avoid confusion with the use of subgame (and subarena) which is common in computer science [Zie98]. It is easy to see that this definition captures the intuitive idea that a good strategy should "enforce victory whenever it can" better than the classical one.

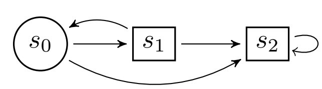

Fig. 1: A game where victory cannot be enforced.

Next, consider games where victory cannot be enforced at any point during the play. Take the Büchi game in Figure 1 [1], whose goal is to visit infinitely often s_0. No matter how many visits to s_0 Player 1 manages to make, he will never reach a point where he can enforce victory. Still, it is intuitively better for him to keep trying (i.e., move to s_1) rather than give up (i.e., move to s_2). To capture this intuition, we resort to the classical game-theoretic notion of *dominance* [OR94]. Given two strategies σ and σ' of Player 1, we say that σ *dominates* σ' if σ is always at least as good as σ', and better than σ' in at least one case. Dominance induces a strict partial order on strategies, whose maximal elements are called *admissible* strategies. In Section 3,

[1] Player-1 states are represented by circles and Player-2 states by squares.

we compare the above notions, and we prove that a strategy is admissible if and only if it is simultaneously strongly winning and cooperatively strongly winning (i.e., strongly winning with the help of Player 2).

To the best of our knowledge, the only paper dealing with admissibility in a context similar to ours is [Ber07], which provides existence results for general multi-player games of infinite duration, and does not address the issue of the memory requirements of the strategies.

Memory. A useful measure for the complexity of a strategy consists in evaluating how much memory it needs regarding the history of the game. In the simplest case, a strategy requires no memory at all: its decisions are based solely on the current state of the game. Such strategies are called *positional* or *memoryless* [GZ05]. In other cases, a strategy may require the amount of memory that can be provided by a finite automaton (*finite memory*), or more [DJW97].

The memory measure of a strategy is particularly important for the applications that we target in this paper. Since we are interested in actually implementing strategies in hardware or software, the simplest the strategy, the easiest and most efficient it is to implement.

We devote Section 4 to studying the memory requirements for various types of "good" strategies. In particular, we prove that all goals that have positional winning strategies also have positional *strongly* winning strategies. On the other hand, admissible strategies may require an unbounded amount of memory on some of those goals. We then provide necessary and sufficient conditions for a goal to have positional admissible strategies, building on the results of [GZ05].

We also prove that for *prefix-independent* goals, all positional winning strategies are automatically strongly winning. Additionally, prefix-independent goals admitting positional winning strategies also admit positional admissible strategies, as we show by presenting a simple algorithm which computes positional admissible strategies for these goals.

2 Definitions

We treat games that are played by two players on a finite graph, for an infinite number of turns. The aim of the first player is to obtain an infinite trace that belongs to a fixed set of accepting traces. In the literature, such games are termed *two-player*, *turn-based*, and *qualitative*.The following definitions make this framework formal.

A *game* is a tuple $G = (S_1, S_2, \delta, C, F)$ such that: S_1 and S_2 are disjoint finite sets of states; let $S = S_1 \cup S_2$, we have that $\delta \subseteq S \times S$ is the transition relation and $C : S \to \mathbb{N}$ is the coloring function, where $\mathbb{N}$ denotes the set of natural numbers including zero. Finally, $F \subseteq \mathbb{N}^\omega$ is the *goal*, where $\mathbb{N}^\omega$ denotes the set of infinite sequences of natural numbers. We denote by $\neg F$ the complement of F, i.e., $\mathbb{N}^\omega \setminus F$. We assume that games are non-blocking, i.e. each state has at least one successor in δ.

A (finite or infinite) path in G is a (finite or infinite) path in the directed graph (S, δ). With an abuse of notation, we extend the coloring function from states to paths, with the obvious meaning. If a finite path ρ is a prefix of a finite or infinite path ρ', we also say that ρ' *extends* ρ. We denote by $\mathit{first}(\rho)$ the first state of a path ρ and by $\mathit{last}(\rho)$ the last state of a finite path ρ.

Strategies. A *strategy* in G is a function $\sigma : S^* \to S$ such that for all $\rho \in S^*$, $(\mathit{last}(\rho), \sigma(\rho)) \in \delta$. Our strategies are deterministic, or, in game-theoretic terms, *pure*. We denote Stra_G the set of all strategies in G. We do not distinguish *a priori* between strategies of Player 1 and Player 2. However, for sake of clarity, we write σ for a strategy that should intuitively be interpreted as belonging to Player 1, and τ for the (rare) occasions when a strategy of Player 2 is needed.

Consider two strategies σ and τ, and a finite path ρ, and let $n = |\rho|$. We denote by $\mathrm{Outc}_G(\rho, \sigma, \tau)$ the unique infinite path $s_0 s_1 \ldots$ such that *(i)* $s_0 s_1 \ldots s_{n-1} = \rho$, and *(ii)* for all $i \geq n$, $s_i = \sigma(s_0 \ldots s_{i-1})$ if $s_{i-1} \in S_1$ and $s_i = \tau(s_0 \ldots s_{i-1})$ otherwise. We set $\mathrm{Outc}_G(\rho, \sigma) = \bigcup_{\tau \in \mathit{Stra}_G} \mathrm{Outc}_G(\rho, \sigma, \tau)$ and $\mathrm{Outc}_G(\rho) = \bigcup_{\sigma \in \mathit{Stra}_G} \mathrm{Outc}_G(\rho, \sigma)$. For all $s \in S$ and $\rho \in \mathrm{Outc}_G(s, \sigma)$, we say that ρ is *consistent* with σ. Similarly, we say that $\mathrm{Outc}_G(s, \sigma, \tau)$ is consistent with σ and τ. We extend the definition of consistent to finite paths in the obvious way.

A strategy σ is *positional* (or *memoryless*) if $\sigma(\rho)$ only depends on the last state of ρ. Formally, for all $\rho, \rho' \in S^*$, if $\mathit{last}(\rho) = \mathit{last}(\rho')$ then $\sigma(\rho) = \sigma(\rho')$.

Dominance. Given two strategies σ and τ, and a state s, we set $\mathit{val}_G(s, \sigma, \tau) = 1$ if $C(\mathrm{Outc}_G(s, \sigma, \tau)) \in F$, and $\mathit{val}_G(s, \sigma, \tau) = 0$ otherwise. Given two strategies σ and σ', we say that σ' *dominates* σ if: *(i)* for all $\tau \in \mathit{Stra}_G$ and all $s \in S$, $\mathit{val}_G(s, \sigma', \tau) \geq \mathit{val}_G(s, \sigma, \tau)$, and *(ii)* there exists $\tau \in \mathit{Stra}_G$ and $s \in S$ such that $\mathit{val}_G(s, \sigma', \tau) > \mathit{val}_G(s, \sigma, \tau)$.

It is easy to check that dominance is an irreflexive, asymmetric and transitive relation. Hence, it is a *strict partial order* on strategies.

Good strategies. In the following, unless stated otherwise, we consider a fixed game $G = (S_1, S_2, \delta, C, F)$ and we omit the G subscript.

For an infinite sequence $x \in \mathbb{N}^\omega$, we say that x is *accepting* if $x \in F$ and *rejecting* otherwise. We reserve the term "winning" to strategies and finite paths, as explained in the following. Let ρ be a finite path in G, we say that a strategy σ is *winning from* ρ if, for all $\rho' \in \mathrm{Outc}(\rho, \sigma)$, we have that $C(\rho')$ is accepting. We say that ρ is *winning* if there is a strategy σ which is winning from ρ. The above definition extends to states, by considering them as length-1 paths. A state that is not winning is called *losing*.

Further, a strategy σ is *cooperatively winning from* ρ if there exists a strategy τ such that $C(\mathrm{Outc}(\rho, \sigma, \tau))$ is accepting. We say that ρ is *cooperatively winning* if there is a strategy σ which is cooperatively winning from ρ. Intuitively, a path is cooperatively winning if the two players together can extend that path into an infinite path that satisfies the goal. Again, the above definitions extend to states, by considering them as length-1 paths.

We can now present the following set of winning criteria. Each of them is a possible definition of what a "good" strategy is.

– A strategy is *winning* if it is winning from all winning states. This criterion intuitively demands that strategies enforce victory whenever the initial state allows it.
– A strategy is *strongly winning* if it is winning from all winning paths that are consistent with it.
– A strategy is *subgame perfect* if it is winning from all winning paths. This criterion states that a strategy should enforce victory whenever the current history of the game allows it.
– A strategy is *cooperatively winning* (in short, *c-winning*) if it is cooperatively winning from all cooperatively winning states. This criterion essentially asks a strategy to be winning with the help of Player 2.
– A strategy is *cooperatively strongly winning* (in short, *cs-winning*) if it is cooperatively winning from all cooperatively winning paths that are consistent with it.
– A strategy is *cooperatively subgame perfect* (in short, *c-perfect*) if it is cooperatively winning from all cooperatively winning paths.
– A strategy is *admissible* if there is no strategy that dominates it. This criterion favors strategies that are maximal w.r.t. the partial order defined by dominance.

The notions of winning and cooperatively winning strategies are customary to computer scientists [Tho95,AHK97]. The notion of subgame perfect strategy comes from classical game theory [OR94]. The introduction of the notion of strongly winning strategy is motivated by the fact that in the target applications game histories that are inconsistent with the strategy of Player 1 cannot occur. Being strongly winning is strictly weaker than being subgame perfect. In particular, there are games for which there is a positional strongly winning strategy, but no positional subgame perfect strategy. The term "strongly winning" seems appropriate since this notion is a natural strengthening of the notion of winning strategy.

We say that a goal F is *positional* if, for all games G with goal F, there is a positional winning strategy in G.

3 Comparing Winning Criteria

In this section, we compare the winning criteria presented in Section 2. Figure 2 summarizes the relationships between the winning criteria under consideration. We start by stating the following basic properties.

Lemma 1. *The following properties hold:*

1. *all strongly winning strategies are winning, but not vice versa;*
2. *all subgame perfect strategies are strongly winning, but not vice versa;*
3. *all cs-winning strategies are c-winning, but not vice versa;*
4. *all c-perfect strategies are cs-winning, but not vice versa;*
5. *all games have a winning (respectively, strongly winning, subgame perfect, c-winning, cs-winning, c-perfect, admissible) strategy.*

Proof. The containments stated in (1) and (2) are obvious by definition. The fact that those containments are strict is proved by simple examples. Similarly for statements (3) and (4).

Regarding statement (5), the existence of a winning (respectively, strongly winning, subgame perfect, c-winning, cs-winning, c-perfect) strategy is obvious by definition. The existence of an admissible strategy can be derived from Theorem 11 from [Ber07]. □

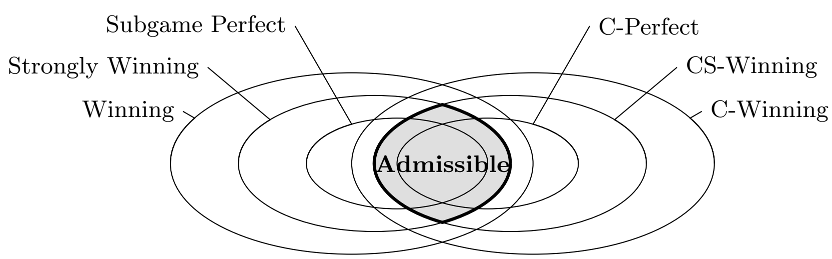

Fig. 2: Comparing winning criteria.

The following result provides a characterization of admissibility in terms of the simpler criteria of strongly winning and cooperatively strongly winning. Such characterization will be useful to derive further properties of admissible strategies. The result can be proved as a consequence of Lemma 9 from [Ber07].

Theorem 1. *A strategy is admissible if and only if it is strongly winning and cooperatively strongly winning.*

4 Memory

In this section, we study the amount of memory required by "good" strategies for achieving different kinds of goals. We are particularly interested in identifying those goals which admit positional good strategies, because positional strategies are the easiest to implement.

4.1 Positional Winning Strategies

In this section, we recall the main result of [GZ05], which provides necessary and sufficient conditions for a goal to be positional w.r.t. both players. Such characterization provides the basis for our characterization of the goals admitting positional admissible strategies, in Section 4.3.

We start with some additional notation. For a goal $F \subseteq \mathbb{N}^\omega$, we define its *preference relation* $\preceq_F$ and its strict version $\prec_F$ as follows: for two sequences $x, y \in \mathbb{N}^\omega$,

$$x \prec_F y \stackrel{\text{def}}{\iff} x \notin F \text{ and } y \in F \qquad x \preceq_F y \stackrel{\text{def}}{\iff} \text{if } x \in F \text{ then } y \in F.$$

Next, define the following relations between two languages $X, Y \subseteq \mathbb{N}^\omega$.

$$X \sqsubset_F^b Y \overset{\text{def}}{\iff} \exists y \in Y \,.\, \forall x \in X \,.\, x \prec_F y \qquad X \sqsubseteq_F^b Y \overset{\text{def}}{\iff} \forall x \in X \,.\, \exists y \in Y \,.\, x \preceq_F y$$
$$X \sqsubset_F^w Y \overset{\text{def}}{\iff} \exists x \in X \,.\, \forall y \in Y \,.\, x \prec_F y \quad X \sqsubseteq_F^w Y \overset{\text{def}}{\iff} \forall y \in Y \,.\, \exists x \in X \,.\, x \preceq_F y.$$

In the above definitions, the superscripts b and w stand for "best" and "worst", respectively. For instance, $X \sqsubset_F^b Y$ intuitively means that the *best* sequence in Y is strictly better than the best sequence in X. In other words, there is an accepting sequence in Y, while all sequences in X are rejecting. Similarly, $X \sqsubseteq_F^b Y$ means that the best sequence in Y is at least as good as the best sequence in X, i.e., if there is an accepting sequence in X, there is an accepting sequence in Y as well. We omit the subscript "F" when the goal is clear from the context.

A language $M \subseteq \mathbb{N}^*$ is *recognizable* if it is accepted by a finite automaton. We denote by Rec the set of all recognizable languages in $\mathbb{N}^*$. For a language $M \subseteq \mathbb{N}^*$, we denote by $[M]$ the language of all infinite words $x \in \mathbb{N}^\omega$ such that all prefixes of x are prefixes of some word in M.

A goal is *monotone* if for all recognizable sets $M, N \in \text{Rec}$,

$$\exists x \in \mathbb{N}^* \,.\, [xM] \sqsubset^b [xN] \implies \forall x \in \mathbb{N}^* \,.\, [xM] \sqsubseteq^b [xN].$$

A goal is *selective* iff, for all $x \in \mathbb{N}^*$ and all recognizable sets $M, N, K \in \text{Rec}$,

$$[x(M \cup N)^* K] \sqsubseteq^b [xM^*] \cup [xN^*] \cup [xK].$$

The following result is an adaptation to our setting of Theorem 2 from [GZ05].

Theorem 2 ([GZ05]). *Given a goal F, both players have a positional winning strategy for all games with goal F, if and only if both F and $\neg F$ are monotone and selective.*

4.2 Positional Strongly Winning and Subgame Pefect Strategies

For a game $G = (S_1, S_2, \delta, C, F)$ and a path $\rho = s_0 \dots s_n$ in G, define $detach(G, \rho)$ as the game obtained from G by adding a copy of the path ρ to it as a chain of new states ending in the original state s_n. Formally, $detach(G, \rho) = (S_1, S_2', \delta', C', F)$, where $S_2' = S_2 \cup \{s_0', s_1', \dots, s_{n-1}'\}$ and $s_0', s_1', \dots, s_{n-1}'$ are new distinct states not belonging to S_2 or to S_1. Then, $(s, t) \in \delta'$ iff either *(i)* $(s, t) \in \delta$, or *(ii)* $s = s_i'$ and $t = s_{i+1}'$, or *(iii)* $s = s_{n-1}'$ and $t = s_n$. Finally, the color labeling is defined by:

$$C'(s) = \begin{cases} C(s_i) & \text{if } s = s_i' \text{ for some } i \in \{0, \dots, n-1\}, \\ C(s) & \text{otherwise.} \end{cases}$$

The key idea of the detach operation consists in converting a path that may need the collaboration of Player 2 to occur, into a path which must occur if the game starts in a certain (new) state. This operation allows us to prove the following result.

Theorem 3. *For a goal F, the following are equivalent:*

1. *F is positional;*
2. *F admits positional strongly winning strategies;*
3. *F admits positional subgame perfect strategies.*

Proof. (Sketch) Since ($3 \implies 2$) and ($2 \implies 1$) are obvious by definition, it remains to prove that $1 \implies 3$. Hence, assume that the goal is positional. Let G be a game and let W be the set of winning paths of G. W may be infinite but it is certainly countable. Consider any ordering of W into $\rho_0, \rho_1, \ldots$. Consider the sequence of games $(G_i)_{i\geq 0}$ defined by $G_0 = G$ and $G_{i+1} = \mathit{detach}(G_i, \rho_i)$. Additionally, consider the sequence of strategies $(\sigma_i)_{i\geq 0}$ defined by: σ_0 is any positional winning strategy in G_0, and

$$\sigma_{i+1} = \begin{cases} \sigma_i & \text{if } \sigma_i \text{ is winning in } G_{i+1}, \\ \text{any positional winning strategy in } G_{i+1} & \text{otherwise.} \end{cases}$$

Due to space constraints, we omit the proof that the sequence $(\sigma_i)_{i\geq 0}$ converges to a subgame-perfect strategy σ^* within a finite number of steps. □

4.3 Positional Admissible Strategies

Since all admissible strategies are winning, admissible strategies require at least as much memory as winning strategies. The following example shows that there are positional goals for which all admissible strategies require an unbounded amount of memory.

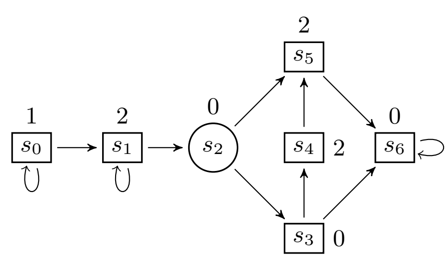

Fig. 3: Admissible strategies may require unbounded memory for a positional goal.

Example 1. Consider the goal informally described as follows: an infinite sequence is accepting if and only if either it contains infinitely many 2's, or it contains at least as many 2's as 1's. Such goal is positional, as it is monotone and selective.

Consider the game in Figure 3, with said goal. In the figure, the color of each state appears next to it or above it. The only choice for Player 1 occurs in s_2, where he can choose between s_3 and s_5. One can easily check that all states are losing. However, for all $n > 0$, the initial prefix $s_0^{n+1} s_1^n$ is winning and requires Player 1 to choose s_5 after s_2. On the other hand, the initial prefix $s_0^{n+2}s_1^n$ is cooperatively winning, and requires Player 1 to choose s_3 after s_2.

In conclusion, any admissible strategy must be able to distinguish $s_0^{n+1}s_1^n$ from $s_0^{n+2}s_1^n$, which requires an unbounded amount of memory. □

As shown in the following, in order to obtain positional admissible strategies, we need the goal to satisfy the following additional property.

Definition 1. *A goal is* strongly monotone *if for all recognizable sets $M, N \in Rec$,*

$$\exists x \in \mathbb{N}^* . [xM] \sqsubset^b [xN] \implies \forall x \in \mathbb{N}^* . [xM] \sqsubseteq^b [xN] \wedge [xM] \sqsubseteq^w [xN].$$

To gain some intuition, consider the goal of Example 1. We show that it is monotone, but not strongly so. Let x, M, and N be such that $[xM] \sqsubset^b [xN]$. This means that all sequences in $[xM]$ are rejecting, i.e., have a number of 2's smaller than the number of 1's. On the other hand, at least one sequence in $[xN]$ is accepting. Thus, there is a word $z \in [N]$ that has more excess 2's (possibly infinitely many) than any word in $[M]$.

Now, consider any $y \in \mathbb{N}^*$. Assume that $y \cdot y' \in [yM]$ is accepting. Then, $y \cdot z$ must also be accepting. This shows that $[yM] \sqsubseteq^b [yN]$ and the goal is monotone.

Assume instead that there is a rejecting sequence in $[yN]$. This does not imply that there is a rejecting sequence in $[yM]$, as would be required by the definition of strong monotonicity. A concrete counter-example is provided by Example 1. Let $x = 1 \cdot 1 \cdot 1 \cdot 2$, $M = 2 \cdot 0^*$, $N = 0^* + (2 \cdot 2 \cdot 0^*)$. Notice that $x = C(s_0 s_0 s_0 s_1)$, $[M] = C(\mathrm{Outc}(s_5))$, and $[N] = C(\mathrm{Outc}(s_3))$. Since $[xM] = 1{\cdot}1{\cdot}1{\cdot}2{\cdot}2{\cdot}0^\omega$ is rejecting, while $1{\cdot}1{\cdot}1{\cdot}2{\cdot}2{\cdot}2{\cdot}0^\omega \in [xN]$ is accepting, we have that $[xM] \sqsubset^b [xN]$. However, with $y = 1{\cdot}1{\cdot}2$, we have that all paths in $[yM]$ are accepting (there is only one) and there is one rejecting path in $[yN]$ (namely, $1{\cdot}1{\cdot}2{\cdot}0^\omega$). So, $[yM] \not\sqsubseteq^w [yN]$ and the goal is not strongly monotone.

The following is the main result of this section, providing a characterization of the goals admitting positional admissible strategies for both players.

Theorem 4. *Given a goal F, both players have a positional admissible strategy for all games with goal F if and only if both F and $\neg F$ are strongly monotone and selective.*

Due to space constraints, we provide the proof for one direction of Theorem 4, and a proof sketch for the other one.

Lemma 2. *Given a goal F, if Player 1 has a positional admissible strategy for all games with goal F, then F is strongly monotone and selective.*

Proof. By Lemma 5 of [GZ05], F is monotone and selective. It remains to prove that it is strongly monotone. Let $x \in \mathbb{N}^*$, $M, N \in \mathrm{Rec}$ such that

$$[xM] \sqsubset^b [xN]. \tag{1}$$

In other words, all paths in $[xM]$ are rejecting and at least one path in $[xN]$ is accepting. Let $y \in \mathbb{N}^*$, we will prove that $[yM] \sqsubseteq^b [yN]$ and $[yM] \sqsubseteq^w [yN]$.

Assume w.l.o.g. that M, N are not empty. Let $\mathcal{A}_x, \mathcal{A}_y, \mathcal{A}_M, \mathcal{A}_N$ be the finite automata recognizing the languages $\{x\}, \{y\}, M, N$, respectively. We can assume w.l.o.g. that these automata are deterministic, and in particular that they have a unique initial state. Let s_x, s_y, s_M, s_N be their respective initial states. Moreover, we can assume that $\mathcal{A}_x$ and $\mathcal{A}_y$ are linearly ordered chains of states, and that there are no edges in $\mathcal{A}_M$ (resp., $\mathcal{A}_N$) that go back to s_M (resp., s_N).

We build a game G as follows. We assign all states of $\mathcal{A}_x, \mathcal{A}_y, \mathcal{A}_M$ to Player 1, and all states of $\mathcal{A}_N$ to Player 2. We remove the two final states of $\mathcal{A}_x$ and $\mathcal{A}_y$ and the two initial states of $\mathcal{A}_M$ and $\mathcal{A}_N$. We connect the four automata as follows. Let t be a brand new state, we connect the penultimate state of $\mathcal{A}_x$ and the penultimate state of $\mathcal{A}_y$ to t. Then, we connect t to all the successors of s_M and s_N. We assign state t to Player 1. A technical difficulty is due to the fact that games are required to be non-blocking, while automata are not. This issue can be overcome using the notion of *essential state*, as proposed in the proof of Lemma 5 of [GZ05].

Let σ^* be a positional admissible strategy for Player 1 in G. By (1), if the game starts in s_x, once in t any cs-winning strategy must choose N. By Theorem 1, we have $\sigma^*(t) = s_N$. Assume that $[yM]$ contains an accepting sequence. Since all states in $\mathcal{A}_M$ have been assigned to Player 1, s_y is a winning state in G. Since σ^* is a winning strategy and $C(\text{Outc}(s_y, \sigma^*)) \subseteq [yN]$, there must be an accepting sequence in $[yN]$ too. Therefore, $[yM] \sqsubseteq^b [yN]$.

Finally, assume that $[yN]$ contains a rejecting sequence. Then, σ^* is not winning from s_y. Since σ^* is a winning strategy, s_y is not a winning state. If we assume that $[yM]$ contains no rejecting sequences, we obtain that $[yM]$, being non-empty, contains at least one accepting sequence. As before, this means that s_y is a winning state, which is a contradiction. Therefore, $[yM] \sqsubseteq^w [yN]$, which concludes the proof. □

Lemma 3. *Given a goal F, if both F and $\neg F$ are strongly monotone and selective, then both players have a positional admissible strategy for all games with goal F.*

Proof. (Sketch) Let $G = (S_1, S_2, \delta, C, F)$, with both F and $\neg F$ strongly monotone and selective. By Theorems 3 and 2, let σ_1 be a positional subgame-perfect strategy for Player 1 in G. Let SW be the set of all states s such that there is a finite path π in G such that: *(i)* $last(\pi) = s$, *(ii)* π is winning, and *(iii)* π can be extended into an infinite rejecting path (i.e., $C(\text{Outc}(\pi)) \setminus F \neq \emptyset$). Let G_1 be the game obtained from G by removing the edges which start in $SW \cap S_1$ and do not belong to σ_1. Let σ^* be a positional cs-winning strategy in G_1. It can be proved that σ^* is subgame-perfect and cs-winning. It follows by Theorem 1 that σ^* is admissible. □

5 Prefix-independent Goals

A goal F is *prefix-independent* iff for all $x \in \mathbb{N}^\omega$ and all $c \in \mathbb{N}$, $c \cdot x \in F$ if and only if $x \in F$. Examples of common prefix-independent goals include Büchi, co-Büchi and parity goals.

Theorem 5. *If a goal F is prefix-independent, then, for all games with goal F, all positional winning strategies are strongly winning, and all positional c-winning strategies are cs-winning.*

The positionality assumption is necessary in the above result. For a prefix-independent goal, it is easy to devise winning strategies that are not positional and not strongly winning. On the other hand, being prefix-independent is not necessary for ensuring that all positional winning strategies are strongly winning. For instance, safety and reachability goals are not prefix-independent, but they ensure said property.

Computing positional admissible strategies. Suppose that we are given a game G with a positional prefix-independent goal F, and that we have an algorithm for computing the set of winning states and a positional winning strategy for all games with goal F. Consider the following procedure, inspired by the proof of Lemma 3.

Procedure 1.

1. Compute the set of winning states *Win* and a positional winning strategy σ for G.
2. Remove from G the edges of Player 1 which start in *Win* and do not belong to σ.
3. In the resulting game, compute and return a positional cooperatively winning strategy.

The following theorem shows that the strategy returned by Procedure 1 is admissible. As far as the complexity of the procedure is concerned, assuming the usual graph-like adjacency-list representation for games, we obtain the same asymptotical complexity as finding a positional winning strategy for F. In particular, step 3 can easily be performed by attributing all states to Player 1 and then running the algorithm for a positional winning strategy.

Theorem 6. *Assume that F is a positional and prefix-independent goal, and that there is an algorithm for computing the set of winning states and a positional winning strategy for all games G with goal F in time $\mathcal{O}(f(|G|))$. Then, one can compute a positional admissible strategy for all games G with goal F using Procedure 1 in time $\mathcal{O}(f(|G|))$.*

This result allows us to easily compute admissible strategies for several common goals such as Büchi, co-Büchi, and parity. However, prefix-independence is not necessary for Procedure 1 to work. For instance, it is easy to prove that the procedure also returns an admissible strategy for reachability and safety goals.

6 Conclusions

We advanced the claim that computer science applications of game theory, especially in the domain of automatic verification and synthesis of controllers, may benefit from considering various winning criteria, such as admissibility, in addition to the classical one. Given the importance of (the lack of) memory for those applications, and considering that admissible strategies may require unboundedly more memory than plain winning strategies (Example 1), with Theorem 4 we characterize the goals that admit positional admissible strategies.

Further investigation and experimentation is needed to verify our claim in a concrete applicative setting. Moreover, it remains to determine how to compute admissible strategies for goals that are not prefix-independent.

References

[AHK97] R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. In *Proc. 38th IEEE Symp. Found. of Comp. Sci.*, pages 100–109. IEEE Computer Society Press, 1997.

[ALW89] M. Abadi, L. Lamport, and P. Wolper. Realizable and unrealizable concurrent program specifications. In *Proc. 16th Int. Colloq. Aut. Lang. Prog.*, volume 372 of *Lect. Notes in Comp. Sci.*, pages 1–17. Springer-Verlag, 1989.

[BBF07] S. Baselice, P.A. Bonatti, and M. Faella. On interoperable trust negotiation strategies. In *POLICY 07: 8th IEEE International Workshop on Policies for Distributed Systems and Networks*. IEEE Computer Society, 2007.

[Ber07] D. Berwanger. Admissibility in infinite games. In *STACS*, volume 4393 of *LNCS*, 2007.

[dAFMR05] L. de Alfaro, M. Faella, R. Majumdar, and V. Raman. Code aware resource management. In *EMSOFT 05: 5th Intl. ACM Conference on Embedded Software*, pages 191–202. ACM Press, 2005.

[dAH01] L. de Alfaro and T.A. Henzinger. Interface theories for component-based design. In *EMSOFT 01: 1st Intl. Workshop on Embedded Software*, volume 2211 of *Lect. Notes in Comp. Sci.*, pages 148–165. Springer-Verlag, 2001.

[DJW97] S. Dziembowski, M. Jurdziński, and I. Walukiewicz. How much memory is needed to win infinite games? In *LICS*, IEEE Computer Society, 1997.

[EJ91] E.A. Emerson and C.S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *Proc. 32nd IEEE Symp. Found. of Comp. Sci.*, pages 368–377. IEEE Computer Society Press, 1991.

[GZ05] H. Gimbert and W. Zielonka. Games where you can play optimally without any memory. In *CONCUR 05: Concurrency Theory. 16th Int. Conf.*, volume 3653 of *LNCS*. Springer, 2005.

[KR01] S. Kremer and J.-F. Raskin. A game-based verification of non-repudiation and fair exchange protocols. In *CONCUR 01: Concurrency Theory. 12th Int. Conf.*, volume 2154 of *LNCS*, pages 551–565. Springer, 2001.

[OR94] M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[PR89] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proceedings of the 16th Annual Symposium on Principles of Programming Languages*, pages 179–190. ACM Press, 1989.

[Tho95] W. Thomas. On the synthesis of strategies in infinite games. In *Proc. of 12th Annual Symp. on Theor. Asp. of Comp. Sci.*, volume 900 of *LNCS*, pages 1–13. Springer-Verlag, 1995.

[Zie98] W. Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science*, 200:135–183, June 1998.

A Appendix

A.1 Proof of Theorem 1

Although Theorem 1 can be seen as a consequence of Lemma 9 from [Ber07], we provide here a direct, independent proof. First, we state a trivial property of winning paths.

Lemma A-1. *If $s_0 s_1 \dots s_n$ is a winning path and $s_{n-1} \in S_1$, then $s_0 s_1 \dots s_{n-1}$ is a winning path.*

Theorem 1. *A strategy is admissible if and only if it is strongly winning and cooperatively strongly winning.*

Proof. For the "if" part, let σ^* be a strategy which is both strongly winning and cs-winning. Assume that there is a strategy σ that is better than σ^* in at least one case. In particular, let s be a state and τ be a strategy of Player 2 such that $val(s, \sigma^*, \tau) = 0$ and $val(s, \sigma, \tau) = 1$. We can build a Player-2 strategy τ' such that $val(s, \sigma^*, \tau') = 1$ and $val(s, \sigma, \tau') = 0$, thus proving that σ does not dominate σ^*.

Let αu be the longest prefix common to both $\mathrm{Outc}(s, \sigma, \tau)$ and $\mathrm{Outc}(s, \sigma^*, \tau)$. Precisely, let $\mathrm{Outc}(s, \sigma, \tau) = \alpha u v \dots$ and $\mathrm{Outc}(s, \sigma^*, \tau) = \alpha u v^* \dots$, where $\alpha \in S^*$, $u, v, v^* \in S$, and $v \neq v^*$. Clearly, it must be $u \in S_1$. We have that αu is not a winning path because σ^* is strongly winning and still $val(s, \sigma^*, \tau) = 0$. By Lemma A-1, $\alpha u v$ is not a winning path either. We also have that αu is c-winning, because $val(s, \sigma, \tau) = 1$. Since σ^* is cs-winning, $\alpha u v^*$ is also c-winning.

Then, we define the Player-2 strategy τ' as follows. We let τ' coincide with τ on all finite paths that are prefixes of αu. On all paths that are extensions of $\alpha u v$, we let τ' behave in such a way to ensure that $val(s, \sigma, \tau') = 0$. This is possible because $\alpha u v$ is not a winning path. On all paths that are extensions of $\alpha u v^*$, we let τ' behave in such a way to ensure that $val(s, \sigma^*, \tau') = 1$. This is possible because $\alpha u v^*$ is a c-winning path and σ^* is cs-winning. We conclude that σ does not dominate σ^*. Therefore, no strategy dominates σ^*.

Next, we prove the "only if" part. Let σ^* be an admissible strategy. By contradiction, assume that σ^* is not strongly winning. Therefore, there exist a winning path ρ and a strategy τ^* such that ρ is consistent with σ^* and τ^* and $val(\rho, \sigma^*, \tau^*) = 0$. Now, let σ be a strategy that is winning from ρ. Define another Player-1 strategy σ' as follows. For all $\pi \in S^*$, $\sigma'(\pi) = \sigma(\pi)$ if π extends ρ, and $\sigma'(\pi) = \sigma^*(\pi)$ otherwise. We show that σ' dominates σ^*, which is a contradiction. Take any Player-2 strategy τ. If σ' and τ together do not give rise to the finite path ρ, σ' behaves exactly like σ^*. If σ' and τ together do give rise to the path ρ, from that point on σ' behaves like σ, and therefore ensures victory. This proves that σ' always performs at least as well as σ^*. Finally, there is a case where σ' performs better than σ^*: we have $val(\mathit{first}(\rho), \sigma', \tau^*) = 1$ and $val(\mathit{first}(\rho), \sigma^*, \tau^*) = 0$, which concludes the proof of the contradiction.

Next, we show that σ^* is cs-winning. Let ρ be a c-winning path and assume by contradiction that, for all strategies τ, $val(\rho, \sigma^*, \tau) = 0$. Let $\sigma^\bullet, \tau^\bullet$ be a pair of strategies such that $val(\rho, \sigma^\bullet, \tau^\bullet) = 1$. Let σ be a strategy that behaves like σ^*, except that for all paths extending ρ it behaves like $\sigma^\bullet$. We prove that σ dominates σ^*. If the path ρ is not formed during the game, σ behaves exactly like σ^*. If the path ρ is formed, σ^* loses with certainty, while σ wins in at least one case, namely against $\tau^\bullet$. This proves that σ dominates σ^*, which is a contradiction. □

A.2 Proof of Theorem 3

We start with two auxiliary lemmas. First, notice that the detach operation adds some states to the game, but the new states are not reachable from any of the old states. As a consequence, if games G and $detach(G, \rho)$ start from an old state, they are indistinguishable to both players. In particular, the detach operation preserves the winning property of paths, as stated by the following lemma.

Notice that in the following we commit a slight abuse of language by identifying positional strategies in G and in $detach(G, \rho)$. This is justified by the fact that each new state introduced in $detach(G, \rho)$ only has one successor, thus giving no more choices to either player.

Lemma A-2. *For all strategies σ, and finite paths ρ, ρ' in G, σ is winning from ρ in G if and only if it is winning from ρ in $detach(G, \rho')$.*

The second auxiliary result states that, if ρ is a winning path in G, then the first state of the detached copy of ρ in $detach(G, \rho)$ is winning.

Lemma A-3. *Let $\rho = s_0 \dots s_n$ be a path in G and let $G' = detach(G, \rho)$. Let s'_0 be the new state added to G' in correspondence to s_0. If ρ is a winning path in G, then s'_0 is a winning state in G'.*

Proof. Let σ be a strategy that is winning from ρ in G, and let $\rho' = s'_0 \dots s'_{n-1} s_n$ be the detachment of ρ added to G'. Define a strategy σ' in G' as follows. For all finite paths π in G', if π extends ρ', i.e. $\pi = \rho'\pi'$, set $\sigma'(\pi) = \sigma(\rho\pi')$. For all other paths in G', σ' chooses an arbitrary move available to Player 1. It is easy to check that the set of infinite paths in G which are consistent with σ and extend ρ is color-equivalent to the set of infinite paths in G' which are consistent with σ' and start in s'_0. Therefore, σ' is winning from s'_0 in G'. □

We are now ready to prove the following.

Theorem 3. *For a goal F, the following are equivalent:*

1. *F is positional;*
2. *F admits positional strongly winning strategies;*
3. *F admits positional subgame perfect strategies.*

Proof. We prove that the sequence $(\sigma_i)_{i \geq 0}$, as defined in the proof sketch on page 7, converges to a strategy σ^* within a finite number of steps. Specifically, we show that once a strategy σ occurs in the sequence and is then replaced by another one, it will not occur at any later point in the sequence. This fact, together with the fact that the number of positional strategies is finite, leads to the convergence of the sequence. Assume by contradiction that there exist indices $0 \leq a < b$ such that $\sigma_a \neq \sigma_{a+1} \neq \sigma_b$ and $\sigma_a = \sigma_b$. Since $\sigma_a \neq \sigma_{a+1}$, σ_a is not winning in G_{a+1}. Therefore, by repeated application of Lemma A-2, σ_a (or equivalently, σ_b) cannot be winning in G_b, which is a contradiction.

Next, we prove that σ^* is subgame perfect in G. Let $\rho = s_0 \dots s_n$ be a winning path in G. There is an index $j > 0$ such that $G_j = \mathit{detach}(G_{j-1}, \rho)$. Let s'_0 be the "copy" of s_0 added to G_j by the detach operation. Since ρ is a winning path in G, by Lemma A-2 it is still winning in G_{j-1}, and by Lemma A-3 s'_0 is a winning state in G_j. Therefore, σ_j is winning from s'_0. Since σ^* is the ultimate value for the sequence of strategies, σ^* is also winning from s'_0. Take an infinite path ρ' in G, which extends ρ and that after ρ is consistent with σ^*. If we replace in ρ' the prefix ρ with $s'_0 \dots s'_{n-1} s_n$, we obtain an infinite path ρ'' in G_j, which is consistent with σ^* (because σ^* is positional) and color-equivalent to ρ'. Since σ^* is winning from s'_0, $C(\rho'') \in F$ and therefore $C(\rho') \in F$, which concludes the proof. □

A.3 Proof of Lemma 3

First, notice that the orderings between languages $X, Y \subseteq \mathbb{N}^\omega$ satisfy the following properties.

$$X \not\sqsubseteq^b Y \iff Y \sqsubset^b X \qquad\qquad X \not\sqsubseteq^w Y \iff Y \sqsubset^w X. \tag{2}$$

Before giving the full proof of Lemma 3, we prove the following auxiliary result.

Lemma A-4. *If F is strongly monotone, then*

$$\exists x \in \mathbb{N}^* \,.\, [xM] \sqsubset^w [xN] \implies \forall x \in \mathbb{N}^* \,.\, [xM] \sqsubseteq^b [xN] \wedge [xM] \sqsubseteq^w [xN].$$

Proof. We prove the counter-positive. Assume there is $x \in \mathbb{N}^*$ such that $[xM] \not\sqsubseteq^b [xN]$. By (2), $[xN] \sqsubset^b [xM]$. By strong monotonicity, for all $y \in \mathbb{N}^*$, $[yN] \sqsubseteq^w [yM]$. So, it cannot be that $[yM] \sqsubset^w [yN]$. Similarly if we assume that there is $x \in \mathbb{N}^*$ such that $[xM] \not\sqsubseteq^w [xN]$. □

Lemma 3. *Given a goal F, if both F and $\neg F$ are strongly monotone and selective, then both players have a positional admissible strategy for all games with goal F.*

Proof. We prove that the strategy σ^*, as defined in the proof sketch on page 10, is subgame-perfect and cs-winning.

Let π be a winning path in G, we prove that σ^* is winning from π. If $C(\mathrm{Outc}(\pi)) \subseteq F$ then all strategies, σ^* in particular, are winning from π, and we are done. Otherwise, let $s = \mathit{last}(\pi)$, we have $s \in SW$. If $s \in S_1$, $\sigma^*(s) = \sigma_1(s)$. Since σ_1 is subgame perfect, $\pi \cdot \sigma^*(s)$ is still a winning path. If $s \in S_2$, we also have that $\pi \cdot \sigma^*(s)$ is a winning path. In conclusion, using σ^* from π, either we eventually build a path ρ such that $C(\mathrm{Outc}(\rho)) \subseteq F$, or we remain forever in SW. In the first case, we have nothing to prove. In the second case, after π the strategy σ^* behaves always like σ_1. Since σ_1 is subgame perfect and π is winning, we end up in an accepting infinite path. This proves that σ^* is subgame perfect in G.

Next, let π be a c-winning path consistent with σ^*, we prove that σ^* is c-winning from π. Since π is consistent with σ^*, π is a path of G_1. Let π' be an infinite extension of π belonging to F. If π' contains no state in SW, π' is a path in G_1, and therefore π is c-winning in G_1. Since σ^* is cs-winning in G_1, σ^* is c-winning from π, as requested. Otherwise, let $\rho = \pi s_1 \ldots s_n$ be the prefix of π' such that s_n is the first occurrence of SW in π', after π. By the definition of SW applied to s_n, let ρ' be the finite path satisfying the above properties *(i)*, *(ii)*, and *(iii)*. Let $x = C(\rho')$, $y = C(\rho)$, $M = \{C(\pi) \mid \pi \text{ is a prefix of a path in } \mathrm{Outc}(s_n)\}$ and $N = \{C(\pi) \mid \pi \text{ is a prefix of a path in } \mathrm{Outc}(s_n, \sigma^*)\}$. Clearly, $M, N \in \mathrm{Rec}$. We have the following properties:

set	contains	reason
$[x \cdot M]$	both accepting and rejecting	definition of SW
$[x \cdot N]$	all accepting	σ^* is subgame perfect
$[y \cdot M]$	at least one accepting	π' is accepting

So, $[x \cdot M] \sqsubset^w [x \cdot N]$. By the strong monotonicity of F and by Lemma A-4, we have that $[yM] \sqsubseteq^b [yN]$, and therefore $[yN]$ contains an accepting sequence. By definition of N, such sequence corresponds to a path in G_1. Therefore, π is c-winning in G_1. □

A.4 Proof of Theorem 5

In order to prove Theorem 5, we first introduce two auxiliary results. A goal F is *shrinkable* iff for all $c\rho \in F$, with $c \in \mathbb{N}$ and $\rho \in \mathbb{N}^\omega$, $\rho \in F$. A goal F is *extensible* iff for all $\rho \in F$ and all $c \in \mathbb{N}$, $c\rho \in F$. It is immediate that a goal is prefix-independent iff it is both shrinkable and extensible.

Lemma A-5. *If a goal F is shrinkable, then, for all games with goal F, all winning paths end in a winning state, and all c-winning paths end in a c-winning state.*

Proof. We prove the statement for winning paths, as the one regarding c-winning paths can be proved along similar lines. Let G be a game with a shrinkable goal F, let $\rho = s_0 \ldots s_n$ be a winning path and let σ be a strategy which is winning from ρ. Consider all infinite paths that extend ρ and are consistent with σ. These paths all satisfy the goal F. If we remove the prefix ρ from these paths, they still all satisfy F, since F is shrinkable. Consider the strategy σ' defined by: for all $\pi \in S^*$,

$$\sigma'(\pi) = \begin{cases} \sigma(s_0 \ldots s_{n-1}\pi) & \text{if } \mathit{first}(\pi) = s_n, \\ \text{arbitrarily defined} & \text{otherwise.} \end{cases}$$

It is immediate that σ' is winning from s_n and therefore s_n is a winning state. □

The following corollary states that if a goal is shrinkable, winning strategies confine the game in the winning region.

Corollary A-1. *If a goal is shrinkable, for all winning strategies σ, and for all finite paths ρ consistent with σ, if first(ρ) is winning then last(ρ) is winning.*

Theorem 5. *If a goal F is prefix-independent, then, for all games with goal F, all positional winning strategies are strongly winning, and all positional c-winning strategies are cs-winning.*

Proof. We prove the statement for positional winning strategies, as the one regarding positional c-winning strategies can be proved along similar lines. Let G be a game with goal F, and let σ be a winning strategy for G. Let $\rho = s_0 \dots s_n$ be a winning path which is consistent with σ. By Lemma A-5, s_n is a winning state. Let $\rho' = s_0 \dots s_n s_{n+1} \dots$ be an infinite path which extends ρ and is consistent with σ. Since σ is positional and winning from s_n, $C(s_n s_{n+1} \dots) \in F$. Since F is extensible, $C(\rho') \in F$. Therefore, σ is winning from ρ. Being ρ generic, we conclude that σ is strongly winning. □

Simple examples show that neither shrinkability nor extensibility alone can replace prefix-independence in the assumptions of the above result.

A.5 Proof of Theorem 6

Theorem 6. *Assume that F is a positional and prefix-independent goal, and that there is an algorithm for computing the set of winning states and a positional winning strategy for all games G with goal F in time $\mathcal{O}(f(|G|))$. Then, one can compute a positional admissible strategy for all games G with goal F using Procedure 1 in time $\mathcal{O}(f(|G|))$.*

Proof. Let G be a game with goal F, where F is positional and prefix-independent. Let *Win* be the set of winning states of G, and consider the application of the above procedure to G. Let σ_1 be the strategy computed at step 1 of the procedure, and σ_3 the output strategy. Clearly, σ_3 is positional. We prove that it is also admissible. According to Theorems 5 and 1, it is sufficient to prove that σ_3 is winning and cooperatively winning.

By Corollary A-1, whenever the game starts in a winning state, σ_1 confines the game in *Win*. Since σ_3 coincides with σ_1 on *Win*, σ_3 is winning.

Next, let G_2 be the game built at Step 2 of the procedure. We know that σ_3 is cooperatively winning in G_2. We now prove that σ_3 is cooperatively winning in G. Let s be a cooperatively winning state in G. We prove that s is cooperatively winning in G_2 as well. Let σ, τ be a pair of strategies in G such that $C(\mathrm{Outc}(s, \sigma, \tau)) \in F$. We consider the following two cases: *(1)* the infinite path $\mathrm{Outc}(s, \sigma, \tau)$ does not visit *Win*; then, strategies σ and τ are valid in G_2 and thus s is cooperatively winning in G_2; *(2)* let ρ be the shortest prefix of $\mathrm{Outc}(s, \sigma, \tau)$ which ends in *Win*; then, consider a new strategy σ' which behaves as follows, for all $\pi \in S^*$:

$$\sigma'(\pi) = \begin{cases} \sigma_1(\pi) & \text{if } \pi \text{ extends } \rho, \\ \sigma(\pi) & \text{otherwise.} \end{cases}$$

Consider strategies σ' and τ playing together in G_2. At first, σ' behaves like σ and so the path ρ is built. Then, σ' behaves like the winning strategy σ_1. Eventually, the infinite path $\rho{\cdot}\rho'$ is built, where $C(last(\rho){\cdot}\rho') \in F$. Since F is extensible, $C(\rho{\cdot}\rho') \in F$ and thus s is cooperatively winning in G_2.

The complexity of Procedure 1 has already been discussed in Section 5. □

A.6 Additional Examples

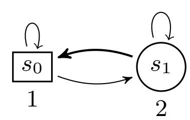

Fig. 4: A game having a winning strategy which is not strongly winning, and a strongly winning strategy which is not subgame perfect. The goal is $\varphi = 1 \wedge \Diamond\Box 2$. The positional strategy consisting of going from s_1 to s_0 (the thick edge) is (trivially) winning but not strongly winning. The strategy that chooses s_1 when the current history contains exactly one occurrence of s_0, and chooses s_0 otherwise is strongly winning but not subgame perfect, due for instance to the winning path $s_0 s_1 s_0 s_1$.

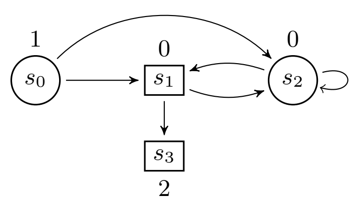

Fig. 5: A game having a c-winning strategy which is not cs-winning, and a cs-winning strategy which is not c-perfect. The goal is $\varphi = 1 \wedge \Diamond\Box 2$. The only c-winning state is s_0. The positional strategy $s_0 \to s_1$, $s_2 \to s_2$ is c-winning but not cs-winning. Consider the strategy that chooses $s_0 \to s_1$ and then *(i)* $s_2 \to s_1$ if s_1 was visited in the current history and *(ii)* $s_2 \to s_2$ if s_1 was not visited in the current history. Such strategy is cs-winning but not c-perfect, due to the c-winning path $s_0 s_2$.

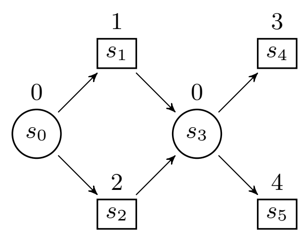

Fig. 6: A game having a positional strongly winning strategy, but no positional subgame perfect strategy. The goal is $\varphi = 0 \wedge \big((\Diamond 1 \wedge \Diamond 3) \vee (\Diamond 2 \wedge \Diamond 4)\big)$. The positional strategy $s_0 \to s_1$, $s_3 \to s_4$ is strongly winning. All subgame perfect strategies need memory in state s_3.