



Subgame Optimal Strategies in Finite Concurrent Games with Prefix-Independent Objectives

Benjamin Bordais^(✉), Patricia Bouyer and Stéphane Le Roux

Université Paris-Saclay, CNRS, ENS Paris-Saclay, LMF, 91190 Gif-sur-Yvette, France
bordais@lsv.fr

Abstract. We investigate concurrent two-player win/lose stochastic games on finite graphs with prefix-independent objectives. We characterize subgame optimal strategies and use this characterization to show various memory transfer results: 1) For a given (prefix-independent) objective, if every game that has a subgame *almost-surely winning* strategy also has a positional one, then every game that has a subgame *optimal* strategy also has a positional one; 2) Assume that the (prefix-independent) objective has a neutral color. If every *turn-based* game that has a subgame almost-surely winning strategy also has a positional one, then every game that has a *finite-choice* (notion to be defined) subgame optimal strategy also has a positional one.

We collect or design examples to show that our results are tight in several ways. We also apply our results to Büchi, co-Büchi, parity, mean-payoff objectives, thus yielding simpler statements.

1 Introduction

Turn-based two-player win/lose (stochastic) games on finite graphs have been intensively studied in the context of model checking in a broad sense [19,1]. These games behave well regarding optimality in various settings. Most importantly for this paper, [14] proved the following results for finite turn-based stochastic games with prefix-independent objectives: (1) every game has deterministic optimal strategies; (2) from every value-1 state, there is an optimal, i.e. almost-surely winning, strategy; (3) if from every value-1 state of every game there is an optimal strategy using some fixed amount of memory, every game has an optimal strategy using this amount of memory. These results are of either of the following generic forms:

- In all games, (from all nice states) there is a nice strategy.
- If from all *nice states* of all games there is a nice strategy, so it is from all *states*.

The concurrent version of these turn-based (stochastic) games has a higher modeling power than the turn-based version: this is really useful in practice since real-world systems are intrinsically concurrent [15]. They are played on a finite graph as follows: at each player state, the two players stochastically and independently choose one among finitely many actions. This yields a Nature state,

which stochastically draws a next player state, from where each player chooses one action again, and so on. Each player state is labelled by a color, and who wins depends on the infinite sequence of colors underlying the (stochastically) generated infinite sequence of player states. Unfortunately, these concurrent games do not behave well in general even for simple winning conditions and simple graph structures, like finite graphs:

- Reachability objectives: there is a game without optimal strategies [13];
- Büchi objectives: there is a game with value 1 while all finite-memory strategies have value 0 [12];
- Co-Büchi objectives: although there are always positional ε -optimal strategies [8], there is a game with optimal strategies but without finite-memory optimal strategies [4];
- Parity [12] and mean-payoff [10] objectives: there is a game with subgame almost-surely-winning strategies, but where all finite-memory strategies have value 0.

In this paper, we focus on concurrent stochastic finite games. Therefore, the generic forms of our results will be more complex, in order to take into account the above-mentioned discrepancies. They will somehow be given as generic statements as follows:

- Every game that has a *nice* strategy also has a *nicer* one.
- If all *special games* that have a nice strategy have a nicer one, so it is for all *games*.

Much of the difficulty consists in fine-tuning the strength of “nice”, “nicer” and “special” above. We present below our main contributions on finite two-player win/lose concurrent stochastic games with prefix-independent objectives:

1. We provide a characterization of subgame optimal strategies, which are strategies that are optimal after every history (Theorem 1): a Player A strategy is subgame optimal iff 1) it is locally optimal and 2) for every Player B deterministic strategy, after every history, if the visited states have the same positive value, Player A wins with probability 1. This characterization is used to prove all the results below.
2. We prove memory transfer results from subgame almost-surely winning strategies to subgame optimal strategies:
 - (a) Theorem 2: If every game that has a subgame *almost-surely winning* strategy also has a positional one, then every game that has a subgame *optimal* strategy also has a positional one.
 - (b) Corollary 1: every Büchi or co-Büchi game that has a subgame optimal strategy has a positional one. (Whereas parity games may require infinite memory [12].)

Note that the transfer result 2a can be generalized from positional to finite memory.

3. We say that a strategy has finite-choice, if it uses only finitely many action distributions. Note that finite-memory (resp. deterministic) strategies clearly have finite choice.

- (a) Theorem 4: In a given game, if there is a finite-choice optimal strategy, there is a finite-choice *subgame* optimal strategy.
- (b) Theorem 5: Assume that the objective has a neutral color. If every *turn-based* game that has a subgame almost-surely winning strategy also has a positional one, then every game that has a *finite-choice* subgame optimal strategy also has a positional one.
- (c) Corollary 2: every parity or mean-payoff game that has a finite-memory subgame optimal strategy also has a positional one.

Note that **3a** and **3b** are false if the word finite-choice is removed [4]. The proof of **3b** invokes **3a**. Flavor (and proofs) of **3b** and **2a** are similar, but both premises and conclusions are weakened in **3b**, as emphasized.

Related works. A large part of this paper is dedicated to the extension to concurrent games of the results from [14] regarding the transfer of memory from almost-surely winning strategies to optimal strategies in turn-based games. Note that the proof technique used in [14] is different and could not be adapted to our more general setting. In their proof, both players agree on a preference over Nature states and play according to this preference. In our proof, we slice the graph into value areas (that is, sets of states with the same value), and show that it is sufficient to play an almost-sure winning strategy in each slice; we then glue these (partial) strategies together to get a subgame-optimal strategy over the whole graph.

The slicing technique was already used in the context of concurrent games in [8]. The authors focus on parity objectives and establishes a memory transfer result from limit-sure winning strategies to almost-optimal strategies. As an application, they show that, for co-Büchi objectives, since positional strategies are sufficient to win limit-surely, they also are to win almost-optimally. Their construction made heavy use of the specific nature of parity objectives.

We also mention [6], where the focus is also on concurrent games with prefix-independent objectives. In particular, the authors establish a (very useful) result: if all states have positive values, then they all have value 1. (Note that a strengthening of this result is presented in this paper (Theorem 3), which also appears as an adaptation of a result proved in [14]). This result is then used in another context with non-zero-sum games.

Finally, some recent works on concurrent games have been done in [2,3,4], where the goal is the following: local interactions of the two players in the player state are given by bi-dimensional tables; those tables can be abstracted as *game forms*, where (output) variables are issues of the local interaction (possibly several issues are labelled by the same variable). The goal of this series of works is to give (intrinsic) properties of these game forms, so that, when used in a graph game, the existence of optimal strategies is ensured. For instance, in [3], a property of games forms, called RM, is given, which ensures that, if one only uses RM game forms in a graph, then for every reachability objective, Player A will always have an optimal strategy for that objective. This property is a characterization of well-behaved game forms regarding reachability objectives

since every game form which is not RM can be embedded into a (small) graph game in such a way that Player A does not have an optimal strategy. This line of works really differs from the target of the current paper.

Structure of the paper. Section 2 presents notations, Section 3 recalls the notion of game forms, Section 4 introduces our formalism, Section 5 exhibits a necessary and sufficient pair of conditions for subgame optimality, Section 6 shows a memory transfer from subgame almost-surely winning to subgame optimal in concurrent games, and Section 7 adapts the results of the previous section to the case of the existence of a subgame finite-choice strategy.

Detailed proofs and additional formal definitions are available in [5].

2 Preliminaries

Consider a non-empty set Q . We denote by Q^* , Q^+ and Q^ω the set of finite sequences, non-empty finite sequences and infinite sequences of elements of Q respectively. For $n \in \mathbb{N}$, we denote by Q^n (resp. $Q^{\leq n}$) the set of sequences of (resp. at most) n elements of Q . For all $\rho = q_1 \cdots q_n \in Q^n$ and $i \leq n$, we denote by ρ_i the element $q_i \in Q$ and by $\rho_{\leq i} \in Q^i$ the finite sequence $q_1 \cdots q_i$. For a subset $S \subseteq Q$, we denote by $Q^* \cdot S^\omega \subseteq Q^\omega$ the set of infinite paths that eventually settle in S and by $(Q^* \cdot S)^\omega \subseteq Q^\omega$ the set of infinite paths visiting infinitely often the set S .

A *discrete probabilistic distribution* over a non-empty finite set Q is a function $\mu : Q \rightarrow [0, 1]$ such that $\sum_{x \in Q} \mu(x) = 1$. The *support* $\text{Supp}(\mu)$ of a probabilistic distribution $\mu : Q \rightarrow [0, 1]$ is the set of non-zeros of the distribution: $\text{Supp}(\mu) = \{q \in Q \mid \mu(q) \in (0, 1]\}$. The set of all distributions over Q is denoted $\mathcal{D}(Q)$.

3 Game forms

We recall the definition of game forms – informally, bi-dimensional tables with variables – and of games in normal forms – game forms whose outcomes are values between 0 and 1.

Definition 1 (Game form and game in normal form). A game form (*GF for short*) is a tuple $\mathcal{F} = \langle \text{Act}_A, \text{Act}_B, \mathcal{O}, \varrho \rangle$ where Act_A (resp. Act_B) is the non-empty finite set of actions available to Player A (resp. B), \mathcal{O} is a non-empty set of outcomes, and $\varrho : \text{Act}_A \times \text{Act}_B \rightarrow \mathcal{O}$ is a function that associates an outcome to each pair of actions. When the set of outcomes \mathcal{O} is equal to $[0, 1]$, we say that \mathcal{F} is a game in normal form. For a valuation $v \in [0, 1]^\mathcal{O}$ of the outcomes, the notation $\langle \mathcal{F}, v \rangle$ refers to the game in normal form $\langle \text{Act}_A, \text{Act}_B, [0, 1], v \circ \varrho \rangle$.

We use game forms to represent interactions between two players. The strategies available to Player A (resp. B) are convex combinations of actions given as the rows (resp. columns) of the table. In a game in normal form, Player A tries to maximize the outcome, whereas Player B tries to minimize it.

Definition 2 (Outcome of a game in normal form). Consider a game in normal form $\mathcal{F} = \langle \text{Act}_A, \text{Act}_B, [0, 1], \varrho \rangle$. The set $\mathcal{D}(\text{Act}_A)$ (resp. $\mathcal{D}(\text{Act}_B)$) is the set of strategies available to Player A (resp. B). For a pair of strategies $(\sigma_A, \sigma_B) \in \mathcal{D}(\text{Act}_A) \times \mathcal{D}(\text{Act}_B)$, the outcome $\text{out}_{\mathcal{F}}(\sigma_A, \sigma_B)$ in \mathcal{F} of the strategies (σ_A, σ_B) is defined as: $\text{out}_{\mathcal{F}}(\sigma_A, \sigma_B) := \sum_{a \in \text{Act}_A} \sum_{b \in \text{Act}_B} \sigma_A(a) \cdot \sigma_B(b) \cdot \varrho(a, b) \in [0, 1]$.

Definition 3 (Value of a game in normal form and optimal strategies). Consider a game in normal form $\mathcal{F} = \langle \text{Act}_A, \text{Act}_B, [0, 1], \varrho \rangle$ and a strategy $\sigma_A \in \mathcal{D}(\text{Act}_A)$ for Player A. The value of the strategy σ_A , denoted $\text{val}_{\mathcal{F}}(\sigma_A)$ is equal to: $\text{val}_{\mathcal{F}}(\sigma_A) := \inf_{\sigma_B \in \mathcal{D}(\text{Act}_B)} \text{out}_{\mathcal{F}}(\sigma_A, \sigma_B)$, and analogously for Player B, with a sup instead of an inf. When $\sup_{\sigma_A \in \mathcal{D}(\text{Act}_A)} \text{val}_{\mathcal{F}}(\sigma_A) = \inf_{\sigma_B \in \mathcal{D}(\text{Act}_B)} \text{val}_{\mathcal{F}}(\sigma_B)$, it defines the value of the game \mathcal{F} , denoted $\text{val}_{\mathcal{F}}$.

A strategy $\sigma_A \in \mathcal{D}(\text{Act}_A)$ ensuring $\text{val}_{\mathcal{F}} = \text{val}_{\mathcal{F}}(\sigma_A)$ is called optimal. The set of all optimal strategies for Player A is denoted $\text{Opt}_A(\mathcal{F}) \subseteq \mathcal{D}(\text{Act}_A)$, and analogously for Player B. Von Neuman's minimax theorem [20] ensures the existence of optimal strategies (for both players).

In the following, strategies in games in normal forms will be called GF-strategies, in order not to confuse them with strategies in concurrent (graph) games.

4 Concurrent games and optimal strategies

4.1 Concurrent arenas and strategies

We introduce the definition of concurrent arenas played on a finite graph.

Definition 4 (Finite stochastic concurrent arena). A colored concurrent arena \mathcal{C} is a tuple $\langle Q, (A_q)_{q \in Q}, (B_q)_{q \in Q}, D, \delta, \text{dist}, K, \text{col} \rangle$ where Q is the non-empty finite set of states, for all $q \in Q$, A_q (resp. B_q) is the non-empty finite set of actions available to Player A (resp. B) at state q , D is the finite set of Nature states, $\delta : \bigcup_{q \in Q} (\{q\} \times A_q \times B_q) \rightarrow D$ is the transition function, $\text{dist} : D \rightarrow \mathcal{D}(Q)$ is the distribution function. Furthermore, K is the non-empty finite set of colors and $\text{col} : Q \rightarrow K$ is the coloring function.

In the following, the arena \mathcal{C} will refer to the tuple $\langle Q, (A_q)_{q \in Q}, (B_q)_{q \in Q}, D, \delta, \text{dist}, K, \text{col} \rangle$, unless otherwise stated. A concurrent game is obtained from a concurrent arena by adding a winning condition: the set of infinite paths winning for Player A (and losing for Player B).

Definition 5 (Finite stochastic concurrent game). A finite concurrent game is a pair $\langle \mathcal{C}, W \rangle$ where \mathcal{C} is a finite concurrent colored arena and $W \subseteq K^\omega$ is Borel. The set W is called the objective, as it corresponds to the set of colored paths winning for Player A.

In this paper, we only consider a specific kind of objectives: prefix-independent ones. Informally, they correspond to objectives W such that an infinite path ρ is in W if and only if any of its suffixes is in W . More formally:

Definition 6 (Prefix-independent objectives). For a non-empty finite set of colors K and $W \subseteq K^\omega$, W is said to be prefix-independent (PI for short) if, for all $\rho \in K^\omega$ and $i \geq 0$, $\rho \in W \Leftrightarrow \rho_{\geq i} \in W$.

In the following, we refer to concurrent games with prefix-independent objectives as PI concurrent games.

Definition 7 (Parity, Büchi, co-Büchi objectives). Let $K \subset \mathbb{N}$ be a finite non-empty set of integers. Consider a concurrent arena \mathcal{C} with K as set of colors. For an infinite path $\rho \in Q^\omega$, we denote by $\text{col}(\rho)_\infty \subseteq \mathbb{N}$ the set of colors seen infinitely often in ρ : $\text{col}(\rho)_\infty := \{n \in \mathbb{N} \mid \forall i \in \mathbb{N}, \exists j \geq i, \text{col}(\rho_j) = n\}$. Then, the parity objective w.r.t. col is the set $W^{\text{Parity}}(\text{col}) := \{\rho \in Q^\omega \mid \max \text{col}(\rho)_\infty \text{ is even}\}$. The Büchi (resp. co-Büchi) objective correspond to the parity objective with $K := \{1, 2\}$ (resp. $K := \{0, 1\}$).

Strategies are then defined as functions that, given the history of the game (i.e. the sequence of states already seen) associate a distribution on the actions available to the Player.

Definition 8 (Strategies). Consider a concurrent game \mathcal{C} . A strategy for Player A is a function $s_A : Q^+ \rightarrow \mathcal{D}(A)$ with $A := \bigcup_{q \in Q} A_q$ such that, for all $\rho = q_0 \cdots q_n \in Q^+$, we have $s_A(\rho) \in \mathcal{D}(A_{q_n})$. We denote by $S_{\mathcal{C}}^A$ the set of all strategies in arena \mathcal{C} for Player A. This is analogous for Player B.

Given two strategies s_A, s_B for both players in an arena \mathcal{C} from a starting state q_0 , we define in the usual manner the probability $\mathbb{P}_{s_A, s_B}^{\mathcal{C}, q_0}$ of a finite path which induces the probability of an arbitrary Borel subset of infinite paths. Values of strategies and of the game are defined below.

Definition 9 (Value of strategies and of the game). Let $\mathcal{G} = \langle \mathcal{C}, W \rangle$ be a PI concurrent game and consider a strategy $s_A \in S_{\mathcal{C}}^A$ for Player A. The function $\chi_{\mathcal{G}}[s_A] : Q \rightarrow [0, 1]$ giving the value of the strategy s_A is such that, for all $q_0 \in Q$, we have $\chi_{\mathcal{G}}[s_A](q_0) := \inf_{s_B \in S_{\mathcal{C}}^B} \mathbb{P}_{s_A, s_B}^{\mathcal{C}, q_0}[W]$. The function $\chi_{\mathcal{G}}[A] : Q \rightarrow [0, 1]$ giving the value for Player A: is such that, for all $q_0 \in Q$, we have $\chi_{\mathcal{G}}[A](q_0) := \sup_{s_A \in S_{\mathcal{C}}^A} \chi_{\mathcal{G}}[s_A](q_0)$. The function $\chi_{\mathcal{G}}[B] : Q \rightarrow [0, 1]$ giving the value of the game for Player B is defined similarly by reversing the supremum and infimum.

By Martin’s result on the determinacy of Blackwell games [17], for all concurrent games $\mathcal{G} = \langle \mathcal{C}, W \rangle$, the value functions for both Players are equal, this defines the value function $\chi_{\mathcal{G}} : Q \rightarrow [0, 1]$ of the game: $\chi_{\mathcal{G}} := \chi_{\mathcal{G}}[A] = \chi_{\mathcal{G}}[B]$.

We define value areas: subsets of states whose values are the same.

Definition 10 (Value area). In a PI concurrent game \mathcal{G} , $V_{\mathcal{G}}$ refers to the set of values appearing in the game: $V_{\mathcal{G}} := \{\chi_{\mathcal{G}}[q] \mid q \in Q\}$. Furthermore, for all $u \in V_{\mathcal{G}}$, $Q_u \subseteq Q$ refers to the set of states whose values are u w.r.t. $\chi_{\mathcal{G}}$: $Q_u := \{q \in Q \mid \chi_{\mathcal{G}}(q) = u\}$.

In concurrent games, game forms appear at each state and describe the interactions of the players at that state. Furthermore, the valuation mapping each

state to its value in the game can be lifted, via a convex combination, into a valuation of the Nature states. This, in turn, induces a natural way to define the game in normal form appearing at each state.

Definition 11 (Local interactions, Lifting valuations). *In a PI concurrent game \mathcal{G} where the valuation $\chi_{\mathcal{G}} : Q \rightarrow [0, 1]$ gives the values of the game, the lift $\nu_{\mathcal{G}} : \mathcal{D} \rightarrow [0, 1]$ is such that, for all $d \in \mathcal{D}$, we have $\nu_{\mathcal{G}}(d) := \sum_{q \in Q} \chi_{\mathcal{G}}(q) \cdot \text{dist}(d)(q)$ (recall that $\text{dist} : \mathcal{D} \rightarrow \mathcal{D}(Q)$ is the distribution function).*

Let $q \in Q$. The local interaction at state q is the game form $\mathcal{F}_q = \langle A_q, B_q, \mathcal{D}, \delta(q, \cdot, \cdot) \rangle$. The game in normal form at state q is then $\mathcal{F}_q^{\text{nf}} := \langle \mathcal{F}_q, \nu_{\mathcal{G}} \rangle$.

The values of the game in normal form $\mathcal{F}_q^{\text{nf}}$ and of the state q are equal.

Proposition 1. *In a PI concurrent game \mathcal{G} , for all states $q \in Q$, we have $\chi_{\mathcal{G}}(q) = \text{out}_{\mathcal{F}_q^{\text{nf}}}$.*

4.2 More on strategies

In this subsection, we define several kinds of strategies. Let us fix a PI concurrent game \mathcal{G} for the rest of this section. First, we consider optimal strategies, i.e. strategies realizing the value of the game. Strategies are positively-optimal if their values are positive from all states whose value is positive.

Definition 12 ((Positively-) optimal strategies). *A Player A strategy $s_A \in S_C^A$ is (resp. positively-) optimal from a state $q \in Q$ if $\chi_{\mathcal{G}}(q) = \chi_{\mathcal{G}}[s_A](q)$ (resp. if $\chi_{\mathcal{G}}(q) > 0 \Rightarrow \chi_{\mathcal{G}}[s_A](q) > 0$). It is (resp. positively-) optimal if this holds from all states $q \in Q$.*

Note that the definition of optimal strategies we consider is sometimes referred to as uniform optimality, as it holds from every state of the game. However, it does not say anything about what happens once some sequence of states have been seen. We would like now to define a notion of strategy that is optimal from any point that can occur after any finite sequence of states has been seen. This correspond to subgame optimal strategies. To define them, we need to introduce the notion of residual strategy.

Definition 13 (Residual and Subgame Optimal Strategies). *For all finite sequences $\rho \in Q^+$, the residual strategy s_A^ρ of a Player A strategy s_A is the strategy $s_A^\rho : Q^+ \rightarrow \mathcal{D}(A)$ such that, for all $\pi \in Q^+$, we have $s_A^\rho(\pi) := s_A(\rho \cdot \pi)$.*

The Player A strategy s_A is subgame optimal if, for all $\rho = \rho' \cdot q \in Q^+$, the residual strategy s_A^ρ is optimal from q , i.e. $\chi_{\mathcal{G}}[s_A^\rho](q) = \chi_{\mathcal{G}}(q)$.

Note that, in particular, subgame optimal strategies are optimal strategies. When such strategies do exist, we want them to be as simple as possible, for instance we want them to be positional, that is that they only depend on the current state of the game.

As for Player B, we will consider a specific kind of strategies, namely deterministic strategies. That is because, once a Player A strategy is fixed we obtain an (infinite) MDP. In such a context, ε -optimal strategies can be chosen among deterministic strategies (see for instance the explanation in [9, Thm. 1]).

Definition 14 (Positional, Deterministic strategies). A Player A strategy s_A is positional if, for all states $q \in Q$ and paths $\rho \in Q^+$ we have $s_A(\rho \cdot q) = s_A(q)$.

A Player B strategy s_B is deterministic if, for all finite sequences $\rho \cdot q \in Q^+$, there exists $b \in B_q$ such that $s_B(\rho \cdot q)(b) = 1$.

5 Necessary and sufficient condition for subgame optimality

In this section, we present a necessary and sufficient pair of conditions for a Player A strategy to be subgame optimal, formally stated in Theorem 1. The arguments given here are somewhat similar to the ones given in Section 4 of [4], which deals with the same question restricted to positional strategies.

The first condition is local: it specifies how a strategy behaves in the games in normal form at each local interaction of the game. As mentioned in Proposition 1, at each state q , the value of the game in normal form \mathcal{F}_q^{nf} is equal to the value of the state q (given by the valuation $\chi_G \in [0, 1]^Q$). This suggests that, for all finite sequences of states $\rho \in Q^+$ ending at that state q , the GF-strategy $s_A(\rho)$ needs to be optimal in the game in normal form \mathcal{F}_q^{nf} for the residual strategy s_A^ρ to be optimal from q . Strategies with such a property are called locally optimal. This is a necessary condition for subgame optimality. (However, it is neither a necessary nor a sufficient condition for optimality, as argued in Section 6).

Definition 15 (Locally optimal strategies). Consider a PI concurrent game \mathcal{G} . A Player A strategy s_A is locally optimal if, for all $\rho = \rho' \cdot q \in Q^+$, the GF-strategy $s_A(\rho)$ is optimal in the game in normal form \mathcal{F}_q^{nf} . That is – recalling that $\nu_G \in [0, 1]^D$ lifts the valuation $\chi_G \in [0, 1]^Q$ to the Nature states – for all $b \in B_q$: $\chi_G(q) \leq \sum_{a \in A} s_A(\rho)(a) \cdot \nu_G \circ \delta(q, a, b) = \text{out}_{\mathcal{F}_q^{nf}}(s_A(\rho), b)$

Lemma 1. In a PI concurrent game, subgame optimal strategies are locally optimal.

Note that this was already shown for positional strategies in [4].

Local optimality does not ensure subgame optimality in general. However, it does ensure that, for all Player B deterministic strategies, the game almost-surely eventually settles in a value area, i.e. in some Q_u for some $u \in V_G$.

Lemma 2. Consider a PI concurrent game \mathcal{G} and a Player A locally optimal strategy s_A . For all Player B deterministic strategies, almost surely the states seen infinitely often have the same value. That is: $\mathbb{P}^{s_A, s_B}[\bigcup_{u \in V_G} Q^* \cdot (Q_u)^\omega] = 1$.

Proof (Sketch). First, if a state of value 1 is reached (i.e. a state in Q_1), then all states that can be seen with positive probability have value 1 (i.e. are in Q_1), since the strategy s_A is locally optimal. Let now $u \in V_G$ be the highest value in V_G that is not 1 and consider the set of infinite paths such that the set Q_u is seen infinitely often but the game does not settle in it, i.e. the set $(Q^* \cdot (Q \setminus Q_u))^\omega \cap (Q^* \cdot Q_u)^\omega \subseteq Q^\omega$. Since the strategy s_A is locally optimal (and

since $V_{\mathcal{G}}$ is finite), one can show that there is a positive probability $p > 0$ such that, the conditional probability of reaching Q_1 knowing that Q_u is left is at least p . Hence, if Q_u is left infinitely often, almost-surely the set Q_1 is seen (and never left). It follows that the probability of the event $(Q^* \cdot (Q \setminus Q_u))^\omega \cap (Q^* \cdot Q_u)^\omega$ is 0. This implies that, almost-surely, if the set Q_u is seen infinitely often, then at some point it is never left. The same arguments can then be used with the highest value in $V_{\mathcal{G}}$ that is less than u , etc. Overall, we obtain that, for all $u \in V_{\mathcal{G}}$, if a set Q_u is seen infinitely often, it is eventually never left almost-surely.

Local optimality ensures that, at each step, the expected values of the states reached does not worsen (and may even improve if Player B does not play optimally). By propagating this property, we obtain that, given a Player A locally optimal strategy and a Player B deterministic strategy, the convex combination of the values u in $V_{\mathcal{G}}$ weighted by the probability of settling in the value area Q_u , from a state q is at least equal to its value $\chi_{\mathcal{G}}(q)$. This is stated in Lemma 3 below.

Lemma 3. *For a PI concurrent game \mathcal{G} , a Player A locally optimal strategy s_A , a Player B deterministic strategy s_B and a state $q \in Q$: $\chi_{\mathcal{G}}(q) \leq \sum_{u \in V_{\mathcal{G}}} u \cdot \mathbb{P}_q^{s_A, s_B}[Q^* \cdot (Q_u)^\omega]$.*

Note that if Player B plays subgame optimally, then this inequality is an equality.

Proof (Sketch). First, let us denote $\mathbb{P}_q^{s_A, s_B}$ by \mathbb{P} . It can be shown by induction that, for all $i \in \mathbb{N}^*$, we have the property $\mathcal{P}(i) : \chi_{\mathcal{G}}(q) \leq \sum_{\pi \cdot q' \in q \cdot Q^i} \chi_{\mathcal{G}}(q') \cdot \mathbb{P}(\pi \cdot q') = \sum_{u \in V_{\mathcal{G}} \setminus \{0\}} u \cdot \mathbb{P}[q \cdot Q^{i-1} \cdot Q_u]$. Furthermore, since by Lemma 2, the game almost-surely settles in a value area, it can be shown that for n large enough, the probability of being in Q_u after n steps (i.e. $\mathbb{P}[q \cdot Q^{n-1} \cdot Q_u]$) is arbitrarily close to the probability of eventually settling in Q_u (i.e. $\mathbb{P}[Q^* \cdot (Q_u)^\omega]$). We can then apply $\mathcal{P}(n)$ to obtain the desired inequality.

Recall that we are considering a pair of conditions to characterize that a strategy is subgame optimal. The first condition is local optimality. To summarize, we have seen that the fact that a strategy is locally optimal ensures that, from any state q , the expected values of the value areas where the game settles is at least $\chi_{\mathcal{G}}(q)$. However, local optimality does not ensure anything as to the probability of W given that the game settles in a specific value area. This is where the second condition comes into play. For the explanations regarding this condition, we will need Lemma 4 below: a consequence of Levy’s 0-1 Law.

Lemma 4. *Let \mathcal{M} be a countable Markov chain with a PI objective. If there is a $q \in Q$ such that $\chi_{\mathcal{M}}(q) < 1$, then $\inf_{q' \in Q} \chi_{\mathcal{M}}(q') = 0$.*

Consider now a Player A subgame optimal strategy s_A and a Player B deterministic strategy. Let us consider what happens if the game eventually settles in Q_u for some $u \in V_{\mathcal{G}} \setminus \{0\}$. Assume towards a contradiction that there is a finite path after which the probability of W given that the play eventually settles in Q_u is less than 1. Then, there is a continuation of this path ending in Q_u for which this probability of W is less than u . Indeed, it was shown that, for a PI objective,

in a countable Markov chain (which is what we obtain once strategies for both players are fixed), if there is a state with a value less than 1, then the infimum of the values in the Markov chain is 0 (this is what is stated in Lemma 4). Following our above towards-a-contradiction-assumption, there would be a finite path from which the Player A strategy s_A is not optimal. This is in contradiction with the fact that it is subgame optimal. Hence, a second necessary condition – in addition to the local optimality assumption – for subgame optimality is: from all finite paths, for all Player B deterministic strategies, for all positive values $u \in V_G \setminus \{0\}$, the probability of W and eventually settling in Q_u is equal to the probability of eventually settling in Q_u . We obtain the theorem below.

Theorem 1. *Consider a concurrent game \mathcal{G} with a PI objective W and a Player A strategy $s_A \in S_C^A$. The strategy s_A is subgame optimal if and only if:*

- *it is locally optimal;*
- *for all $\rho \in Q^+$, for all Player B deterministic strategies s_B , for all values $u \in V_G \setminus \{0\}$, we have $\mathbb{P}_\rho^{s_A, s_B}[W \cap Q^* \cdot (Q_u)^\omega] = \mathbb{P}_\rho^{s_A, s_B}[Q^* \cdot (Q_u)^\omega]$.*

Proof (Sketch). Lemma 1 states that local optimality is necessary and we have informally argued above why the second condition is also necessary for subgame optimality. As for the fact that they are sufficient conditions, this is a direct consequence of Lemmas 2 and 3 and the fact that deterministic strategies can achieve the same values as arbitrary strategies in MDPs (which we obtain once a Player A strategy is fixed), as cited in Subsection 4.2.

One may ask what happens in the special case where the strategy s_A considered is positional. As mentioned above, such a characterization was already presented in [4]¹. Overall, we obtain a similar result except that the second condition is replaced by what happens in the game restricted to the End Components in the Markov Decision Process induced by the positional strategy s_A .

6 From subgame almost-surely winning to subgame optimality

In [14, Thm. 4.5], the authors have proved a transfer result in PI turn-based games: the amount of memory sufficient to play optimally in every state of value 1 of every game is also sufficient to play optimally in every game. This result does not hold on concurrent games as is. First, although there are always optimal strategies in PI turn-based games (as proved in the same paper [14, Thm. 4.3]), there are PI concurrent games without optimal strategies. Second, infinite memory may be required to play optimally in co-Büchi concurrent games whereas almost-surely winning strategies can be found among positional strategies in a turn-based setting. This can be seen in the game of Figure 1 with $\text{col}(q_0) = 0$ and $\text{col}(q_1) = \text{col}(q'_1) = 1$. The green values in the local interaction at state q_0 are the

¹ The proof was only presented for a specific class of objectives.




Fig. 1. A co-Büchi game.




Fig. 2. A parity game.




Fig. 3. A concurrent game with $A_{q_0} = \{a_1, a_2\}$.

values of the game if they are reached (the game ends immediately). If a green value is not reached, the objective of Player A is to see only finitely often states q_1 and q'_1 . It has already been argued in [4] that the value of this game is $1/2$ and that there is an optimal strategy for Player A but it requires infinite memory. To play optimally, Player A must play the top row with probability $1 - \varepsilon_k$ and the middle row with probability ε_k for $\varepsilon_k > 0$ that goes (fast) to 0 when k goes to ∞ (where k denotes the number of steps). The ε_k must be chosen so that, if Player B always plays the left column with probability 1, then the state q_1 is seen finitely often with probability 1. Furthermore, as soon as the state q'_1 is visited, Player A switches to a positional strategy playing the bottom row with probability ε'_k small enough (where k denotes the number of steps before the state q'_1 was seen) and the two top rows with probability $(1 - \varepsilon'_k)/2$.

Hence, the transfer of memory from almost-surely winning to optimal does not hold in concurrent games even if it is assumed that optimal strategies exist. However, one can note that although the strategy described above is optimal, it is not subgame optimal. Indeed, when the strategy switches, the value of the residual strategy is $1/2 - \varepsilon'_k < 1/2$. In fact, there is no subgame optimal strategy in that game. Actually, if we assume that, not only optimal but subgame optimal strategies exist, then the transfer of memory will hold.

The aim of this section is twofold: first, we identify a necessary and sufficient condition for the existence of subgame optimal strategies². Second, we establish the above-mentioned memory transfer that relates the amount of memory to play subgame optimally and to be almost-surely winning. Before stating the main theorem of this section, let us first introduce the definition of positionally subgame almost-surely winnable objective, i.e. objectives for which subgame almost-surely winning strategies can be found among positional strategies.

Definition 16 (Positionally subgame almost-surely winnable objective). Consider a PI objective $W \subseteq \mathcal{K}^\omega$. It is said to be a positionally subgame almost-surely winnable objective (PSAW for short) if the following holds: in all concur-

² Note that this is different from what we did in the previous section: there, we established a necessary and sufficient condition for a specific strategy to be subgame optimal. Here, given a game, we consider necessary and sufficient conditions on the game for the existence of a subgame optimal strategy.

rent games $\mathcal{G} = \langle \mathcal{C}, W \rangle$ where there is a subgame almost-surely winning strategy, there is a positional one.

Theorem 2. Consider a non-empty finite set of colors K and a PI objective $\emptyset \subsetneq W \subseteq K^\omega$. Consider a concurrent game \mathcal{G} with objective W . Then, the three following assertions are equivalent:

- a. there exists a subgame optimal strategy;
- b. there exists an optimal strategy that is locally optimal;
- c. there exists a positively-optimal strategy that is locally optimal.

Furthermore, if this holds and if the objective W is PSAW, then there exists a subgame optimal positional strategy.

First, note that the equivalence is stated in terms of existence of strategies, not on the strategies themselves. In particular, any subgame optimal strategy is both optimal and locally optimal, however, an optimal strategy that is locally optimal is not necessarily a subgame optimal strategy. Second, it is straightforward that point a implies point b (from Theorem 1) and that point b implies point c (by definition of positively-optimal strategies). In the remainder of this section, we explain informally the constructions leading to the proof of this theorem, i.e. to the proof that point c implies point a . The transfer of memory is a direct consequence of the way this theorem is proven. We fix a PI concurrent game $\mathcal{G} = \langle \mathcal{C}, W \rangle$ for the rest of the section.

The idea is as follows. As stated in Theorem 1, subgame optimal strategies are locally optimal and win the game almost-surely if the game settles in a value area Q_u for some positive $u \in V_G \setminus \{0\}$. Our idea is therefore to consider subgame almost-surely winning strategies in the derived game \mathcal{G}_u : a “restriction” of the game \mathcal{G} to Q_u (more details will be given later). We can then glue together these subgame almost-surely winning strategies – defined for all $u \in V_G \setminus \{0\}$ – into a subgame optimal strategy. However, there are some issues:

- 1. the state values in the game \mathcal{G}_u should be all equal to 1;
- 2. furthermore, there must exist a subgame almost-surely winning strategy in \mathcal{G}_u ;
- 3. this subgame almost-surely winning strategy in \mathcal{G}_u should be locally optimal when considered in the whole game \mathcal{G} .

Note that the method we use here is different from what the authors of [14] did to prove the transfer of memory in turn-based games.

Let us first deal with issue 3. One can ensure that the almost-surely winning strategies in the game \mathcal{G}_u are all locally optimal in \mathcal{G} by properly defining the game \mathcal{G}_u . More specifically, this is done by enforcing that the only Player A possible strategies in \mathcal{G}_u are locally optimal in the game \mathcal{G} . To do so, we construct the game \mathcal{G}_u whose state space is Q_u (plus gadget states) but whose set of actions $A_{\mathcal{F}_q^{\text{nf}}}$, at a state $q \in Q_u$, is such that the set of strategies $\mathcal{D}(A_{\mathcal{F}_q^{\text{nf}}})$ corresponds exactly to the set of optimal strategies in the original game in normal form $\mathcal{F}_q^{\text{nf}}$, while keeping the set of actions $A_{\mathcal{F}_q^{\text{nf}}}$ for Player A finite. This is possible thanks

$$\begin{array}{cccc}
 a_1 \begin{bmatrix} q_0 & q_1 \\ q_0 & q_2 \end{bmatrix} & a_1 \begin{bmatrix} \frac{1}{2} & \frac{1}{4} \\ \frac{1}{2} & \frac{3}{4} \end{bmatrix} & \frac{a_1+a_2}{2} \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{3}{4} \end{bmatrix} & \frac{a_1+a_2}{2} \begin{bmatrix} q_0 & \frac{q_1+q_2}{2} \\ q_0 & q_2 \end{bmatrix}
 \end{array}$$

Fig. 4. The local interaction \mathcal{F}_{q_0} at normal form state q_0 . **Fig. 5.** The game $\mathcal{F}_{q_0}^{\text{nf}}$ at the state q_0 . **Fig. 6.** The game $\mathcal{F}_{q_0}^{\text{opt, nf}}$ with only optimal strategies. **Fig. 7.** The game $\mathcal{F}_{q_0}^{\text{opt}}$ with only optimal strategies.

to Proposition 2 below: in every game in normal form $\mathcal{F}_q^{\text{nf}}$ at state $q \in Q_u$, there exists a finite set $A_{\mathcal{F}_q^{\text{nf}}}$ of optimal strategies such that the optimal strategies in $\mathcal{F}_q^{\text{nf}}$ are exactly the convex combinations of strategies in $A_{\mathcal{F}_q^{\text{nf}}}$. This is a well known result, argued for instance in [18].

Proposition 2. *Consider a game in normal form $\mathcal{F}^{\text{nf}} = \langle A, B, [0, 1], \delta \rangle$ with $|A| = n$ and $|B| = k$. There exists a set $A_{\mathcal{F}^{\text{nf}}} \subseteq \text{Opt}_A(\mathcal{F}^{\text{nf}})$ of optimal strategies such that $|A_{\mathcal{F}^{\text{nf}}}| \leq n + k$ and $\mathcal{D}(A_{\mathcal{F}^{\text{nf}}}) = \text{Opt}_A(\mathcal{F}^{\text{nf}})$.*

Proof (Sketch). One can write a system of $n + k$ inequalities (with some additional equalities) whose set of solutions is exactly the set of optimal GF-strategies $\text{Opt}_A(\mathcal{F}^{\text{nf}})$. The result then follows from standard system of inequalities arguments as the space of solutions is in fact a polytope with at most $n + k$ vertices.

We illustrate this construction: a part of a concurrent game is depicted in Figure 3 and the change of the interaction of the players at state q_0 is depicted in Figures 4, 5, 6 and 7.

The game \mathcal{G}_u has the same objective W as the game \mathcal{G} . Since we want all the states to have value 1 in \mathcal{G}_u (recall issue 1), we will build the game \mathcal{G}_u such that any edge leading to a state not in Q_u in \mathcal{G} now leads to a PI concurrent game \mathcal{G}_W (with the same objective W) where all states have value 1. The game \mathcal{G}_W is (for instance) a clique with all colors in K where Player A plays alone.

An illustration of this construction can be found in Figures 8 and 9. The blue dotted arrows are the ones that need to be redirected when the game is changed. With such a definition, we have made some progress w.r.t. the issue 1 cited previously (regarding the values being equal to 1): the values of all states of the game \mathcal{G}_u are positive (for positive u).

Lemma 5. *Consider the game \mathcal{G}_u for some positive $u \in V_G \setminus \{0\}$ and assume that, in \mathcal{G} , there exists a positively-optimal strategy that is locally optimal. Then, for all states q in \mathcal{G}_u , the value of the state q in \mathcal{G}_u is positive: $\chi_{\mathcal{G}_u}(q) > 0$.*

Proof (Sketch). Consider a state $q \in Q_u$ and a Player A locally optimal strategy s_A in \mathcal{G} that is positively-optimal from q . Then, the strategy s_A (restricted to Q_u^+) can be seen as a strategy in \mathcal{G}_u (it has to be defined in \mathcal{G}_W , but this can be done straightforwardly). Note that this is only possible because the strategy s_A is locally optimal (due to the definition of \mathcal{G}_u). For a Player B strategy s_B in \mathcal{G}_u , consider what happens with strategies s_A and s_B in both games \mathcal{G}_u and \mathcal{G} . Either




Fig. 8. The depiction of a PI concurrent game with its value areas.




Fig. 9. The PI concurrent game after the modifications described above.

the game stays indefinitely in Q_u , and what happens in \mathcal{G}_u and \mathcal{G} is identical. Or it eventually leaves Q_u , leading to states of value 1 in \mathcal{G}_u . Hence, the value of the game \mathcal{G}_u from q with strategies s_A and s_B is at least the value of the game \mathcal{G} from q with the same strategies. Thus, the value of the state q is positive in \mathcal{G}_u .

As it turns out, Lemma 5 suffices to deal with both issues 1 and 2 at the same time. Indeed, as stated in Theorem 3 below, it is a general result that in a PI concurrent game, if all states have positive values, then all states have value 1 and there is a subgame almost-surely winning strategy.

Theorem 3. *Consider a PI concurrent game \mathcal{G} and assume that all state values are greater than or equal to $c > 0$, i.e. for all $q \in Q$, $\chi_{\mathcal{G}}(q) \geq c$. Then, there is a subgame almost-surely winning strategy in \mathcal{G} .*

Remark 1. This theorem can be seen as a strengthening of Theorem 1 from [6]. Indeed, this Theorem 1 states that if all states have positive values, then they all have value 1 (this is then generalized to games with countably-many states). Theorem 3 is stronger since it ensures the existence of (subgame) almost-surely winning strategies. Although a detailed proof is provided in the complete version of this paper [5], note that this theorem was already stated and proven in [14] in the context of PI turn-based games. Nevertheless their arguments could have been used *verbatim* for concurrent games as well. In [5], we give a proof using the same construction (namely, reset strategies) but we argue differently why the construction proves the theorem.

We can now glue together pieces of strategies s_A^u defined in all games \mathcal{G}_u into a single strategy $s_A[(s_A^u)_{u \in V_{\mathcal{G}} \setminus \{0\}}]$. Informally, the glued strategy mimics the strategy on Q_u^+ and switches strategy when a value area is left and another one is reached.

Definition 17 (Gluing strategies). *Consider a PI concurrent game \mathcal{G} and for all values $u \in V_{\mathcal{G}} \setminus \{0\}$, a strategy s_A^u in the game \mathcal{G}_u . Then, we glue these*

strategies into the strategy $s_A[(s_A^u)_{u \in V_G \setminus \{0\}}] : Q^+ \rightarrow \mathcal{D}(A)$ simply written s_A such that, for all ρ ending at state $q \in Q$:

$$s_A(\rho) := \begin{cases} s_A^u(\pi) & \text{if } u = \chi_G(q) > 0 \text{ for } \pi \text{ the longest suffix of } \rho \text{ in } Q_u^+ \\ \text{is arbitrary} & \text{if } \chi_G(q) = 0 \end{cases}$$

As stated in Lemma 6 below, the construction described in Definition 17 transfers almost-surely winning strategies in \mathcal{G}_u into a subgame optimal strategy in \mathcal{G} .

Lemma 6. *For all $u \in V_G \setminus \{0\}$, let s_A^u be a subgame almost-surely winning strategy in \mathcal{G}_u . The glued strategy $s_A[(s_A^u)_{u \in V_G \setminus \{0\}}]$, denoted s_A , is subgame optimal in \mathcal{G} .*

Proof (Sketch). We apply Theorem 1. First, the strategy s_A is locally optimal in all Q_u for $u > 0$ by the strategy restriction done to define the game \mathcal{G}_u (only optimal strategies are considered at each game in normal form $\mathcal{F}_q^{\text{nf}}$ at states $q \in Q_u$). Furthermore, any strategy is optimal in a game in normal form of value 0 (which is the case of the game in normal forms of states in Q_0). Second, if the game eventually settles in a value area Q_u for some $u > 0$, from then on the strategy s_A mimics the strategy s_A^u , which is subgame almost-surely winning in \mathcal{G}_u . Hence, the probability of W given that the game eventually settles in Q_u is 1. This holds for all $u \in V_G \setminus \{0\}$, so the second condition of Theorem 1 holds.

We now have all the ingredients to prove Theorem 2.

Proof (Of Theorem 2). We consider the PI concurrent game \mathcal{G} and assume that there is a positively-optimal strategy that is locally optimal. Then, by Lemma 5, for all positive values $u \in V_G \setminus \{0\}$, all states in \mathcal{G}_u have positive values. It follows, by Theorem 3, that there exists a subgame almost-surely winning strategy in every game \mathcal{G}_u for $u \in V_G \setminus \{0\}$. We then obtain a subgame optimal strategy by gluing these strategies together, given by Lemma 6.

The second part of the theorem, dealing with transfer of positionality from subgame almost-surely winning to subgame optimal follows from the fact that if all strategies s_A^u are positional for all $u \in V_G \setminus \{0\}$, then so is the glued strategy $s_A[(s_A^u)_{u \in V_G \setminus \{0\}}]$.

We now apply the result of Theorem 2 to two specific classes of objectives: Büchi and co-Büchi objectives. Note that this result is already known for Büchi objectives, proven in [4].

Corollary 1. *Consider a concurrent game with a Büchi (resp. co-Büchi) objective and assume that there is a positively-optimal strategy that is locally optimal. Then there is a subgame optimal positional strategy.*

Note that it is also possible to prove a memory transfer from subgame almost-surely winning to subgame optimal for an arbitrary memory skeleton, instead of only positional strategies. This adds only a few minor difficulties.

Application to the turn-based setting. The aim of Section 6 was to extend an already existing result on turn-based games in the context of concurrent games. This required an adaptation of the assumptions. However, it is in fact possible to retrieve the original result on turn-based games from Theorem 2 in a fairly straightforward manner. It amounts to show that, in all finite turn-based games \mathcal{G} , for all values $u \in V_{\mathcal{G}} \setminus \{0\}$, there is a locally optimal strategy that is positively-optimal from all states in Q_u .

7 Finite-choice strategies

In this section, we introduce a new kind of strategies, namely finite-choice strategies. Let us first motivate why we consider such strategies. Consider again the co-Büchi game of Figure 1. Recall that the optimal strategy we described first plays the top row with increasing probability and the middle row with decreasing probability and then, once Player B plays the second column, switches to a positional strategy playing the bottom row with positive, yet small enough probability. Note that switching strategy is essential. Indeed, if Player A does not switch, Player B could at some point opt for the middle column and see indefinitely the state q'_1 with very high probability. In fact, what happens in that case is rather counter-intuitive: once Player B switches, there is infinitely often a positive probability to reach the outcome of value 1. However, the probability to ever reaching this outcome can be arbitrarily small, if Player B waits long enough before playing the middle row. This happens because the probability ε_k to visit that outcome goes (fast) to 0 when k goes to ∞ . In fact, such an optimal strategy has “infinite choice” in the sense that it may prescribe infinitely many different probability distribution.

In this section, we consider *finite-choice strategies*, i.e. strategies that can use only finitely many GF-strategies at each state.

Definition 18 (Finite-choice strategy). *Let \mathcal{G} be a concurrent game. A Player A strategy s_A in \mathcal{G} has finite choice if, for all $q \in Q$, the set $S_q^{s_A} := \{s_A(\rho \cdot q) \mid \rho \in Q^+\} \subseteq \mathcal{D}(A_q)$ is finite.*

Note that positional (even finite-memory) and deterministic strategies are examples of finite-choice strategies.

Interestingly, we can link finite-choice strategies with the existence of subgame optimal strategies. In general it does not hold that if there are optimal strategies, then there exists subgame optimal strategies (as exemplified in the game of Figure 1). However, in Theorem 4 below, we state that if we additionally assume that the optimal strategy considered has finite choice, then there is a subgame optimal strategy (that has also finite choice).

Theorem 4. *Consider a PI concurrent game \mathcal{G} . If there is a finite-choice optimal strategy, then there is a finite-choice subgame optimal strategy.*

Proof (Sketch). Consider such an optimal finite-choice strategy s_A . In particular, note that there is a constant $c > 0$ such that for all $\rho \cdot q \in Q^+$, for all $a \in A_q$ we

have: $s_A(\rho \cdot q)(q) > 0 \Rightarrow s_A(\rho \cdot q)(q) \geq c$. We build a subgame optimal strategy s'_A in the following way: for all $\rho = \rho' \cdot q \in Q^+$, if the residual strategy s^ρ_A is optimal, then $s'_A(\rho) := s_A(\rho)$, otherwise $s'_A(\rho) := s_A(q)$ (i.e. we reset the strategy). Straightforwardly, the strategy s'_A has finite choice. We want to apply Theorem 1 to prove that it is subgame optimal. One can see that it is locally optimal (by the criterion chosen for resetting the strategy). Consider now some $\rho \in Q^+$ ending at state $q \in Q$ and another state $q' \in Q$. Assume that the residual strategy s^ρ_A is optimal but that the residual strategy $s^{\rho \cdot q'}_A$ is not. Then, similarly to why local optimality is necessary for subgame optimality (see Proposition 1), one can show that any Player B action b leading to q' from ρ with positive probability is such that $\chi_G(q) < \text{out}_{\mathcal{F}^{\text{gr}}_q}(s_A(\rho), b)$. Hence, there is positive probability from ρ , if Player B opts for the action b , to reach a state of value different from $u = \chi_G(q)$. And if this happens infinitely often, a state of value different from u will be reached almost-surely³. In other words, if a value area is never left, almost-surely, the strategy s'_A only resets finitely often.

Consider now some $\rho \in Q^+$, a Player B deterministic strategy s_B and a value $u \in V_G \setminus \{0\}$. From what we argued above, the probability of the event $Q^* \cdot (Q_u)^\omega$ (resp. $W \cap Q^* \cdot (Q_u)^\omega$) is the same if we intersect it with the fact that the strategy s'_A only resets finitely often. Furthermore, if the strategy does not reset anymore from some point on, and all states have the same value $u > 0$, then it follows that the probability of W is 1 (since W is PI). We can then conclude by applying Theorem 1.

Finite-choice strategies are interesting for another reason. In the previous section, we applied the memory transfer from Theorem 2 to the Büchi and co-Büchi objectives. We did not apply it to other objectives – in particular to the parity objective. Indeed, in general, contrary to the case of turn-based games, infinite-memory is necessary to be almost-surely winning in parity games. This happens in Figure 2 (already described in [12]) where the objective of Player A is to see q_1 infinitely often, while seeing q_2 only finitely often. Let us describe a Player A subgame almost-surely winning strategy. The top row is played with probability $1 - \varepsilon_k$ and the bottom row is played with probability $\varepsilon_k > 0$ with ε_k going to 0 when k goes to ∞ (the (ε_k) used in the game in Figure 1 works here as well) where k denotes the number of times the state q_0 is seen. Such a strategy is subgame almost-surely winning and does not have finite choice. In fact, it can be shown that all Player A finite-choice strategies have value 0 in that game.

Interestingly, the transfer of memory of Theorem 2 is adapted in Theorem 5 with the memory that is sufficient in turn-based games – for those PI objectives that have a “neutral color” – if we additionally assume that the subgame optimal strategy considered has finite choice. First, let us define what is meant by “neutral color”, then we define the turn-based version of PSAW.

³ This holds because the strategy s_A has finite choice: the probability to see a state of different value is bounded below by the product of c and the smallest positive probability among all Nature states.

Definition 19 (Objective with a neutral color). Consider a set of colors K and a PI objective $W \subseteq K^\omega$. It has a neutral color if there is some (neutral) color $k \in K$ such that, for all $\rho = \rho_0 \cdot \rho_1 \cdots \in K^\omega$, we have $\rho \in W \Leftrightarrow \rho_0 \cdot k \cdot \rho_1 \cdot k \cdots \in W$.

Definition 20 (PSAW objective in turn-based games). Consider a PI objective $W \subseteq K^\omega$. It is positionally subgame almost-surely winnable in turn-based games (PSAWT for short) if in all turn-based games $\mathcal{G} = \langle C, W \rangle$ where there is a subgame almost-surely winning strategy, there is a positional one.

Theorem 5. Consider a PSAWT PI objective $W \subseteq K^\omega$ with a neutral color and a concurrent game \mathcal{G} with objective W . Assume there is a subgame optimal strategy that has finite choice. Then, there is a positional one.

Proof (Sketch). A finite-choice strategy s_A plays only among a finite number of GF-strategies at each state. The idea is therefore to modify the game \mathcal{G}_u of the previous subsection into a game \mathcal{G}'_u by transforming it into a (finite) turn-based game. At each state, Player A chooses first her GF-strategy. She can choose among only a finite number of them: she has at her disposal, at a state q , only optimal GF-strategies in $S_q^{S_A}$ (recall Definition 18). We consider the objective W in that new arena where Player B states are colored with a neutral color. The existence, in \mathcal{G} , of a subgame optimal strategy that has finite choice ensures that all states in \mathcal{G}'_u have positive values. We can then conclude as for Theorem 2: a subgame optimal strategy can be obtained by gluing together subgame almost-surely winning strategies in the (turn-based) games \mathcal{G}'_u (that can be chosen positional by assumption).

As an application, one can realize that the parity, mean-payoff and generalized Büchi objectives have a neutral color and are PSAWT ([11,16,7]). Hence, for these objectives, if there exists an optimal strategy that has finite choice, then there is one that is positional.

Corollary 2. Consider a concurrent game \mathcal{G} with a parity (resp. mean-payoff, resp. generalized Büchi) objective. Assume that there is an optimal strategy that has finite choice in \mathcal{G} . Then, there is a positional one.

References

1. Roderick Bloem, Krishnendu Chatterjee, and Barbara Jobstmann. *Handbook of Model Checking*, chapter Graph games and reactive synthesis, pages 921–962. Springer, 2018.
2. Benjamin Bordais, Patricia Bouyer, and Stéphane Le Roux. From local to global determinacy in concurrent graph games. In Mikolaj Bojanczyk and Chandra Chekuri, editors, *41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2021, December 15-17, 2021, Virtual Conference*, volume 213 of *LIPICs*, pages 41:1–41:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

3. Benjamin Bordaïs, Patricia Bouyer, and Stéphane Le Roux. Optimal strategies in concurrent reachability games. In Florin Manea and Alex Simpson, editors, *30th EACSL Annual Conference on Computer Science Logic, CSL 2022, February 14-19, 2022, Göttingen, Germany (Virtual Conference)*, volume 216 of *LIPICs*, pages 7:1–7:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
4. Benjamin Bordaïs, Patricia Bouyer, and Stéphane Le Roux. Playing (almost-)optimally in concurrent büchi and co-büchi games. *CoRR*, abs/2203.06966, 2022.
5. Benjamin Bordaïs, Patricia Bouyer, and Stéphane Le Roux. Sub-game optimal strategies in concurrent games with prefix-independent objectives. *CoRR*, abs/2301.10697, 2023.
6. Krishnendu Chatterjee. Concurrent games with tail objectives. *Theor. Comput. Sci.*, 388(1-3):181–198, 2007.
7. Krishnendu Chatterjee, Luca de Alfaro, and Thomas A. Henzinger. Trading memory for randomness. In *1st International Conference on Quantitative Evaluation of Systems (QEST 2004), 27-30 September 2004, Enschede, The Netherlands*, pages 206–217. IEEE Computer Society, 2004.
8. Krishnendu Chatterjee, Luca de Alfaro, and Thomas A. Henzinger. The complexity of quantitative concurrent parity games. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22-26, 2006*, pages 678–687. ACM Press, 2006.
9. Krishnendu Chatterjee, Laurent Doyen, Hugo Gimbert, and Thomas A. Henzinger. Randomness for free. *Inf. Comput.*, 245:3–16, 2015.
10. Krishnendu Chatterjee and Rasmus Ibsen-Jensen. Qualitative analysis of concurrent mean-payoff games. *Inf. Comput.*, 242:2–24, 2015.
11. Krishnendu Chatterjee, Marcin Jurdzinski, and Thomas A. Henzinger. Quantitative stochastic parity games. In J. Ian Munro, editor, *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2004, New Orleans, Louisiana, USA, January 11-14, 2004*, pages 121–130. SIAM, 2004.
12. Luca de Alfaro and Thomas A. Henzinger. Concurrent omega-regular games. In *15th Annual IEEE Symposium on Logic in Computer Science, Santa Barbara, California, USA, June 26-29, 2000*, pages 141–154. IEEE Computer Society, 2000.
13. Hugh Everett. Recursive games. *Annals of Mathematics Studies – Contributions to the Theory of Games*, 3:67–78, 1957.
14. Hugo Gimbert and Florian Horn. Solving simple stochastic tail games. In Moses Charikar, editor, *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 847–862. SIAM, 2010.
15. Marta Kwiatkowska, Gethin Norman, Dave Parker, and Gabriel Santos. Automatic verification of concurrent stochastic systems. *Formal Methods in System Design*, 58:188–250, 2021.
16. Thomas M Liggett and Steven A Lippman. Stochastic games with perfect information and time average payoff. *Siam Review*, 11(4):604–607, 1969.
17. Donald A. Martin. The determinacy of blackwell games. *The Journal of Symbolic Logic*, 63(4):1565–1581, 1998.
18. Lloyd S Shapley and RN Snow. Basic solutions of discrete games. *Contributions to the Theory of Games*, 1(24):27–27, 1950.
19. Wolfgang Thomas. Infinite games and verification. In *Proc. 14th International Conference on Computer Aided Verification (CAV’02)*, volume 2404 of *Lecture Notes in Computer Science*, pages 58–64. Springer, 2002. Invited Tutorial.
20. John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton Univ. Press, Princeton, 1944.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

