Qualitative Concurrent Parity Games*

Krishnendu Chatterjee[§] Luca de Alfaro[§] Thomas A. Henzinger[†‡]

[§] CE, University of California, Santa Cruz,USA

[†] EECS, University of California, Berkeley,USA

[‡] Computer and Communication Sciences, EPFL, Switzerland

{c_krish,tah}@eecs.berkeley.edu, luca@soe.ucsc.edu

April 2008

Technical Report No. UCSC-CRL-08-02

School of Engineering, University of California, Santa Cruz, CA, USA

Abstract

We consider 2-player games played on a finite state space for an infinite number of rounds. The games are *concurrent*: in each round, the two players choose their moves independently and simultaneously; the current state and the two moves determine the successor state. We consider ω-regular winning conditions specified as *parity* objectives on the resulting infinite state sequence. Both players are allowed to use randomization when choosing their moves. We study the computation of the *limit-winning* set of states, consisting of the states where the sup-inf value of the game for player 1 is 1: in other words, a state is *limit-winning* if player 1 can ensure a probability of winning arbitrarily close to 1. We show that the limit-winning set can be computed in $\mathcal{O}(n^{2m+2})$ time, where n is the size of the game structure and $2m$ is the number of parities; membership of a state in the limit-winning set can be decided in NP $\cap$ coNP. While this complexity is the same as for the simpler class of *turn-based* parity games, where in each state only one of the two players has a choice of moves, our algorithms are considerably more involved than those for turn-based games. This is because concurrent games violate two of the most fundamental properties of turn-based parity games. First, in concurrent games limit-winning strategies require randomization; and second, they require infinite memory.

1 Introduction

Concurrent games are played by two players on a finite state space for an infinite number of rounds. In each round, the two players independently choose moves, and the current state and the two chosen moves determine the successor state. In *deterministic* concurrent games, the successor state is unique; in *probabilistic* concurrent games, the successor state is given by a probability distribution. The outcome of the game is an infinite sequence of states. We consider ω-regular objectives; that is, given an ω-regular set W of infinite state sequences, player 1 wins if the outcome of the game lies in W. Otherwise, player 2 wins, i.e., the game is zero-sum. Such games occur in the synthesis and verification of reactive systems [Chu62, RW87, PR89, ALW89, Dil89, AHK97].

*This research was suppored in part by the NSF grants CCR-0132780, CNS-0720884, and CCR-0225610, and by the Swiss National Science Foundation.

The player-1 *value* $v_1(s)$ of the game at a state s is the limit probability with which player 1 can ensure that the outcome of the game lies in W; that is, the value $v_1(s)$ is realized by the supremum of player-1 strategies (called the *optimal* player-1 strategy) against the infimum of player-2 strategies. Symmetrically, the player-2 *value* $v_2(s)$ is the limit probability with which player 2 can ensure that the outcome of the game lies outside W. The *qualitative* analysis of games ask for a computation of the set of states whose values are 0 or 1; the *quantitative* analysis asks for a precise computation of values.

Traditionally, the special case of *turn-based* games has received most attention. In turn-based games, in each round, only one of the two players has a choice of moves. This situation arises in the modeling of reactive systems if the interaction between components (e.g., plant and controller) is asynchronous; that is, if a scheduler decides in each round which of the components gets to proceed. In turn-based deterministic games, all values are 0 or 1 and can be computed using combinatorial algorithms [Tho90]; in turn-based probabilistic games, values can be computed by iterative approximation [Con93].

In this paper we focus on the more general *concurrent* situation, where in each round, both players choose their moves simultaneously and independently. Such concurrency is necessary for modeling the synchronous interaction of components [dAHM00, dAHM01]. The concurrent probabilistic games fall into a class of stochastic games studied in game theory [Sha53], and the ω-regular objectives, which arise from the safety and liveness specifications of reactive systems, fall into a low level ($\Sigma_3 \cap \Pi_3$) of the Borel hierarchy. It follows from a classical result of Martin [] that these games are determined, i.e., for each state s we have $v_1(s) + v_2(s) = 1$. We study the problem of computing the *limit-winning* sets of states for both players, i.e., the states s with $v_1(s) = 1$ and the states s with $v_1(s) = 0$.

Concurrent games differ from turn-based games in that optimal strategies require, in general, randomization. A player that uses a *pure* strategy must, in each round, choose a move based on the current state and the history (i.e., past state sequence) of the game. By contrast, a player that uses a *randomized* strategy may choose not a single move, but a probability distribution over moves. The move to be played is then selected at random, according to the chosen distribution. Randomized strategies are not helpful for achieving a value of 1 in turn-based probabilistic games, but they can be helpful in concurrent games, even if the game itself is deterministic. To see this, consider the concurrent deterministic game called MatchBit. In each round, both players simultaneously and independently choose a bit (0 or 1), and the objective of player 1 is satisfied if the two chosen bits match in any round. For every pure strategy of player 1, there is a corresponding pure strategy for player 2 that prevents player 1 from winning (the strategy for player 2 always chooses a different bit than the one chosen by the strategy for player 1). However, if both players choose their bits truly simultaneously and independently, then it is extremely "likely" that the chosen bits will match in some round. This intuition is captured mathematically by randomization: if player 1 chooses her bits at random, with uniform probability, then player 1 wins with probability 1/2 at each round, and she can win the game with probability 1.

For deterministic games that are played with pure strategies, once the strategies of both players are fixed, the unique outcome of the game either does or does not satisfy the given objective W. For fixed randomized strategies, there are many possible outcomes: a given strategy pair may result in player 1 winning with probability 1 (as in the game MatchBit), or with probability greater than 7/8, etc. While there are several "qualitative" interpretations of winning in games with randomized strategies (winning with certainty, winning with probability 1 [dAHK98]), the

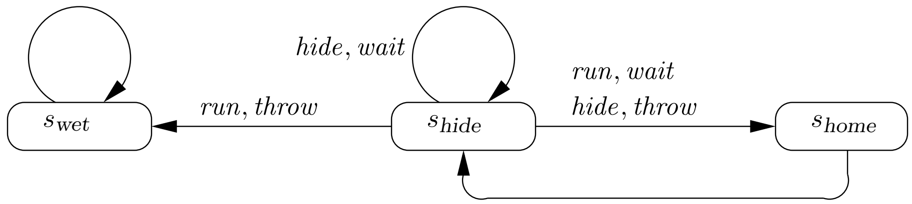

Figure 1: Game SKIRMISH

classical game-theoretic notion of qualitative winning (i.e., game value 1) corresponds to "winning with probability arbitrarily close to 1": a state s is limit-winning for player 1 (i.e., $v_1(s) = 1$) iff for each real $\varepsilon > 0$, there is a player-1 strategy that guarantees a win with probability $1-\varepsilon$ against any strategy of player 2. Limit-winning is illustrated by the concurrent deterministic game SKIRMISH, which is derived from a game of [KS81]: player 1 is hiding (state s_{hide}); her goal is to run and reach home (state s_{home}) without being hit by a snowball; player 2 is armed with a single snowball. There are three states, s_{wet}, s_{hide}, and s_{home}. The only state with more than one move for either player is the initial state s_{hide}, where player 1 must choose between *hide* and *run*, and player 2 must choose between *wait* and *throw*. The effects of the moves are shown in Figure 1. To see that from the state s_{hide} player 1 can limit-win, given any $\varepsilon > 0$, suppose that player 1 chooses *run* with probability ε and *hide* with probability $1 - \varepsilon$. While player 1 can win this game with probability arbitrarily close to 1, she cannot win with probability 1: if she never chooses *run*, she risks that player 2 always chooses *wait*, confining her in s_{hide}; on the other hand, if in any round she chooses *run* with positive probability, then the strategy of player 2 that chooses *throw* in that round causes her to lose with positive probability.

We consider objectives in *parity* form, which is a normal form for all ω-regular objectives. Specifically, every finite-state game with an ω-regular winning condition can be reduced to another finite-state game with a parity winning condition, because every ω-regular set can be defined by a deterministic parity automaton [Mos84, Tho90]. Concurrent probabilistic games generalize Markov chains, Markov decision processes, deterministic as well as probabilistic turn-based games, and concurrent deterministic games. Previously, algorithms for computing the limit-winning sets of parity games have been known only for (1) all varieties of turn-based deterministic games [BL69, GH82, EJ91, Tho95], (2) concurrent parity games that are played with pure strategies (these games can be solved like turn-based games) [AHK97], and (3) concurrent reachability games played with randomized strategies (reachability objectives are a very special case of ω-regular objectives, as in the MATCHBIT and SKIRMISH games) [dAHK07]. We provide an algorithm for computing, given a concurrent probabilistic parity game, the set of limit-winning states for player 1. The limit-winning set can be computed in time $\mathcal{O}(n^{2m+1})$, where n is the size of the game structure and $2m$ is the number of parities in the winning condition.

Also, given a concurrent probabilistic parity game and a state s, we show that whether s is a limit-winning state for player 1 can be decided in NP $\cap$ coNP. While this complexity is the same as for solving turn-based deterministic parity games [EJ91], our algorithms and correctness proofs are considerably more involved than those for turn-based games. This has several reasons.

First, in sharp contrast with turn-based games [BL69], the optimal (i.e., limit-winning) strategies for a concurrent (deterministic or probabilistic) game may require both randomization and an infinite amount of memory about the history of the game. Randomization is required for reachability objectives, and infinite memory is required for Büchi objectives. Consider again the game

SKIRMISH, together with the Büchi winning condition that the state s_{home} be visited infinitely often. The limit-winning states are s_{home} and s_{hide}. As explained above, for every ε, from s_{hide} player 1 can reach s_{home} with probability at least $1-\varepsilon$. However, if player 1 uses the same probability ε to choose *run* in every visit to s_{hide}, by always choosing *throw* player 2 can ensure that the probability of infinitely many visits to s_{home} is 0. The proof that there are no finite-memory (rather than memoryless) winning strategies follows a similar argument. On the other hand, for $\varepsilon > 0$, an infinite-memory strategy that ensures winning with probability at least $1-\varepsilon$ can be constructed as follows: for $k \geq 0$, let $\varepsilon_k = 1 - (1-\varepsilon)^{-1/2^{k+1}}$, so that $\prod_{k=0}^{\infty}(1-\varepsilon_k) = 1-\varepsilon$; then, at s_{hide}, choose *run* with probability ε_k, where k is the number of prior visits to s_{home}. Thus, the construction of winning strategies for concurrent games often hinges on the analysis of the limit behavior of infinite-memory randomized strategies. In the paper, we provide a complete characterization of the types of winning and spoiling strategies needed for the various subclasses of concurrent games.

Second, the fact that both players can choose from several available moves at a state breaks the standard recursive divide-and-conquer approach to the solution of turn-based parity games [McN93, Tho95]. For example, the set of states from which player 1 cannot reach a goal no longer forms a proper subgame. Our algorithms are instead presented in symbolic fixpoint form, using μ-calculus notation, which, as first remarked in [EJ91], offers a powerful tool for writing and analyzing algorithms that traverse state spaces. The fixpoint solution also suggests a way for implementing the algorithms symbolically, potentially enabling the analysis of systems with very large state spaces [BCM+90].

A preliminary version of this article appeared in [dAH00]. The preliminary version contained several gaps that are filled in the present article. In addition to computing limit-winning states, [dAH00] also gave algorithms for computing the smaller set of *almost-sure winning* states, where there exists an optimal strategy that allows player 1 to win with probability 1. Following the techniques of [dAH00], algorithms for computing almost-sure winning states can be derived from algorithms for computing limit-winning states. However, we omit the discussion of almost-sure winning from the present article, because it would increase the length of what is already a very long presentation.

2 Definitions

In this section we define game structures, strategies, objectives, winning modes and other preliminary definitions.

2.1 Game structures

Probability distributions. For a finite set A, a *probability distribution* on A is a function $\delta : A \mapsto [0,1]$ such that $\sum_{a \in A} \delta(a) = 1$. We denote the set of probability distributions on A by $\mathcal{D}(A)$. Given a distribution $\delta \in \mathcal{D}(A)$, we denote by $\mathrm{Supp}(\delta) = \{x \in A \mid \delta(x) > 0\}$ the *support* of the distribution δ.

Concurrent game structures. A concurrent (two-player) *game structure* $\mathcal{G} = \langle S, A, \Gamma_1, \Gamma_2, \delta \rangle$ consists of the following components.

- A finite state space S.
- A finite set A of moves.

- Two move assignments $\Gamma_1, \Gamma_2 \colon S \mapsto 2^A \setminus \emptyset$. For $i \in \{1,2\}$, assignment Γ_i associates with each state $s \in S$ the nonempty set $\Gamma_i(s) \subseteq A$ of moves available to player i at state s. For technical convenience, we assume that $\Gamma_i(s) \cap \Gamma_j(t) = \emptyset$ unless $i = j$ and $s = t$, for all $i, j \in \{1,2\}$ and $s, t \in S$. The moves can be trivially renamed if this assumption is not met.

- A probabilistic transition function $\delta \colon S \times A \times A \mapsto \mathcal{D}(S)$, which associates with every state $s \in S$ and moves $a_1 \in \Gamma_1(s)$ and $a_2 \in \Gamma_2(s)$ a probability distribution $\delta(s, a_1, a_2) \in \mathcal{D}(S)$ for the successor state.

Plays. At every state $s \in S$, player 1 chooses a move $a_1 \in \Gamma_1(s)$, and simultaneously and independently player 2 chooses a move $a_2 \in \Gamma_2(s)$. The game then proceeds to the successor state t with probability $\delta(s, a_1, a_2)(t)$, for all $t \in S$. For all states $s \in S$ and moves $a_1 \in \Gamma_1(s)$ and $a_2 \in \Gamma_2(s)$, we indicate by $\mathit{Dest}(s, a_1, a_2) = \mathrm{Supp}(\delta(s, a_1, a_2))$ the set of possible successors of s when moves a_1, a_2 are selected. A *path* or a *play* of $\mathcal{G}$ is an infinite sequence $\omega = \langle s_0, s_1, s_2, \ldots \rangle$ of states in S such that for all $k \geq 0$, there are moves $a_1^k \in \Gamma_1(s_k)$ and $a_2^k \in \Gamma_2(s_k)$ such that $s_{k+1} \in \mathit{Dest}(s_k, a_1^k, a_2^k)$. We denote by Ω the set of all paths.

2.2 Strategies

A *strategy* for a player is a recipe that describes how to extend a play. Formally, a strategy for player $i \in \{1,2\}$ is a mapping $\pi_i \colon S^+ \mapsto \mathcal{D}(A)$ that associates with every nonempty finite sequence $x \in S^+$ of states, representing the past history of the game, a probability distribution $\pi_i(x)$ used to select the next move. Thus, the choice of the next move can be history-dependent and randomized. The strategy π_i can prescribe only moves that are available to player i; that is, for all sequences $x \in S^*$ and states $s \in S$, we require that $\mathrm{Supp}(\pi_i(x \cdot s)) \subseteq \Gamma_i(s)$. We denote by Π_i the set of all strategies for player $i \in \{1,2\}$.

Given a state $s \in S$ and two strategies $\pi_1 \in \Pi_1$ and $\pi_2 \in \Pi_2$, we define $\mathit{Outcomes}(s, \pi_1, \pi_2) \subseteq \Omega$ to be the set of paths that can be followed by the game, when the game starts from s and the players use the strategies π_1 and π_2. Formally, $\langle s_0, s_1, s_2, \ldots \rangle \in \mathit{Outcomes}(s, \pi_1, \pi_2)$ if $s_0 = s$ and if for all $k \geq 0$ there exist moves $a_1^k \in \Gamma_1(s_k)$ and $a_2^k \in \Gamma_2(s_k)$ such that

$$\pi_1(s_0, \ldots, s_k)(a_1^k) > 0, \qquad \pi_2(s_0, \ldots, s_k)(a_2^k) > 0, \qquad s_{k+1} \in \mathit{Dest}(s_k, a_1^k, a_2^k).$$

Once the starting state s and the strategies π_1 and π_2 for the two players have been chosen, the game is reduced to an ordinary stochastic process. Hence, the probabilities of events are uniquely defined, where an *event* $\mathcal{A} \subseteq \Omega$ is a measurable set of paths[1]. For an event $\mathcal{A} \subseteq \Omega$, we denote by $\mathrm{Pr}_s^{\pi_1,\pi_2}(\mathcal{A})$ the probability that a path belongs to $\mathcal{A}$ when the game starts from s and the players use the strategies π_1 and π_2.

Types of strategies. We classify strategies according to their use of *memory* and *randomization*:

- A strategy π is *deterministic* if for all $x \in S^+$ there exists $a \in A$ such that $\pi(x)(a) = 1$. Thus, deterministic strategies are equivalent to functions $S^+ \mapsto A$. We denote by Π_i^D the set of deterministic strategies for player i.

[1] To be precise, we should define events as measurable sets of paths *sharing the same initial state,* and we should replace our events with families of events, indexed by their initial state [KSK66]. However, our (slightly) improper definition leads to more concise notation.

- Strategies in general require *memory* to remember the history of plays. An equivalent definition of strategies is as follows: let $\mathcal{M}$ be a set called *memory* to remember the history of plays. A strategy with memory can be described as a pair of functions: (a) a *memory update* function $\pi_u : S \times \mathcal{M} \mapsto \mathcal{M}$, that given the memory $\mathcal{M}$ with the information about the history and the current state updates the memory; and (b) a *next move* function $\pi_n : S \times \mathcal{M} \mapsto \mathcal{D}(A)$ that given the memory and the current state specifies the next move of the player. A strategy is *finite-memory* if the memory $\mathcal{M}$ is finite. We denote by Π_i^F the set of finite-memory strategies for player i.

- A *memroyless* strategy is independent of the history of play and only depends on the current state. Formally, for a memoryless strategy π we have $\pi(x \cdot s) = \pi(s)$ for all $s \in S$ and all $x \in S^*$. Thus memoryless strategies are equivalent to functions $S \mapsto \mathcal{D}(A)$. We denote by Π_i^M the set of memoryless strategies for player i. A strategy is *deterministic memoryless* if it is both deterministic and memoryless. The deterministic memoryless strategy neither use memory, nor use randomization and are equivalent to functions $S \mapsto A$. We denote by $\Pi_i^{DM} = \Pi_i^D \cap \Pi_i^M$ the set of deterministic memoryless strategies for player i.

In the tables listing strategy types, we indicate with Π^H (for *history-dependent*) a generic strategy, in which the distributions chosen may depend in an arbitrary way on the past history of the game.

2.3 Objectives

We specify objectives for the players by providing the set of *winning plays* $\Phi \subseteq \Omega$ for each player. In this paper we study only zero-sum games [RF91, FV97], where the objectives of the two players are complementary. A general class of objectives are the Borel objectives [Kec95]. A *Borel objective* $\Phi \subseteq S^\omega$ is a Borel set in the Cantor topology on S^ω. In this paper we consider *ω-regular objectives* [Tho90], which lie in the first $2^1\!/_2$ levels of the Borel hierarchy (i.e., in the intersection of Σ_3 and Π_3). The ω-regular objectives, and subclasses thereof, can be specified in the following forms. For a play $\omega = \langle s_0, s_1, s_2, \ldots \rangle \in \Omega$, we define $\mathit{Inf}(\omega) = \{s \in S \mid s_k = s \text{ for infinitely many } k \geq 0\}$ to be the set of states that occur infinitely often in ω.

- *Reachability and safety objectives.* Given a set $T \subseteq S$ of "target" states, the reachability objective requires that some state of T be visited. The set of winning plays is thus $\mathrm{Reach}(T) = \{\omega = \langle s_0, s_1, s_2, \ldots \rangle \in \Omega \mid s_k \in T \text{ for some } k \geq 0\}$. Given a set $F \subseteq S$, the safety objective requires that only states of F be visited. Thus, the set of winning plays is $\mathrm{Safe}(F) = \{\omega = \langle s_0, s_1, s_2, \ldots \rangle \in \Omega \mid s_k \in F \text{ for all } k \geq 0\}$.

- *Büchi and co-Büchi objectives.* Given a set $B \subseteq S$ of "Büchi" states, the Büchi objective requires that B is visited infinitely often. Formally, the set of winning plays is $\text{Büchi}(B) = \{\omega \in \Omega \mid \mathit{Inf}(\omega) \cap B \neq \emptyset\}$. Given $C \subseteq S$, the co-Büchi objective requires that all states visited infinitely often are in C. Formally, the set of winning plays is $\text{co-Büchi}(C) = \{\omega \in \Omega \mid \mathit{Inf}(\omega) \subseteq C\}$.

- *Rabin-chain (parity) objectives.* For $c, d \in \mathbb{N}$, we let $[c..d] = \{c, c+1, \ldots, d\}$. Let $p : S \mapsto [0..d]$ be a function that assigns a *priority* $p(s)$ to every state $s \in S$, where $d \in \mathbb{N}$. The *Even parity objective* requires that the maximum priority visited infinitely often is even. Formally, the set of winning plays is defined as $\mathrm{Parity}(p) = \{\omega \in \Omega \mid \max\big(p(\mathit{Inf}(\omega))\big) \text{ is even }\}$. The dual

Odd parity objective is defined as $\text{coParity}(p) = \{\omega \in \Omega \mid \max\big(p(\mathit{Inf}(\omega))\big) \text{ is odd }\}$. Note that for a priority function $p : S \mapsto \{1, 2\}$, an even parity objective $\text{Parity}(p)$ is equivalent to the Büchi objective $\text{Büchi}(p^{-1}(2))$, i.e., the Büchi set consists of the states with priority 2. Hence Büchi and co-Büchi objectives are simpler and special cases of parity objectives.

Given a set $U \subseteq S$ we use shorthand notations $\Box U, \Diamond U, \Box\Diamond U$ and $\Diamond\Box U$ to denote $\text{Safe}(U), \text{Reach}(U), \text{Büchi}(U)$ and $\text{co-Büchi}(U)$, respectively.

2.4 Winning modes

Given an objective Φ, for all initial states $s \in S$, the set of paths Φ is measurable for all choices of the strategies of the player [Var85]. Given an initial state $s \in S$ and an objective Φ, we consider the following *winning modes* for player 1:

Limit. We say that player 1 *wins limit surely* if the player has a strategy to win with probability arbitrarily close to 1, or $\sup_{\pi_1 \in \Pi_1} \inf_{\pi_2 \in \Pi_2} \Pr_s^{\pi_1,\pi_2}(\Phi) = 1$.

Bounded. We say that player 1 *wins boundedly* if the player has a strategy to win with probability bounded away from 0, or $\sup_{\pi_1 \in \Pi_1} \inf_{\pi_2 \in \Pi_2} \Pr_s^{\pi_1,\pi_2}(\Phi) > 0$.

Analogous definitions apply for player 2. We abbreviate the winning modes by *limit* and *bounded*, respectively. We call these winning modes the *qualitative* winning modes. Using a notation derived from *alternating temporal logic* [AHK97], given a player $i \in \{1, 2\}$, a winning mode $\zeta \in \{\mathit{limit}, \mathit{bounded}\}$ and an objective Φ, we denote by $\langle\langle i \rangle\rangle_\zeta(\Phi)$ the set of states from which player i can win in mode ζ the game with objective Φ.

2.5 Winning and spoiling strategies

Given an objective Φ, the *winning strategies* are the strategies that enable player 1 to win the game whenever possible. We define limit-winning strategies as follows.

- A *limit-winning family of strategies for* Φ is a family $\{\pi_1(\varepsilon) \mid \varepsilon > 0\}$ of strategies for player 1 such that for all reals $\varepsilon > 0$, all states $s \in \langle\langle 1 \rangle\rangle_{\mathit{limit}}(\Phi)$, and all strategies π_2 of player 2, we have $\Pr_s^{\pi_1(\varepsilon),\pi_2}(\Phi) \geq 1 - \varepsilon$.

The *spoiling strategies* for an objective Φ are the strategies that enable player 2 to prevent player 1 from winning the game whenever it cannot be won. We define limit-spoiling strategies for player 2 as follows.

- A *limit-spoiling strategy for* Φ is a strategy π_2 for player 2 such that there exists a real $q > 0$ such that for all states $s \notin \langle\langle 1 \rangle\rangle_{\mathit{limit}}(\Phi)$ and all strategies π_1 of player 1, we have $\Pr_s^{\pi_1,\pi_2}(\Phi) \leq 1 - q$.

We will show that limit-winning or limit-spoiling strategies always exist. In the following sections, we consider objectives that consist in safety and reachability, Büchi, co-Büchi, and Rabin-chain (or parity) objectives [Mos84, Tho90]. We call games with such objectives *safety, reachability, Büchi, co-Büchi,* and *Rabin-chain (parity) games,* respectively. We remark that the ability of solving games with Rabin-chain objectives suffices for solving games with respect to arbitrary ω-regular objectives. In fact, we can encode a general ω-regular objective as a deterministic Rabin-chain automaton. By taking the synchronous product of the automaton and the original game, we obtain an (enlarged) game with a Rabin-chain objective [Tho95, LW95, KPBV95, BLV96]. The set of winning states of the original structure can be computed by computing the set of winning states of this enlarged game.

2.6 Mu-calculus, complementation, and levels

Consider a mu-calculus expression $Y = \mu X \,.\, \phi(X)$ over a finite set S, where $\phi : 2^S \mapsto 2^S$ is monotonic. The least fixpoint $Y = \mu X \,.\, \phi(X)$ of $X = \phi(X)$ is equal to the limit $Y = \lim_{k\to\infty} X_k$, where $X_0 = \emptyset$, and $X_{k+1} = \phi(X_k)$. For every state $s \in Y$, we define the *level* $k \geq 0$ of s in $\mu X \,.\, \phi(X)$ to be the integer such that $s \notin X_k$ and $s \in X_{k+1}$. The greatest fixpoint $Y = \nu X \,.\, \phi(X)$ of $X = \phi(X)$ is equal to the limit $Y = \lim_{k\to\infty} X_k$, where $X_0 = S$, and $X_{k+1} = \phi(X_k)$. For every state $s \notin Y$, we define the *level* $k \geq 0$ of s in $\nu X \,.\, \phi(X)$ to be the integer such that $s \in X_k$ and $s \notin X_{k+1}$. The *height* of a mu-calculus expression $Y = \lambda X \,.\, \phi(X)$, where $\lambda \in \{\mu, \nu\}$, is the maximal level of any state in Y, i.e., the integer h such that $X_h = \lim_{k\to\infty} X_k$. An expression of height h can be computed in $h+1$ iterations. Given a mu-calculus expression $Y = \lambda X \,.\, \phi(X)$, where $\lambda \in \{\mu, \nu\}$, the complement $\neg Y = S \setminus Y$ of λ is given by $\neg Y = \overline{\lambda} X \,.\, \neg\phi(\neg X)$, where $\overline{\lambda} = \mu$ if $\lambda = \nu$, and $\overline{\lambda} = \nu$ if $\lambda = \mu$.

Distributions and one-step transitions. Given a state $s \in S$, we denote by $\chi_1^s = \mathcal{D}(\Gamma_1(s))$ and $\chi_2^s = \mathcal{D}(\Gamma_2(s))$ the sets of probability distributions over the moves at s available to player 1 and 2, respectively. Moreover, for $s \in S$, $X \subseteq S$, $\xi_1 \in \chi_1^s$, and $\xi_2 \in \chi_2^s$ we denote by

$$P_s^{\xi_1,\xi_2}(X) = \sum_{a\in\Gamma_1(s)} \sum_{b\in\Gamma_2(s)} \sum_{t\in X} \xi_1(a) \cdot \xi_2(b) \cdot \delta(s,a,b)(t)$$

the one-step probability of a transition into X when players 1 and 2 play at s with distributions ξ_1 and ξ_2, respectively. Given a state s and distributions $\xi_1 \in \chi_1^s$ and $\xi_2 \in \chi_2^s$ we denote by $\mathit{Dest}(s, \xi_1, \xi_2) = \{t \in S \mid P_2^{\xi_1,\xi_2}(t) > 0\}$ the set of states that have positive probability of transition from s when the players play ξ_1 and ξ_2 at s.

3 Safety Games

Given a safety objective $\Box U$, where $U \subseteq S$ is a subset of states, we present symbolic algorithms to compute $\langle\langle 1 \rangle\rangle_{\mathit{limit}}(\Box U)$. We introduce the necessary predecessors operators for the symbolic algorithm.

The Pre and Epre operators. The *controllable predecessor* operators Pre_1 and Epre_2, defined for all $s \in S$ and $X \subseteq S$ by:

$$\begin{aligned} \mathrm{Pre}_1(X) &= \{s \in S \mid \exists \xi_1 \in \chi_1^s \,.\, \forall \xi_2 \in \chi_2^s \,.\, P_s^{\xi_1,\xi_2}(X) = 1\}; \\ \mathrm{Epre}_2(X) &= \{s \in S \mid \exists \xi_2 \in \chi_2^s \,.\, \forall \xi_1 \in \chi_1^s \,.\, P_s^{\xi_1,\xi_2}(X) > 0\} \,. \end{aligned}$$

Intuitively, if $s \in \mathrm{Pre}_1(X)$ then player 1 can enforce reaching X with probability 1 in one step, and if $s \in \mathrm{Epre}_1(X)$, then player 1 can enforce reaching X with positive probability in one step. Given a subset $Z \subseteq S$, a move $a \in \Gamma_1(s)$ *risks* moving into Z if there exists $b \in \Gamma_2(s)$ such that $\mathit{Dest}(s, a, b) \cap Z \neq \emptyset$. Observe that if player 2 plays all moves uniformly at random and player 1 plays a move that risks to move into Z, then the game reaches Z with positive probability. In the above definitions, the operators Pre and Epre are defined in terms of distributions. The following lemma provides a definition in terms of moves.

Lemma 1 *For all $X \subseteq S$ we have*

$$\begin{array}{rrcl} 1. & s \in \mathit{Pre}_1(X) & \mathit{iff} & \exists a \in \Gamma_1(s) \,.\, \forall b \in \Gamma_2(s) \,.\, \mathit{Dest}(s,a,b) \subseteq X; \\ 2. & s \in \mathit{Epre}_2(\neg X) & \mathit{iff} & \forall a \in \Gamma_1(s) \,.\, \exists b \in \Gamma_2(s) \,.\, \mathit{Dest}(s,a,b) \cap \neg X \neq \emptyset; \\ 3. & \mathit{Epre}_2(\neg X) & = & \neg \mathit{Pre}_1(X). \end{array}$$

Proof. Assume that there exists $a \in \Gamma_1(s)$ such that for all $b \in \Gamma_2(s)$ we have $\mathit{Dest}(s,a,b) \subseteq X$, and let ξ_1 be a distribution that deterministically chooses one such a. This distribution ξ_1 is a witness to the existential quantifier in the definition of Pre_1. This shows that if $\exists a \in \Gamma_1(s) \,.\, \forall b \in \Gamma_2(s) \,.\, \mathit{Dest}(s,a,b) \subseteq X$, then $s \in \mathrm{Pre}_1(X)$.

Conversely, assume that for all $a \in \Gamma_1(s)$ there exists $b \in \Gamma_2(s)$ such that $\mathit{Dest}(s,a,b) \cap \neg X \neq \emptyset$. Let ξ_2 be the distribution that plays uniformly at random all the moves in the set $\Gamma_2(s)$. This distribution is a witness to the existential quantifier in the definition of Epre_2, and thus, $s \in \mathrm{Epre}_2(\neg X)$. Observe that the distribution ensures that $\neg X$ is reached with probability at least $\frac{1}{|\Gamma_2(s)|} \cdot \min\{\delta(s,a,b)(t) \mid a \in \Gamma_1(s), b \in \Gamma_2(s), \delta(s,a,b)(t) > 0\}$, so that the probability is bounded away from 0.

The result follows by noting that the two above assumptions are complementary, and $\mathrm{Pre}_1(X) \cap \mathrm{Epre}_2(\neg X) = \emptyset$. ∎

In view of these results, if $s \in \mathrm{Pre}_1(X)$ we denote by $\xi^{\mathrm{Pre}}_{s,1}(X) \in \chi^s_1$ a distribution that deterministically chooses a move from the set $\{a \in \Gamma_1(s) \mid \forall b \in \Gamma_2(s) \,.\, \mathit{Dest}(s,a,b) \subseteq X\}$. Similarly, if $s \in \mathrm{Epre}_2(\neg X)$, we define the distribution $\xi^{\mathrm{Epre}}_{s,2}(\neg X) \in \chi^s_2$ as the distribution that plays uniformly at random all the moves in the set $\Gamma_2(s)$.

Lemma 2 *For all $U \subseteq S$ we have*

$$\nu X \,.\, [\mathit{Pre}_1(X) \cap U] \subseteq \langle\!\langle 1 \rangle\!\rangle_{\mathit{limit}}\big(\Box U\big) \,. \tag{1}$$

Proof. Let $W = \nu X \,.\, [\mathrm{Pre}_1(X) \cap U]$. Define a deterministic memoryless strategy π_1 as follows: $\pi_1(s) = \xi^{\mathrm{Pre}}_{s,1}(W)$ for $s \in W$, and an arbitrary deterministic memoryless strategy for $s \in S \setminus W$. The strategy π_1 ensures that for all π_2 and for all $s \in W$ we have $\mathrm{Pr}^{\pi_1,\pi_2}_s(\Box U) = 1$; clearly π_1 is a limit-winning strategy for all $s \in W$. ∎

By Lemma 1 we have $\mathrm{Epre}_2(X) = \neg \mathrm{Pre}_1(\neg X)$, and complementing the mu-calculus expression (1) we obtain

$$V = \mu X \,.\, [\neg U \cup \mathrm{Epre}_2(X)] \,. \tag{2}$$

Lemma 3 *For all $U \subseteq S$, let $V = \mu X \,.\, [\neg U \cup \mathit{Epre}_2(X)]$. We have*

$$V \subseteq \big\{s \in S \bigm| \exists q > 0 \,.\, \exists \pi_2 \,.\, \forall \pi_1 \,.\, \mathrm{Pr}^{\pi_1,\pi_2}_s(\Diamond \neg U) \geq q > 0\big\}.$$

Proof. Let V_k be the set of states of level $k \geq 0$ in $V = \mu X.(\neg U \cup \mathrm{Epre}_2(X))$, so that $V = \bigcup_{k \geq 0} V_k$. Then a memoryless limit-spoiling strategy for the safety objective $\Box U$ can be obtained by playing at each $s \in V_k$ the distribution $\xi^{\mathrm{Epre}}_{s,2}(\bigcup_{i=0}^{k-1} V_i)$, for all $k > 0$. The distributions played at $V_0 = \neg U$ and $S \setminus V$ are irrelevant. ∎

From Lemma 2 and Lemma 3, we obtain the following theorem, which summarizes the results on safety games. See [dAHK07] for an example to illustrate that limit-spoiling strategies require randomization.

Theorem 1 [dAHK07] *For all concurrent game structures, for all safety objectives $\Box U$ for player 1, where $U \subseteq S$, the following assertions hold.*

1. *We have $\langle\langle 1 \rangle\rangle_{limit}\big(\Box U\big) = \nu X \,.\, [Pre_1(X) \cap U]$.*
2. *The set of limit-winning states can be computed in linear time in the size of the game.*
3. *Deterministic memoryless limit-winning strategy exists for player 1.*
4. *The most restrictive class of strategies in which limit-spoiling strategies are guaranteed to exist for player 2 is the class of memoryless strategies.*

4 Reachability Games

In this section we consider reachability games with objective $\Diamond U$ for player 1. We will present symbolic algorithm to compute limit-winning states for player 1. To solve these games, we introduce new predecessor operators.

4.1 Predecessor operators for reachability games

To solve reachability games, we define two new operators, Lpre and Fpre. For $s \in S$ and $X, Y \subseteq S$, these two-argument predecessor operators are defined as follows:

$$\mathrm{Lpre}_1(Y, X) = \{s \in S \mid \forall \alpha > 0 \,.\, \exists \xi_1 \in \chi_1^s \,.\, \forall \xi_2 \in \chi_2^s \,.\, \big[P_s^{\xi_1,\xi_2}(X) > \alpha \cdot P_s^{\xi_1,\xi_2}(\neg Y)\big]\}; \tag{3}$$

$$\mathrm{Fpre}_2(X, Y) = \{s \in S \mid \exists \beta > 0 \,.\, \exists \xi_2 \in \chi_2^s \,.\, \forall \xi_1 \in \chi_1^s \,.\, \big[P_s^{\xi_1,\xi_2}(Y) \geq \beta \cdot P_s^{\xi_1,\xi_2}(\neg X)\big]\} \,. \tag{4}$$

The operator $\mathrm{Lpre}_1(Y, X)$ states that player 1 can choose distributions to ensure that the probability to progress to X can be made arbitrarily large as compared to the probability of escape from Y. The operator $\mathrm{Fpre}_2(X, Y)$ states that player 2 can choose distributions to ensure that the probability to progress to Y can be made greater than a positive constant times the probability of escape from X. The definitions of the predecessor operators Lpre_2 and Fpre_1 can be obtained from these definitions simply by exchanging the subscripts 1 and 2.

The above definitions (3) and (4) is not computational. We now present symbolic algorithms to compute $\mathrm{Lpre}_1(Y, X)$ and $\mathrm{Fpre}_2(X, Y)$; these algorithms will also lead to a proof of the duality of these operators (Lemma 4). Given a state $s \in S$ and two subsets $X, Y \subseteq S$, to decide whether $s \in \mathrm{Lpre}_1(Y, X)$ we evaluate a μ-calculus expressions over the set $\Gamma_s = \Gamma_1(s) \cup \Gamma_2(s)$. For $X, Y \subseteq S$, $A \subseteq \Gamma_s$, $i \in \{1, 2\}$, and $s \in S$ we define two predicates, $\mathsf{Stay}_i(s, Y, A)$ and $\mathsf{Cover}_i(s, X, A)$ by:

$$\mathsf{Stay}_i(s, Y, A) = \{a \in \Gamma_i(s) \mid \forall b \in \Gamma_{\neg i}(s) \setminus A \,.\, \big[Dest(s, a, b) \subseteq Y\big]\} \tag{5}$$

$$\mathsf{Cover}_i(s, X, A) = \{b \in \Gamma_{\neg i}(s) \mid \exists a \in \Gamma_i(s) \cap A \,.\, Dest(s, a, b) \cap X \neq \emptyset\} \,; \tag{6}$$

where $\neg i = 2$ if $i = 1$, and $\neg i = 1$ if $i = 2$. The set $\mathsf{Stay}_i(s, Y, A) \subseteq \Gamma_i(s)$ consists of the set of player i moves a such that for all moves b for the other player that are not in A, the next state is in Y with probability 1. The set $\mathsf{Cover}_i(s, X, A) \subseteq \Gamma_{\neg i}(s)$ consists of the other player moves b such that there is a move a for player i in A such that the next state is in X with positive probability. The duals of the predicates are as follows:

$$\begin{aligned}
\neg\mathsf{Stay}_i(s, \neg Y, \neg A) &= \Gamma_{\neg i}(s) \cup \{a \in \Gamma_i(s) \mid \exists b \in \Gamma_{\neg i}(s) \cap A \,.\, \mathit{Dest}(s, a, b) \cap Y \neq \emptyset\} \\
&= \Gamma_{\neg i}(s) \cup \mathsf{Cover}_{\neg i}(s, Y, A) && (7)\\
\neg\mathsf{Cover}_i(s, \neg X, \neg A) &= \Gamma_i(s) \cup \{b \in \Gamma_{\neg i}(s) \mid \forall \Gamma_i(s) \setminus A \,.\, \mathit{Dest}(s, a, b) \subseteq X\} \\
&= \Gamma_i(s) \cup \mathsf{Stay}_{\neg i}(s, X, A). && (8)
\end{aligned}$$

Lemma 4 *Given $X \subseteq Y \subseteq S$, $s \in S$, the following assertions hold.*

1. $s \in \mathit{Lpre}_1(Y, X)$ *iff* $\Gamma_2(s) \subseteq \mu A \,.\, \big[\mathsf{Stay}_1(s, Y, A) \cup \mathsf{Cover}_1(s, X, A)\big]$; (9)
2. $s \in \mathit{Fpre}_2(\neg X, \neg Y)$ *iff* $\Gamma_2(s) \setminus \mu A \,.\, \big[\mathsf{Stay}_1(s, Y, A) \cup \mathsf{Cover}_1(s, X, A)\big] \neq \emptyset$; (10)
3. $\mathit{Fpre}_2(\neg X, \neg Y) = \neg \mathit{Lpre}_1(Y, X)$. (11)

Proof. We first prove that if $\Gamma_2(s) \subseteq \mu A \,.\, \big[\mathsf{Stay}_1(s, Y, A) \cup \mathsf{Cover}_1(s, X, A)\big]$, then $s \in \mathrm{Lpre}_1(Y, X)$. For all $0 < \varepsilon < 1$, we define the distribution $\xi^{\mathrm{Lpre}}_{s,1}[\varepsilon](Y, X) \in \chi^s_1$ as the distribution that plays each move $a \in \Gamma_1(s)$ with probability proportional to ε^j, where j is the level of a in $\mu A \,.\, [\mathsf{Stay}_1(s, Y, A) \cup \mathsf{Cover}_1(s, X, A)]$ (the details of the construction are available in [dAHK07]). Since $\Gamma_2(s) \subseteq \mu A \,.\, [\mathsf{Stay}_1(s, Y, A) \cup \mathsf{Cover}_1(s, X, A)]$, for every move $b \in \Gamma_2(s)$ we denote by $\ell(b)$ the level of b in the μ-calculus expression. Consider a move b with $\ell(b) = k$. Given player 2 plays move b we have:

- the probability of going to X is proportional to at least ε^{k-1}, since for some move $a \in \Gamma_1(s)$ of level at most $k - 1$ we have $\mathit{Dest}(s, a, b) \cap X \neq \emptyset$; and
- the probability of leaving Y is at most proportional to ε^k, since for all moves $a \in \Gamma_1(s)$ of level $k - 1$ we have $\mathit{Dest}(s, a, b) \subseteq Y$.

It follows that for all distributions $\xi_2 \in \chi^s_2$, the ratio of the probability of going to X as compared to leaving Y is proportional to at least $1/\varepsilon$. Since $\varepsilon > 0$ is arbitrary, the result follows. Also observe that as $\varepsilon \to 0$, the distribution converges to a distribution with support $\mathsf{Stay}_1(s, Y, \emptyset)$. Given $\varepsilon > 0$, the construction also ensures that X is reached with probability at least ε^M, where $M = |\Gamma_s|$.

To complete the proof we show that if $\Gamma_2(s) \setminus \mu A \,.\, [\mathsf{Stay}_1(s, Y, A) \cup \mathsf{Cover}_1(s, X, A)] \neq \emptyset$, then $\mathrm{Fpre}_2(\neg X, \neg Y)$ can be satisfied. Let $A^* = \mu A \,.\, [\mathsf{Stay}_1(s, Y, A) \cup \mathsf{Cover}_1(s, X, A)]$ and $B = \Gamma_s \setminus A^*$. Note that $B \cap \Gamma_2(s) \neq \emptyset$. Then the following assertions hold: (a) for all $a \in A^* \cap \Gamma_1(s)$ and for all $b \in B \cap \Gamma_2(s)$ we have $\mathit{Dest}(s, a, b) \cap X = \emptyset$, i.e., $\mathit{Dest}(s, a, b) \subseteq \neg X$; and (b) for all $a \in B \cap \Gamma_1(s)$, there exists $b \in B \cap \Gamma_2(s)$ such that $\mathit{Dest}(s, a, b) \cap \neg Y \neq \emptyset$. We define the distribution $\xi^{\mathrm{Fpre}}_{s,2}(\neg X, \neg Y) \in \chi^s_2$ as the distribution that plays uniformly at random all the moves in $B \cap \Gamma_2(s)$. The distribution $\xi^{\mathrm{Fpre}}_{s,2}(\neg X, \neg Y) \in \chi^s_2$ is a witness of ξ_2 in the characterization of (4) with $\beta = \frac{1}{|B|} \cdot \min\{\delta(s, a, b)(t) \mid s \in S, a \in \Gamma_1(s), b \in \Gamma_2(s), t \in \mathit{Dest}(s, a, b)\}$. ∎

Remark 1 *The proof of Lemma 4 can be easily extended as follows: given $X \subseteq Y \subseteq S$, and $s \in S$, let $C = \mu A \,.\, \big[\mathsf{Stay}_1(s, Y, A) \cup \mathsf{Cover}_1(s, X, A)\big]$; then for all $\alpha > 0$, there exists $\xi_1 \in \chi^s_1$ such that for all $\xi_2 \in \chi^s_2$ with $\mathrm{Supp}(\xi_2) \subseteq C$ we have $\big[P^{\xi_1,\xi_2}_s(X) > \alpha \cdot P^{\xi_1,\xi_2}_s(\neg Y)\big]$. In other words, if player 2's moves are restricted to $C \cap \Gamma_2(s)$, then player 1 can ensure that $s \in \mathit{Lpre}_1(Y, X)$.*

The algorithm for the computation of $\mathrm{Lpre}_1(Y, X)$ is illustrated in Figure 2.

Example 1 A state s is called an *absorbing state* in a concurrent game graph if for all $a_1 \in \Gamma_1(s)$ and $a_2 \in \Gamma_2(s)$ we have $\delta(s, a_1, a_2)(s) = 1$. In other words, at s for all choices of moves of the players the next state is always s.

Consider the game shown in Fig. 3, originally due to [KS81]. The states s_1 and s_2 are absorbing states. The transition function is defined as follows:

$$\delta(s_0, a, c)(s_0) = 1; \qquad \delta(s_0, b, d)(s_2) = 1; \qquad \delta(s_0, a, d)(s_1) = \delta(s_0, b, c)(s_1) = 1.$$

The objective of player 1 is to reach s_1, i.e., $\Diamond\{s_1\}$. For $\varepsilon > 0$, consider the memoryless strategy $\pi_\varepsilon \in \Pi_1^M$ that plays move a with probability $1-\varepsilon$, and move b with probability ε. The game starts at s_0, and in each round if player 2 plays move c, then the play reaches s_1 with probability ε and stays in s_0 with probability $1-\varepsilon$; whereas if player 2 plays move d, then the game reaches state s_1 with probability $1-\varepsilon$ and state s_2 with probability ε. Hence it is easy to argue that player 1, using strategy π_ε, can reach s_1 with probability at least $1-\varepsilon$, regardless of the strategy adopted by player 2. Hence for all $\varepsilon > 0$, there exists a strategy π_ε for player 1, such that against all strategies π_2, we have $\mathrm{Pr}_{s_0}^{\pi_\varepsilon,\pi_2}(\Diamond\{s_1\}) \geq 1-\varepsilon$; and thus $s_0 \in \langle\langle 1\rangle\rangle_{limit}\big(\Diamond\{s_1\}\big)$. Thus, player 1 can win with probability arbitrarily close to 1.

To see that player 1 cannot win with probability 1, consider the strategy π_2 obtained by the initial randomization of the two memoryless strategies: π_2^1 that plays move c and d with probability $1/2$, and π_2^2 that plays move c deterministically. Assume player 2 plays according to π_2, and consider any strategy π_1 for player 1. If the strategy π_1 always deterministically plays a at s_0, then the game never reaches s_1. If at any round j, at s_0 the strategy π_1 plays b with positive probability, then the play reaches s_2 with positive probability. Hence it follows that player 1 cannot ensure that $\Diamond\{s_1\}$ is satisfied with probability 1. ∎

The following technical lemma will lead to the algorithm for reachability games, and will play a key role in several of the arguments.

Lemma 5 (Basic Lpre principle) *Let $X \subseteq S$, $Y = X \cup \{s\}$, and $Y \subseteq Z \subseteq S$. Let $s \in \mathrm{Lpre}_1(Z, X)$. For all events $\mathcal{A} \subseteq \Box(Z \setminus Y)$, the following assertion holds:*

> *Assume that for all $\eta > 0$ there exists $\pi_1^\eta \in \Pi_1^M$ (resp. $\pi_1^\eta \in \Pi_1$) such that for all $\pi_2 \in \Pi_2$ and for all $z \in Z \setminus Y$ we have*
>
> $$\mathrm{Pr}_z^{\pi_1^\eta,\pi_2}(\mathcal{A} \cup \Diamond Y) \geq 1-\eta, \qquad (i.e., \lim_{\eta\to 0} \mathrm{Pr}_z^{\pi_1^\eta,\pi_2}(\mathcal{A} \cup \Diamond Y) = 1).$$
>
> *Then, for all $\varepsilon > 0$ there exists $\pi_1^\varepsilon \in \Pi_1^M$ (resp. $\pi_1^\varepsilon \in \Pi_1$) such that for all $\pi_2 \in \Pi_2$ we have*
>
> $$\mathrm{Pr}_s^{\pi_1^\varepsilon,\pi_2}(\mathcal{A} \cup \Diamond X) \geq 1-\varepsilon, \qquad (i.e., \lim_{\varepsilon\to 0} \mathrm{Pr}_z^{\pi_1^\varepsilon,\pi_2}(\mathcal{A} \cup \Diamond X) = 1).$$

Proof. The situation is depicted in Fig 4.(a). Since $s \in \mathrm{Lpre}_1(Z, X)$, given $\varepsilon > 0$, player 1 can play the distribution $\xi_{s,1}^{\mathrm{Lpre}}[\varepsilon](Z, X)$ to ensure that the probability of going to $\neg Z$ is at most ε times the probability to going to X; we denote the probabilities as $\gamma \cdot \varepsilon$ and γ, respectively. Observe that $\gamma > \varepsilon^l$, where $l = |\Gamma_s|$. Let α denote the probability of the event $\mathcal{A}$ and since $\mathcal{A} \subseteq \mathcal{A} \cup \Diamond X$, the worst-case analysis for the result correspond to the case when $\alpha = 0$, and the simplified situation is

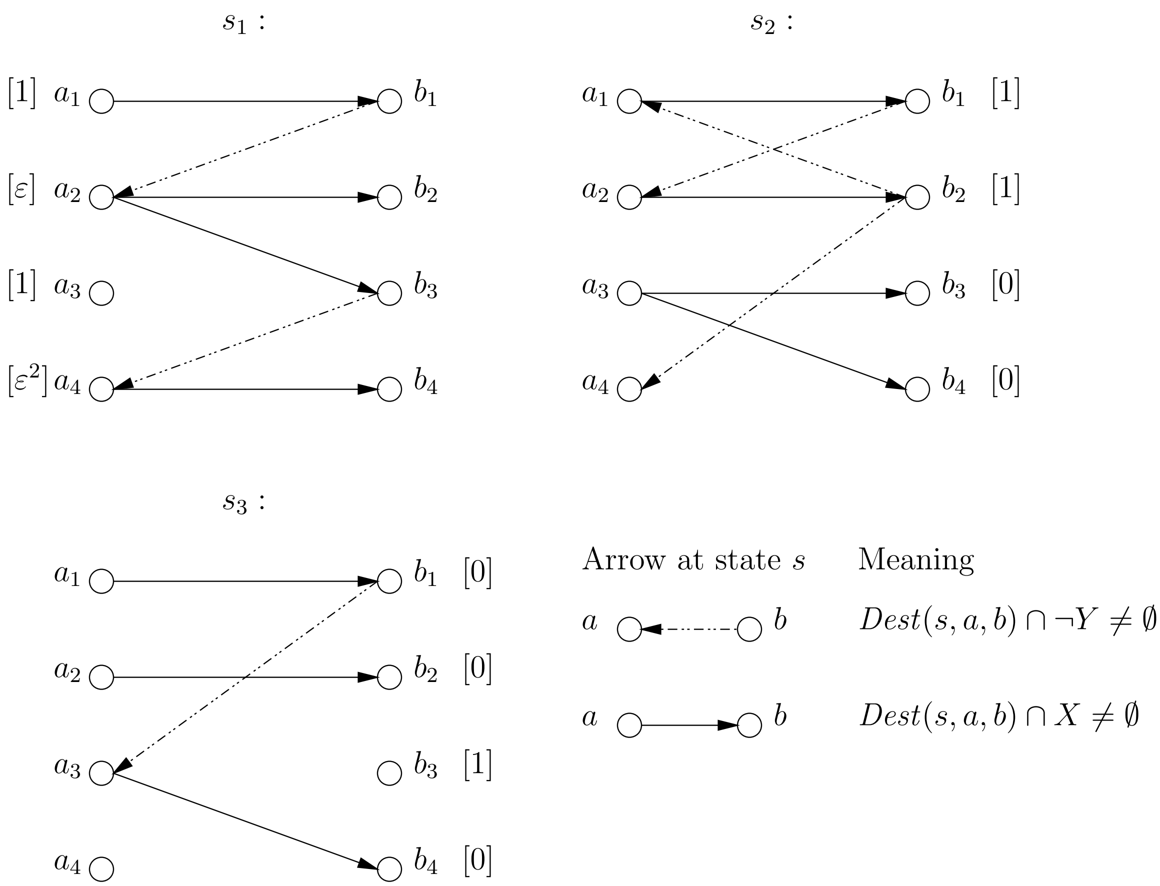

Figure 2: Example of computation of predicate $\mathrm{Lpre}_1(Y, X)$. The algorithm for deciding $\mathrm{Lpre}_1(Y, X)$ works as follows. Initially, all moves are marked "uncovered", and $i = 0$. Then, for $i \geq 0$, the two following steps are applied alternatively, until no more moves can be covered:
(a) all moves of player 1 that have incoming edges only from already covered moves are marked as covered;
(b) all moves of player 2 that have at least one incoming edge from a covered move are marked as covered.
A state satisfies $\mathrm{Lpre}_1(Y, X)$ if all the moves of player 2 can be marked as covered.
In the figure, at state s_1, it is possible to mark as covered all moves of player 2, and we have $s_1 \in \mathrm{Lpre}_1(Y, X)$. In this case, we also label each move of player 1 with ε^i, where i is the round (beginning from 0) at which it has been covered. These labels can be used to construct a family of distributions for player 1 that satisfy (3).
At state s_2, we can only cover moves b_3 and b_4 of player 2; hence, $s_2 \notin \mathrm{Lpre}_1(Y, X)$. Similarly, at s_3 we cannot cover move b_3, hence $s_3 \notin \mathrm{Lpre}_1(Y, X)$. In these cases, we have $s_2 \in \mathrm{Fpre}_2(\neg X, \neg Y)$ and $s_3 \in \mathrm{Fpre}_2(\neg X, \neg Y)$. If we label with 0 the moves of player 2 that have been covered, and with 1 those that have not, we can use the resulting labels to obtain a distribution for player 2 satisfying the characterization for Fpre_2 obtained by exchanging the roles of players 1 and 2 in (4). The labels associated with the moves are called *move weights.*

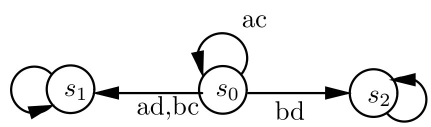

Figure 3: Reachability games

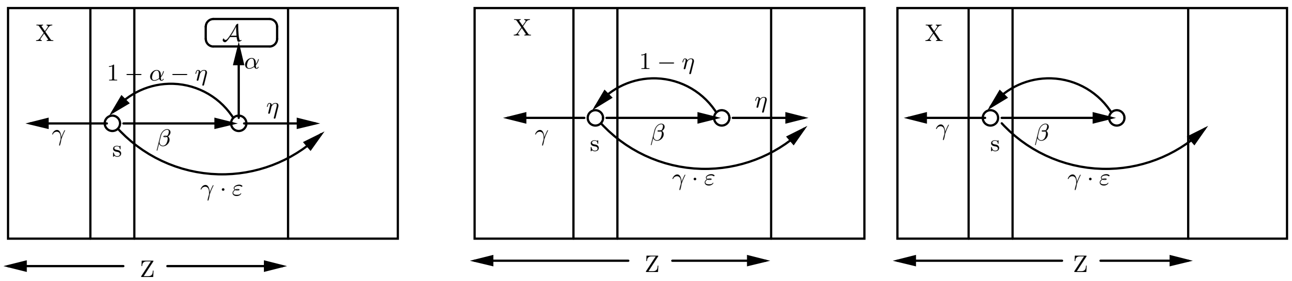

(a) (b) (c)

Figure 4: Basic Lpre principle; in the figures $\beta = 1 - \gamma - \gamma \cdot \varepsilon$

shown as Fig 4.(b). We first present an informal argument and then present rigorous calculations. Once we let $\eta \to 0$, then we only have an edge from $Z \setminus Y$ to Y and the situation is shown in Fig 4.(c). If q is the probability to reach X, then the probability to reach $\neg Z$ is $q \cdot \varepsilon$ and we have $q + q\varepsilon = 1$, i.e., $q = \frac{1}{1+\varepsilon}$, and given $\varepsilon' > 0$ we can chose ε to ensure that $q \geq 1 - \varepsilon'$.

We now present detailed calculations. Given $\varepsilon' > 0$ we construct a strategy $\pi_1^{\varepsilon'}$ as follows: let $\varepsilon = \frac{\varepsilon'}{2(1-\varepsilon')}$ and $\eta = \varepsilon^{l+1} > 0$; and fix the strategy π_1^{η} for states in $Z \setminus Y$ and the distribution $\xi_{s,1}^{\text{Lpre}}[\varepsilon](Z, X)$ at s. Observe that by choice we have $\eta \leq \gamma \cdot \varepsilon$. Let $q = \Pr_s^{\pi_1^{\varepsilon'},\pi_2}(\mathcal{A} \cup \Diamond X)$. Then we have $q \geq \gamma + \beta \cdot (1-\eta) \cdot q$; since the set $Z \setminus Y$ is reached with probability at most β and then again Y is reached with probability at least $1 - \eta$. Thus we have

$$
\begin{array}{rcll}
q & \geq & \gamma + (1-\gamma-\gamma\cdot\varepsilon)\cdot(1-\eta)\cdot q; & \\
q & \geq & \dfrac{\gamma}{\gamma + \gamma\cdot\varepsilon + \eta - \eta\cdot\gamma - \eta\cdot\gamma\cdot\varepsilon} & \\
 & \geq & \dfrac{\gamma}{\gamma + \gamma\cdot\varepsilon + \eta} & \\
 & \geq & \dfrac{\gamma}{\gamma + \gamma\cdot\varepsilon + \gamma\cdot\varepsilon} & \quad (\text{since } \eta \leq \gamma\cdot\varepsilon) \\
 & \geq & \frac{1}{1+2\varepsilon} \geq 1 - \varepsilon'. &
\end{array}
$$

The desired result follows. ∎

4.2 Winning sets for reachability games

The characterizations of the predecessor operators given above lead to algorithms for the symbolic computation of the limit-winning states of reachability games.

Lemma 6 *For all $U \subseteq S$ we have*

$$\nu Y \,.\, \mu X \,.\, [Lpre_1(Y,X) \cup U] \subseteq \langle\langle 1 \rangle\rangle_{limit}(\Diamond U). \tag{12}$$

Proof. We exhibit limit-winning strategies as follows. Let $W = \nu Y \,.\, \mu X \,.\, [\mathrm{Lpre}_1(Y,X) \cup U]$, so that $W = \mu X \,.\, [\mathrm{Lpre}_1(W,X) \cup U]$. The computation of W can be obtained as a sequence $T_0 \subset T_1 \subset T_2 \subset \cdots \subset T_{m-1} \subset T_m = W$ such that $T_0 = U$ and for $i \geq 1$ we have $T_{i+1} \setminus T_i = \{s\}$ and $s \in \mathrm{Lpre}_1(W,T_i)$. By induction on i, from $i = m$ down to $i = 1$, we can show that for $s \in T_{i+1} \setminus T_i$, for all $\varepsilon > 0$, there exists $\pi_1^{\varepsilon} \in \Pi_1^M$ such that for all π_2, we have $\mathrm{Pr}_s^{\pi_1^{\varepsilon},\pi_2}(\Diamond T_i) \geq 1 - \varepsilon$. The result is an application of the basic Lpre principle (Lemma 5), with $Z = W$, $X = T_i$, $Y = T_{i+1}$ and $\mathcal{A} = \emptyset$. It is easy to show that for all $\varepsilon > 0$, there exists $\pi_1^{limit}[\varepsilon] \in \Pi_1^M$ such that for all strategies π_2 for player 2, for all states $s \in W$, we have $\mathrm{Pr}_s^{\pi_1^{limit}[\varepsilon],\pi_2}(\Diamond U) \geq 1 - \varepsilon$. The result follows. ∎

We complement the mu-calculus expression (12) and exhibit spoiling strategies for player 2 for the complementary set. Note that by Lemma 4 we have $\neg\mathrm{Lpre}_1(\neg Y, \neg X) = \mathrm{Fpre}_2(X,Y)$.

Lemma 7 *For all $U \subseteq S$ we have*

$$\begin{aligned} \neg \nu Y \,.\, \mu X \,.\, [Lpre_1(Y,X) \cup U] &= \mu Y \,.\, \nu X \,.\, [(\neg Lpre_1(\neg Y, \neg X)) \cap \neg U] \\ &= \mu Y \,.\, \nu X \,.\, [Fpre_2(X,Y) \cap \neg U] \subseteq \neg \langle\langle 1 \rangle\rangle_{limit}(\Diamond U). \end{aligned} \tag{13}$$

Proof. We define the spoiling strategies for limit reachability as follows. We have $\neg W = \mu Y \,.\, \nu X \,.\, (\mathrm{Fpre}_2(X,Y) \cap \neg U)$, from (13). The set $\neg W$ is obtained as follows: $Y_0 \subset Y_1 \subset Y_2 \cdots \subset Y_m = \neg W$ such that for $Y_0 = \emptyset$ and for all $i \geq 0$, for all $s \in Y_{i+1} \setminus Y_i$ we have $s \in \mathrm{Fpre}_2(Y_{i+1}, Y_i)$. A spoiling strategy for player 2 can be obtained by playing, at all $k \geq 1$ and $s \in Y_k$, with distribution $\xi_{s,2}^{\mathrm{Fpre}}(Y_k, Y_{k-1})$. The strategy ensures for all $k \geq 1$, and $s \in Y_k$, that against all strategies π_1, either Y_k is not left or else Y_{k-1} is reached with positive bounded probability. This proves that the strategy constructed is a spoiling strategy. The detailed proofs that the above strategy is a spoiling strategy can be found in [dAHK07]. ∎

From Lemma 6 and 7 we obtain the following result.

Theorem 2 *For all concurrent game structures, for all reachability objectives $\Diamond U$ for player 1, where $U \subseteq S$, the following assertions hold.*

1. *We have*
$$\langle\langle 1 \rangle\rangle_{limit}(\Diamond U) = \nu Y \,.\, \mu X \,.\, [Lpre_1(Y,X) \cup U] \tag{14}$$
2. *The sets of limit-winning states can be computed using the relations (14) in quadratic time.*
3. *The most restrictive class of strategies in which limit-winning strategies are guaranteed to exist for player 1 is the class of memoryless strategies.*
4. *The most restrictive class of strategies in which limit-spoiling strategies are guaranteed to exist for player 2 is the class of memoryless strategies.*

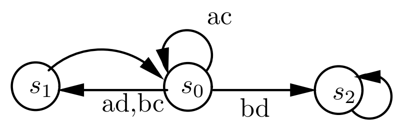

Figure 5: Büchi games

5 Büchi Games

In this section we consider Büchi games, where the objective for player 1 is $\Box\Diamond B_2$ for $B_2 \subseteq S$. We first present an example that shows that in case of Büchi games, limit-winning strategies may require infinite memory.

Example 2 Consider the game shown in Fig. 5. The transition function at state s_0 is same the as the one in Fig 3. The state s_2 is an absorbing state, and from the state s_1 the next state is always s_0. The objective of player 1 is to visit s_1 infinitely often, i.e., $\Box\Diamond\{s_1\}$. For $\varepsilon > 0$, we construct a strategy π_ε as follows: for $i \geq 0$, construct a sequence of ε_i, such that $\varepsilon_i > 0$, and $\prod_i (1 - \varepsilon_i) \geq (1 - \varepsilon)$. At state s_0, between the i-th and the $i+1$-th visits to s_1, we use a strategy that reaches s_1 with probability $1 - \varepsilon_i$; such a strategy can be constructed as in the solution of reachability games (see the discussion for Fig. 3 and, for a more rigorous treatment, [dAHK07]). The overall strategy π_ε constructed in this fashion ensures that against any strategy π_2, the state s_1 is visited infinitely often with probability $1 - \varepsilon$. However, the strategy π_ε needs to count the number of visits to s_1, and therefore requires infinite memory. ∎

The following lemma shows that the fact that the infinite memory requirement, in general, cannot be avoided.

Lemma 8 *Limit-winning strategies for Büchi games may require infinite memory.*

Proof. Consider again the game described in Example 2 and illustrated in Fig. 5. Example 2 shows the existence of an infinite-memory limit-winning strategy. We show now that all finite-memory strategies visit s_2 infinitely often with probability 0. Let π_1 be an arbitrary finite-memory strategy for player 1, and let M be the (finite) memory set used by the strategy. On the set $\{s_0, s_1, s_2\} \times M$, strategy π_1 is memoryless. Consider now a strategy π_2 for player 2 constructed as follows. From a state $(s_0, m) \in \{s_0, s_1, s_2\} \times M$, if player 1 plays a with probability 1, then player 2 plays c with probability 1, ensuring that the successor is (s_0, m') for some $m' \in M$. If player 1 plays b with positive probability, then player 2 plays c and d uniformly at random, ensuring that (s_2, m') is reached with positive probability, for some $m' \in M$. Under π_1, π_2 the game is reduced to a Markov chain, and since the set $\{s_2\} \times M$ is absorbing, and since all states in $\{s_0\} \times M$ either stay safe in $\{s_0\} \times M$ or reach $\{s_2\} \times M$ in one step with positive probability, and all states in $\{s_1\} \times M$ reaches $\{s_0\} \times M$ in one step, the closed recurrent classes must be either entirely contained in $\{s_0\} \times M$, or in $\{s_2\} \times M$. This shows that, under π_1, π_2, player 1 achieves the Büchi goal $\Box\Diamond\{s_1\}$ with probability 0. ∎

We now present symbolic algorithms to compute limit-winning states in Büchi games using the same predecessor operators of reachability games. The algorithms are based on the *basic principle of repeated reachability.*

Basic principle of repeated reachability. We say that an objective is *infinitary* if it is independent of all finite prefixes. Formally, an objective $\mathcal{A}$ is infinitary if, for all $u, v \in S^*$ and $\omega \in S^\omega$, we have $u\omega \in \mathcal{A}$ iff $v\omega \in \mathcal{A}$. Observe that parity objectives are defined based on the states that appear infinitely often along a play, and hence independent of all finite prefixes, so that, parity objectives are infinitary objectives.

Lemma 9 *Given sets $T \subseteq S$, $B \subseteq S$, and an infinitary objective $\mathcal{A}$, let*

$$W \subseteq \langle\langle 1 \rangle\rangle_{limit}\big(\Diamond T \cup \Diamond(B \cap Pre_1(W)) \cup \mathcal{A}\big).$$

Then

$$W \subseteq \langle\langle 1 \rangle\rangle_{limit}\big(\Diamond T \cup \Box\Diamond B \cup \mathcal{A}\big).$$

Proof. Let $Z = B \cap \mathrm{Pre}_1(W)$. For all states $s \in W \setminus (B \cup T)$, for all $\varepsilon > 0$, there is a player 1 strategy $\pi_1^{limit}[\varepsilon]$ that ensures that against all player 2 strategies π_2 we have

$$\mathrm{Pr}_s^{\pi_1^{limit}[\varepsilon],\pi_2}\big(\Diamond(T \cup Z) \cup \mathcal{A}\big) \geq 1 - \varepsilon.$$

For all states in Z player 1 can ensure that the successor state is in W (since $\mathrm{Pre}_1(W)$ holds in Z). Given $\varepsilon > 0$, fix a sequence $\varepsilon_1, \varepsilon_2, \ldots$ such that for all $i \geq 1$ we have $\varepsilon_i > 0$ and $\prod_{i=1}^{\infty}(1-\varepsilon_i) \geq (1-\varepsilon)$ (for example let $(1-\varepsilon_i) = (1-\varepsilon)^{\frac{1}{2^i}}$, i.e., $\varepsilon_i = 1-(1-\varepsilon)^{\frac{1}{2^i}}$). Consider a strategy π_1^* as follows: for states $s \in Z$ play a memoryless strategy for $\mathrm{Pre}_1(W)$ to ensure that the next state is in W; for states $s \in W \setminus (Z \cup T)$ play a strategy $\pi_1^{limit}[\varepsilon_{j+1}]$ between the j-th and $j+1$-th visit to $Z \cup T$. Let us denote by $\Diamond_k Z \cup \Diamond T$ to be the set of paths that visits Z at least k-times or visits T at least once. Observe that $\lim_{k\to\infty}\big(\Diamond_k Z \cup \Diamond T\big) \subseteq \Box\Diamond B \cup \Diamond T$. Hence for all $s \in W$ and for all $\pi_2 \in \Pi_2$ we have

$$\begin{aligned}
\mathrm{Pr}_s^{\pi_1^*,\pi_2}(\Box\Diamond B \cup \Diamond T \cup \mathcal{A}) \;&\geq\; \mathrm{Pr}_s^{\pi_1^*,\pi_2}\big(\Diamond Z \cup \Diamond T \cup \mathcal{A}\big) \cdot \prod_{k=1}^{\infty} \mathrm{Pr}_s^{\pi_1^*,\pi_2}\big(\Diamond_{k+1} Z \cup \Diamond T \cup \mathcal{A} \mid \Diamond_k Z \cup \Diamond T \cup \mathcal{A}\big) \\
&=\; \mathrm{Pr}_s^{\pi_1^{limit}[\varepsilon_1],\pi_2}\big(\Diamond Z \cup \Diamond T \cup \mathcal{A}\big) \\
&\quad\; \cdot \textstyle\prod_{k=1}^{\infty} \mathrm{Pr}_s^{\pi_1^{limit}[\varepsilon_{k+1}],\pi_2}\big(\Diamond_{k+1} Z \cup \Diamond T \mid \Diamond_k Z \cup \Diamond T \cup \mathcal{A}\big) \\
&\geq\; \textstyle\prod_{k=1}^{\infty}(1-\varepsilon_k) \geq 1 - \varepsilon.
\end{aligned}$$

Hence we have $\mathrm{Pr}_s^{\pi_1^*,\pi_2}(\Box\Diamond B \cup \Diamond T \cup \mathcal{A}) \geq 1 - \varepsilon$, for all $s \in W$ and for all $\pi_2 \in \Pi_2$. Since $\varepsilon > 0$ is arbitrary, it follows that $W \subseteq \langle\langle 1 \rangle\rangle_{limit}\big(\Box\Diamond B \cup \Diamond T \cup \mathcal{A}\big)$. ∎

Lemma 10 *For $T \subseteq S$, $B_2 \subseteq S$, and $B_1 = S \setminus B_2$, we have*

$$\nu Y_0 \,.\, \mu X_0 \,.\, \big[T \cup (B_1 \cap Lpre_1(Y_0, X_0)) \cup (B_2 \cap Pre_1(Y_0))\big] \subseteq \langle\langle 1 \rangle\rangle_{limit}\big(\Box\Diamond B_2 \cup \Diamond T\big) \;. \tag{15}$$

Proof. Let

$$W = \nu Y_0 \,.\, \mu X_0 \,.\, \big[T \cup (B_1 \cap \mathrm{Lpre}_1(Y_0, X_0)) \cup (B_2 \cap \mathrm{Pre}_1(Y_0))\big].$$

Then we can rewrite W as follows:

$$W = \nu Y_0 \,.\, \mu X_0 \,.\, \big[T \cup (B_2 \cap \mathrm{Pre}_1(W)) \cup (B_1 \cap \mathrm{Lpre}_1(Y_0, X_0))\big].$$

The result follows from an application of Lemma 9. By the correctness of limit-reachability (Lemma 6), we treat $T \cup (B_2 \cup \mathrm{Pre}_1(W))$ as the target set U for reachability and obtain that that $W \subseteq \langle\langle 1 \rangle\rangle_{limit}\big(\Diamond T \cup (B_2 \cap \mathrm{Pre}_1(W)\big)$. The result then follows from Lemma 9, with $B = B_2$ and $\mathcal{A} = \emptyset$. Also observe that the witness strategy $\pi_1^{limit}[\varepsilon]$, for $\varepsilon > 0$, for Lemma 9, in the present case can be memoryless strategy to ensure that $T \cup (B_2 \cap \mathrm{Pre}_1(W))$ is reached with probability at least $1 - \varepsilon$ (the construction of such a strategy is similar to Lemma 6). ∎

We note that since $\mathrm{Fpre}_2(X_0, Y_0) = \neg\mathrm{Lpre}_1(\neg Y_0, \neg X_0)$ (by Lemma 4) and $\mathrm{Epre}_2(Y_0) = \neg\mathrm{Pre}_1(\neg Y_0)$ (by Lemma 1), we have:

$$\begin{aligned}
&\neg\ \nu Y_0 \,.\, \mu X_0 \,.\, \big[(B_1 \cap \mathrm{Lpre}_1(Y_0, X_0)) \cup (B_2 \cap \mathrm{Pre}_1(Y_0))\big] \\
&= \mu Y_0 \,.\, \nu X_0 \,.\, \big[(B_2 \cup \mathrm{Fpre}_2(X_0, Y_0)) \cap (B_1 \cup \mathrm{Epre}_2(Y_0))\big].
\end{aligned} \tag{16}$$

The following lemma complements the result of Lemma 10.

Lemma 11 *For $B_2 \subseteq S$ and $B_1 = S \setminus B_2$ we have*

$$\begin{aligned}
&\mu Y_0 \,.\, \nu X_0 \,.\, \big[(B_2 \cup \mathit{Fpre}_2(X_0, Y_0)) \cap (B_1 \cup \mathit{Epre}_2(Y_0))\big] \\
&= \mu Y_0 \,.\, \nu X_0 \,.\, \big[(B_1 \cap \mathit{Fpre}_2(X_0, Y_0)) \cup (B_2 \cap \mathit{Epre}_2(Y_0))\big] \\
&\subseteq \neg\langle\langle 1 \rangle\rangle_{limit}\big(\Box\Diamond B_2\big)\ .
\end{aligned} \tag{17}$$

Proof. We exhibit the existence of memoryless spoiling strategies. Let

$$V = \mu Y_0 \,.\, \nu X_0 \,.\, \big[(B_1 \cap \mathrm{Fpre}_2(X_0, Y_0)) \cup (B_2 \cap \mathrm{Epre}_2(Y_0))\big].$$

Then we analyze the computation of V as follows: the set V is obtained as a increasing sequence $\emptyset = T_0 \subseteq T_1 \subseteq T_2 \ldots \subseteq T_m = V$ of states, such that the states in $T_{i+1} \setminus T_i$ are obtained as follows: (a) either a set of B_2 states such that $\mathrm{Epre}_2(T_i)$ holds; or (b) a set of B_1 states such that $\mathrm{Fpre}_2(T_{i+1}, T_i)$ holds. For state $s \in (T_{i+1} \setminus T_i) \cap B_2$ the distribution $\xi_{s,2}^{\mathrm{Epre}}(T_i)$ ensures that T_i is reached with positive probability. For state $s \in (T_{i+1} \setminus T_i) \cap B_1$ the distribution $\xi_{s,2}^{\mathrm{Fpre}}(T_{i+1}, T_i)$ ensures that (i) either T_i is reached with positive probability, or (ii) the game stays in $T_{i+1} \setminus T_i$, leading to $\Box(T_{i+1} \setminus T_i)$ and thus $\Diamond\Box B_1$. It follows that for all $s \in T_{i+1}$, there is a memoryless strategy π_2 such that for all π_1 we have $\mathrm{Pr}_s^{\pi_1,\pi_2}(\Diamond T_i \cup \Diamond\Box B_1) \geq q > 0$ for some $q > 0$. By induction, we obtain that for all $s \in T_{i+1}$ there is a memoryless strategy π_2 such that for all π_1 we have $\mathrm{Pr}_s^{\pi_1,\pi_2}(\Diamond T_0 \cup \Box\Diamond B_1) \geq q > 0$ for some $q > 0$. Since $T_0 = \emptyset$, with $T_{i+1} = V$ we obtain that $V \subseteq \neg\langle\langle 1 \rangle\rangle_{limit}\big(\Box\Diamond B_2\big)$ ∎

Strategy constructions. Note that the strategies constructed in Lemma 10 require in general infinite memory (for counting the number of visits to $W \cap B_2$) for limit-winning. In the witness strategies for limit-winning, the moves that are played with positive bounded probabilities at a state s are those in the set $\mathsf{Stay}_1(s, W, \emptyset)$; all other moves are played with probabilities that tend to 0 as $\varepsilon \to 0$. Hence, given any distribution $\xi_1 \in \chi_1^s$ such that $\mathrm{Supp}(\xi_1) = \mathsf{Stay}_1(s, W, \emptyset)$, there exists a sequence of limit-winning strategies that in the limit, as $\varepsilon \to 0$, converges to a memoryless strategy $\pi_1 \in \Pi_1^M$ such that for all $s \in W$ we have $\pi_1(s) = \xi_1$. However, the limit strategy π_1 is not limit-winning in general. We formalize this in the lemma below.

Lemma 12 *Given a Büchi objective $\Box\Diamond B_2$, let W be the set of limit-winning states. Let $\pi_1 \in \Pi_1^M$ be any memoryless strategy such that for all $s \in W$ we have $\mathrm{Supp}(\pi_1(s)) = \mathsf{Stay}_1(s, W, \emptyset)$. Then there exists a sequence of $(\varepsilon_i)_{i\geq 0}$ such that (a) $\varepsilon_0 > \varepsilon_1 > \varepsilon_2 > \ldots$ and $\lim_{i\to\infty} \varepsilon_i = 0$, (b) for all $i \geq 0$ there exists a strategy $\pi_1^{\varepsilon_i}$ such that for all $\pi_2 \in \Pi_2$ and for all $s \in W$ we have $\mathrm{Pr}_s^{\pi_1^{\varepsilon_i},\pi_2}(\Box\Diamond B_2) \geq 1-\varepsilon_i$, and (c) the strategies $\pi_1^{\varepsilon_i}$ converges to π_1 as $i \to \infty$, i.e., $\lim_{i\to\infty} \pi_1^{\varepsilon_i} = \pi_1$.*

Observe that by (16) the expression (15) complements to (17) with $T = \emptyset$. Thus from Lemma 10 and Lemma 11, along with the characterizations of the predecessor operators, we have the following theorem, that summarizes the results on Büchi games. The quadratic complexity is an immediate consequence of the quadratic complexity of reachability games [dAHK07].

Theorem 3 *For all concurrent game structures, for all Büchi objectives $\Box\Diamond B_2$ for player 1, where $B_2 \subseteq S$ and $B_1 = S \setminus B_2$, the following assertions hold.*

1. *We have*

$$\langle\!\langle 1 \rangle\!\rangle_{limit}\big(\Box\Diamond B_2\big) = \nu Y_0 \,.\, \mu X_0 \,.\, \big[(B_1 \cap Lpre_1(Y_0, X_0)) \cup (B_2 \cap Pre_1(Y_0))\big] \,. \tag{18}$$

2. *The set of limit-winning states can be computed using the relations (18) in time quadratic in the size of the game.*

3. *Limit-winning strategies for player 1 require infinite-memory in general.*

4. *The most restrictive class of strategies in which limit-spoiling strategies are guaranteed to exist for player 2 is the class of memoryless strategies.*

6 co-Büchi Games

The winning condition of a coBüchi game is a formula $\Diamond\Box B_0$, where $B_0 \subseteq S$ is a subset of states. To solve coBüchi games, we begin by introducing the required predecessor operators.

6.1 Predecessor operators in limit co-Büchi games

While the operators Lpre and Pre suffice for solving Büchi games, coBüchi games, and general parity games, can be solved using predecessor operators that are best understood as the combination of simpler predecessor operators. We use the operators $\uplus_*$ and $\Cap_*$ to combine predecessor operators; the operators $\uplus_*$ and $\Cap_*$ are different from the usual union $\cup$ and intersection $\cap$. Roughly, for two predecessor operators α and β, the predecessor operator $\alpha \Cap_* \beta$ requires that the distributions of player 1 and 2 satisfy the conjunction of the conditions stipulated by α and β; similarly, $\uplus_*$ corresponds to disjunction. We first introduce the operator Lpre $\uplus_*$ Pre. For all $s \in S$ and $X_1, Y_0, Y_1 \subseteq S$, we define

$$\mathrm{Lpre}_1(Y_1, X_1) \uplus_* \mathrm{Pre}_1(Y_0) = \left\{ s \in S \mid \forall \alpha > 0 \,.\, \exists \xi_1 \in \chi_1^s . \forall \xi_2 \in \chi_2^s . \left[\begin{array}{c} P_s^{\xi_1,\xi_2}(X_1) > \alpha \cdot P_s^{\xi_1,\xi_2}(\neg Y_1) \\ \bigvee \\ P_s^{\xi_1,\xi_2}(Y_0) = 1 \end{array} \right] \right\} .$$

Note that the above formula corresponds to a disjunction of the predicates for Lpre_1 and Pre_1. However, it is important to note that the distribution ξ_1 that player 1 needs to use to satisfy the predicate is *the same.* In other words, $\mathrm{Lpre}_1(Y_1, X_1) \uplus \mathrm{Pre}_1(Y_0)$ is *not* equivalent to $\mathrm{Lpre}_1(Y_1, X_1) \cup \mathrm{Pre}_1(Y_0)$, because in the latter union, player 1 could use one distribution to satisfy $\mathrm{Lpre}_1(Y_1, X_1)$, and a different one to satisfy $\mathrm{Pre}_1(Y_0)$, whereas in $\mathrm{Lpre}_1(Y_1, X_1) \uplus \mathrm{Pre}_1(Y_0)$, the same ξ_1 needs to satisfy both predicates.

Similarly, we introduce the operator $(\mathrm{Fpre}_2 \Cap \mathrm{Epre}_2) \uplus \mathrm{Lpre}_2$ as follows: for $s \in S$ and $X_1, Y_0, Y_1 \subseteq S$, we define

$$\big(\mathrm{Fpre}_2(X_1, Y_1) \Cap \mathrm{Epre}_2(Y_1)\big) \quad \uplus \quad \mathrm{Lpre}_2(X_1, Y_0) \tag{19}$$
$$= \left\{ s \in S \mid \exists \beta > 0. \forall \alpha \geq 0. \exists \xi_2 \in \chi_2^s. \forall \xi_1 \in \chi_1^s. \left[\begin{array}{c} \big(P_s^{\xi_1,\xi_2}(Y_1) \geq \beta \cdot P_s^{\xi_1,\xi_2}(\neg X_1) \\ \bigwedge \\ P_s^{\xi_1,\xi_2}(Y_1) > 0\big) \\ \bigvee \\ P_s^{\xi_1,\xi_2}(Y_0) > \alpha \cdot P_s^{\xi_1,\xi_2}(\neg X_1) \end{array} \right] \right\}.$$

To decide whether $s \in \mathrm{Lpre}_1(Y_1, X_1) \uplus \mathrm{Pre}_1(Y_0)$, for $s \in S$ and $X_1, Y_0, Y_1 \subseteq S$, we provide a mu-calculus expression over the set $\Gamma_s = \Gamma_1(s) \cup \Gamma_2(s)$. We also prove that the above predecessor operators are dual.

Computation of Lpre $\uplus$ Pre. We now give an algorithm for computing Lpre $\uplus$ Pre. Let $X_1 \subseteq Y_0 \subseteq Y_1 \subseteq S$ and $s \in S$. To understand the algorithm, recall the algorithm for $\mathrm{Lpre}_1(Y_1, X_1)$: we compute $C = \mu A. [\mathsf{Stay}_1(s, Y_1, A)) \cup \mathsf{Cover}_1(s, X_1, A)]$, and we require $\Gamma_2(s) \subseteq C$. To compute $\mathrm{Lpre}_1(Y_1, X_1) \uplus \mathrm{Pre}_1(Y_0)$, we add a νB quantifier, we add a constraint of $\mathsf{Stay}_1(s, Y_0, B)$, and we require that the set of moves obtained by the formula is non-empty. Formally, we consider the formula

$$C' = \nu B \,.\, \mu A \,.\, \left[\begin{array}{c} (\mathsf{Stay}_1(s, Y_1, A) \cap \mathsf{Stay}_1(s, Y_0, B)) \\ \cup \\ \mathsf{Cover}_1(s, X_1, A) \end{array} \right]$$

we have $C' \cap \Gamma_1(s) \neq \emptyset$ iff $s \in (\mathrm{Lpre}_1(Y_1, X_1) \uplus \mathrm{Pre}_1(Y_0))$. The basic intuition is as follows: if player 2 plays moves outside C', then the condition $\mathsf{Stay}_1(s, Y_0, B)$ ensures that Y_0 is not left, and as $Y_0 \subseteq Y_1$, also Y_1 is not left. Observe that $C' \subseteq C$, and by Remark 1 if we restrict player 2 to play moves in C', then player 1 can ensure $\mathrm{Lpre}_1(Y_1, X_1)$. Combining the above arguments we show that player 1 can ensure $\mathrm{Lpre}_1(Y_1, X_1) \uplus \mathrm{Pre}_2(Y_0)$. We now formally present the proof.

Lemma 13 *The following assertions hold.*

1. *For all $X_1 \subseteq Y_0 \subseteq Y_1 \subseteq S$ and $s \in S$, if*

$$\nu B \,.\, \mu A \,.\, \left[\begin{array}{c} (\mathsf{Stay}_1(s, Y_1, A) \cap \mathsf{Stay}_1(s, Y_0, B)) \\ \cup \\ \mathsf{Cover}_1(s, X_1, A) \end{array} \right] \cap \Gamma_1(s) \neq \emptyset,$$

then $s \in Lpre_1(Y_1, X_1) \uplus Pre_1(Y_0)$.

2. *For all $Y_1 \subseteq Y_0 \subseteq X_1 \subseteq S$ and $s \in S$, if*

$$\Gamma_1(s) \subseteq \mu B \,.\, \nu A \,.\, \left[\begin{array}{c} (\mathsf{Cover}_2(s, Y_1, A) \cup \mathsf{Cover}_2(s, Y_0, B)) \\ \cup \\ \mathsf{Stay}_2(s, X_1, A) \end{array} \right],$$

then $s \in \big(Fpre_2(X_1, Y_1) \Cap Epre_2(Y_1)\big) \uplus Lpre_2(X_1, Y_0)$.

Proof. We prove the two cases below.

1. Let

$$A^* = \nu B \,.\, \mu A \,.\, \left[\begin{array}{c} (\mathsf{Stay}_1(s, Y_1, A) \cap \mathsf{Stay}_1(s, Y_0, B)) \\ \cup \\ \mathsf{Cover}_1(s, X_1, A) \end{array} \right]$$

and let $A^* \cap \Gamma_1(s) \neq \emptyset$. By replacing B with A^* we have $A^* = \mu A \,.\, [(\mathsf{Stay}_1(s, Y_1, A) \cap \mathsf{Stay}_1(s, Y_0, A^*)) \cup \mathsf{Cover}_1(s, X_1, A)]$. For $\varepsilon > 0$, consider a distribution $\xi_1[\varepsilon] \in \chi_s^1$ that plays each move $a \in A^* \cap \Gamma_1(s)$ with probability proportional to ε^k, where k is the level of the move a in the expression $\mu A \,.\, [(\mathsf{Stay}_1(s, Y_1, A) \cap \mathsf{Stay}_1(s, Y_0, A^*)) \cup \mathsf{Cover}_1(s, X_1, A)]$. We analyze the following cases.

   - Consider a move $b \in \Gamma_2(s)$ such that $b \notin A^*$. Observe that $A^* \cap \Gamma_1(s) \subseteq \mathsf{Stay}_1(s, Y_0, A^*)$. Hence for all $a \in A^* \cap \Gamma_1(s)$ we have $Dest(s, a, b) \subseteq Y_0$, and since $Y_0 \subseteq Y_1$ we also have $Dest(s, a, b) \subseteq Y_1$.
   - For $b \in A^*$, let the level of b in $\mu A.\, [(\mathsf{Stay}_1(s, Y_1, A) \cap \mathsf{Stay}_1(s, Y_0, A^*)) \cup \mathsf{Cover}_1(s, X_1, A)]$ be k. Then we have: (a) there exists $a \in A^* \cap \Gamma_1(s)$ such that level of a is at most $k-1$ and $Dest(s, a, b) \cap X_1 \neq \emptyset$, i.e., $P_s^{\xi_1, b}(X)$ is at least proportional to ε^{k-1}; and (b) for all $a \in A^* \cap \Gamma_1(s)$ such that level of a is at least $k-1$, we have $Dest(s, a, b) \subseteq Y_1$, i.e., $P_s^{\xi_1, b}(\neg Y_1)$ is at most proportional to ε^k. It follows that the probability of the ratio of going to X as compared to leaving Y_1 is at least proportional to $\frac{1}{\varepsilon}$.

   Thus playing $\xi_1[\varepsilon]$ player 1 ensures that (a) if player 2 plays a distribution ξ_2 such that $\mathrm{Supp}(\xi_2) \subseteq \Gamma_2(s) \setminus A^*$, then $\mathrm{Pre}_1(Y_0)$ holds, and (b) if $\mathrm{Supp}(\xi_2) \cap A^* \neq \emptyset$, then $\mathrm{Lpre}_1(Y_1, X_1)$ holds. Hence the distributions $\xi_1[\varepsilon]$, for $\varepsilon > 0$, shows that $s \in \mathrm{Lpre}_1(Y_1, X_1) \uplus \mathrm{Pre}_1(Y_0)$. We denote such a distribution by $\xi_{s,1}^{\mathrm{Lpre} \uplus \mathrm{Pre}}[\varepsilon](Y_1, Y_0, X_1)$ that is a witness that $s \in \mathrm{Lpre}_1(Y_1, X_1) \uplus \mathrm{Pre}_1(Y_0)$.

2. Let

$$\Gamma_1(s) \subseteq \mu B \,.\, \nu A \,.\, \left[\begin{array}{c} (\mathsf{Cover}_2(s, Y_1, A) \cup \mathsf{Cover}_2(s, Y_0, B)) \\ \cup \\ \mathsf{Stay}_2(s, X_1, A) \end{array} \right]$$

   For $\varepsilon > 0$, consider a distribution $\xi_2[\varepsilon] \in \chi_s^2$ that plays each move $b \in \Gamma_2(s)$ with probability proportional to ε^k, where k is the level of the move b in the μ-calculus expression above. Let B_k denote the set of moves with level at most k. Consider $a \in \Gamma_1(s)$ and let the level of a be k, i.e., $a \in B_k \setminus B_{k-1}$. Observe that we can write B_{k-1} and B_k as follows:

$$B_{k-1} = [(\mathsf{Cover}_2(s, Y_1, B_{k-1}) \cup \mathsf{Cover}_2(s, Y_0, B_{k-2})) \cup \mathsf{Stay}_2(s, X_1, B_{k-1})]\,;$$

$$B_k = [(\mathsf{Cover}_2(s, Y_1, B_k) \cup \mathsf{Cover}_2(s, Y_0, B_{k-1})) \cup \mathsf{Stay}_2(s, X_1, B_k)]\,.$$

   By the description of B_{k-1} above and the property of $\mathsf{Stay}_2(s, X_1, B_{k-1})$ it follows that for all moves $b \in B_{k-1} \cap \Gamma_2(s)$ and for all moves $a' \in \Gamma_1(s) \setminus B_{k-1}$ we have $Dest(s, a', b) \subseteq X_1$. Hence for all $a \in (B_k \setminus B_{k-1}) \cap \Gamma_1(s)$ and $b \in \Gamma_2(s) \cap B_{k-1}$ we have $Dest(s, a, b) \subseteq X_1$, so the probability of reaching $\neg X_1$ is at most $\alpha \cdot \varepsilon^k$ for some constant $\alpha > 0$. Furthermore, there must exist at least one $b \in \Gamma_2(s)$ such that:

- Either $b \in B_{k-1}$ and $Dest(s,a,b) \cap Y_0 \neq \emptyset$ (due to $\mathsf{Cover}_2(s, Y_0, B_{k-1})$). In this case, the probability of a transition to Y_0 is proportional to at least ε^{k-1}. In this case, the ratio probability of reaching Y_0 as compared to leaving X_1 is least proportional to $\frac{1}{\varepsilon}$ and $\mathrm{Lpre}_2(X_1, Y_0)$ holds.
- Or $b \in B_k$ and $Dest(s,a,b) \cap Y_1 \neq \emptyset$ (due to $\mathsf{Cover}_2(s, Y_1, B_k)$). In this case, the probability of a transition to Y_1 is proportional to at least ε^k. In this case, the probability of reaching Y_1 is at least proportional to the probability to reach $\neg X_1$ and $\mathrm{Fpre}_2(X_1, Y_1) \cap\!\!\!\!* \, \mathrm{Epre}_2(Y_1)$ holds.

The result follows from the above case analysis. For $\varepsilon > 0$, we denote by $\xi_{s,2}^{(\mathrm{Fpre} \cap\!\!\!\!* \, \mathrm{Epre}) \cup\!\!\!\!* \, \mathrm{Lpre}}[\varepsilon](X_1, Y_0, Y_1)$ the distribution that witnesses that $s \in (\mathrm{Fpre}_2(X_1, Y_1) \cap\!\!\!\!* \, \mathrm{Epre}_2(Y_1)) \cup\!\!\!\!* \, \mathrm{Lpre}_2(X_1, Y_0)$. ∎

Remark 2 *Similar to Remark 1, the proof of Lemma 13 can be easily extended to show the following: given $Y_1 \subseteq Y_0 \subseteq X_1 \subseteq S$ and $s \in S$, let*

$$C = \mu B \,.\, \nu A \,.\, \left[\begin{array}{c} (\mathsf{Cover}_2(s, Y_1, A) \cup \mathsf{Cover}_2(s, Y_0, B)) \\ \cup \\ \mathsf{Stay}_2(s, X_1, A) \end{array} \right].$$

If player 1's moves are restricted to the set $C \cap \Gamma_1(s)$, then player 2 can ensure that $s \in \big(Fpre_2(X_1, Y_1) \cap\!\!\!\!* \, Epre_2(Y_1)\big) \cup\!\!\!\!* \, Lpre_2(X_1, Y_0)$.

Note 1 Observe that the distribution $\xi_{s,1}^{\mathrm{Lpre} \cup\!\!\!\!* \, \mathrm{Pre}}[\varepsilon](Y_1, Y_0, X_1)$ constructed in Lemma 13 satisfy the following: as $\varepsilon \to 0$, the distribution converges to a distribution with support $\Gamma_1(s) \cap \big(\mathsf{Stay}_1(s, Y_1, \emptyset) \cap \mathsf{Stay}_1(s, Y_0, \emptyset)\big)$.

Lemma 14 *For all $X_1 \subseteq Y_0 \subseteq Y_1 \subseteq S$ and $s \in S$ the following assertions hold.*

1. $Lpre_1(Y_1, X_1) \cup\!\!\!\!* \, Pre_1(Y_0) = \neg\big((Fpre_1(\neg X_1, \neg Y_1) \cap\!\!\!\!* \, Epre_2(\neg Y_1)) \cup\!\!\!\!* \, Lpre_2(\neg X_1, \neg Y_0)\big)$ (20)

2. $s \in Lpre_1(Y_1, X_1) \cup\!\!\!\!* \, Pre_1(Y_0)$ *iff* $\nu B \,.\, \mu A \,.\, \left[\begin{array}{c} (\mathsf{Stay}_1(s, Y_1, A) \cap \mathsf{Stay}_1(s, Y_0, B)) \\ \cup \\ \mathsf{Cover}_1(s, X_1, A) \end{array} \right] \cap \Gamma_1(s) \neq \emptyset.$

Proof. The result follows from Lemma 13 and the following duality of the predicates over moves.

$$\neg \left[\begin{array}{c} (\mathsf{Stay}_1(s, \neg Y_1, \neg A) \cap \mathsf{Stay}_1(s, \neg Y_0, \neg B)) \\ \cup \\ \mathsf{Cover}_2(s, \neg X_1, \neg A) \end{array} \right] = \left[\begin{array}{c} (\mathsf{Cover}_2(s, Y_1, A) \cup \mathsf{Cover}_2(s, Y_0, B) \cup \Gamma_2(s)) \\ \cap \\ (\mathsf{Stay}_2(s, X_1, A) \cup \Gamma_1(s)) \end{array} \right]$$

From the observations that $\mathsf{Cover}_2(s, Y_1, A) \subseteq \Gamma_1(s)$, $\mathsf{Cover}_2(s, Y_0, B) \subseteq \Gamma_1(s)$, and $\mathsf{Stay}_2(s, X_1, A) \subseteq \Gamma_2(s)$, it follows that

$$\left[\begin{array}{c} (\mathsf{Cover}_2(s, Y_1, A) \cup \mathsf{Cover}_2(s, Y_0, B) \cup \Gamma_2(s)) \\ \cap \\ (\mathsf{Stay}_2(s, X_1, A) \cup \Gamma_1(s)) \end{array} \right] = \mathsf{Cover}_s(s, Y_1, A) \cup \mathsf{Cover}_2(s, Y_0, B) \cup \mathsf{Stay}_2(s, X_1, A).$$

The above equalities, the complementation of μ-calculus formulas and Lemma 13 prove the desired results. ∎

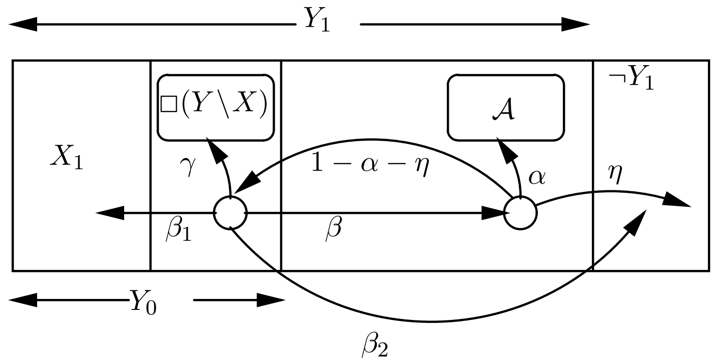

Figure 6: Pictorial description of the basic Lpre $\uplus$ Pre principle.

Lemma 15 (Basic Lpre $\uplus$ Pre principle) *Let $X_1 \subseteq Y_0 \subseteq Y_1 \subseteq S$ such that all $s \in Y_0 \setminus X_1$ satisfies that $s \in Lpre_1(Y_1, X_1) \uplus Pre_1(Y_0)$. For all events $\mathcal{A} \subseteq \Box(Y_1 \setminus Y_0)$, the following assertion hold: if for all $\eta > 0$, there exists $\pi_1^\eta \in \Pi_1^M$ such that for all $\pi_2 \in \Pi_2$ and for all $z \in Y_1 \setminus Y_0$ we have*

$$\mathrm{Pr}_z^{\pi_1^\eta,\pi_2}(\mathcal{A} \cup \Diamond Y_0) \geq 1-\eta, \qquad (i.e., \lim_{\eta \to 0} \mathrm{Pr}_z^{\pi_1^\eta,\pi_2}(\mathcal{A} \cup \Diamond Y_0) = 1),$$

then for all $\varepsilon > 0$, there exists $\pi_1^\varepsilon \in \Pi_1^M$ such that for all $\pi_2 \in \Pi_2$ and for all $s \in (Y_1 \setminus X_1)$ we have

$$\mathrm{Pr}_s^{\pi_1^\varepsilon,\pi_2}(\mathcal{A} \cup \Diamond X_1 \cup \Box(Y_0 \setminus X_1)) \geq 1-\varepsilon, \qquad (i.e., \lim_{\varepsilon \to 0} \mathrm{Pr}_z^{\pi_1^\varepsilon,\pi_2}(\mathcal{A} \cup \Diamond X_1 \cup \Box(Y_0 \setminus X_1)) = 1).$$

Proof. Since all $s \in (Y_0 \setminus X_1)$ satisfies that $s \in \mathrm{Lpre}_1(Y_1, X_1) \uplus \mathrm{Pre}_1(Y_0)$, given $\varepsilon > 0$, player 1 can play the distribution $\xi_{s,1}^{\mathrm{Lpre} \uplus \mathrm{Pre}}[\varepsilon](Y_1, Y_0, X_1)$ at states in $(Y_0 \setminus X_1)$ to ensure that the following conditions hold at all states in $(Y_0 \setminus X_1)$: (a) the probability β_1 to reach X_1 is at least proportional to $\beta \cdot \varepsilon^{l+1} + \frac{\beta_2}{\varepsilon}$, where β is the probability to reach $(Y_1 \setminus Y_0)$, β_2 is the probability to reach $\neg Y_1$, and $l = |\Gamma_s|$; and (b) if $\beta_1 = 0$, then $\beta = 0$ and $\beta_2 = 0$. The situation is pictorially depicted in Fig 6.

Since $\mathcal{A} \cup \Box(Y_0 \setminus X_1) \subseteq \mathcal{A} \cup \Box(Y_0 \setminus X_1) \cup \Diamond X_1$, the worst case analysis for the result correspond to the case with $\alpha = 0$ and $\gamma = 0$. The simplified case is shown in Fig 7. Similar to Lemma 4 once we let $\eta \to 0$, then the situation simplifies to the case where there is only an edge from $(Y_1 \setminus Y_0)$ to Y_0. Similar to Lemma 4 (via a more tedious calculation), it can be shown that for all $\varepsilon' > 0$, we can choose $\varepsilon > 0$ and $\eta > 0$, such that the strategy $\pi_1^{\varepsilon'}$ obtained by fixing the strategy π_1^η for states in $(Y_1 \setminus Y_0)$ and the distribution $\xi_{s,1}^{\mathrm{Lpre} \uplus \mathrm{Pre}}[\varepsilon](Y_1, Y_0, X_1)$ for states $s \in (Y_0 \setminus X_1)$ satisfies that for all $s \in (Y_1 \setminus X_1)$ and for all π_2 we have $\mathrm{Pr}_s^{\pi_1^{\varepsilon'},\pi_2}(\mathcal{A} \cup \Diamond X_1 \cup \Box(Y_0 \setminus X_1)) \geq 1-\varepsilon'$. ∎

6.2 Limit winning sets in co-Büchi games

We now present the computation of the limit-winning set in coBüchi games.

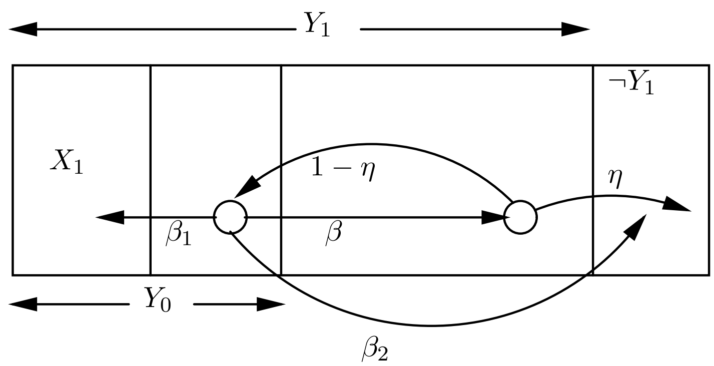

Figure 7: Pictorial description of the basic Lpre $\uplus\!\!\!\ast$ Pre principle, simplified to yield the worst case analysis.

Lemma 16 *For $T \subseteq S$, $B_0 \subseteq S$, and $B_1 = S \setminus B_0$, let*

$$W = \nu Y_1 \,.\, \mu X_1 \,.\, \nu Y_0 \,.\, \left[\begin{array}{c} T \\ \cup \\ B_0 \cap (Lpre_1(Y_1, X_1) \uplus\!\!\!\ast\; Pre_1(Y_0)) \\ \cup \\ B_1 \cap Lpre_1(Y_1, X_1) \end{array} \right].$$

Then we have $W \subseteq \langle\!\langle 1 \rangle\!\rangle_{limit}\big(\Diamond\Box B_0 \cup \Diamond T\big)$.

Proof. We first reformulate the algorithm for computing W in an equivalent form. We have

$$W = \mu X_1 \,.\, \nu Y_0 \,.\, \left[\begin{array}{c} T \\ \cup \\ B_0 \cap (\mathrm{Lpre}_1(W, X_1) \uplus\!\!\!\ast\; \mathrm{Pre}_1(Y_0)) \\ \cup \\ B_1 \cap \mathrm{Lpre}_1(W, X_1) \end{array} \right]. \tag{21}$$

This mu-calculus formula computes W as the limit of a sequence of sets $W_0 = T$, W_1, W_2, $\ldots$. At each iteration, both states satisfying B_0 (corresponding to the Lpre $\uplus\!\!\!\ast$ Pre operator) and states satisfying B_1 (corresponding to the Lpre operator) can be added. The fact that both types of states can be added complicates the analysis of the algorithm. To simplify the correctness proof, we formulate an alternative algorithm for the computation of W; an iteration will add either a single B_1 state, or a set of B_0 states.

To obtain the simpler algorithm, notice that the set Y_0 does not appear as an argument of the $\mathrm{Lpre}_1(W, X_1)$ operator. Hence, each B_1-state can be added without regards to B_0-states that are not already in W. Moreover, since the νY_0 operator applies only to B_0-states, B_1-states can be added one at a time. From these considerations, we can reformulate the algorithm for the computation of W as follows.

The algorithm computes W as an increasing sequence $T = T_0 \subset T_1 \subset T_2 \subset \cdots \subset T_m = W$ of states, where $m \geq 0$. Let $L_i = T_i \setminus T_{i-1}$ and the sequence is computed by computing T_i as follows, for $0 < i \leq m$:

1. either the set L_i consists of states such that $s \in L_i$ implies that

$$s \in \mathrm{Lpre}_1(W, T_{i-1}) \circledast \mathrm{Pre}_1(L_i \cup T_{i-1}) \cap B_0 \ ,$$

   i.e., $s \in T_i \setminus T_{i-1}$ implies that $s \in \mathrm{Lpre}_1(W, T_{i-1}) \circledast \mathrm{Pre}_1(T_i)$;

2. or the set $L_i = \{s\}$ is a singleton such that $s \in \mathrm{Lpre}_1(W, T_{i-1}) \cap B_1$.

The proof that $W \subseteq \langle\langle 1 \rangle\rangle_{limit}(\Diamond\Box B_0 \cup \Diamond T)$ is based on an induction on the sequence $T = T_0 \subset T_1 \subset T_2 \subset \cdots \subset T_m = W$. For $1 \le i \le m$, let $V_i = W \setminus T_{m-i}$, so that V_1 consists of the last block of states that has been added, V_2 to the two last blocks, and so on until $V_m = W$. We prove by induction on $i \in \{1, \ldots, m\}$, from $i = 1$ to $i = m$, that for all $s \in V_i$, for all $\eta > 0$, there exists a memoryless strategy π_1^η for player 1 such that for all $\pi_2 \in \Pi_2$ we have

$$\mathrm{Pr}_s^{\pi_1^\eta, \pi_2}\big(\Diamond T_{m-i} \cup (\Diamond\Box B_0 \cap \Box V_i)\big) \ge 1 - \eta.$$

Since the base case is a simplified version of the induction step, we focus on the latter. There are two cases, depending on whether $V_i \setminus V_{i-1}$ is composed of B_0 or of B_1-states.

1. If $V_i \setminus V_{i-1} \subseteq B_0$, then all $s \in V_i \setminus V_{i-1}$ satisfies that $s \in \mathrm{Lpre}_1(W, T_{m-i}) \circledast \mathrm{Pre}_1(T_{m-i+1}) = \mathrm{Lpre}_1(W, T_{m-i}) \circledast \mathrm{Pre}_1(V_i \setminus V_{i-1})$. Observe that $\Box(V_i \setminus V_{i-1}) \subseteq \Box(B_0 \cap V_i) \subseteq \Diamond\Box B_0 \cap \Box V_i$, since $V_i \setminus V_{i-1} \subseteq B_0$. The result then follows from the above observation and by an application of the basic Lpre $\circledast$ Pre principle (Lemma 15), with $Y_1 = W$, $X_1 = T_{m-i}$, $Y_0 = X \cup (V_i \setminus V_{i-1})$ (i.e., $Y_0 \setminus X_1 = V_i \setminus V_{i-1}$), $Y_1 \setminus Y_0 = V_{i-1}$ and $\mathcal{A} = \Box\Diamond B_0 \cap \Box V_{i-1}$.

2. If $V_i \setminus V_{i-1} \subseteq B_1$, then $V_i \setminus V_{i-1} = \{s\}$ for some $s \in S$ and $s \in \mathrm{Lpre}_1(W, T_{m-i})$. The result then follows from the application of the basic Lpre principle (Lemma 5) with $Z = W$, $X = T_{m-i}$, $Z \setminus Y = V_{i-1}$ and $\mathcal{A} = \Box\Diamond B_0 \cap \Box V_{i-1}$.

This completes the inductive proof. With $i = m$ we obtain that for all $\eta > 0$, there exists a memoryless strategy π_1^η such that for all states $s \in V_m = W$ and for all π_2 we have $\mathrm{Pr}_s^{\pi_1^\eta, \pi_2}(\Diamond T_0 \cup \Diamond\Box B_0) \ge 1 - \eta$. Since $T_0 = T$, the desired result follows. ∎

The following lemma complements the result of Lemma 16.

Lemma 17 *For $T \subseteq S$, $B_0 \subseteq S$, and $B_1 = S \setminus B_0$, let*

$$Z = \mu Y_1 \,.\, \nu X_1 \,.\, \mu Y_0 \,.\, \left[\begin{array}{c} B_0 \cap \big((\mathit{Fpre}_2(X_1, Y_1) \circledast \mathit{Epre}_2(Y_1)) \circledast \mathit{Lpre}_2(X_1, Y_0)\big) \\ \cup \\ B_1 \cap \mathit{Fpre}_2(X_1, Y_1) \end{array}\right].$$

Then we have $Z \subseteq \neg\langle\langle 1 \rangle\rangle_{limit}(\Diamond\Box B_0)$.

Proof. For $k \ge 0$, let Z_k be the set of states of level k in the above mu-calculus expression. We will construct a spoiling strategy for player 2, and show by induction on k that $Z_k \cap \langle\langle 1 \rangle\rangle_{limit}(\Diamond\Box B_0) = \emptyset$. The base case, for $k = 0$, corresponds to the set

$$Z_0 = \nu X_1 \,.\, \mu Y_0 \,.\, \left[\begin{array}{c} B_0 \cap \big((\mathrm{Fpre}_2(X_1, \emptyset) \circledast \mathrm{Epre}_2(\emptyset)) \circledast \mathrm{Lpre}_2(X_1, Y_0)\big) \\ \cup \\ B_1 \cap \mathrm{Fpre}_2(X_1, \emptyset) \end{array}\right]$$

$$= \nu X_1 \,.\, \mu Y_0 \,.\, \left[\begin{array}{c} B_0 \cap \mathrm{Lpre}_2(X_1, Y_0) \\ \cup \\ B_1 \cap \mathrm{Pre}_2(X_1) \end{array}\right] = \langle\langle 2 \rangle\rangle_{limit}(\Box\Diamond B_1)\ .$$

For the induction step, consider the case for $k > 0$:

$$Z_k = \nu X_1 \,.\, \mu Y_0 \,.\, \left[\begin{array}{c} B_0 \cap \big((\mathrm{Fpre}_2(X_1, Z_{k-1}) \mathbin{\overset{*}{\cap}} \mathrm{Epre}_2(Z_{k-1})) \mathbin{\overset{*}{\cup}} \mathrm{Lpre}_2(X_1, Y_0) \big) \\ \cup \\ B_1 \cap \mathrm{Fpre}_2(X_1, Z_{k-1}) \end{array} \right] .$$

For $X = Z_k$, we have:

$$X = \mu Y_0 \,.\, \left[\begin{array}{c} B_0 \cap \big((\mathrm{Fpre}_2(X, Z_{k-1}) \mathbin{\overset{*}{\cap}} \mathrm{Epre}_2(Z_{k-1})) \mathbin{\overset{*}{\cup}} \mathrm{Lpre}_2(X, Y_0) \big) \\ \cup \\ B_1 \cap \mathrm{Fpre}_2(X, Z_{k-1}) \end{array} \right] . \tag{22}$$

We can reformulate the above mu-calculus expression for the computation of X in the following equivalent form. We let $T = X \setminus Z_{k-1}$ be the set of states added to Z_{k-1} to yield X, and we compute T as the limit $\emptyset \subseteq T_0 \subseteq \cdots \subseteq T_m \subseteq T$ of an increasing sequence of sets of states, for some $m \geq 0$. The first set, T_0, is obtained by adding the B_1-states mentioned in (22), noting that these sets do not depend on Y:

$$T_0 = (B_1 \cap \mathrm{Fpre}_2(X, Z_{k-1})) \setminus Z_{k-1} \,.$$

The sets of states $T_1, \ldots, T_m$ are singleton sets. For $1 \geq i \geq m$, set T_i contains a single state s_i, where

$$s_i \in (\mathrm{Fpre}_2(X, Z_{k-1}) \mathbin{\overset{*}{\cap}} \mathrm{Epre}_2(Z_{k-1})) \mathbin{\overset{*}{\cup}} \mathrm{Lpre}_2(X, \bigcup_{j=0}^{i-1} T_j) \,.$$

Given $\varepsilon > 0$, consider a sequence $\varepsilon_1, \varepsilon_2, \ldots$ such that $\varepsilon_i > 0$ for all i and $\prod_{i=1}^{\infty}(1 - \varepsilon_i) \geq 1 - \varepsilon$. We construct a strategy $\pi_2[\varepsilon]$ for player 2 as follows: the strategy is played in rounds and it proceeds from round j to $j+1$ for a visit to a state in T_0, and in round j at all $s \in T_0$, $\pi_2[\varepsilon]$ plays the distribution $\xi_{s,2}^{\mathrm{Fpre}}(X, Z_{k-1})$; at state $s_i \in T_i$, player 2 plays with distribution $\xi_{s,2}^{(\mathrm{Fpre} \mathbin{\overset{*}{\cap}} \mathrm{Epre}) \mathbin{\overset{*}{\cup}} \mathrm{Lpre}}[\varepsilon_j](X, \bigcup_{j=0}^{i-1} T_j, Z_{k-1})$. Given the strategy $\pi_2[\varepsilon]$ we analyze two cases to prove the inductive case.

- If player 1 plays in such a way as to cause a positive probability of transition into Z_{k-1}, the induction hypothesis leads to the conclusion.
- Otherwise, if player 1 plays such that the probability of reaching Z_{k-1} is 0, then operator $\mathrm{Fpre}_2(X, Z_{k-1})$ simplifies to $\mathrm{Pre}_2(X)$, and the operator $(\mathrm{Fpre}_2(X, Z_{k-1}) \mathbin{\overset{*}{\cap}} \mathrm{Epre}_2(Z_{k-1})) \mathbin{\overset{*}{\cup}} \mathrm{Lpre}_2(X, \bigcup_{j=0}^{i-1} T_j)$ simplifies to $\mathrm{Lpre}_2(X, \bigcup_{j=0}^{i-1} T_j)$. Hence the strategy ensures that player 2 limit-win the $\Box\Diamond T_0$ objective (i.e., satisfies the objective $\Box\Diamond T_0$ with probability at least $1 - \varepsilon$). Since $T_0 \subseteq B_1$, player 2 satisfies $\Box\Diamond B_1$ with probability at least $1 - \varepsilon$.

In both cases, the inductive case is proved. The result follows. ∎

By the complementation of the predecessor operators (Lemma 14), the μ-calculus expression for W of Lemma 16 with $T = \emptyset$, complements to give the μ-calculus expression for Z of Lemma 17. From this we obtain Theorem 4, summarizing the result on coBüchi games.

Theorem 4 *For all concurrent game structures, for all co-Büchi objectives $\Diamond\Box B_0$ for player 1, where $B_0 \subseteq S$ and $B_1 = S \setminus B_0$, the following assertions hold.*

1. *We have*

$$\langle\langle 1 \rangle\rangle_{limit}\big(\Diamond\Box B_0\big) = \nu Y_1 \,.\, \mu X_1 \,.\, \nu Y_0 \,.\, \left[\begin{array}{c} B_0 \cap (Lpre_1(Y_1, X_1) \uplus\!\!\!\!* \;\; Pre_1(Y_0)) \\ \cup \\ B_1 \cap Lpre_1(Y_1, X_1) \end{array}\right]. \tag{23}$$

2. *The set of limit-winning states can be computed using the relations (23) in time $\mathcal{O}(|S|^3 \cdot \sum_{s \in S} |\Gamma_1(s) \cup \Gamma_2(s)|^2)$, where S is the set of states.*

3. *The most restrictive class of strategies in which limit-winning strategies are guaranteed to exist for player 1 is the class of memoryless strategies.*

4. *Limit-spoiling strategies for player 2 require infinite-memory in general.*

The time complexity of item (2) of Theorem 4 is obtained as follows: the triple-nested fixed point ensures that in at most $|S|^3$ iterations the μ-calculus formula of (23) converges to the fixpoint. By Lemma 4 and Lemma 13 it follows that whether a state $s \in \mathrm{Lpre}_1(Y_1, X_1)$ or $s \in \mathrm{Lpre}_1(Y_1, X_1) \uplus\!\!\!\!* \;\; \mathrm{Pre}_1(Y_0)$ can be decided in time $\mathcal{O}(|\Gamma_1(s) \cup \Gamma_2(s)|^2)$. This yields the result of item (2) of Theorem 4.

7 Parity Games

In this section we consider parity (Rabin-chain) games with objective $\mathrm{Parity}(p)$ for player 1, where $p : S \mapsto [0..2n-1]$ or $p : S \mapsto [1..2n]$ is a function that maps states with priorities from 0 to $2n-1$ or 1 to $2n$. We denote by $m = 2n$ the number of priorities and denote by $B_i = p^{-1}(i)$ the set of states with priority i. We will use the following notation in our proofs: for $\bowtie \in \{\leq, <, \geq, >\}$ we denote by $B_{\bowtie n} = \bigcup_{i \bowtie n} B_i$. We first introduce the predecessor operators for limit parity games. Given a parity function $p : S \mapsto [0..2n-1]$, the parity function $p+1 : S \mapsto [1..2n]$ is defined as follows: $p+1(s) = p(s)+1$ and observe that $\mathrm{coParity}(p) = \mathrm{Parity}(p+1)$. Similarly, given a parity function $p : S \mapsto [1..2n]$, the parity function $p-1 : S \mapsto [0..2n-1]$ is defined as follows: $p-1(s) = p(s)-1$ and observe that $\mathrm{coParity}(p) = \mathrm{Parity}(p-1)$.

7.1 Predecessor operators for Rabin-chain games

We introduce the predecessor operators for limit parity games. We first introduce two limit predecessor operators as follows:

$$\begin{aligned} &\mathrm{LPreOdd}_1(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i}) \\ &\;= \mathrm{Lpre}_1(Y_n, X_n) \uplus\!\!\!\!* \;\; \mathrm{Lpre}_1(Y_{n-1}, X_{n-1}) \uplus\!\!\!\!* \;\; \cdots \uplus\!\!\!\!* \;\; \mathrm{Lpre}_1(Y_{n-i}, X_{n-i}); \\[1ex] &\mathrm{LPreEven}_1(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i}, Y_{n-i-1}) \\ &\;= \mathrm{Lpre}_1(Y_n, X_n) \uplus\!\!\!\!* \;\; \mathrm{Lpre}_1(Y_{n-1}, X_{n-1}) \uplus\!\!\!\!* \;\; \cdots \uplus\!\!\!\!* \;\; \mathrm{Lpre}_1(Y_{n-i}, X_{n-i}) \uplus\!\!\!\!* \;\; \mathrm{Pre}_1(Y_{n-i-1}). \end{aligned}$$

The formal definitions of the above operators are as follows:

$$\mathrm{LPreOdd}_1(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i}) =$$

$$\left\{ s \in S \mid \forall \alpha > 0 \,.\, \exists \xi_1 \in \chi_1^s . \forall \xi_2 \in \chi_2^s . \begin{bmatrix} P_s^{\xi_1,\xi_2}(X_n) > \alpha \cdot P_s^{\xi_1,\xi_2}(\neg Y_n) \\ \bigvee \\ P_s^{\xi_1,\xi_2}(X_{n-1}) > \alpha \cdot P_s^{\xi_1,\xi_2}(\neg Y_{n-1}) \\ \bigvee \\ \vdots \\ \bigvee \\ P_s^{\xi_1,\xi_2}(X_{n-i}) > \alpha \cdot P_s^{\xi_1,\xi_2}(\neg Y_{n-i}) \end{bmatrix} \right\}.$$

$$\mathrm{LPreEven}_1(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i}, Y_{n-i-1}) =$$

$$\left\{ s \in S \mid \forall \alpha > 0 \,.\, \exists \xi_1 \in \chi_1^s . \forall \xi_2 \in \chi_2^s . \begin{bmatrix} P_s^{\xi_1,\xi_2}(X_n) > \alpha \cdot P_s^{\xi_1,\xi_2}(\neg Y_n) \\ \bigvee \\ P_s^{\xi_1,\xi_2}(X_{n-1}) > \alpha \cdot P_s^{\xi_1,\xi_2}(\neg Y_{n-1}) \\ \bigvee \\ \vdots \\ \bigvee \\ P_s^{\xi_1,\xi_2}(X_{n-i}) > \alpha \cdot P_s^{\xi_1,\xi_2}(\neg Y_{n-i}) \\ \bigvee \\ P_s^{\xi_1,\xi_2}(Y_{n-i-1}) = 1 \end{bmatrix} \right\}.$$

Observe that the above definition can be inductively written as follows:

1. We have $\mathrm{LPreOdd}_1(0, Y_n, X_n) = \mathrm{Lpre}_1(Y_n, X_n)$ and for $i \geq 1$ we have

$$\begin{aligned} &\mathrm{LPreOdd}_1(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i}) \\ &= \mathrm{Lpre}_1(Y_n, X_n) \mathbin{\cup\!\!\!\!\ast} \mathrm{LPreOdd}_1(i-1, Y_{n-1}, X_{n-1}, \ldots, Y_{n-i}, X_{n-i}) \end{aligned}$$

2. We have $\mathrm{LPreEven}_1(0, Y_n, X_n, Y_{n-1}) = \mathrm{Lpre}_1(Y_n, X_n) \mathbin{\cup\!\!\!\!\ast} \mathrm{Pre}_1(Y_{n-1})$ and for $i \geq 1$ we have

$$\begin{aligned} &\mathrm{LPreEven}_1(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i}, Y_{n-i-1}) \\ &= \mathrm{Lpre}_1(Y_n, X_n) \mathbin{\cup\!\!\!\!\ast} \mathrm{LPreEven}_1(i-1, Y_{n-1}, X_{n-1}, \ldots, Y_{n-i}, X_{n-i}, Y_{n-i-1}) \end{aligned}$$

The operators $\mathrm{LPreOdd}_2$ and $\mathrm{LPreEven}_2$ can be obtained from $\mathrm{LPreOdd}_1$ and $\mathrm{LPreEven}_1$ by exchanging the subscripts 1 and 2. We now introduce two positive predecessor operators as follows:

$$\begin{aligned} &\mathrm{FPreOdd}_2(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i}) \\ &= (\mathrm{Fpre}_2(X_n, Y_n) \mathbin{\cap\!\!\!\!\ast} \mathrm{Epre}_2(Y_n)) \mathbin{\cup\!\!\!\!\ast} \mathrm{Lpre}_2(X_n, Y_{n-1}) \mathbin{\cup\!\!\!\!\ast} \cdots \mathbin{\cup\!\!\!\!\ast} \mathrm{Lpre}_2(X_{n-i+1}, Y_{n-i}) \mathbin{\cup\!\!\!\!\ast} \mathrm{Pre}_2(X_{n-i}) \end{aligned}$$

$$\begin{aligned} &\mathrm{FPreEven}_2(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i}, Y_{n-i-1}) \\ &= (\mathrm{Fpre}_2(X_n, Y_n) \mathbin{\cap\!\!\!\!\ast} \mathrm{Epre}_2(Y_n)) \mathbin{\cup\!\!\!\!\ast} \mathrm{Lpre}_2(X_n, Y_{n-1}) \\ &\qquad \mathbin{\cup\!\!\!\!\ast} \cdots \mathbin{\cup\!\!\!\!\ast} \mathrm{Lpre}_2(X_{n-i+1}, Y_{n-i}) \mathbin{\cup\!\!\!\!\ast} \mathrm{Lpre}_2(X_{n-i}, Y_{n-i-1}) \end{aligned}$$

The formal definitions of the above operators are as follows:

$$\mathrm{FPreOdd}_2(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i}) =$$

$$\left\{ s \in S \mid \exists \beta > 0. \forall \alpha \geq 0. \exists \xi_2 \in \chi_2^s . \forall \xi_1 \in \chi_1^s . \left[\begin{array}{c} \big(P_s^{\xi_1,\xi_2}(Y_n) \geq \beta \cdot P_s^{\xi_1,\xi_2}(\neg X_n) \\ \bigwedge \\ P_s^{\xi_1,\xi_2}(Y_n) > 0\big) \\ \bigvee \\ P_s^{\xi_1,\xi_2}(Y_{n-1}) > \alpha \cdot P_s^{\xi_1,\xi_2}(\neg X_n) \\ \bigvee \\ P_s^{\xi_1,\xi_2}(Y_{n-2}) > \alpha \cdot P_s^{\xi_1,\xi_2}(\neg X_{n-1}) \\ \bigvee \\ \vdots \\ \bigvee \\ P_s^{\xi_1,\xi_2}(Y_{n-i}) > \alpha \cdot P_s^{\xi_1,\xi_2}(\neg X_{n-i+1}) \\ \bigvee \\ P_s^{\xi_1,\xi_2}(X_{n-i}) = 1 \end{array} \right] \right\}.$$

$$\mathrm{FPreEven}_2(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i}, Y_{n-i-1}) =$$

$$\left\{ s \in S \mid \exists \beta > 0. \forall \alpha \geq 0. \exists \xi_2 \in \chi_2^s . \forall \xi_1 \in \chi_1^s . \left[\begin{array}{c} \big(P_s^{\xi_1,\xi_2}(Y_n) \geq \beta \cdot P_s^{\xi_1,\xi_2}(\neg X_n) \\ \bigwedge \\ P_s^{\xi_1,\xi_2}(Y_n) > 0\big) \\ \bigvee \\ P_s^{\xi_1,\xi_2}(Y_{n-1}) > \alpha \cdot P_s^{\xi_1,\xi_2}(\neg X_n) \\ \bigvee \\ P_s^{\xi_1,\xi_2}(Y_{n-2}) > \alpha \cdot P_s^{\xi_1,\xi_2}(\neg X_{n-1}) \\ \bigvee \\ \vdots \\ \bigvee \\ P_s^{\xi_1,\xi_2}(Y_{n-i}) > \alpha \cdot P_s^{\xi_1,\xi_2}(\neg X_{n-i+1}) \\ \bigvee \\ P_s^{\xi_1,\xi_2}(Y_{n-i-1}) > \alpha \cdot P_s^{\xi_1,\xi_2}(\neg X_{n-i}) \end{array} \right] \right\}.$$

The above definitions can be alternatively written as follows

$$\begin{aligned} \mathrm{FPreOdd}_2(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i}) \quad &= \\ (\mathrm{Fpre}_2(X_n, Y_n) \mathbin{\cap\!\!\!\!*} \mathrm{Epre}_2(Y_n)) \quad \mathbin{\cup\!\!\!\!*} \quad &\mathrm{LPreEven}_2(i-1, X_n, Y_{n-1}, \ldots, X_{n-i+1}, Y_{n-i}, X_{n-i}); \end{aligned}$$

$$\begin{aligned} \mathrm{FPreEven}_2(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i}, Y_{n-i-1}) \quad &= \\ (\mathrm{Fpre}_2(X_n, Y_n) \mathbin{\cap\!\!\!\!*} \mathrm{Epre}_2(Y_n)) \quad \mathbin{\cup\!\!\!\!*} \quad &\mathrm{LPreOdd}_2(i, X_n, Y_{n-1}, \ldots, X_{n-i}, Y_{n-i-1}). \end{aligned}$$

We will prove the duality of $\mathrm{LPreOdd}_1$ and $\mathrm{FPreOdd}_2$ and $\mathrm{LPreEven}_1$ and $\mathrm{FPreEven}_2$ later. We first show how to characterize the limit winning sets and its complement for parity games using the above predecessor operators.

7.2 Limit winning in parity games

We will prove the following result by induction.

1. *Case 1.* For a parity function $p : S \mapsto [0..2n-1]$ the following assertions hold.

   (a) For all $T \subseteq S$ we have $W \subseteq \langle\langle 1 \rangle\rangle_{limit}\big(\text{Parity}(p) \cup \Diamond T\big)$, where W is defined as follows:

$$\nu Y_n.\mu X_n.\nu Y_{n-1}.\mu X_{n-1}.\cdots \nu Y_1.\mu X_1.\nu Y_0. \left[\begin{array}{c} T \\ \cup \\ B_{2n-1} \cap \text{LPreOdd}_1(0, Y_n, X_n) \\ \cup \\ B_{2n-2} \cap \text{LPreEven}_1(0, Y_n, X_n, Y_{n-1}) \\ \cup \\ B_{2n-3} \cap \text{LPreOdd}_1(1, Y_n, X_n, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-4} \cap \text{LPreEven}_1(1, Y_n, X_n, Y_{n-1}, X_{n-1}, Y_{n-2}) \\ \vdots \\ B_1 \cap \text{LPreOdd}_1(n-1, Y_n, X_n, \ldots, Y_1, X_1) \\ \cup \\ B_0 \cap \text{LPreEven}_1(n-1, Y_n, X_n, \ldots, Y_1, X_1, Y_0) \end{array}\right]$$

   We refer to the above expression as the limit-expression for case 1.

   (b) We have $Z \subseteq \neg\langle\langle 1 \rangle\rangle_{limit}\big(\text{Parity}(p)\big)$, where Z is defined as follows

$$\mu Y_n.\nu X_n.\mu Y_{n-1}.\nu X_{n-1}.\cdots \mu Y_1.\nu X_1.\mu Y_0. \left[\begin{array}{c} B_{2n-1} \cap \text{FPreOdd}_2(0, Y_n, X_n) \\ \cup \\ B_{2n-2} \cap \text{FPreEven}_2(0, Y_n, X_n, Y_{n-1}) \\ \cup \\ B_{2n-3} \cap \text{FPreOdd}_2(1, Y_n, X_n, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-4} \cap \text{FPreEven}_2(1, Y_n, X_n, Y_{n-1}, X_{n-1}, Y_{n-2}) \\ \vdots \\ B_1 \cap \text{FPreOdd}_2(n-1, Y_n, X_n, \ldots, Y_1, X_1) \\ \cup \\ B_0 \cap \text{FPreEven}_2(n-1, Y_n, X_n, \ldots, Y_1, X_1, Y_0) \end{array}\right]$$

   We refer to the above expression as the positive-expression for case 1.

2. *Case 2.* For a parity function $p : S \mapsto [1..2n]$ the following assertions hold.

(a) For all $T \subseteq S$ we have $W \subseteq \langle\langle 1 \rangle\rangle_{limit}\big(\text{Parity}(p) \cup \Diamond T\big)$, where W is defined as follows:

$$
\nu Y_{n-1}.\mu X_{n-1}.\cdots \nu Y_1.\mu X_1.\nu Y_0.\mu X_0 \begin{bmatrix} T \\ \cup \\ B_{2n} \cap \text{Pre}_1(Y_{n-1}) \\ \cup \\ B_{2n-1} \cap \text{LPreOdd}_1(0, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-2} \cap \text{LPreEven}_1(0, Y_{n-1}, X_{n-2}, Y_{n-2}) \\ \cup \\ B_{2n-3} \cap \text{LPreOdd}_1(1, Y_{n-1}, X_{n-1}, Y_{n-2}, X_{n-2}) \\ \vdots \\ B_2 \cap \text{LPreEven}_1(n-2, Y_{n-1}, X_{n-1}, \ldots, Y_1, X_1, Y_0) \\ \cup \\ B_1 \cap \text{LPreOdd}_1(n-1, Y_{n-1}, X_{n-1}, \ldots, Y_0, X_0) \end{bmatrix}
$$

We refer to the above expression as the limit-expression for case 2.

(b) We have $Z \subseteq \neg\langle\langle 1 \rangle\rangle_{limit}\big(\text{Parity}(p)\big)$, where Z is defined as follows

$$
\mu Y_{n-1}.\nu X_{n-1}.\cdots \mu Y_1.\nu X_1.\mu Y_0.\nu X_0 \begin{bmatrix} B_{2n} \cap \text{Epre}_2(Y_{n-1}) \\ \cup \\ B_{2n-1} \cap \text{FPreOdd}_2(0, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-2} \cap \text{FPreEven}_2(0, Y_{n-1}, X_{n-2}, Y_{n-2}) \\ \cup \\ B_{2n-3} \cap \text{FPreOdd}_2(1, Y_{n-1}, X_{n-1}, Y_{n-2}, X_{n-2}) \\ \vdots \\ B_2 \cap \text{FPreEven}_2(n-2, Y_{n-1}, X_{n-1}, \ldots, Y_1, X_1, Y_0) \\ \cup \\ B_1 \cap \text{FPreOdd}_2(n-1, Y_{n-1}, X_{n-1}, \ldots, Y_0, X_0) \end{bmatrix}
$$

We refer to the above expression as the positive-expression for case 2.

Proof structure. The base case follows from the coBüchi and Büchi case (see Lemma 16 and Lemma 17 for correctness of case 1 and Lemma 10 and Lemma 11 for correctness of case 2). The proof of induction proceeds in four steps as follows:

1. *Step 1.* We assume the correctness of case 1 and case 2, and then extend the result to parity objective with parity function $p : S \mapsto [0..2n]$, i.e., we add a max even priority. The result is obtained as follows: for the correctness of the limit-expression we use the correctness of case 1 and for complementation we use the correctness of case 2.

2. *Step 2.* We assume the correctness of step 1 and extend the result to parity objectives with parity function $p : S \mapsto [1..2n+1]$, i.e., we add a max odd priority. The result is obtained as follows: for the correctness of the limit-expression we use the correctness of case 2 and for complementation we use the correctness of step 1.

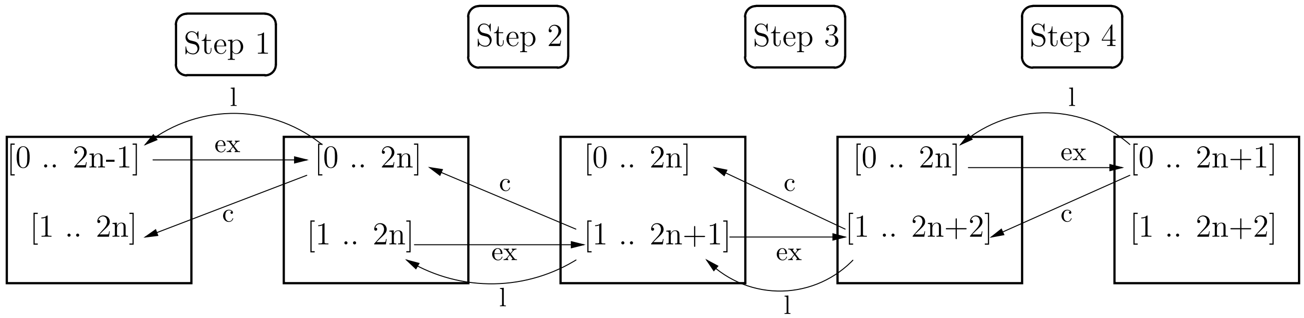

Figure 8: Pictorial description of the proof structure. The back-arrows marked as l denote dependence for proving correctness for the limit-expression; the back-arrows marked as c denote dependence for proving correctness of the complementation; and the back arrows marked as *ex* denote extension of a result. In each box, the upper line is the parity for case 1, and the lower line is the parity for case 2.

3. *Step 3.* We assume correctness of step 2 and extend the result to parity objectives with parity function $p : S \mapsto [1..2n+2]$. This step adds a max even priority and the proof will be similar to step 1. The result is obtained as follows: for the correctness of the limit-expression we use the correctness of step 2 and for complementation we use the correctness of step 1.

4. *Step 4.* We assume correctness of step 3 and extend the result to parity objectives with parity function $p : S \mapsto [0..2n+1]$. This step adds a max odd priority and the proof will be similar to step 2. The result is obtained as follows: for the correctness of the limit-expression we use the correctness of step 1 and for complementation we use the correctness of step 3.

A pictorial view of the proof structure is shown in Fig 8.

Correctness of step 1. We now proceed with the proof of step 1 and by inductive hypothesis we will assume that case 1 and case 2 hold.

Lemma 18 *For a parity function $p : S \mapsto [0..2n]$, and for all $T \subseteq S$, we have $W \subseteq$*

$\langle\langle 1\rangle\rangle_{limit}\big(Parity(p) \cup \Diamond T\big)$, *where* W *is defined as follows:*

$$\nu Y_n.\mu X_n.\nu Y_{n-1}.\mu X_{n-1}.\cdots \nu Y_1.\mu X_1.\nu Y_0. \left[\begin{array}{c} T \\ \cup \\ B_{2n} \cap Pre_1(Y_n) \\ \cup \\ B_{2n-1} \cap LPreOdd_1(0, Y_n, X_n) \\ \cup \\ B_{2n-2} \cap LPreEven_1(0, Y_n, X_n, Y_{n-1}) \\ \cup \\ B_{2n-3} \cap LPreOdd_1(1, Y_n, X_n, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-4} \cap LPreEven_1(1, Y_n, X_n, Y_{n-1}, X_{n-1}, Y_{n-2}) \\ \vdots \\ B_1 \cap LPreOdd_1(n-1, Y_n, X_n, \ldots, Y_1, X_1) \\ \cup \\ B_0 \cap LPreEven_1(n-1, Y_n, X_n, \ldots, Y_1, X_1, Y_0) \end{array}\right]$$

Proof. The formula is obtained from the limit-expression for case 1 by just adding the expression $B_{2n} \cap \mathrm{Pre}_1(Y_n)$. To prove the result we first rewrite W as follows:

$$\nu Y_n.\mu X_n.\nu Y_{n-1}\mu X_{n-1}\cdots \nu Y_1.\mu X_1.\nu Y_0. \left[\begin{array}{c} T \cup (B_{2n} \cap \mathrm{Pre}_1(W)) \\ \cup \\ B_{2n-1} \cap \mathrm{LPreOdd}_1(0, Y_n, X_n) \\ \cup \\ B_{2n-2} \cap \mathrm{LPreEven}_1(0, Y_n, X_n, Y_{n-1}) \\ \cup \\ B_{2n-3} \cap \mathrm{LPreOdd}_1(1, Y_n, X_n, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-4} \cap \mathrm{LPreEven}_1(1, Y_n, X_n, Y_{n-1}, X_{n-1}, Y_{n-2}) \\ \vdots \\ B_1 \cap \mathrm{LPreOdd}_1(n-1, Y_n, X_n, \ldots, Y_1, X_1) \\ \cup \\ B_0 \cap \mathrm{LPreEven}_1(n-1, Y_n, X_n, \ldots, Y_1, X_1, Y_0) \end{array}\right]$$

Treating $T \cup (B_{2n} \cap \mathrm{Pre}_1(W))$, as the set T for the limit-expression for case 1, we obtain from the inductive hypothesis that

$$W \subseteq \langle\langle 1\rangle\rangle_{limit}\big(\mathrm{Parity}(p) \cup \Diamond(T \cup (B_{2n} \cap \mathrm{Pre}_1(W)))\big).$$

By Lemma 9, with $B = B_{2n}$ and $\mathcal{A} = \mathrm{Parity}(p)$ we obtain that

$$W \subseteq \langle\langle 1\rangle\rangle_{limit}\big(\mathrm{Parity}(p) \cup \Diamond T \cup \Box\Diamond B_{2n}\big).$$

Since B_{2n} is the maximal priority and it is even we have $\Box\Diamond B_{2n} \subseteq \mathrm{Parity}(p)$. Hence $W \subseteq \langle\langle 1\rangle\rangle_{limit}\big(\mathrm{Parity}(p) \cup \Diamond T\big)$ and the result follows. ∎

Lemma 19 *For a parity function $p : S \mapsto [0..2n]$, we have $Z \subseteq \neg\langle\langle 1\rangle\rangle_{limit}\big(Parity(p)\big)$, where Z is defined as follows*

$$\mu Y_n.\nu X_n.\mu Y_{n-1}.\nu X_{n-1}.\cdots \mu Y_1.\nu X_1.\mu Y_0. \begin{bmatrix} B_{2n} \cap Epre_2(Y_n) \\ \cup \\ B_{2n-1} \cap FPreOdd_2(0, Y_n, X_n) \\ \cup \\ B_{2n-2} \cap FPreEven_2(0, Y_n, X_n, Y_{n-1}) \\ \cup \\ B_{2n-3} \cap FPreOdd_2(1, Y_n, X_n, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-4} \cap FPreEven_2(1, Y_n, X_n, Y_{n-1}, X_{n-1}, Y_{n-2}) \\ \vdots \\ B_1 \cap FPreOdd_2(n-1, Y_n, X_n, \ldots, Y_1, X_1) \\ \cup \\ B_0 \cap FPreEven_2(n-1, Y_n, X_n, \ldots, Y_1, X_1, Y_0) \end{bmatrix}$$

Proof. For $k \geq 0$, let Z_k be the set of states of level k in the above μ-calculus expression. We will show that in Z_k player 2 can ensure that either Z_{k-1} is reached with positive probability or else coParity(p) is satisfied with probability arbitrarily close to 1. Since $Z_0 = \emptyset$, it would follow by induction that $Z_k \cap \langle\langle 1\rangle\rangle_{limit}\big(\text{Parity}(p)\big) = \emptyset$ and the desired result will follow.

We simplify the computation of Z_k given Z_{k-1} and allow that Z_k is obtained from Z_{k-1} in the following two ways.

1. Add a set states satisfying $B_{2n} \cap \text{Epre}_2(Z_{k-1})$, and if such a non-emptyset is added, then clearly player 2 can ensure from Z_k that Z_{k-1} is reached with positive probability. Thus the inductive case follows.

2. Add a set of states satisfying the following condition:

$$\nu X_n.\mu Y_{n-1}.\nu X_{n-1}.\cdots \mu Y_1.\nu X_1.\mu Y_0. \begin{bmatrix} B_{2n-1} \cap \text{FPreOdd}_2(0, Z_{k-1}, X_n) \\ \cup \\ B_{2n-2} \cap \text{FPreEven}_2(0, Z_{k-1}, X_n, Y_{n-1}) \\ \cup \\ B_{2n-3} \cap \text{FPreOdd}_2(1, Z_{k-1}, X_n, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-4} \cap \text{FPreEven}_2(1, Z_{k-1}, X_n, Y_{n-1}, X_{n-1}, Y_{n-2}) \\ \vdots \\ B_1 \cap \text{FPreOdd}_2(n-1, Z_{k-1}, X_n, \ldots, Y_1, X_1) \\ \cup \\ B_0 \cap \text{FPreEven}_2(n-1, Z_{k-1}, X_n, \ldots, Y_1, X_1, Y_0) \end{bmatrix}$$

   If the probability of reaching to Z_{k-1} is not positive, then the following conditions hold:

   - $\text{FPreOdd}_2(i, Z_{k-1}, X_n, Y_{n-1}, \ldots, Y_{n-i}, X_{n-i})$ simplifies to the predecessor operator $\text{LPreEven}_2(i-1, X_{n-1}, Y_{n-2}, \ldots, Y_{n-i}, X_{n-i})$ and

- $\text{FPreEven}_2(i, Z_{k-1}, X_n, Y_{n-1}, \ldots, Y_{n-i}, X_{n-i}, Y_{n-i-1})$ simplifies to the predecessor operator $\text{LPreOdd}_2(i, X_{n-1}, Y_{n-2}, \ldots, Y_{n-i}, X_{n-i}, Y_{n-i-1})$.

Hence if we rule out the possibility of positive probability to reach Z_{k-1}, then the above μ-calculus expression simplifies to

$$Z^* = \nu X_n.\mu Y_{m-1}\nu X_{m-1} \cdots \mu Y_1.\nu X_1.\mu Y_0. \begin{bmatrix} B_{2n-1} \cap \text{Pre}_2(X_n) \\ \cup \\ B_{2n-2} \cap \text{LPreOdd}_2(0, X_n, Y_{n-1}) \\ \cup \\ B_{2n-3} \cap \text{LPreEven}_2(1, X_n, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-4} \cap \text{LPreOdd}_2(1, X_n, Y_{n-1}, X_{n-1}, Y_{n-2}) \\ \vdots \\ B_1 \cap \text{LPreEven}_2(n-2, X_n, \ldots, Y_1, X_1) \\ \cup \\ B_0 \cap \text{LPreOdd}_2(n-1, X_n, \ldots, Y_1, X_1, Y_0) \end{bmatrix}.$$

We now consider the parity function $p+1 : S \mapsto [1..2n]$, and by applying the correctness of the limit-expression for case 2 (inductive hypothesis) with the roles of player 1 and player 2 exchanged we have $Z^* \subseteq \langle\langle 2 \rangle\rangle_{limit}\big(\text{coParity}(p)\big)$ (since $\text{Parity}(p+1) = \text{coParity}(p)$). Hence the desired claim follows.

The result follows from the above case analysis. ∎

Correctness of step 2. We now prove correctness of step 2 and we will rely on the correctness of step 1 and the inductive hypothesis. Since correctness of step 1 follows from the inductive hypothesis, we obtain the correctness of step 2 from the inductive hypothesis.

Lemma 20 *For a parity function $p : S \mapsto [1..2n+1]$, and for all $T \subseteq S$ we have $W \subseteq \langle\langle 1 \rangle\rangle_{limit}\big(Parity(p) \cup \Diamond T\big)$, where W is defined as follows:*

$$\nu Y_n.\mu X_n.\nu Y_{n-1}.\mu X_{n-1}.\cdots \nu Y_0.\mu X_0 \begin{bmatrix} T \\ \cup \\ B_{2n+1} \cap \mathit{LPreOdd}_1(0, Y_n, X_n) \\ \cup \\ B_{2n} \cap \mathit{LPreEven}_1(0, Y_n, X_n, Y_{n-1}) \\ \cup \\ B_{2n-1} \cap \mathit{LPreOdd}_1(1, Y_n, X_n, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-2} \cap \mathit{LPreEven}_1(1, Y_n, X_n, Y_{n-1}, X_{n-2}, Y_{n-2}) \\ \cup \\ B_{2n-3} \cap \mathit{LPreOdd}_1(2, Y_n, X_n, Y_{n-1}, X_{n-1}, Y_{n-2}, X_{n-2}) \\ \vdots \\ B_2 \cap \mathit{LPreEven}_1(n-1, Y_n, X_n, Y_{n-1}, X_{n-1}, \ldots, Y_1, X_1, Y_0) \\ \cup \\ B_1 \cap \mathit{LPreOdd}_1(n, Y_n, X_n, Y_{n-1}, X_{n-1}, \ldots, Y_0, X_0) \end{bmatrix}$$

Proof. We first explain how the μ-calculus expression is obtained from the limit-expression for case 2: we add a $\nu Y_n.\mu X_n$ (adding a quantifier alternation of the μ-calculus formula), and every LPreOdd and LPreEven predecessor operators are modified by adding $\mathrm{Lpre}_1(Y_n, X_n) \circledast\!\!\cup$ with the respective predecessor operators, and we add $B_{2n+1} \cap \mathrm{LPreOdd}_1(0, Y_n, X_n)$.

We first reformulate the algorithm for computing W in an equivalent form.

$$W = \mu X_n.\nu Y_{n-1}.\mu X_{n-1}.\cdots \nu Y_0.\mu X_0 \left[\begin{array}{c} T \\ \cup \\ B_{2n+1} \cap \mathrm{LPreOdd}_1(0, W, X_n) \\ \cup \\ B_{2n} \cap \mathrm{LPreEven}_1(0, W, X_n, Y_{n-1}) \\ \cup \\ B_{2n-1} \cap \mathrm{LPreOdd}_1(1, W, X_n, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-2} \cap \mathrm{LPreEven}_1(1, W, X_n, Y_{n-1}, X_{n-2}, Y_{n-2}) \\ \cup \\ B_{2n-3} \cap \mathrm{LPreOdd}_1(2, W, X_n, Y_{n-1}, X_{n-1}, Y_{n-2}, X_{n-2}) \\ \vdots \\ B_2 \cap \mathrm{LPreEven}_1(n-1, W, X_n, Y_{n-1}, X_{n-1}, \ldots, Y_1, X_1, Y_0) \\ \cup \\ B_1 \cap \mathrm{LPreOdd}_1(n, W, X_n, Y_{n-1}, X_{n-1}, \ldots, Y_0, X_0) \end{array}\right].$$

This mu-calculus formula computes W as the limit of a sequence of sets $W_0 = T$, W_1, W_2, $\ldots$. At each iteration, both states in B_{2n+1} and states satisfying $B_{\leq 2n}$ can be added. The fact that both types of states can be added complicates the analysis of the algorithm. To simplify the correctness proof, we formulate an alternative algorithm for the computation of W; an iteration will add either a single B_{2n+1} state, or a set of $B_{\leq 2n}$ states.

To obtain the simpler algorithm, notice that the set of variables $Y_{n-1}, X_{n-1}, \ldots, Y_0, X_0$ does not appear as an argument of the $\mathrm{LPreOdd}_1(0, W, X_n) = \mathrm{Lpre}_1(W, X_n)$ operator. Hence, each B_{2n+1}-state can be added without regards to $B_{\leq 2n}$-states that are not already in W. Moreover, since the $\nu Y_{n-1}.\mu X_{n-1}.\ldots\nu Y_0.\mu X_0$ operator applies only to $B_{\leq 2n}$-states, B_{2n+1}-states can be added one at a time. From these considerations, we can reformulate the algorithm for the computation of W as follows.

The algorithm computes W as an increasing sequence $T = T_0 \subset T_1 \subset T_2 \subset \cdots \subset T_m = W$ of states, where $m \geq 0$. Let $L_i = T_i \setminus T_{i-1}$ and the sequence is computed by computing T_i as follows, for $0 < i \leq m$:

1. either the set $L_i = \{s\}$ is a singleton such that $s \in \mathrm{Lpre}_1(W, T_{i-1}) \cap B_{2n+1}$.

2. or the set L_i consists of states in $B_{\leq 2n}$ such that L_i is a subset of the following expression

$$\nu Y_{n-1}.\mu X_{n-1}.\cdots \nu Y_0.\mu X_0 \begin{bmatrix} B_{2n} \cap \mathrm{LPreEven}_1(0, W, T_{i-1}, Y_{n-1}) \\ \cup \\ B_{2n-1} \cap \mathrm{LPreOdd}_1(1, W, T_{i-1}, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-2} \cap \mathrm{LPreEven}_1(1, W, T_{i-1}, Y_{n-1}, X_{n-2}, Y_{n-2}) \\ \cup \\ B_{2n-3} \cap \mathrm{LPreOdd}_1(2, W, T_{i-1}, Y_{n-1}, X_{n-1}, Y_{n-2}, X_{n-2}) \\ \vdots \\ B_2 \cap \mathrm{LPreEven}_1(n-1, W, T_{i-1}, Y_{n-1}, X_{n-1}, \dots, Y_1, X_1, Y_0) \\ \cup \\ B_1 \cap \mathrm{LPreOdd}_1(n, W, T_{i-1}, Y_{n-1}, X_{n-1}, \dots, Y_0, X_0) \end{bmatrix}$$

The proof that $W \subseteq \langle\langle 1 \rangle\rangle_{limit}\big(\mathrm{Parity}(p) \cup \Diamond T\big)$ is based on an induction on the sequence $T = T_0 \subset T_1 \subset T_2 \subset \cdots \subset T_m = W$. For $1 \leq i \leq m$, let $V_i = W \setminus T_{m-i}$, so that V_1 consists of the last block of states that has been added, V_2 to the two last blocks, and so on until $V_m = W$. We prove by induction on $i \in \{1, \dots, m\}$, from $i = 1$ to $i = m$, that for all $s \in V_i$, for all $\eta > 0$, there exists a strategy π_1^η for player 1 such that for all $\pi_2 \in \Pi_2$ we have

$$\mathrm{Pr}_s^{\pi_1^\eta, \pi_2}\big(\Diamond T_{m-i} \cup \mathrm{Parity}(p)\big) \geq 1 - \eta.$$

Since the base case is a simplified version of the induction step, we focus on the latter. There are two cases, depending on whether $V_i \setminus V_{i-1}$ is composed of B_{2n+1} or of $B_{\leq 2n}$-states.

- If $V_i \setminus V_{i-1} \subseteq B_{2n+1}$, then $V_i \setminus V_{i-1} = \{s\}$ for some $s \in S$ and $s \in \mathrm{Lpre}_1(W, T_{m-i})$. The result then follows from the application of the basic Lpre principle (Lemma 5) with $Z = W$, $X = T_{m-i}$, $Z \setminus Y = V_{i-1}$ and $\mathcal{A} = \mathrm{Parity}(p)$.

- If $V_i \setminus V_{i-1} \subseteq B_{\leq 2n}$, then we analyze the predecessor operator that $s \in V_i \setminus V_{i-1}$ satisfies. The predecessor operator are essentially the predecessor operator of the limit-expression for case 2 modified by the addition of the operator $\mathrm{Lpre}_1(W, T_{m-i}) \uplus_*$. If player 2 plays such the $\mathrm{Lpre}_1(W, T_{m-i})$ part of the predecessor operator gets satisfied, then the analysis reduces to the previous case, and player 1 can ensure that T_{m-i} is reached with probability close to 1. Once we rule out the possibility of $\mathrm{Lpre}_1(W, T_{m-i})$, then the μ-calculus expression simplifies to the limit-expression of case 2, i.e.,

$$\nu Y_{n-1}.\mu X_{n-1}.\cdots \nu Y_0.\mu X_0 \begin{bmatrix} B_{2n} \cap \mathrm{Pre}_1(Y_{n-1}) \\ \cup \\ B_{2n-1} \cap \mathrm{LPreOdd}_1(0, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-2} \cap \mathrm{LPreEven}_1(0, Y_{n-1}, X_{n-2}, Y_{n-2}) \\ \cup \\ B_{2n-3} \cap \mathrm{LPreOdd}_1(1, Y_{n-1}, X_{n-1}, Y_{n-2}, X_{n-2}) \\ \vdots \\ B_2 \cap \mathrm{LPreEven}_1(n-2, Y_{n-1}, X_{n-1}, \dots, Y_1, X_1, Y_0) \\ \cup \\ B_1 \cap \mathrm{LPreOdd}_1(n-1, Y_{n-1}, X_{n-1}, \dots, Y_0, X_0) \end{bmatrix}$$

This ensures that if we rule out $\text{Lpre}_1(W, T_{m-i})$ form the predecessor operators, then by inductive hypothesis (limit-expression for case 2) we have $L_i \subseteq \langle\langle 1 \rangle\rangle_{limit}(\text{Parity}(p))$, and if $\text{Lpre}_1(W, T_{m-i})$ is satisfied then T_{m-i} is ensured to reach with probability arbitrary close to 1. Hence player 1 can ensure that for all $\eta > 0$, for all $s \in V_i$, there is a strategy π_1^η for player 1 such that for all π_2 for player 2 we have

$$\text{Pr}_s^{\pi_1^\eta, \pi_2}(\Diamond T_{m-i} \cup \text{Parity}(p)) \geq 1 - \eta.$$

This completes the inductive proof. With $i = m$ we obtain that for all $\eta > 0$, there exists a strategy π_1^η such that for all states $s \in V_m = W$ and for all π_2 we have $\text{Pr}_s^{\pi_1^\eta, \pi_2}(\Diamond T_0 \cup \text{Parity}(p)) \geq 1 - \eta$. Since $T_0 = T$, the desired result follows. ∎

Lemma 21 *For a parity function $p : S \mapsto [1..2n+1]$ we have $Z \subseteq \neg\langle\langle 1 \rangle\rangle_{limit}(\text{Parity}(p))$, where Z is defined as follows:*

$$\mu Y_n.\nu X_n.\mu Y_{n-1}.\nu X_{n-1}.\cdots \mu Y_0.\nu X_0 \left[\begin{array}{c} B_{2n+1} \cap \mathit{FPreOdd}_2(0, Y_n, X_n) \\ \cup \\ B_{2n} \cap \mathit{FPreEven}_2(0, Y_n, X_n, Y_{n-1}) \\ \cup \\ B_{2n-1} \cap \mathit{FPreOdd}_2(1, Y_n, X_n, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-2} \cap \mathit{FPreEven}_2(1, Y_n, X_n, Y_{n-1}, X_{n-2}, Y_{n-2}) \\ \cup \\ B_{2n-3} \cap \mathit{FPreOdd}_2(2, Y_n, X_n, Y_{n-1}, X_{n-1}, Y_{n-2}, X_{n-2}) \\ \vdots \\ B_2 \cap \mathit{FPreEven}_2(n-1, Y_n, X_n, Y_{n-1}, X_{n-1}, \ldots, Y_1, X_1, Y_0) \\ \cup \\ B_1 \cap \mathit{FPreOdd}_2(n, Y_n, X_n, Y_{n-1}, X_{n-1}, \ldots, Y_0, X_0) \end{array} \right]$$

Proof. For $k \geq 0$, let Z_k be the set of states of level k in the above μ-calculus expression. We will show that in Z_k player 2 can ensure that either Z_{k-1} is reached with positive probability or else $\text{coParity}(p)$ is satisfied with probability arbitrarily close to 1. Since $Z_0 = \emptyset$, it would follow by induction that $Z_k \cap \langle\langle 1 \rangle\rangle_{limit}(\text{Parity}(p)) = \emptyset$ and the desired result will follow.

We obtain of Z_k from Z_{k-1} as follows:

$$\nu X_n.\mu Y_{n-1}.\nu X_{n-1}.\cdots \mu Y_0.\nu X_0 \left[\begin{array}{c} B_{2n+1} \cap \mathrm{FPreOdd}_2(0, Z_{k-1}, X_n) \\ \cup \\ B_{2n} \cap \mathrm{FPreEven}_2(0, Z_{k-1}, X_n, Y_{n-1}) \\ \cup \\ B_{2n-1} \cap \mathrm{FPreOdd}_2(1, Z_{k-1}, X_n, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-2} \cap \mathrm{FPreEven}_2(1, Z_{k-1}, X_n, Y_{n-1}, X_{n-2}, Y_{n-2}) \\ \cup \\ B_{2n-3} \cap \mathrm{FPreOdd}_2(2, Z_{k-1}, X_n, Y_{n-1}, X_{n-1}, Y_{n-2}, X_{n-2}) \\ \vdots \\ B_2 \cap \mathrm{FPreEven}_2(n-1, Z_{k-1}, X_n, Y_{n-1}, X_{n-1}, \ldots, Y_1, X_1, Y_0) \\ \cup \\ B_1 \cap \mathrm{FPreOdd}_2(n, Z_{k-1}, X_n, Y_{n-1}, X_{n-1}, \ldots, Y_0, X_0) \end{array}\right]$$

If player 1 risks into moving to Z_{k-1} with positive probability, then the inductive case is proved as Z_{k-1} is reached with positive probability. If the probability of reaching to Z_{k-1} is not positive, then the following conditions hold:

- $\mathrm{FPreOdd}_2(i, Z_{k-1}, X_n, Y_{n-1}, \ldots, Y_{n-i}, X_{n-i})$ simplifies to the predecessor operator $\mathrm{LPreEven}_2(i-1, X_{n-1}, Y_{n-2}, \ldots, Y_{n-i}, X_{n-i})$ and
- $\mathrm{FPreEven}_2(i, Z_{k-1}, X_n, Y_{n-1}, \ldots, Y_{n-i}, X_{n-i}, Y_{n-i-1})$ simplifies to the predecessor operator $\mathrm{LPreOdd}_2(i, X_{n-1}, Y_{n-2}, \ldots, Y_{n-i}, X_{n-i}, Y_{n-i-1})$.

Hence if we rule out the possibility of positive probability to reach Z_{k-1}, then the above μ-calculus expression simplifies to

$$Z^* = \nu X_n.\mu Y_{n-1}.\nu X_{n-1}.\cdots \mu Y_0.\nu X_0 \left[\begin{array}{c} B_{2n+1} \cap \mathrm{Pre}_2(X_n) \\ \cup \\ B_{2n} \cap \mathrm{LPreOdd}_2(0, X_n, Y_{n-1}) \\ \cup \\ B_{2n-1} \cap \mathrm{LPreEven}_2(0, X_n, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-2} \cap \mathrm{LPreOdd}_2(1, X_n, Y_{n-1}, X_{n-2}, Y_{n-2}) \\ \cup \\ B_{2n-3} \cap \mathrm{LPreEven}_2(1, X_n, Y_{n-1}, X_{n-1}, Y_{n-2}, X_{n-2}) \\ \vdots \\ B_2 \cap \mathrm{LPreOdd}_2(n-1, X_n, Y_{n-1}, X_{n-1}, \ldots, Y_1, X_1, Y_0) \\ \cup \\ B_1 \cap \mathrm{LPreEven}_2(n-1, X_n, Y_{n-1}, X_{n-1}, \ldots, Y_0, X_0) \end{array}\right].$$

We now consider the parity function $p-1 : S \mapsto [0..2n]$ and by applying the correctness of the limit-expression for step 1 (Lemma 18) with the roles of player 1 and player 2 exchanged we have $Z^* \subseteq \langle\langle 2 \rangle\rangle_{limit}\big(\mathrm{coParity}(p)\big)$ (since $\mathrm{coParity}(p) = \mathrm{Parity}(p-1)$). Hence the result follows. ∎

Correctness of step 3. The correctness of step 3 is similar to correctness of step 1.

Lemma 22 *For a parity function $p : S \mapsto [1..2n+2]$, and for all $T \subseteq S$, we have $W \subseteq \langle\langle 1\rangle\rangle_{limit}\big(Parity(p) \cup \Diamond T\big)$, where W is defined as follows:*

$$\nu Y_n.\mu X_n.\nu Y_{n-1}.\mu X_{n-1}.\cdots \nu Y_0.\mu X_0 \begin{bmatrix} T \\ \cup \\ B_{2n+2} \cap Pre_1(Y_n) \\ \cup \\ B_{2n+1} \cap LPreOdd_1(0, Y_n, X_n) \\ \cup \\ B_{2n} \cap LPreEven_1(0, Y_n, X_n, Y_{n-1}) \\ \cup \\ B_{2n-1} \cap LPreOdd_1(1, Y_n, X_n, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-2} \cap LPreEven_1(1, Y_n, X_n, Y_{n-1}, X_{n-2}, Y_{n-2}) \\ \cup \\ B_{2n-3} \cap LPreOdd_1(2, Y_n, X_n, Y_{n-1}, X_{n-1}, Y_{n-2}, X_{n-2}) \\ \vdots \\ B_2 \cap LPreEven_1(n-1, Y_n, X_n, Y_{n-1}, X_{n-1}, \ldots, Y_1, X_1, Y_0) \\ \cup \\ B_1 \cap LPreOdd_1(n, Y_n, X_n, Y_{n-1}, X_{n-1}, \ldots, Y_0, X_0) \end{bmatrix}$$

Proof. Similar to step 1 (Lemma 18), we add a max even priority. The proof of the result is essentially similar to the proof of Lemma 18, the only modification is instead of the correctness of the limit-expression of case 1 we need to consider the correctness of the limit-expression for step 2 (i.e., Lemma 20 for parity function $p : S \mapsto [1..2n+1]$). ∎

Lemma 23 *For a parity function $p : S \mapsto [1..2n+2]$ we have $Z \subseteq \neg\langle\langle 1\rangle\rangle_{limit}\big(Parity(p)\big)$, where Z is defined as follows:*

$$\mu Y_n.\nu X_n.\mu Y_{n-1}.\nu X_{n-1}.\cdots \mu Y_0.\nu X_0 \begin{bmatrix} B_{2n+2} \cap Epre_2(Y_n) \\ \cup \\ B_{2n+1} \cap FPreOdd_2(0, Y_n, X_n) \\ \cup \\ B_{2n} \cap FPreEven_2(0, Y_n, X_n, Y_{n-1}) \\ \cup \\ B_{2n-1} \cap FPreOdd_2(1, Y_n, X_n, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-2} \cap FPreEven_2(1, Y_n, X_n, Y_{n-1}, X_{n-2}, Y_{n-2}) \\ \cup \\ B_{2n-3} \cap FPreOdd_2(2, Y_n, X_n, Y_{n-1}, X_{n-1}, Y_{n-2}, X_{n-2}) \\ \vdots \\ B_2 \cap FPreEven_2(n-1, Y_n, X_n, Y_{n-1}, X_{n-1}, \ldots, Y_1, X_1, Y_0) \\ \cup \\ B_1 \cap FPreOdd_2(n, Y_n, X_n, Y_{n-1}, X_{n-1}, \ldots, Y_0, X_0) \end{bmatrix}$$

Proof. The proof of the result is essentially similar to the proof of Lemma 19, the only modification is instead of the correctness of the limit-expression of case 2 we need to consider the correctness of the limit-expression for step 1 (i.e., Lemma 18). This is because in the proof, after we rule out states in B_{2n+2} and analyze the sub-formula as in Lemma 18, we consider parity function $p - 1 : S \mapsto [0..2n]$ and then invoke the correctness of Lemma 18. ∎

Correctness of step 4. The correctness of step 4 is similar to correctness of step 2.

Lemma 24 *For a parity function $p : S \mapsto [0..2n+1]$, and for all $T \subseteq S$, we have $W \subseteq \langle\langle 1 \rangle\rangle_{limit}\big(Parity(p) \cup \Diamond T\big)$, where W is defined as follows:*

$$\nu Y_{n+1}.\mu X_{n+1}.\cdots \nu Y_1.\mu X_1.\nu Y_0. \begin{bmatrix} T \\ \cup \\ B_{2n+1} \cap LPreOdd_1(0, Y_{n+1}, X_{n+1}) \\ \cup \\ B_{2n} \cap LPreEven_1(0, Y_{n+1}, X_{n+1}, Y_n) \\ \cup \\ B_{2n-1} \cap LPreOdd_1(1, Y_{n+1}, X_{n+1}, Y_n, X_n) \\ \cup \\ B_{2n-2} \cap LPreEven_1(1, Y_{n+1}, X_{n+1}, Y_n, X_n, Y_{n-1}) \\ \cup \\ B_{2n-3} \cap LPreOdd_1(2, Y_{n+1}, X_{n+1}, Y_n, X_n, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-4} \cap LPreEven_1(2, Y_{n+1}, X_{n+1}, Y_n, X_n, Y_{n-1}, X_{n-1}, Y_{n-2}) \\ \vdots \\ B_1 \cap LPreOdd_1(n, Y_{n+1}, X_{n+1}, Y_n, X_n, \ldots, Y_1, X_1) \\ \cup \\ B_0 \cap LPreEven_1(n, Y_{n+1}, X_{n+1}, Y_n, X_n, \ldots, Y_1, X_1, Y_0) \end{bmatrix}$$

Proof. Similar to step 2 (Lemma 20), we add a max odd priority. The proof of the result is essentially similar to the proof of Lemma 20, the only modification is instead of the correctness of the limit-expression of case 2 we need to consider the correctness of the limit-expression for step 1 (i.e., Lemma 18 for parity function $p : S \mapsto [0..2n]$). ∎

Lemma 25 *For a parity function $p : S \mapsto [0..2n+1]$ we have $Z \subseteq \neg\langle\langle 1 \rangle\rangle_{limit}\big(Parity(p)\big)$, where Z*

is defined as follows:

$$\mu Y_{n+1}.\nu X_{n+1}.\cdots\mu Y_1.\nu X_1.\mu Y_0.\left[\begin{array}{c} B_{2n+1}\cap \mathit{FPreOdd}_2(0,Y_{n+1},X_{n+1}) \\ \cup \\ B_{2n}\cap \mathit{FPreEven}_2(0,Y_{n+1},X_{n+1},Y_n) \\ \cup \\ B_{2n-1}\cap \mathit{FPreOdd}_2(1,Y_{n+1},X_{n+1},Y_n,X_n) \\ \cup \\ B_{2n-2}\cap \mathit{FPreEven}_2(1,Y_{n+1},X_{n+1},Y_n,X_n,Y_{n-1}) \\ \cup \\ B_{2n-3}\cap \mathit{FPreOdd}_2(2,Y_{n+1},X_{n+1},Y_n,X_n,Y_{n-1},X_{n-1}) \\ \cup \\ B_{2n-4}\cap \mathit{FPreEven}_2(2,Y_{n+1},X_{n+1},Y_n,X_n,Y_{n-1},X_{n-1},Y_{n-2}) \\ \vdots \\ B_1\cap \mathit{FPreOdd}_2(n,Y_{n+1},X_{n+1},Y_n,X_n,\ldots,Y_1,X_1) \\ \cup \\ B_0\cap \mathit{FPreEven}_2(n,Y_{n+1},X_{n+1},Y_n,X_n,\ldots,Y_1,X_1,Y_0) \end{array}\right]$$

Proof. The proof of the result is essentially similar to the proof of Lemma 21, the only modification is instead of the correctness of the limit-expression of step 1 (Lemma 18) we need to consider the correctness of the limit-expression for step 3 (i.e., Lemma 22). This is because in the proof, while we analyze the sub-formula as in Lemma 22, we consider parity function $p+1 : S \mapsto [1..2n+2]$ and then invoke the correctness of Lemma 22. ∎

7.3 Duality of predecessor operators

In this section we prove the duality of the predecessor operators of subsection 7.1. We present the details of result for $\mathrm{Lpre}_1(Y_1,X_1) \uplus_{*} \mathrm{Lpre}_1(Y_0,X_0)$ and its complement; the general result follows similarly by an induction.

Lemma 26 *Given $X_1 \subseteq X_0 \subseteq Y_0 \subseteq Y_1 \subseteq S$, and $s \in S$, let*

$$B = \mu W_2 \,.\, \nu W_1 \,.\, \mu W_0 \,.\left[\begin{array}{c} \mathsf{Stay}_1(s,Y_1,W_0)\cap \mathsf{Stay}_1(s,Y_0,W_1) \\ \cup \\ \mathsf{Cover}_1(s,X_1,W_0)\cap \mathsf{Cover}_1(s,X_0,W_2) \end{array}\right].$$

If $\Gamma_2(s) \subseteq B$, then $s \in Lpre_1(Y_1,X_1) \uplus_{*} Lpre_2(Y_0,X_0)$.

Proof. We first analyze the computation of B. The set B of moves is obtained as follows:

$$\emptyset = W_2^0 \subseteq W_2^1 \subseteq W_2^2 \subseteq \cdots \subseteq W_2^k = B;$$

and the set W_2^{i+1} is obtained from W_2^i as follows:

$$W_2^{i+1} = \nu W_1 \,.\, \mu W_0 \,.\left[\begin{array}{c} \mathsf{Stay}_1(s,Y_1,W_0)\cap \mathsf{Stay}_1(s,Y_0,W_1) \\ \cup \\ \mathsf{Cover}_1(s,X_1,W_0)\cap \mathsf{Cover}_1(s,X_0,W_2^i) \end{array}\right].$$

Alternatively we have

$$W_2^{i+1} = \mu W_0 \,.\, \left[\begin{array}{c} \mathsf{Stay}_1(s, Y_1, W_0) \cap \mathsf{Stay}_1(s, Y_0, W_2^{i+1}) \\ \cup \\ \mathsf{Cover}_1(s, X_1, W_0) \cap \mathsf{Cover}_1(s, X_0, W_2^{i}) \end{array} \right].$$

Equivalently, we can characterize the computation of W_2^{i+1} as follows:

$$\begin{array}{rcl} W_2^{i+1,0} &=& W_2^i \,\cup\, \mathsf{Cover}_1(s, X_0, W_2^i) \quad \text{(we call moves in } W_2^{i+1,0} \setminus W_2^i \text{ as “openers”);} \\ W_2^{i+1,j+1} &=& W_2^{i+1,j} \,\cup\, (\mathsf{Stay}_1(s, Y_1, W_2^{i+1,j}) \cap \mathsf{Stay}_1(s, Y_0, W_2^{i+1})) \cup \mathsf{Cover}_1(s, X_1, W_2^{i+1,j}). \end{array}$$

Properties. We now describe the following key properties of the moves.

1. for all $b \in W_2^{i+1,0} \setminus W_2^i$, there exists $a \in W_2^i$ such that $Dest(s, a, b) \cap X_0 \neq \emptyset$ (by property of $\mathsf{Cover}_1(s, X_0, W_2^i)$).
2. for all $b \in W_2^{i+1,j+1} \setminus W_2^{i+1,j}$, there exists $a \in W_2^{i+1,j}$ such that $Dest(s, a, b) \cap X_1 \neq \emptyset$ (by property of $\mathsf{Cover}_1(s, X_1, W_2^{i+1,j})$).
3. for all $a \in W_2^{i+1}$, for all $b \in \Gamma_2(s) \setminus W_2^{i+1}$ we have $Dest(s, a, b) \subseteq Y_0 \subseteq Y_1$ (by property of $\mathsf{Stay}_1(s, Y_0, W_2^{i+1})$).
4. for all $a \in W_2^{i+1,j+1}$, for all $b \in \Gamma_2(s) \setminus W_2^{i+1,j}$ we have $Dest(s, a, b) \subseteq Y_1$ (by property of $\mathsf{Stay}_1(s, Y_1, W_2^{i+1,j})$).

For all $0 < \varepsilon < 1$, consider a distribution for player 1 $\xi_1[\varepsilon]$ that plays moves in $W_2^{i+1,j+1} \setminus W_2^{i+1,j}$ with probability proportional to $\varepsilon^{|W_2^{i+1,j}|+1}$. For a move $b \in \Gamma_2(s)$ we have the following case analysis.

1. If $b \in W_2^{i+1,0}$ for some i (i.e., b is an opener move), then (a) there exists $a \in W_2^i$ with $Dest(s, a, b) \cap X_0 \neq \emptyset$ (property 1), i.e., the probability of going to X_0 is at least proportional to $\varepsilon^{|W_2^i|}$; and (b) for all $a \in W_2^i$ we have $Dest(s, a, b) \subseteq Y_0 \subseteq Y_1$ (property 3), i.e., the probability of leaving Y_0 (or Y_1) is at most proportional to $\varepsilon^{|W_2^i|+1}$.
2. If $b \in W_2^{i+1,j+1}$ for some i, j, then (a) there exists $a \in W_2^{i+1,j}$ with $Dest(s, a, b) \cap X_1 \neq \emptyset$ (property 2), i.e., the probability of going to X_1 is at least proportional to $\varepsilon^{|W_2^{i+1,j}|}$; and (b) for all $a \in W_2^{i+1,j}$ we have $Dest(s, a, b) \subseteq Y_1$ (property 4), i.e., the probability of leaving Y_1 is at most proportional to $\varepsilon^{|W_2^{i+1,j}|+1}$.

It follows from above that given the distribution $\xi_1[\varepsilon]$ for player 1, and for a distribution ξ_2 for player 2, (a) if ξ_2 plays only openers with positive probability then the probability of going to X_0 as compared to leaving Y_0 is proportional to $\frac{1}{\varepsilon}$; (b) otherwise, the probability of going to X_1 as compared to leaving Y_1 is proportional to $\frac{1}{\varepsilon}$. Since $\varepsilon > 0$ is arbitrary, it follows that $s \in \mathrm{Lpre}_1(Y_1, X_1) \uplus \mathrm{Lpre}_1(Y_0, X_0)$. ∎

Lemma 27 *Given $Y_1 \subseteq Y_0 \subseteq X_0 \subseteq X_1 \subseteq S$, and $s \in S$, let*

$$C = \nu W_2 \,.\, \mu W_1 \,.\, \nu W_0 \,.\, \left[\begin{array}{c} \mathsf{Cover}_2(s, Y_1, W_0) \cup \mathsf{Cover}_1(s, Y_0, W_1) \\ \cup \\ \mathsf{Stay}_2(s, X_1, W_0) \cap \mathsf{Stay}_2(s, X_0, W_2) \end{array} \right].$$

If $\Gamma_2(s) \cap B \neq \emptyset$, then $s \in \big(Fpre_2(X_1, Y_1) \circledast Epre_2(Y_1)\big) \uplus Lpre_2(X_1, Y_0) \uplus Pre_2(X_0)$.

Proof. We first observe that for moves $a \in \Gamma_1(s) \setminus C$, and $b \in C$ we have $Dest(s,a,b) \subseteq X_0 \subseteq X_1$ (this follows from the condition $\mathsf{Stay}_2(s, X_0, W_2)$). Hence by playing a distribution with support subset of C player 2 can ensure that for moves of player 1 outside C the set X_0 and also X_1 is not left. We have

$$
\begin{aligned}
C &= \mu W_1 \,.\, \nu W_0 \,.\, \left[\begin{array}{c} \mathsf{Cover}_2(s, Y_1, W_0) \cup \mathsf{Cover}_1(s, Y_0, W_1) \\ \cup \\ \mathsf{Stay}_2(s, X_1, W_0) \cap \mathsf{Stay}_2(s, X_0, C) \end{array} \right] \\
&\subseteq \mu W_1 \,.\, \nu W_0 \,.\, \left[\begin{array}{c} \mathsf{Cover}_2(s, Y_1, W_0) \cup \mathsf{Cover}_1(s, Y_0, W_1) \\ \cup \\ \mathsf{Stay}_2(s, X_1, W_0) \end{array} \right]
\end{aligned}
$$

Compare the above μ-calculus with the formula of Lemma 13, and by Remark 2 it follows that player 2 can play distribution with support subset of C such that if player 1 plays a distribution with support subset of C, then $\big(\mathrm{Fpre}_2(X_1, Y_1) \mathbin{\Cap_{*}} \mathrm{Epre}_2(Y_1)\big) \mathbin{\Cup_{*}} \mathrm{Lpre}_2(X_1, Y_0)$ can be ensured. It follows that $s \in \big(\mathrm{Fpre}_2(X_1, Y_1) \mathbin{\Cap_{*}} \mathrm{Epre}_2(Y_1)\big) \mathbin{\Cup_{*}} \mathrm{Lpre}_2(X_1, Y_0) \mathbin{\Cup_{*}} \mathrm{Pre}_2(X_0)$. ∎

From Lemmas 26 and 27 we obtain the following lemma.

Lemma 28 *Given $X_1 \subseteq X_0 \subseteq Y_0 \subseteq Y_1 \subseteq S$, and $s \in S$, the following assertions hold.*

1. *Let*

$$
B = \mu W_2 \,.\, \nu W_1 \,.\, \mu W_0 \,.\, \left[\begin{array}{c} \mathsf{Stay}_1 s, (Y_1, W_0) \cap \mathsf{Stay}_1(s, Y_0, W_1) \\ \cup \\ \mathsf{Cover}_1(s, X_1, W_0) \cap \mathsf{Cover}_1(s, X_0, W_2) \end{array} \right].
$$

   We have $s \in Lpre_1(Y_0, X_0) \mathbin{\Cup_{*}} Lpre_1(Y_1, X_1)$ iff $\Gamma_2(s) \subseteq B$.

2. *We have*

$$
\big(Fpre_2(\neg X_1, \neg Y_1) \mathbin{\Cap_{*}} Epre_2(\neg Y_1)\big) \mathbin{\Cup_{*}} Lpre_2(\neg X_1, \neg Y_0) \mathbin{\Cup_{*}} Pre_2(\neg X_0) = \neg(Lpre_1(Y_0, X_0) \mathbin{\Cup_{*}} Lpre_1(Y_1, X_1)).
$$

Generalization. In the proof of Lemma 26 the basic idea is as follows: the computation of the set B is obtained as chunks of sets W_2^j, within the chunks of a W_2^j the condition for $\mathrm{Lpre}_1(Y_1, X_1)$ was satisfied, and across the chunks of W_2^j the condition for $\mathrm{Lpre}_1(Y_0, X_0)$ is satisfied. In the general case for $\mathrm{LPreOdd}_1(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i})$ we have a similar μ-calculus formula for a set B of moves that is obtained as chunks of W_{2i}^j, such that within the chunks of W_{2i}^j the condition for $\mathrm{LPreOdd}_1(i-1, Y_n, X_n, \ldots, Y_{n-i-1}, X_{n-i-1})$ is satisfied and across the chunks the condition for $\mathrm{Lpre}_1(Y_{n-i}, X_{n-i})$ is satisfied. The proof structures is essentially similar to Lemma 26. Lemmas 26 and 27 present the basic arguments of the inductive step to generalize Lemma 28 to obtain the following result.

Lemma 29 *Given $X_n \subseteq X_{n-1} \subseteq \cdots \subseteq X_{n-i} \subseteq Y_{n-i} \subseteq Y_{n-i+1} \subseteq \cdots \subseteq Y_n$ and $s \in S$, the following assertions hold:*

1. *Let*

$$
B = \mu W_{2i} \,.\, \nu W_{2i-1} \,.\, \cdots \,.\, \mu W_0 \,.\, \left[\begin{array}{c} \big(\mathsf{Stay}_1(s, Y_n, W_0) \cap \bigcap_{j=1}^{i} \mathsf{Stay}_1(s, Y_{n-j}, W_{2j-1})\big) \\ \cup \\ \bigcap_{j=0}^{i} \mathsf{Cover}_1(s, X_{n-j}, W_{2j}) \end{array} \right].
$$

   We have $s \in LPreOdd_1(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i})$ iff $\Gamma_2(s) \subseteq B$.

2. *We have*

$$FPreOdd_2(i, \neg Y_n, \neg X_n, \ldots, \neg Y_{n-i}, \neg X_{n-i}) = \neg LPreOdd_1(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i}).$$

The basic principle to obtain Lpre $\circledast$ Pre from Lpre also extends to obtain $\mathrm{LPreEven}_1$ from LPreOdd. This gives us the following lemma.

Lemma 30 *Given $X_n \subseteq X_{n-1} \subseteq \cdots \subseteq X_{n-i} \subseteq Y_{n-i-1} \subseteq Y_{n-i} \subseteq Y_{n-i+1} \subseteq \cdots \subseteq Y_n$ and $s \in S$, the following assertions hold:*

1. *Let*

$$B = \nu W_{2i+1} \,.\, \mu W_{2i} \,.\, \nu W_{2i-1} \,.\, \cdots \,.\, \mu W_0 \,.\, \left[\begin{array}{c} \big(\mathsf{Stay}_1(s, Y_n, W_0) \cap \bigcap_{j=1}^{i} \mathsf{Stay}_1(s, Y_{n-j}, W_{2j-1}) \\ \cap\ \mathsf{Stay}_1(s, Y_{n-i-1}, W_{2i+1})\big) \\ \cup \\ \bigcap_{j=0}^{i} \mathsf{Cover}_1(s, X_{n-j}, W_{2j}) \end{array} \right].$$

We have $s \in LPreEven_1(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i}, Y_{n-i-1})$ iff $\Gamma_1(s) \cap B \neq \emptyset$.

2. *We have*

$$FPreEven_2(i, \neg Y_n, \neg X_n, \ldots, \neg Y_{n-i}, \neg X_{n-i}, \neg Y_{n-i-1}) = \neg LPreEven_1(i, Y_n, X_n, \ldots, Y_{n-i}, X_{n-i}, Y_{n-i-1}).$$

Characterization of limit winning set. From Lemmas 18—25, and the duality of predecessor operators (Lemmas 29 and 30) we obtain the following result characterizing the limit-winning sets for Rabin-chain objectives.

Theorem 5 *For all concurrent game structures $\mathcal{G}$, for all parity objectives $Parity(p)$ for player 1, the following assertions hold.*

1. *If $p : S \mapsto [0..2n-1]$, then $\langle\langle 1 \rangle\rangle_{limit}\big(Parity(p)\big) = W$, where W is defined as follows*

$$\nu Y_n.\mu X_n.\cdots \nu Y_1.\mu X_1.\nu Y_0. \left[\begin{array}{c} B_{2n-1} \cap LPreOdd_1(0, Y_n, X_n) \\ \cup \\ B_{2n-2} \cap LPreEven_1(0, Y_n, X_n, Y_{n-1}) \\ \cup \\ B_{2n-3} \cap LPreOdd_1(1, Y_n, X_n, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-4} \cap LPreEven_1(1, Y_n, X_n, Y_{n-1}, X_{n-1}, Y_{n-2}) \\ \vdots \\ B_1 \cap LPreOdd_1(n-1, Y_n, X_n, \ldots, Y_1, X_1) \\ \cup \\ B_0 \cap LPreEven_1(n-1, Y_n, X_n, \ldots, Y_1, X_1, Y_0) \end{array} \right] \tag{24}$$

	Complexity	Winning	Spoiling
Safety	$\mathcal{O}(\lvert\mathcal{G}\rvert)$	Π_1^{DM}	Π_2^{M}
Reachability	$\mathcal{O}(\lvert\mathcal{G}\rvert^2)$	Π_1^{M}	Π_2^{M}
Büchi	$\mathcal{O}(\lvert\mathcal{G}\rvert^2)$	Π_1^{H}	Π_2^{M}
coBüchi	$\mathcal{O}(\lvert\mathcal{G}\rvert^4)$	Π_1^{M}	Π_2^{H}
parity	$\mathcal{O}(\lvert\mathcal{G}\rvert^{2m+2})$	Π_1^{H}	Π_2^{H}

Table 1: Complexity, and types of limit-winning and limit-spoiling strategies for Rabin-chain games; $|\mathcal{G}|$ is the size of the game, and $2m$ is the number parities.

2. *If $p : S \mapsto [1..2n]$, then $\langle\langle 1\rangle\rangle_{limit}\big(Parity(p)\big) = W$, where W is defined as follows*

$$\nu Y_{n-1}.\mu X_{n-1}.\cdots\nu Y_0.\mu X_0 \left[\begin{array}{c} B_{2n} \cap Pre_1(Y_{n-1}) \\ \cup \\ B_{2n-1} \cap LPreOdd_1(0, Y_{n-1}, X_{n-1}) \\ \cup \\ B_{2n-2} \cap LPreEven_1(0, Y_{n-1}, X_{n-2}, Y_{n-2}) \\ \cup \\ B_{2n-3} \cap LPreOdd_1(1, Y_{n-1}, X_{n-1}, Y_{n-2}, X_{n-2}) \\ \vdots \\ B_2 \cap LPreEven_1(n-2, Y_{n-1}, X_{n-1}, \ldots, Y_1, X_1, Y_0) \\ \cup \\ B_1 \cap LPreOdd_1(n-1, Y_{n-1}, X_{n-1}, \ldots, Y_0, X_0) \end{array}\right] \tag{25}$$

3. *The set of limit-winning states can be computed using the relations (24) and (25) in time $\mathcal{O}(|S|^{2n+1} \cdot \sum_{s\in S} |\Gamma_1(s) \cup \Gamma_2(s)|^{2n})$, where $|S|$ is the number of states of the game.*

4. *Limit-winning strategies for player 1 require infinite-memory in general.*

5. *Limit-spoiling strategies for player 2 require infinite-memory in general.*

The time complexity of item Theorem 5 is obtained as follows. A μ-calculus formula of nesting depth $2n$ converges in at most $|S|^{2n}$ iterations to the fixpoint. The convergence of μ-calculus formula, and the fact that where a state $s \in \text{LPreOdd}_1$ (resp. $s \in \text{LPreEven}_1$) with $2n$-arguments (resp. $2n + 1$ arguments) can be decided in time $\mathcal{O}(|\Gamma_1(s) \cup \Gamma_2(s)|^{2n})$ gives the desired time complexity of Theorem 5. Table 1 summarizes the types of the limit-winning and limit-spoiling strategies that can be found, along with the complexity of computing the set of limit-winning states.

7.4 Strategy complexity and computational complexity

Strategy constructions. The limit-winning strategies for parity objectives requires infinite-memory (for counting the number of rounds). However, similar to the construction of limit-winning strategies for Büchi objectives, there exists limit-winning strategies that in the limit converge to a memoryless strategy. However, the memoryless strategy to which the sequence of limit-winning strategies converges is not limit-winning in general (this fact also holds for reachability objectives, this follows from Example 1). We formalize this in a lemma below. Example 2 shows that limit-winning strategies require infinite-memory in general.

Lemma 31 *Let $\mathcal{G}$ be a concurrent game structure with a parity objective Parity(p), and let $W = \langle\langle 1\rangle\rangle_{limit}(Parity(p))$ be the limit-winning states for player 1. For all $s \in W$, there exist $A(s) \subseteq \Gamma_1(s)$ such that the following condition hold.*

> *For any memoryless strategy $\pi_1^m \in \Pi_1^M$ such that $\mathrm{Supp}(\pi_1(s)) = A(s)$ for all $s \in W$, we have that there exists a sequence of $(\varepsilon_i)_{i\geq 0}$ such that:*
>
> - *$\varepsilon_0 > \varepsilon_1 > \varepsilon_2 > \ldots$ and $\lim_{i\to\infty} \varepsilon_i = 0$,*
> - *for all $i \geq 0$ there exists a strategy $\pi_1^{\varepsilon_i}$ such that for all $\pi_2 \in \Pi_2$ and for all $s \in W$ we have $\mathrm{Pr}_s^{\pi_1^{\varepsilon_i},\pi_2}(Parity(p)) \geq 1 - \varepsilon_i$, and*
> - *the strategies $\pi_1^{\varepsilon_i}$ converges to π_1^m as $i \to \infty$, i.e., $\lim_{i\to\infty} \pi_1^{\varepsilon_i} = \pi_1^m$.*

Witness of limit-winning strategies. The witness strategy for a limit-winning strategy as constructed in lemmas of subsection 7.2 can be described in two parts: a *ranking function* of the states, and a *ranking function* of the actions at a state. These ranking functions are described by μ-calculus formulas. At the round k of a play, the witness strategy π_1 plays at a state s in the limit-winning set, the actions with least rank with positive-bounded probabilities, and the other actions with vanishingly small probabilities as $\varepsilon \to 0$. Hence, the strategy π_1 can be described as

$$\pi_1 = (1 - \varepsilon_k)\pi_1^\ell + \varepsilon_k \cdot \pi_1^d(\varepsilon_k),$$

where π_1^ℓ is a memoryless strategy such that, at each state s, $\mathrm{Supp}(\pi_1^\ell(s))$ is the set of actions with least rank at s (as stated in Lemma 31). The μ-calculus formula and the ranking of the states and the actions can be described as polynomial witness and can be verified in polynomial time. This shows that whether a state is limit-winning for a parity objective can be decided in NP. The existence of polynomial witness for limit-spoiling strategies for player 2 and polynomial time verification procedure is similar, which shows that the problem is also in coNP. This gives us Theorem 6.

Theorem 6 *Given a concurrent game structure $\mathcal{G}$, a parity objective Parity(p), and a state $s \in S$, whether $s \in \langle\langle 1\rangle\rangle_{limit}(Parity(p))$ can be decided in NP $\cap$ coNP.*

7.5 Properties: duality and bounded-limit law

The duality of concurrent games with parity objectives is stated in the corollary. The result follows from the characterization of limit-winning sets as μ-calculus formulas, and the limit-spoiling states as the complementary μ-calculus formulas.

Corollary 1 *For all concurrent game structures $\mathcal{G}$ and all parity objectives Parity(p), we have $\langle\langle 1\rangle\rangle_{limit}(Parity(p)) = S \setminus \langle\langle 2\rangle\rangle_{bounded}(coParity(p))$.*

The following corollary establishes relation between the bounded and limit-winning states in concurrent games with parity objectives.

Corollary 2 (Limit-law for parity games) *Let $\mathcal{G}$ be a concurrent games structure with a parity objective Parity(p). The following assertions hold.*

1. *If $\langle\langle 2\rangle\rangle_{limit}\big(coParity(p)\big) = \emptyset$, then $\langle\langle 1\rangle\rangle_{limit}\big(Parity(p)\big) = S$.*
2. *If $\langle\langle 1\rangle\rangle_{limit}\big(Parity(p)\big) = S$, then $\langle\langle 1\rangle\rangle_{limit}\big(Parity(p)\big) = S$.*
3. *If $\langle\langle 1\rangle\rangle_{bounded}\big(Parity(p)\big) \neq \emptyset$, then $\langle\langle 1\rangle\rangle_{limit}\big(Parity(p)\big) \neq \emptyset$.*

Proof. We prove the results as follows.

1. Consider the proof of Lemma 19, Lemma 21, Lemma 23, and Lemma 25. Let Z_1 be the set of states obtained in the first iteration of the compuatation of the respective μ-calculus expression. It follows from the proof of the lemmas that $Z_1 \subseteq \langle\langle 2\rangle\rangle_{limit}\big(\text{coParity}(p)\big)$. If $\langle\langle 1\rangle\rangle_{limit}\big(\text{Parity}(p)\big) \neq S$, then we have $Z_0 \neq \emptyset$, and hence the result follows.
2. If $\langle\langle 1\rangle\rangle_{limit}\big(\text{Parity}(p)\big) = S$, then by duality (Corollary 1) we have $\langle\langle 2\rangle\rangle_{bounded}\big(\text{coParity}(p)\big) = \emptyset$, and hence $\langle\langle 2\rangle\rangle_{limit}\big(\text{coParity}(p)\big) = \emptyset$. The result then follows from part 1.
3. If $\langle\langle 1\rangle\rangle_{limit}\big(\text{Parity}(p)\big) = \emptyset$, then by part 1 we have $\langle\langle 2\rangle\rangle_{limit}\big(\text{coParity}(p)\big) = S$ (by exchanging roles of palyer 1 and player 2 in part 1). Then we have $\langle\langle 1\rangle\rangle_{bounded}\big(\text{Parity}(p)\big) = \emptyset$. ∎

Independence from precise probabilities. Observe that the computation of all the predecessor operators only depends on the supports of the transition function, and does not depend on the precise transition probabilities. Hence the computation of the limit-winning sets is independent of the precise transition probabilities, and depends only on the support. We formalize this in the following result.

Theorem 7 *Let $G_1 = (S, A, \Gamma_1, \Gamma_2, \delta_1)$ and $G_2 = (S, A, \Gamma_1, \Gamma_2, \delta_2)$ be two concurrent game structures with the same set S of states, same set A of moves, and same move assignment functions Γ_1 and Γ_2. If for all $s \in S$, for all $a_1 \in \Gamma_1(s)$ and $a_2 \in \Gamma_2(s)$ we have $\text{Supp}(\delta_1(s, a_1, a_2)) = \text{Supp}(\delta_2(s, a_1, a_2))$, then for all parity objectives Parity(p), the set of limit-winning states for Parity(p) in G_1 and G_2 coincide.*

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 specifications of reactive systems. In *ICALP'89*, LNCS 372, pages 1–17. Springer, 1989.

[BCM+90] J.R. Burch, E.M. Clarke, K.L. McMillan, D.L. Dill, and L.J. Hwang. Symbolic model checking: 10^{20} states and beyond. In *Proc. 5th IEEE Symp. Logic in Comp. Sci.*, pages 428–439. IEEE Computer Society Press, 1990.

[BL69] J.R. Büchi and L.H. Landweber. Solving sequential conditions by finite-state strategies. *Trans. Amer. Math. Soc.*, 138:295–311, 1969.

[BLV96] N. Buhrke, H. Lescow, and J. Vöge. Strategy construction in infinite games with strett and rabin chain winning conditions. In *TACAS 96*, volume 1055 of *Lect. Notes in Comp. Sci.*, pages 207–225. Springer, 1996.

[Chu62] A. Church. Logic, arithmetic, and automata. In *Proceedings of the International Congress of Mathematicians*, pages 23–35. Institut Mittag-Leffler, 1962.

[Con93] A. Condon. On algorithms for simple stochastic games. In *Advances in Computational Complexity Theory*, volume 13 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 51–73. American Mathematical Society, 1993.

[dAH00] L. de Alfaro and T.A. Henzinger. Concurrent omega-regular games. In *LICS'00*, pages 141–154. IEEE, 2000.

[dAHK98] L. de Alfaro, T.A. Henzinger, and O. Kupferman. Concurrent reachability games. In *Proc. 39th IEEE Symp. Found. of Comp. Sci.*, pages 564–575, 1998.

[dAHK07] L. de Alfaro, T.A. Henzinger, and O. Kupferman. Concurrent reachability games. *Theoretical Computer Science*, 386(3):188–217, 2007.

[dAHM00] L. de Alfaro, T.A. Henzinger, and F.Y.C. Mang. The control of synchronous systems. In *CONCUR'00*, Lecture Notes in Computer Science 1877, pages 458–473. Springer, 2000.

[dAHM01] L. de Alfaro, T.A. Henzinger, and F.Y.C. Mang. The control of synchronous systems, part ii. In *CONCUR'01*, Lecture Notes in Computer Science 2154, pages 566–580. Springer, 2001.

[Dil89] D.L. Dill. *Trace Theory for Automatic Hierarchical Verification of Speed-independent Circuits*. The MIT Press, 1989.

[EJ91] E.A. Emerson and C.S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *32nd Symp. on Foundations of Computer Science (FOCS)*, pages 368–377, 1991.

[FV97] J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer-Verlag, 1997.

[GH82] Y. Gurevich and L. Harrington. Trees, automata, and games. In *Proc. 14th ACM Symp. Theory of Comp.*, 1982.

[Kec95] A. Kechris. *Classical Descriptive Set Theory*. Springer, 1995.

[KPBV95] S. Krishnan, A. Puri, R. Brayton, and P. Varaiya. Rabin index, chain automata and applications to automata and games. In *Proc. 7^{th} Intl. Conference on Computer Aided Verification*, volume 939 of *Lect. Notes in Comp. Sci.* Springer-Verlag, 1995.

[KS81] P.R. Kumar and T.H. Shiau. Existence of value and randomized strategies in zero-sum discrete-time stochastic dynamic games. *SIAM J. Control and Optimization*, 19(5):617–634, 1981.

[KSK66] J.G. Kemeny, J.L. Snell, and A.W. Knapp. *Denumerable Markov Chains*. D. Van Nostrand Company, 1966.

[LW95] H. Lescow and T. Wilke. On polynomial-size programs winning finite-state games. In *Proc. 7^{th} Intl. Conference on Computer Aided Verification*, volume 939 of *Lect. Notes in Comp. Sci.*, pages 239–252. Springer-Verlag, 1995.

[McN93] R. McNaughton. Infinite games played on finite graphs. *Ann. Pure Appl. Logic*, 65:149–184, 1993.

[Mos84] A.W. Mostowski. Regular expressions for infinite trees and a standard form of automata. In *Computation Theory*, volume 208 of *Lect. Notes in Comp. Sci.*, pages 157–168. Springer-Verlag, 1984.

[PR89] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL'89*, pages 179–190. ACM Press, 1989.

[RF91] T.E.S. Raghavan and J.A. Filar. Algorithms for stochastic games — a survey. *ZOR — Methods and Models of Op. Res.*, 35:437–472, 1991.

[RW87] P.J. Ramadge and W.M. Wonham. Supervisory control of a class of discrete-event processes. *SIAM Journal of Control and Optimization*, 25(1):206–230, 1987.

[Sha53] L.S. Shapley. Stochastic games. *Proc. Nat. Acad. Sci. USA*, 39:1095–1100, 1953.

[Tho90] W. Thomas. Automata on infinite objects. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, chapter 4, pages 135–191. Elsevier Science Publishers (North-Holland), Amsterdam, 1990.

[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 *Lect. Notes in Comp. Sci.*, pages 1–13. Springer-Verlag, 1995.

[Var85] M.Y. Vardi. Automatic verification of probabilistic concurrent finite-state systems. In *Proc. 26th IEEE Symp. Found. of Comp. Sci.*, pages 327–338, 1985.