Quantitative Solution of Omega-Regular Games[1]

Luca de Alfaro
Department of Computer Engineering
University of California at Santa Cruz, Santa Cruz, CA 95064
E-mail: luca@soe.ucsc.edu

and

Rupak Majumdar
Department of Electrical Engineering and Computer Sciences
University of California at Berkeley, Berkeley, CA 94720, USA
rupak@eecs.berkeley.edu

Version: November 12, 2001

We consider two-player games played for an infinite number of rounds, with ω-regular winning conditions. The games may be concurrent, in that the players choose their moves simultaneously and independently, and probabilistic, in that the moves determine a probability distribution for the successor state. We introduce *quantitative game μ-calculus,* and we show that the maximal probability of winning such games can be expressed as the fixpoint formulas in this calculus. We develop the arguments both for deterministic and for probabilistic concurrent games; as a special case, we solve probabilistic turn-based games with ω-regular winning conditions, which was also open. We also characterize the optimality, and the memory requirements, of the winning strategies. In particular, we show that while memoryless strategies suffice for winning games with safety and reachability conditions, Büchi conditions require the use of strategies with infinite memory. The existence of optimal strategies, as opposed to ε-optimal, is only guaranteed in games with safety winning conditions.

Key Words: Automata, games, μ-calculus, probabilistic algorithm, temporal logic.

1. INTRODUCTION

We consider two-player games played on finite state spaces for an infinite number of rounds. In each round, depending on the current state of the game, the moves of one or both players determine the next state [30]; we consider games in which the set of available moves is finite. Such games offer a model for systems composed of interacting components, and they have been studied under a wide range of winning conditions. The winning conditions are often codified by associating a *reward* with each state and choice of moves, and by studying the maximal discounted, total, or average reward that player 1 can obtain in such a game; a survey of algorithms for solving games with respect to such winning conditions is e.g. [27, 13]. Here, we consider winning conditions consisting in ω-regular automata acceptance conditions

[1]A preliminary version of this paper appeared in the Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC 2001), pp. 675–683. This research was supported in part by the AFOSR MURI grant F49620-00-1-0327 and the NSF Theory grant CCR-9988172.

defined over the state space of the game [4, 14, 31]. Given a game with an ω-regular winning condition and a starting state s, we study the maximal probability with which player 1 can ensure that the condition holds from s; we call this maximal probability the *value* of the game at s for player 1. The determinacy result of [20] ensures that, at all states and for all ω-regular winning conditions, the value of the game for player 1 is equal to one minus the value of the game with complementary condition for player 2.

We distinguish between *turn-based* and *concurrent* games, and between *deterministic* and *probabilistic* games. Systems in which the interaction between the components is asynchronous give rise to *turn-based* games, where in each round only one of the two players can choose among several moves. On the other hand, synchronous interaction leads to *concurrent games,* where in each round both players can choose simultaneously and independently among several moves. The games are *deterministic* if the current state and the moves uniquely determine the successor state, and are *probabilistic* if the current state and the moves determine a probability distribution for the successor state. For any ω-regular winning condition, the value of a deterministic turn-based game at a state is either 0 or 1; moreover, player 1 can achieve this value by playing according to a *deterministic* strategy, that select a move based on the current state and on the history of the game [4, 14]. In contrast, the value of a concurrent game at a state may be strictly between 0 and 1; furthermore, achieving this value may require the use of *randomized* strategies, that select not a move, but a probability distribution over moves. To see this, consider the concurrent game MatchOneBit. The game starts at state s_0, where both players simultaneously and independently choose a bit (0 or 1); if the bits match, the game proceeds to state s_{win}, otherwise, it proceeds to state s_{lose}. Once at s_{win} (resp. s_{lose}) the state is confined there forever. Consider the *safety* condition $\Box\{s_0, s_{win}\}$, requiring that s_{lose} is not entered. For every deterministic strategy of player 1, player 2 has another (complementary) deterministic strategy that ensures a transition to s_{lose}; hence, if player 1 could only use deterministic strategies, he would win with probability 0. However, if player 1 uses a randomized strategy that chooses both bits at random with uniform probability, then the game enters state s_{win} with probability 1/2, regardless of the strategy of player 2; indeed, the value of the game at s_0 is 1/2.

The value of deterministic turn-based games with ω-regular winning conditions can be computed with the algorithms of [4, 14, 11, 31]. The algorithms of [11] are based on the use of game μ-calculus, obtained by replacing the predecessor operator Pre of classical μ-calculus [17] by the *controllable predecessor* operator Cpre: for a set of states U, the set $\mathrm{Cpre}(U)$ consists of the states from which player 1 can force the game into U in one step. A richer version of game μ-calculus was used in [8] to provide qualitative solutions for concurrent probabilistic games with ω-regular conditions. There, multi-argument predecessor operators are used to compute the set of states from which player 1 can win with probability 1, or arbitrarily close to 1.

We introduce *quantitative game μ-calculus,* and use it to provide a uniform framework for understanding and solving concurrent games with ω-regular winning conditions. In quantitative game μ-calculus, sets of states are replaced by functions from states to the interval $[0, 1]$, and the controllable predecessor operator Cpre is replaced by a quantitative version Ppre. Given a function f from states to the interval $[0, 1]$, the function $g = \mathrm{Ppre}(f)$ associates with each state the maximal expected value of f that player 1 can ensure in one step. The operator Ppre can be evaluated using results about matrix games [35, 26]. Related quantitative predecessor operators for one-player or turn-based structures were considered in [16, 24, 15, 21, 22]. We show that the values of concurrent games with ω-regular conditions can be obtained simply by replacing Cpre by Ppre in the solutions of [11]. The result is surprising because concurrent games differ from deterministic turn-based games in several fundamental respects. First, concurrent games require in general the use of randomized strategies, as remarked above. Second, even for the simple winning condition of reachability, optimal strategies may not exist: one can only guarantee the existence of ε-optimal strategies for all $\varepsilon > 0$ [12]. Third, whereas finite-memory strategies suffice for winning deterministic turn-based games, in concurrent games both ε-optimal strategies, and optimal strategies if they exist, may need an infinite amount of memory [8]. Fourth, the standard recursive structure of proofs for deterministic turn-based games [23, 31] breaks down, as both players can choose a distribution over moves at each state.

We develop the arguments both for deterministic and for probabilistic concurrent games. Hence, as a special case we solve probabilistic turn-based games with ω-regular winning conditions, which was also an open problem. The quantitative game μ-calculus solution formulas provide the value also of games with countable, rather than finite, state space. We also characterize the optimality, and the memory requirements, of the winning strategies. In particular, we show that while memoryless strategies suffice for winning games with safety and reachability conditions, Büchi and Rabin-chain conditions require the use of strategies with infinite memory. The existence of optimal strategies, as opposed to ε-optimal, is only guaranteed in games with safety winning conditions.

The solutions formulas we present in this paper also solve the model-checking problem for the probabilistic temporal logics pCTL and pCTL* over concurrent games. The logics pCTL and pCTL*, originally proposed over Markov chains [1] and Markov decision processes [3], can express the maximal and minimal probability with which linear time temporal logic (LTL) formulas are satisfied. These logics can be immediately generalized to concurrent games, by considering the maximal probability with which a player can ensure that the formula holds. Since LTL formulas can be translated into deterministic Rabin-chain automata [28, 29, 34], our results characterize the validity of pCTL and pCTL* formulas over concurrent games.

As remarked by [11] in the context of deterministic turn-based games, the use of μ-calculus for solving games helps in the formulation of the correctness arguments. In order to argue the correctness of a solution formula, we need to show that player 1 has an optimal (or ε-optimal) strategy that realizes the value given by the formula, and that player 2 has a "spoiling" strategy that is optimal (or ε-optimal) for the game with the complementary condition. Since the operator Ppre in the solution formula refers to player 1, an optimal strategy for player 1 can be constructed from the fixpoint of the formula. On the other hand, the derivation of spoiling strategies for player 2 is not immediate: indeed, even for games with safety or reachability conditions, the standard argument involves the consideration of discounted versions of the games (see, e.g., [13]). In contrast, by writing the solution formula in game μ-calculus, we place the burden of the argument on the syntactic complementation of the solution formula. Specifically, for a winning condition Ψ, we characterize the maximal probabilities of winning the game by a μ-calculus formula ϕ, and from ϕ we construct an optimal (or ε-optimal) strategy for player 1. The syntactic complement $\neg\phi$ of ϕ gives the maximal probabilities for player 2 to win the dual game with condition $\neg\Psi$. From $\neg\phi$, we can again construct an optimal (or ε-optimal) strategy for player 2 for the game with condition $\neg\Psi$. The two constructions are enough to conclude the correctness of our solution formulas.

The iterative interpretation of quantitative game μ-calculus leads to algorithms for the computation of approximate solutions. By representing value functions symbolically, these algorithms may be used for the approximate analysis of games with very large state spaces [5, 10]. Unfortunately, except for safety and reachability conditions, the alternance of least and greatest fixpoint operators in the solution formulas leads to approximation schemes that do not converge monotonically to the value of a game. This situation contrasts with the one for Markov decision processes, where monotonically-converging approximation schemes are available, and where the maximal winning probability can be computed in polynomial time by reduction to linear programming [7]. We show that this discrepancy is no accident, since the basic device for solving Markov decision processes with ω-regular conditions, viz., a reduction to reachability, fails for games.

2. CONCURRENT GAMES

For a countable set A, a *probability distribution* on A is a function $p\colon A \mapsto [0,1]$ such that $\sum_{a\in A} p(a) = 1$. We denote the set of probability distributions on A by $\mathcal{D}(A)$. A (two-player) *concurrent game structure* $\mathcal{G} = \langle S, \mathit{Moves}, \Gamma_1, \Gamma_2, p\rangle$ consists of the following components:

- A finite state space S.
- A finite set *Moves* of moves.
- Two move assignments $\Gamma_1, \Gamma_2 \colon S \mapsto 2^{\mathit{Moves}} \setminus \emptyset$. For $i \in \{1,2\}$, assignment Γ_i associates with each state $s \in S$ the non-empty set $\Gamma_i(s) \subseteq \mathit{Moves}$ of moves available to player i at state s.
- A probabilistic transition function p, that gives the probability $p(t \mid s, a_1, a_2)$ of a transition from s to t for all $s, t \in S$ and all moves $a_1 \in \Gamma_1(s)$ and $a_2 \in \Gamma_2(s)$.

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 $p(t \mid s, a_1, a_2)$, for all $t \in S$. We assume that the players act *non-cooperatively*, i.e., each player chooses her strategy independently and secretly from the other player, and is only interested in maximizing her own reward. A *path* of $\mathcal{G}$ is an infinite sequence $\overline{s} = s_0, s_1, s_2, \ldots$ of states in S such that for all $k \geq 0$, there are moves $a_1^k \in \Gamma_1(s_k)$ and $a_2^k \in \Gamma_2(s_k)$ with $p(s_{k+1} \mid s_k, a_1^k, a_2^k) > 0$. We denote by Ω the set of all paths.

We distinguish the following special classes of concurrent game structures.

- A concurrent game structure $\mathcal{G}$ is *deterministic* if for all $s \in S$ and all $a_1 \in \Gamma_1(s)$, $a_2 \in \Gamma_2(s)$, there is a $t \in S$ such that $p(t \mid s, a_1, a_2) = 1$.
- A concurrent game structure $\mathcal{G}$ is *turn-based* if at every state at most one player can choose among multiple moves; that is, if for every state $s \in S$ there exists at most one $i \in \{1,2\}$ with $|\Gamma_i(s)| > 1$.

For brevity, we refer to concurrent turn-based game structures simply as turn-based game structures.

2.1. Randomized strategies

A *strategy* for player $i \in \{1, 2\}$ is a mapping $\pi_i : S^+ \mapsto \mathcal{D}(\mathit{Moves})$ that associates with every nonempty finite sequence $\sigma \in S^+$ of states, representing the past history of the game, a probability distribution $\pi_1(\sigma)$ 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 $\sigma \in S^*$ and states $s \in S$, if $\pi_i(\sigma s)(a) > 0$, then $a \in \Gamma_i(s)$. We denote by Π_i the set of all strategies for player $i \in \{1, 2\}$. A strategy π is *deterministic* if for all $\sigma \in S^+$ there exists $a \in \mathit{Moves}$ such that $\pi(\sigma)(a) = 1$. Thus, deterministic strategies are equivalent to functions $S^+ \mapsto \mathit{Moves}$. A strategy π is *finite-memory* if the distribution chosen at every state $s \in S$ depends only on s itself, and on a finite number of bits of information about the past history of the game. A strategy π is *memoryless* if $\pi(\sigma s) = \pi(s)$ for all $s \in S$ and all $\sigma \in S^*$.

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[2]. For an event $\mathcal{A} \subseteq \Omega$, we denote by $\Pr_s^{\pi_1,\pi_2}(\mathcal{A})$ 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. Similarly, for a measurable function f that associates a number in $\mathbb{R} \cup \{\infty\}$ with each path, we denote by $\mathrm{E}_s^{\pi_1,\pi_2}\{f\}$ the expected value of f when the game starts from s and the strategies π_1 and π_2 are used. We denote by Θ_i the random variable representing the i-th state of a path; formally, Θ_i is a variable that assumes value s_i on the path $s_0, s_1, s_2, \ldots$.

2.2. Winning conditions

Given a concurrent game structure $\mathcal{G} = \langle S, \mathit{Moves}, \Gamma_1, \Gamma_2, p \rangle$, we consider *winning conditions* expressed by linear-time temporal logic (LTL) formulas, whose atomic propositions correspond to subsets of the set S of states [19]. We focus on winning conditions that correspond to safety or reachability properties, as well as winning conditions that correspond to the accepting criteria of Büchi, co-Büchi, and Rabin-chain automata [25, 11]. We call games with such winning conditions *safety, reachability, Büchi, co-Büchi,* and *Rabin-chain games,* respectively. The ability to solve games with Rabin-chain conditions suffices for solving games with arbitrary LTL (or ω-regular) winning conditions: in fact, it suffices to encode the ω-regular condition as a deterministic Rabin-chain automaton, solving then the game consisting of the synchronous product of the original game with the Rabin-chain automaton [25, 31].

Given an LTL winning condition Ψ, by abuse of notation we denote equally by Ψ the set of paths $\overline{s} \in \Omega$ that satisfy Ψ; this set is measurable for any choice of strategies for the two players [33]. Hence, the probability that a path satisfies Ψ starting from state $s \in S$ under strategies π_1, π_2 for the two players is $\Pr_s^{\pi_1,\pi_2}(\Psi)$. Given a state $s \in S$ and a winning condition Ψ, we are interested in finding the maximal probability with which player $i \in \{1, 2\}$ can ensure that Ψ holds from s. We call such probability the *value of the game Ψ at s for player $i \in \{1, 2\}$*. This

[2] To be precise, we should define events as measurable sets of paths *sharing the same initial state.* However, our (slightly) improper definition leads to more concise notation.

value for player 1 is given by the function $\langle 1\rangle\Psi : S \mapsto [0,1]$, defined for all $s \in S$ by

$$\langle 1\rangle\Psi(s) = \sup_{\pi_1\in\Pi_1} \inf_{\pi_2\in\Pi_2} \Pr_s^{\pi_1,\pi_2}(\Psi).$$

The value for player 2 is given by the function $\langle 2\rangle\Psi$, defined symmetrically. Concurrent games satisfy a *quantitative* version of determinacy [20], stating that for all LTL conditions Ψ and all $s \in S$, we have

$$\langle 1\rangle\Psi(s) = 1 - \langle 2\rangle\neg\Psi(s).$$

A strategy π_1 for player 1 is *optimal* if for all $s \in S$ we have

$$\inf_{\pi_2\in\Pi_2} \Pr_s^{\pi_1,\pi_2} = \langle 1\rangle\Psi(s).$$

For $\varepsilon > 0$, a strategy π_1 for player 1 is ε-*optimal* if for all $s \in S$ we have

$$\inf_{\pi_2\in\Pi_2} \Pr_s^{\pi_1,\pi_2} \geq \langle 1\rangle\Psi(s) - \varepsilon.$$

Note that the quantitative determinacy of concurrent games is equivalent to the existence of ε-optimal strategies for all $\varepsilon > 0$ at all states $s \in S$. For the special case of deterministic turn-based games, it is known that the value of any ω-regular game at a state is either 0 or 1, and finite-memory deterministic optimal strategies always exist; the value of the game can be computed with the algorithms of [4, 14, 11].

2.3. Predecessor operators

Let $\mathcal{F}$ be the space of all functions $S \mapsto [0,1]$ that map states into the interval $[0,1]$. Given two functions $f, g \in \mathcal{F}$, we write $f > g$ (resp. $f \geq g$) if $f(s) > g(s)$ (resp. $f(s) \geq g(s)$) at all $s \in S$, and we define $f \wedge g$ and $f \vee g$ by

$$(f \wedge g)(s) = \min\{f(s), g(s)\}$$
$$(f \vee g)(s) = \max\{f(s), g(s)\}$$

for all $s \in S$. We denote by $\mathbf{0}$ and $\mathbf{1}$ the constant functions that map all states into 0 and 1, respectively. For all $f \in \mathcal{F}$, we denote by $\mathbf{1} - f$ the function defined by $(\mathbf{1} - f)(s) = 1 - f(s)$ for all $s \in S$. Given a subset $Q \subseteq S$ of states, by abuse of notation we denote also by Q the *indicator function* of Q, defined by $Q(s) = 1$ if $s \in Q$ and $Q(s) = 0$ otherwise. We denote by $\neg Q = S \setminus Q$ the complement of the subset Q in S, and again we denote equally by $\neg Q$ the indicator function of $\neg Q$. We denote by $\mathcal{F}_I \subseteq \mathcal{F}$ the set of indicator functions. The *quantitative predecessor operators* $\mathrm{Ppre}_1, \mathrm{Ppre}_2 : \mathcal{F} \mapsto \mathcal{F}$ are defined for every $f \in \mathcal{F}$ by

$$\mathrm{Ppre}_1(f)(s) = \sup_{\pi_1\in\Pi_1} \inf_{\pi_2\in\Pi_2} \mathrm{E}_s^{\pi_1,\pi_2}\{f(\Theta_1)\}$$

and symmetrically for Ppre_2. Intuitively, the value $\mathrm{Ppre}_i(f)(s)$ is the maximum expectation for the next value of f that player $i \in \{1,2\}$ can achieve. Given $f \in \mathcal{F}$ and $i \in \{1,2\}$, the function $\mathrm{Ppre}_i(f)$ can be computed by solving the following *matrix game* at each $s \in S$:

$$\mathrm{Ppre}_1(f)(s) = \mathrm{val}_1\Big[\sum_{t\in S} f(t)p(t \mid s, a_1, a_2)\Big]_{a_1\in\Gamma_1(s), a_2\in\Gamma_2(s)},$$

where $\mathrm{val}_1 A$ denotes the value obtained by player 1 in the matrix game A. The existence of solutions to the above matrix games, and the existence of optimal randomized strategies for players 1 and 2, is guaranteed by the minmax theorem [35]. The matrix games may be solved using traditional linear programming algorithms (see, e.g., [26]). From properties of matrix games we have the following facts.

PROPOSITION 1.

1. *For $i \in \{1,2\}$, the operator $Ppre_i$ is monotonic and continuous, that is, for all $f, g \in \mathcal{F}$, if $f \geq g$ then $Ppre_i(f) \geq Ppre_i(g)$; and for all $f_1 \leq f_2 \leq \ldots$ in $\mathcal{F}$, we have $\lim_n Ppre_i(f_n) = Ppre_i(\lim_n f_n)$.*
2. *For $i \in \{1,2\}$, $f \in \mathcal{F}$ and $c \in \mathcal{F}$ a constant, $Ppre_i(f+c) = Ppre_i(f) + c$.*
3. *The operators $Ppre_1$ and $Ppre_2$ are dual: for all $f \in \mathcal{F}$, we have $Ppre_1(f) = \mathbf{1} - Ppre_2(\mathbf{1} - f)$.*

2.4. Quantitative game μ-calculus

We write the solutions of games with respect to ω-regular winning conditions in *quantitative game μ-calculus.* The formulas of the quantitative game μ-calculus are generated by the grammar

$$\phi \ ::= Q \mid x \mid \phi \vee \phi \mid \phi \wedge \phi \mid \mathrm{Ppre}_1(\phi) \mid \mathrm{Ppre}_2(\phi) \mid \mu x.\phi \mid \nu x.\phi,$$

for proposition $Q \subseteq S$ and variables x from some fixed set X. Hence, as for LTL, the propositions of quantitative μ-calculus formulas correspond to subsets of states of the game. As usual, a formula ϕ is *closed* if every variable x in ϕ occurs in the scope of a fixpoint quantifier μx or νx.

Let $\mathcal{E} : X \mapsto \mathcal{F}$ be a variable valuation that associates a function $\mathcal{E}(x) \in \mathcal{F}$ with each variable $x \in X$. We write $\mathcal{E}[x \mapsto f]$ for the valuation that agrees with $\mathcal{E}$ on all variables, except that $x \in X$ is mapped to $f \in \mathcal{F}$. Given a valuation $\mathcal{E}$, every formula ϕ of quantitative game μ-calculus defines a function $[\![\phi]\!]_{\mathcal{E}} \in \mathcal{F}$:

$$\begin{aligned}
[\![f]\!]_{\mathcal{E}} &= f \\
[\![x]\!]_{\mathcal{E}} &= \mathcal{E}(x) \\
[\![\mathrm{Ppre}_1(\phi)]\!]_{\mathcal{E}} &= \mathrm{Ppre}_1([\![\phi]\!]_{\mathcal{E}}) \\
[\![\mathrm{Ppre}_2(\phi)]\!]_{\mathcal{E}} &= \mathrm{Ppre}_2([\![\phi]\!]_{\mathcal{E}}) \\
[\![\phi_1 \{ {}^{\vee}_{\wedge} \} \phi_2]\!]_{\mathcal{E}} &= ([\![\phi_1]\!]_{\mathcal{E}} \{ {}^{\vee}_{\wedge} \} [\![\phi_2]\!]_{\mathcal{E}}) \\
[\![\{ {}^{\mu}_{\nu} \} x.\phi]\!]_{\mathcal{E}} &= \{ {}^{\sup}_{\inf} \} \{ f \in \mathcal{F} \mid f = [\![\phi]\!]_{\mathcal{E}[x \mapsto f]} \}.
\end{aligned}$$

The existence and uniqueness of the above fixpoints for the μ and ν operators is a consequence of the monotonicity and continuity of all the operators, and in particular of Ppre_1 and Ppre_2. As usual, the fixpoints can be evaluated in an iterative fashion: we have $[\![\mu x.\phi]\!]_{\mathcal{E}} = \lim_{n\to\infty} x_n$, where $x_0 = \mathbf{0}$, and $x_{n+1} = [\![\phi]\!]_{\mathcal{E}[x \mapsto x_n]}$ for $n \geq 0$. Similarly, for the greatest fixpoint operator ν we have $[\![\nu x.\phi]\!]_{\mathcal{E}} = \lim_{n\to\infty} x_n$, where $x_0 = \mathbf{1}$, and $x_{n+1} = [\![\phi]\!]_{\mathcal{E}[x \mapsto x_n]}$ for $n \geq 0$. A quantitative game μ-calculus formula suggests a way to implement approximation algorithms for large state spaces, using a subset $\mathcal{F}' \subseteq \mathcal{F}$ of base functions that have compact representations [2, 6, 10]. We note that the solution algorithms presented in this paper apply also to games with countable (rather than finite) state space and finite set of moves (see Theorem 4); in this case, however, the iterative evaluation of the fixpoints needs to be based on transfinite induction.

The quantitative game μ-calculus defined by (1) suffices for writing the solution formulas of games with ω-regular winning conditions. In some intermediate lemmas, however, we use with slight abuse of notation an extended version of the calculus, in which we have one symbol f for every function $f \in \mathcal{F}$. Obviously, such functions are interpreted as themselves: for all valuations $\mathcal{E}$, we have $[\![f]\!]_{\mathcal{E}} = f$.

2.5. Complementation and correctness

We solve concurrent games with LTL winning condition Ψ by providing a quantitative game μ-calculus formula ϕ such that $\langle 1 \rangle \Psi = [\![\phi]\!]$. To prove this equality, we exploit the *complementation* of μ-calculus expressions. The complement of a closed μ-calculus formula ϕ is a formula $\neg\phi$ such that $\mathbf{1} - [\![\phi]\!] = [\![\neg\phi]\!]$; the complement can be obtained by recursively applying the following transformations, which rely on the duality of Ppre_1 and Ppre_2:

$$\begin{aligned}
&\neg Q && \Rightarrow S \setminus Q \\
&\neg\neg\phi && \Rightarrow \phi \\
&\neg(\mathrm{Ppre}_1(\phi)) && \Rightarrow \mathrm{Ppre}_2(\neg\phi) \\
&\neg(\phi_1 \left\{{\vee \atop \wedge}\right\} \phi_2) && \Rightarrow (\neg\phi_1) \left\{{\wedge \atop \vee}\right\} (\neg\phi_2) \\
&\neg \left\{{\mu \atop \nu}\right\} x.\phi && \Rightarrow \left\{{\nu \atop \mu}\right\} x.\neg\phi[\neg x/x]
\end{aligned}$$

where $\phi[\neg x/x]$ denotes the result of replacing x with $\neg x$ in ϕ. Note that since the formula ϕ is closed, by applying the above transformations to $\neg\phi$ we obtain again a formula of the syntactic form (1). In fact, the above transformations push the $\neg$ operator to the leaves of the syntax tree (1), which consist either in subsets $Q \subseteq S$ or in variables $x \in X$. The subsets are simply complemented. Since ϕ is closed, each variable $x \in X$ in ϕ appears in the scope of a μx or νx quantifier; the transformation rules for μ and ν, together with the rule for double negation elimination, ensure that once all transformations have been applied, no $\neg$ operator remains as prefix to a variable.

Our proofs of $\langle 1 \rangle \Psi = [\![\phi]\!]$ consist in two steps.

- First, from ϕ we construct for all $\varepsilon > 0$ a strategy π_1^ε for player 1 that ensures winning with probability at least $[\![\phi]\!] - \varepsilon$, proving $[\![\phi]\!] \geq \langle 1 \rangle \Psi$.

- Second, we complement ϕ, and we consider the winning condition $\neg\Psi$. From $\neg\phi$ we construct for all $\varepsilon > 0$ a strategy π_2^ε that enables player 2 to win the game with goal $\neg\Psi$ with probability at least $[\![\neg\phi]\!] - \varepsilon$; this shows $[\![\neg\phi]\!] \geq \langle 2 \rangle \neg\Psi$, or equivalently $[\![\phi]\!] \leq \langle 1 \rangle \Psi$.

Even in the cases where solution formulas for concurrent games are known, such as for the reachability winning condition (see [13], Chapter 4.4), this approach yields simpler arguments than the classical one, where the ε-optimal strategies for both players have to be constructed from the solution formula ϕ for player 1 alone, and where it is usually necessary to consider discounted versions of the games.

3. REACHABILITY AND SAFETY GAMES

Concurrent reachability and safety games can be solved by reducing them to positive stochastic games [32, 13]. We present the solution algorithms, reformulating them in quantitative game μ-calculus. As mentioned above, by relying on the complementation of quantitative game μ-calculus, we are able to prove the correctness of the solutions without resorting to the consideration of discounted versions of the same games.

A *concurrent reachability game* consists in a concurrent game structure $\mathcal{G} = \langle S, \mathit{Moves}, \Gamma_1, \Gamma_2, p\rangle$ together with a winning condition $\Diamond U$, where $U \subseteq S$. Intuitively, the winning condition consists in reaching the subset U of states. The solution of such a reachability game is given by

$$\langle 1\rangle \Diamond U = [\![\mu x.(U \vee \mathrm{Ppre}_1(x))]\!]. \tag{1}$$

This solution can be computed iteratively as the limit $\langle 1\rangle \Diamond U = \lim_{k\to\infty} x_k$, where $x_0 = \mathbf{0}$ and $x_{k+1} = U \vee \mathrm{Ppre}_1(x_k)$ for $k \geq 0$. This iteration scheme gives an approximation scheme to solve the reachability game. In Markov decision processes, one can reduce the reachability question to a linear programming problem which can then be solved exactly. This gives an alternative to value iteration. Unfortunately, for concurrent games we cannot reduce the problem to linear programming, because the maximal probability of winning in a game where all probabilities are rationals may still be irrational (see e.g. [27]).

EXAMPLE 1. Consider a concurrent game with three states s, t, and u, and winning condition $\Diamond\{u\}$. The transition relation is as follows: from state t, player 1 has two choices a_1 and b_1, and player 2 the choices a_2 and b_2. The transition probabilities are: $\Pr(u|t,a_1,a_2) = \frac{1}{2}$, $\Pr(t|t,a_1,a_2) = \frac{1}{2}$, $\Pr(u|t,b_1,a_2) = \Pr(u|t,a_1,b_2) = 0$, $\Pr(s|t,b_1,a_2) = \Pr(s|t,a_1,b_2) = 1$, $\Pr(u|t,b_1,b_2) = \frac{3}{4}$, and $\Pr(t|t,b_1,b_2) = \frac{1}{4}$. The states s and u are absorbing: the game never leaves s or u once it reaches these states. The maximal probability of winning the game $\Diamond\{u\}$ is given by the least fixpoint of $x = \mathrm{Ppre}_1(x) \vee \{u\}$; for state t, we have

$$x(t) = \mathrm{val}_1 \begin{bmatrix} \frac{1}{2} + \frac{1}{2}x(t) & 0 \\ 0 & \frac{3}{4} + \frac{1}{4}x(t) \end{bmatrix}$$

which has the solution $x(t) = (-3 + 2\sqrt{6})/5$. ∎

To prove (1), we show separately the two inequalities

$$\langle 1\rangle \Diamond U \geq [\![\mu x.(U \vee \mathrm{Ppre}_1(x))]\!]$$
$$\langle 1\rangle \Diamond U \leq [\![\mu x.(U \vee \mathrm{Ppre}_1(x))]\!].$$

The first inequality is a consequence of the following lemma; the second inequality, as mentioned in Section 2.5, will follow from results on safety games.

LEMMA 1. *Let $w = [\![\mu x.(U \vee \mathit{Ppre}_1(x))]\!]$. For all $\varepsilon > 0$ player 1 has a strategy π_1^ε such that $\Pr_s^{\pi_1^\varepsilon,\pi_2}(\Diamond U) > w(s) - \varepsilon$ for all $\pi_2 \in \Pi_2$ and all $s \in S$.*

Proof. The proof follows a classical argument (see, e.g., [12, 13]). For $n \geq 0$, consider the n-step version of the game, whose winning condition $\Diamond_n U$ requires reaching U in at most n steps. For $n = 0$, we have trivially $\langle 1\rangle \Diamond_0 U = U$, and player 1 has a trivial strategy π_1^0. Let $x_0 = \mathbf{0}$, and $x_{n+1} = U \vee \mathrm{Ppre}_1(x_n)$ for $n \geq 0$. By induction on n, we show that $\langle 1 \rangle \Diamond_n U \geq x_n$ for all $n \geq 0$. Consider the $n+1$ step game. For $s \in U$, player 1 has a trivial strategy. For $s \in S \setminus U$, the strategy π_1^{n+1} plays the first round according to an optimal distribution associated with the matrix game for $\mathrm{Ppre}_1(x_n)$, and then follows the strategy π_1^n from the resulting state. Let π_2 be any strategy of player 2. Then

$$\begin{aligned} \mathrm{E}_s^{\pi_1,\pi_2}\{\Diamond_{n+1} U\} &\geq \sum_{t \in S} \mathrm{E}_t^{\pi_1^n,\pi_2[t]}\{\Diamond_n U\} \mathrm{Pr}_s^{\pi_1,\pi_2}(\Theta_1 = t) \\ &\geq \mathrm{Ppre}_1(x_n)(s) = x_{n+1}(s), \end{aligned}$$

where $\pi_2[t]$ is the strategy that behaves like π_2 after a transition to t has occurred. The result then follows from $w = \lim_{n\to\infty} x_n$, and from the fact that $\Diamond_n U$ implies $\Diamond U$ for all $n \geq 0$: given any $\varepsilon > 0$, there is $N > 0$ such that $\max\{x(s) - x_N(s) \mid s \in S\} < \varepsilon$. Then, by following strategy π_1^N we get $\mathrm{Pr}^{\pi_1^N,\pi_2}(\Diamond U) \geq x - \varepsilon$. ∎

A *concurrent safety game* consists in a concurrent game structure $\mathcal{G} = \langle S, \mathit{Moves}, \Gamma_1, \Gamma_2, p \rangle$ together with a winning condition $\Box U$, where $U \subseteq S$. Intuitively, the winning condition consists in staying forever in the subset U of states. The complement of the reachability condition $\Diamond U$ is the safety condition $\Box \neg U$, and the complement of the quantitative game μ-calculus formula $\mu x.(U \vee \mathrm{Ppre}_1(x))$ is

$$\nu x.(\neg U \wedge \mathrm{Ppre}_2(x)),$$

where $\neg U$ is an abbreviation for $S \setminus U$. We will show that the solution of concurrent safety games is given by

$$\langle 1 \rangle \Box U = [\![\nu x.(U \wedge \mathrm{Ppre}_1(x))]\!], \tag{2}$$

which is dual to (1). To this end, we prove the following lemma.

LEMMA 2. *Let $w = [\![\nu x.(U \wedge \mathit{Ppre}_1(x))]\!]$. Player 1 has a strategy π_1 such that $\mathrm{Pr}_s^{\pi_1,\pi_2}(\Box U) \geq w(s)$ for all $\pi_2 \in \Pi_2$ and all $s \in S$.*

The lemma can be proved using standard arguments about positive reward games [13]. We present here a more direct proof, that will lead to the arguments for Büchi and co-Büchi games.

Proof. Let π_1 be a strategy for player 1 that at all $s \in U$ plays according to an optimal distribution of the matrix game corresponding to $\mathrm{Ppre}_1(w)(s)$, and at all $s \in S \setminus U$ plays arbitrarily. Fix a state $s_0 \in S$ and an arbitrary strategy $\pi_2 \in \Pi_2$. The process $\{H_n\}_{n \geq 0}$ defined by $H_n = w(\Theta_n)$ is a submartingale [36]: in fact, from $w(s) = \mathrm{Ppre}_1(w)(s)$ for $s \in U$ and from the optimality of π_1 follows that

$$\mathrm{E}_{s_0}^{\pi_1,\pi_2}\{H_{n+1} \mid H_0, H_1, \ldots, H_n\} \geq H_n$$

for all $n \geq 0$. Hence, we have $\mathrm{E}_{s_0}^{\pi_1,\pi_2}\{H_n\} \geq H_0 = w(s_0)$. Moreover, since $w(s) \leq 1$ at all $s \in S$, by inspection we have $\mathrm{E}_{s_0}^{\pi_1,\pi_2}\{H_n\} \leq \mathrm{Pr}_{s_0}^{\pi_1,\pi_2}(\Box_n U)$, where $\Box_n U$ is the event of staying in U for at least n steps. Combining these two inequalities we obtain $w(s_0) \leq \mathrm{Pr}_{s_0}^{\pi_1,\pi_2}(\Box_n U)$, and the result follows from $\mathrm{Pr}_{s_0}^{\pi_1,\pi_2}(\Box U) = \lim_{n\to\infty} \mathrm{Pr}_{s_0}^{\pi_1,\pi_2}(\Box_n U)$. ∎

The following theorem summarizes the properties of concurrent reachability and safety games.

THEOREM 1. *The following assertions hold.*

1. *Concurrent reachability and safety games can be solved according to (1) and (2).*

2. *Concurrent reachability games have memoryless ε-optimal strategies; there are deterministic concurrent reachability games without optimal strategies.*

3. *Concurrent safety games have memoryless optimal strategies; there are deterministic concurrent safety games without memoryless deterministic optimal strategies.*

Part 1 is classical [12, 13], except for the notation; the result also follows from the combination of Lemmas 1 and 2. The existence of memoryless ε-optimal strategies for concurrent reachability games follows from results on positive stochastic games (see, e.g., [13], pp. 196). The proof of Lemma 1 constructs an ε-optimal strategy for player 1, but the strategy is in general not memoryless. The existence of deterministic concurrent reachability games without optimal strategies is demonstrated by Example 2 below, adapted from [12, 18]. The existence of memoryless optimal strategies for concurrent safety games is classical; it also follows from the proof of Lemma 2. The existence of deterministic concurrent safety games without optimal deterministic strategies is demonstrated by the game MATCHONEBIT described in the introduction: in fact, randomized strategies are necessary for one-step matrix games [26].

EXAMPLE 2. Consider the following game, adapted from [12, 18] (see also [9] for an intuitive interpretation of the game). The state space of the game is $S = \{s, t, u\}$; the only state where players can choose among more than one move is s. We have $\Gamma_1(s) = \{a, b\}$, and $\Gamma_2(s) = \{c, d\}$. The game has a deterministic transition function: $p(s \mid s, a, c) = p(t \mid s, a, d) = p(t \mid s, b, c) = p(u \mid s, b, d) = 1$, all other transition probabilities are 0. We have $\langle 1 \rangle \Diamond \{t\}(s) = 1$. In fact, player 1 can play moves a and b with probability $1 - \varepsilon$ and ε respectively to ensure a winning probability of $(1 - \varepsilon)$ from s, for $\varepsilon > 0$. However, player 1 has no optimal strategy: if he decides to play move b at the nth round, player 2 can play move d at the n-th round, so that the probability of reaching t is always less than 1. ∎

4. BÜCHI AND CO-BÜCHI GAMES

A *concurrent Büchi game* consists in a concurrent game structure $\mathcal{G} = \langle S, Moves, \Gamma_1, \Gamma_2, p \rangle$ together with a winning condition $\Box\Diamond U$, where $U \subseteq S$. Intuitively, the winning condition consists in visiting the subset U of states infinitely often. The solution of a concurrent Büchi game is given by

$$\langle 1 \rangle \Box\Diamond U = [\![\nu y.\mu x.((\neg U \wedge \mathrm{Ppre}_1(x)) \vee (U \wedge \mathrm{Ppre}_1(y)))]\!] \,. \tag{3}$$

The proof of (3) is based on two lemmas. The first lemma generalizes the result about concurrent reachability games. Given a function $g \in \mathcal{F}$ and a subset U of states, we let $g(\Diamond U)$ be the function that associates with each path $s_0, s_1, s_2, \ldots$ the value $g(s_i)$, for $i = \min\{k \mid s_k \in U\} < \infty$, and the value 0 if $s_k \notin U$ for all $k \geq 0$. Hence, $g(\Diamond U)$ is the value of g at the state where the path first enters U, if such a state exists, and is 0 otherwise. The following lemma can be proved similarly to Lemma 1.

Lemma 3. *For $g \in \mathcal{F}$ and $U \subseteq S$, let*

$$w = [\![\mu x.((\neg U \wedge \mathit{Ppre}_1(x)) \vee (U \wedge g))]\!].$$

Then, for all $\varepsilon > 0$ player 1 has a strategy π_1^ε that ensures $\mathrm{E}_s^{\pi_1^\varepsilon,\pi_2}\{g(\Diamond U)\} \geq w(s) - \varepsilon$ at all $s \in S$.

We call the above game a $g(\Diamond U)$-game; the strategy π_1^ε is an ε-optimal strategy for it. The following lemma shows that the fixpoint (3) is a lower bound for the maximal probability of winning a concurrent Büchi game. The upper-bound result will follow from results on concurrent co-Büchi games.

Lemma 4. *Let*

$$w = [\![\nu y.\mu x.((\neg U \wedge \mathit{Ppre}_1(x)) \vee (U \wedge \mathit{Ppre}_1(y)))]\!].$$

For all $\varepsilon > 0$ player 1 has a strategy π_1^ε such that $\mathrm{Pr}_s^{\pi_1^\varepsilon,\pi_2}(\Box\Diamond U) > w(s) - \varepsilon$ for all $\pi_2 \in \Pi_2$ and all $s \in S$.

Proof. From ε, construct a positive sequence $\{\varepsilon_i\}_{i\geq 0}$ with $\sum_{i=0}^\infty \varepsilon_i < \varepsilon$. The strategy π_1^ε is as follows. In $S \setminus U$ the strategy π_1^ε initially coincides with a ε_0-optimal strategy for the game $w(\Diamond U)$. Upon reaching U, the strategy π_1^ε plays according to an optimal distribution of the matrix game corresponding $\mathrm{Ppre}_1(w)$, until U is left. In the following $\neg U$-phase, π_1^ε coincides with a ε_1-optimal strategy for the game $w(\Diamond U)$; and so forth. Fix a state $s_0 \in S$ and a strategy $\pi_2 \in \Pi_2$. Define the process $\{H_n\}_{n\geq 0}$, where H_n is the value of w at the n-th visit of U. From Lemma 3 and from the construction of π_1^ε, we have $\mathrm{E}_{s_0}^{\pi_1^\varepsilon,\pi_2}\{H_1\} \geq w(s_0) - \varepsilon_0$, and for $n \geq 0$,

$$\mathrm{E}_{s_0}^{\pi_1^\varepsilon,\pi_2}\{H_{n+1} \mid H_1, H_2, \ldots, H_n\} \geq H_n - \varepsilon_n.$$

By taking expectations on both sides, and by induction, this leads to

$$\mathrm{E}_{s_0}^{\pi_1^\varepsilon,\pi_2}\{H_{n+1}\} \geq w(s_0) - \textstyle\sum_{i=0}^n \varepsilon_i$$

for all $n \geq 0$. Denoting by $[\Box\Diamond]_n U$ the event of visiting U at least n times, we have $\mathrm{Pr}_{s_0}^{\pi_1^\varepsilon,\pi_2}([\Box\Diamond]_n U) \geq \mathrm{E}_{s_0}^{\pi_1^\varepsilon,\pi_2}\{H_n\}$. Combining these two results we obtain

$$\mathrm{Pr}_{s_0}^{\pi_1^\varepsilon,\pi_2}([\Box\Diamond]_n U) \geq w(s_0) - \varepsilon,$$

and the result then follows from

$$\lim_{n\to\infty} \mathrm{Pr}_{s_0}^{\pi_1^\varepsilon,\pi_2}([\Box\Diamond]_n U) = \mathrm{Pr}_{s_0}^{\pi_1^\varepsilon,\pi_2}(\Box\Diamond U).$$

∎

A *concurrent co-Büchi game* consists in a concurrent game structure $\mathcal{G} = \langle S, \mathit{Moves}, \Gamma_1, \Gamma_2, p\rangle$ together with a winning condition $\Diamond\Box U$, where $U \subseteq S$. Intuitively, the winning condition consists in eventually staying forever in the subset U of states. The solution of a concurrent co-Büchi game is given by

$$\langle 1\rangle \Diamond\Box U = [\![\mu x.\nu y.((\neg U \wedge \mathrm{Ppre}_1(x)) \vee (U \wedge \mathrm{Ppre}_1(y)))]\!]\ . \tag{4}$$

Again, the proof of the above fixpoint equation is based on two lemmas. Let $[\Box U]$ be the function that associates with each path value 1 if the path always stays in U, and value 0 otherwise. The first lemma generalizes Lemma 2.

LEMMA 5. *For $g \in \mathcal{F}$ and $U \subseteq S$, let*

$$w = [\![\nu y.((U \wedge Ppre_1(y)) \vee (\neg U \wedge g))]\!].$$

Then the strategy π_1 of player 1 that plays at each $s \in S$ according to an optimal distribution of the matrix game corresponding to $Ppre_1(w)(s)$ is such that $\mathrm{E}_s^{\pi_1,\pi_2}\{[\Box U] + g(\Diamond\neg U)\} \geq w$ for all $s \in S$ and $\pi_2 \in \Pi_2$.

The proof is similar to that of Lemma 2. The following lemma shows that the fixpoint of (3) is a lower bound for the maximal probability of winning the concurrent co-Büchi game.

LEMMA 6. *Let*

$$w = [\![\mu x.\nu y.((\neg U \wedge Ppre_1(x)) \vee (U \wedge Ppre_1(y)))]\!].$$

For all $\varepsilon > 0$ player 1 has a strategy π_1^ε such that $\mathrm{Pr}_s^{\pi_1^\varepsilon,\pi_2}(\Diamond\Box U) \geq w(s) - \varepsilon$ for all $\pi_2 \in \Pi_2$ and all $s \in S$.

Proof. Denote by $[\Diamond\Box U]_n$ the event of visiting $\neg U$ at most n times. Let $x_0 = \mathbf{0}$, and for $n > 0$,

$$x_n = [\![\nu y.((\neg U \wedge \mathrm{Ppre}_1(x_{n-1})) \vee (U \wedge \mathrm{Ppre}_1(y)))]\!].$$

By induction on $n \geq 0$, we show that player 1 has a strategy π_1^n such that $\mathrm{Pr}_s^{\pi_1^n,\pi_2}([\Diamond\Box U]_n) \geq x_n(s)$ for all $s \in S$ and all $\pi_2 \in \Pi_2$; the result will then follow by taking the limit $n \to \infty$, noting that $\lim_{n\to\infty} x_n = w$. The base case is trivial. For $n > 0$, the strategy π_1^n plays according to an optimal distribution of the matrix game corresponding to $\mathrm{Ppre}_1(x_n)$ as long as U is not left. At the first visit to $\neg U$, the strategy π_1^n plays one round according to an optimal distribution of the matrix game corresponding to $\mathrm{Ppre}_1(x_{n-1})$, and switches thereafter to the strategy π_1^{n-1}. By definition of π_1^n, from the previous lemma we have

$$\begin{aligned}\mathrm{Pr}_s^{\pi_1^n,\pi_2}([\Diamond\Box U]_n) &\geq \mathrm{Pr}_s^{\pi_1^n,\pi_2}(\Box U) + \mathrm{E}_s^{\pi_1^n,\pi_2}\{x_{n-1}(\Diamond\neg U)\}\\ &\geq x_n.\end{aligned}$$

The result then follows by taking the limit $n \to \infty$. ∎

The following theorem summarizes the results about concurrent Büchi and co-Büchi games.

THEOREM 2. *The following assertions hold.*

1. *Concurrent Büchi and co-Büchi games can be solved according to (3) and (4).*
2. *There are deterministic concurrent Büchi games without optimal strategies, and without finite-memory ε-optimal strategies.*
3. *There are deterministic concurrent co-Büchi games without optimal strategies.*

Part 1 follows from Lemmas 4 and 6, and from quantitative game μ-calculus complementation. Part 2 follows from the lack of optimal strategies for reachability (see Example 2), and from the fact that Büchi games are equivalent to iterated reachability games (see [8] for an example). Part 3 is a consequence of the lack of optimal strategies for concurrent reachability games.

5. RABIN-CHAIN GAMES

A *concurrent Rabin-chain game* consists in a concurrent game structure $\mathcal{G} = \langle S, Moves, \Gamma_1, \Gamma_2, p\rangle$ together with a winning condition

$$\mathcal{R} = \bigvee_{i=0}^{k-1} (\Box\Diamond U_{2i} \wedge \neg\Box\Diamond U_{2i+1}) ,$$

where $k > 0$ and $\emptyset = U_{2k} \subseteq U_{2k-1} \subseteq U_{2k-2} \subseteq \cdots \subseteq U_0 = S$. A more intuitive characterization of this winning condition can be obtained by defining, for $0 \leq i \leq 2k-1$, the set C_i of states of *color* i by $C_i = U_i \setminus U_{i+1}$. The total number of colors is $N = 2k$. Given a path $\overline{s}$, let $\mathit{Infi}(\overline{s}) \subseteq S$ be the set of states that occur infinitely often along $\overline{s}$, and let

$$\mathit{MaxCol}(\overline{s}) = \max\{i \in \{0, \ldots, N-1\} \mid C_i \cap \mathit{Infi}(\overline{s}) \neq \emptyset\}$$

be the largest color appearing along the path. Then,

$$\mathcal{R} = \{\overline{s} \in \Omega \mid \mathit{MaxCol}(\overline{s}) \text{ is even}\}.$$

The solution $\langle 1\rangle\mathcal{R}$ for a Rabin-chain condition with N colors is given by

$$\langle 1\rangle\mathcal{R} = [\![\eta_{N-1}x_{N-1}.\ldots.\mu x_1.\nu x_0.(\bigvee_{i=0}^{N-1}(C_i \wedge \mathrm{Ppre}_1(x_i)))]\!] \tag{5}$$

where $\eta_n = \nu$ if n is even, and $\eta_n = \mu$ if n is odd (compare with [11]). The proof of (5) is based on the following inductive decomposition, inspired by the one of [11]. We denote by $C_{\leq n} = \bigcup_{i=0}^{n} C_i$ (resp. $C_{>n} = \bigcup_{i=n+1}^{N-1} C_i$ and $C_{<n} = \bigcup_{i=0}^{n-1} C_i$) the set of states colored by colors less than or equal to n (resp., greater than n, and smaller than n). Let $z \in \mathcal{F}$, and for $n \geq 0$ define J_n by $J_{-1}(z) = z$, and

$$J_n(z) = \eta_n x.J_{n-1}((C_n \wedge \mathrm{Ppre}_1(x)) \vee (C_{>n} \wedge z)).$$

We can show by induction on n that $[\![J_n(z)]\!]$ is the function that gives the maximal expectation of either winning the concurrent Rabin-chain game while visiting only states in $C_{\leq n}$, or of the value $z(\Diamond C_{>n})$ if $C_{\leq n}$ is exited. Denote by $[\mathcal{R} \wedge \Box C_{\leq n}]$ the random function that has value 1 over a path exactly when the path satisfies condition $\mathcal{R}$ while visiting only states in $C_{\leq n}$. The lemma below makes this statement precise.

LEMMA 7. *For all $\varepsilon > 0$, all $n \in \{0, \ldots, N-1\}$, all $z \in \mathcal{F}$, and all states $s \in S$, there is a strategy $\pi_1 \in \Pi_1$ for player 1 such that for all strategies $\pi_2 \in \Pi_2$ of player 2, we have*

$$\mathrm{E}_s^{\pi_1,\pi_2}\{[\mathcal{R} \wedge \Box C_{\leq n}] + z(\Diamond C_{>n})\} \geq [\![J_n(z)]\!](s) - \varepsilon.$$

Proof. The proof of the lemma is similar to the proof of the lemmas for the Büchi and co-Büchi conditions; we sketch the inductive step for n odd (i.e., $\eta_n = \mu$). From ε, construct a positive sequence $\{\varepsilon_i\}_{i\geq 0}$ with sum less than ε. Let $x_0 = \mathbf{0}$, and for $k > 0$, let

$$x_k = [\![J_{n-1}(C_n \wedge \mathrm{Ppre}(x_{k-1}) \vee C_{>n} \wedge z)]\!].$$

By induction on k, we show that player 1 has a strategy π_1^k such that

$$\mathrm{E}_s^{\pi_1^k,\pi_2}\{[\mathcal{R} \wedge \Box C_{\leq n}] + z(\Diamond C_{>n})\} \geq x_k(s) - \textstyle\sum_{i=0}^k \varepsilon_i$$

for all $s \in S$ and $\pi_2 \in \Pi_2$. The base case is obvious. The strategy π_1^k for player 1 coincides with an ε_k-optimal strategy in the game $J_{n-1}(C_n \wedge \mathrm{Ppre}(x_{k-1}) \vee C_{>n} \wedge z)$ while the game remains in $C_{<n}$; when C_n is hit for the first time, it plays an optimal strategy in the matrix game $\mathrm{Ppre}_1(x_{k-1})$, and thereafter switches to the inductively constructed strategy π_1^{k-1}. Define the shorthand

$$W = [\mathcal{R} \wedge \Box C_{\leq n}] + z(\Diamond C_{>n}),$$

which represents winning while never leaving $C_{\leq n}$, or reaching $C_{>n}$ and getting reward z. For all $s \in S$ and $\pi_2 \in \Pi_2$, we have

$$\begin{aligned}
&\mathrm{E}_s^{\pi_1^k,\pi_2}\{W\} \\
&\geq \left[\!\!\left[J_{n-1} \left(\begin{array}{l} C_n \wedge \left(\lambda r.\mathrm{E}_r^{\pi_1^k[r],\pi_2[r]}\{W\}\right) \\ \vee \\ C_{>n} \wedge z \end{array} \right) \right]\!\!\right] (s) - \varepsilon_k \\
&\geq \left[\!\!\left[J_{n-1} \left(\begin{array}{l} C_n \wedge \left(\lambda r.\sum_{t \in S} \mathrm{E}_t^{\pi_1^{k-1},\pi_2[r,t]}\{W\} \mathrm{Pr}_r^{\pi_1^k[r],\pi_2[r]}\{\Theta_1 = t\}\right) \\ \vee \\ C_{>n} \wedge z \end{array} \right) \right]\!\!\right] (s) - \varepsilon_k \\
&\geq \left[\!\!\left[J_{n-1} \left(\begin{array}{l} C_n \wedge \left(\lambda r.\sum_{t \in S} \left(x_{k-1} - \sum_{i=0}^{k-1} \varepsilon_i\right) \mathrm{Pr}_r^{\pi_1^k[r],\pi_2[r]}\{\Theta_1 = t\}\right) \\ \vee \\ C_{>n} \wedge z \end{array} \right) \right]\!\!\right] (s) - \varepsilon_k \\
&\geq \left[\!\!\left[J_{n-1} \left(C_n \wedge \mathrm{Ppre}_1 \left(x_{k-1} - \textstyle\sum_{i=0}^{k-1} \varepsilon_i \right) \ \vee\ C_{>n} \wedge z \right) \right]\!\!\right] (s) - \varepsilon_k \\
&\geq [\![J_{n-1}(C_n \wedge \mathrm{Ppre}_1(x_{k-1}) \ \vee\ C_{>n} \wedge z)]\!] (s) - \sum_{i=0}^k \varepsilon_i \\
&= x_k(s) - \sum_{i=0}^k \varepsilon_i,
\end{aligned}$$

where $\lambda r.f(r)$ is the usual λ-calculus notation for the function that maps each r into $f(r)$. The first inequality follows by induction on n, and by a case analysis on the possible ways of leaving $C_{<n}$. The strategies $\pi_1^k[r]$ and $\pi_2[r]$ behave like π_1 and π_2 after the path from s to r. The second inequality follows then by an analysis of a single step from r. The strategy $\pi_2[r,t]$ is the strategy that behaves as π_2 after a path from s to r and t; by definition of π_1^k, we have $\pi_1^k[r,t] = \pi_1^{k-1}$. The third inequality follows by induction hypothesis on k, remembering the definition of W. The fourth inequality follows by using the definition of Ppre_1, and the fifth inequality follows by pulling out the constant from the Ppre_1 and the expectation. Finally, the claim follows by taking $k \to \infty$. A similar argument works for n even (i.e., $\eta_n = \nu$). Take $y_0 = \mathbf{1}$, and for $k > 0$, let

$$y_k = [\![J_{n-1}(C_n \wedge \mathrm{Ppre}(x_{k-1}) \ \vee\ C_{>n} \wedge z)]\!].$$

We construct similarly a strategy π_1^k that plays an ε_k-optimal strategy in the game $J_{n-1}(C_n \wedge \mathrm{Ppre}(y_{k-1}) \ \vee\ C_{>n} \wedge z)$ while the game remains in $C_{<n}$; when C_n is

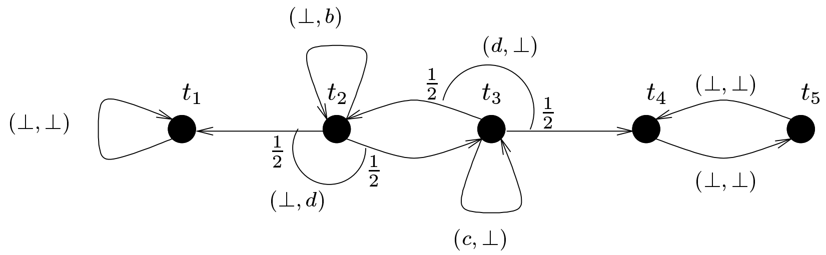

FIG. 1: A game that disproves the reduction to reachability. A label (a, b) of an edge (or of a probabilistic bundle of edges) indicates that the edge is followed when player 1 chooses move a and player 2 chooses move b.

hit for the first time, it plays an optimal strategy in the matrix game $\mathrm{Ppre}_1(y_{k-1})$, and thereafter switches to the inductively constructed strategy π_1^{k-1}. The details are as above. ∎

The value of the game with condition $\mathcal{R}$ is then $[\![J_{N-1}(\mathbf{0})]\!]$. Both the lower and the upper bounds for the value of the game follow from the lemma, because Rabin-chain games are self-dual (the complement of a concurrent Rabin-chain game is again a concurrent Rabin-chain game). We can now summarize the results for concurrent Rabin-chain games.

THEOREM 3. *The following assertions hold.*

1. *Concurrent Rabin-chain games can be solved according to (5).*
2. *There are deterministic concurrent Rabin-chain games without optimal strategies and without finite-memory ε-optimal strategies.*

Part 1 follows from Lemma 7. Again, the lack of optimal strategies, and of finite memory ε-optimal strategies follows from the result proved for Büchi games (which are special cases).

Finally, the next theorem states that if the state space is countable, rather than finite, the quantitative game μ-calculus solutions presented in this paper still define the value of the game.

THEOREM 4. *Consider a concurrent game structure $\mathcal{G} = \langle S, Moves, \Gamma_1, \Gamma_2, p\rangle$, where S is countable. Then, formulas (1), (2), (3), (4), and (5) provide the solutions for concurrent reachability, safety, Büchi, co-Büchi, and Rabin-chain games, respectively.*

This theorem can be proved by the same arguments used for finite concurrent games, using transfinite induction rather than ordinary induction when arguing about the least and greatest fixpoints of the calculus.

6. DISCUSSION

The solution formulas for concurrent games that have been presented in this paper lead to algorithms for the computation of approximate solutions of the games. In the case of safety and reachability games, the solution formulas (1) and (2) contain a single fixpoint operator. By computing these fixpoints in iterative fashion, we obtain approximation schemes that converge monotonically to the solution.

The speed of convergence of such schemes has not been characterized. On the other hand, for Büchi, co-Büchi, and Rabin-chain games, the alternance of fixpoint operators in the solution formulas yields approximation schemes that contain nested iterations, and it is not known how to obtain monotonically converging approximation schemes. Specifically, from (5) the solution of a Rabin-chain game with $2k$ colors has the fixpoint prefix $\mu x_{2k-1}.\nu x_{2k-2}.\ldots\mu x_1.\nu x_0$. Denote by $w(n_{2k-1}, n_{2k-2}, \ldots, n_0)$ the approximation of (5) computed by approximating the fixpoint ηx_i by n_i iterations, for $i \in \{0, \ldots, 2k-1\}$. It is not known how to select a sequence

$$(n_{2k-1}^{(0)}, \ldots, n_0^{(0)}), (n_{2k-1}^{(1)}, \ldots, n_0^{(1)}), (n_{2k-1}^{(2)}, \ldots, n_0^{(2)}), \ldots$$

such that for all $i \in \{0, \ldots, 2k-1\}$ we have $\lim_{j\to\infty} n_i^{(j)} = \infty$, and such that

$$\lim_{j\to\infty} w(n_{2k-1}^{(j)}, \ldots, n_0^{(j)}) = \langle 1 \rangle \mathcal{R}$$

with monotonic convergence.

This situation is in contrast with the situation for Markov decision processes. In a Markov decision process, the problem of computing the maximal probability of satisfying a Büchi, co-Büchi, or Rabin-chain condition Ψ can be solved in polynomial time, by reducing it to the problem of computing a maximal reachability probability [7]. From Ψ, we can first compute the subset $T_\Psi = \{s \in S \mid \langle 1 \rangle \Psi(s) = 1\}$ of states where the maximal probability of Ψ is 1. Then, we have $\langle 1 \rangle \Psi = \langle 1 \rangle \Diamond T_\Psi$, indicating that the maximal probability of satisfying Ψ is equal to the maximal probability of reaching T_Ψ. In concurrent games, given a Büchi, co-Büchi, or Rabin-chain condition Ψ, we can compute the set T_Ψ with the algorithms of [8], setting $T_\Psi = \langle\langle 1 \rangle\rangle_{limit} \Psi$. If the equality $\langle 1 \rangle \Psi = \langle 1 \rangle \Diamond T_\Psi$ held for concurrent games, it would provide monotonic approximation schemes for computing the value of the game (the problem would still not be reducible to linear programming, since the values may be irrational, as mentioned earlier). However, the following example demonstrates that the equality does not hold for games.

Example 3. Consider the game depicted in Figure 1. Let $U = \{t_1, t_2, t_4\}$, and consider the co-Büchi winning condition $\Diamond\Box U$. The set of states R_1 (resp. R_2) where player 1 (resp. 2) can ensure winning (resp. losing) with probability 1 are given by

$$R_1 = T_{\Diamond\Box U} = \{s \in S \mid \langle 1 \rangle \Diamond\Box U(s) = 1\} = \{t_1\}$$
$$R_2 = \{s \in S \mid \langle 2 \rangle \Box\Diamond\neg U(s) = 1\} = \{t_4, t_5\}.$$

For $i \in \{1, 2\}$, the maximal probability for player i of reaching R_i from outside R_i is zero: $\langle 1 \rangle \Diamond R_1(t_k) = 0$ for $k \neq 1$, and $\langle 2 \rangle \Diamond R_2(t_k) = 0$ for $k \notin \{4, 5\}$. Nevertheless, we can verify that $\langle 1 \rangle \Diamond\Box U(t_2) = 2/3$, and $\langle 1 \rangle \Diamond\Box U(t_3) = 1/3$. ∎

REFERENCES

[1] A. Aziz, V. Singhal, F. Balarin, R. Brayton, and A. Sangiovanni-Vincentelli. It usually works: The temporal logic of stochastic systems. In *Computer Aided Verification*, volume 939 of *Lect. Notes in Comp. Sci.* Springer-Verlag, 1995.

[2] D. Bertsekas. *Dynamic Programming and Optimal Control.* Athena Scientific, 1995. Volumes I and II.

[3] A. Bianco and L. de Alfaro. Model checking of probabilistic and nondeterministic systems. In *Found. of Software Tech. and Theor. Comp. Sci.*, volume 1026 of *Lect. Notes in Comp. Sci.*, pages 499–513. Springer-Verlag, 1995.

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

[5] J. Burch, K. McMillan, E. Clarke, and D. Dill. Sequential circuit verification using symbolic model checking. In *Proc. of the 27th ACM/IEEE Design Automation Conference*, pages 46–51, Orlando, FL, USA, June 1990.

[6] E. Clarke, M. Fujita, and X. Zhao. *Multi-Terminal Binary Decision Diagrams and Hybrid Decision Diagrams, Representations of Discrete Functions*, pages 93–108. Kluwer Academic Publishers, 1996.

[7] C. Courcoubetis and M. Yannakakis. Markov decision processes and regular events. In *Proc. 17th Int. Colloq. Aut. Lang. Prog.*, volume 443 of *Lect. Notes in Comp. Sci.*, pages 336–349. Springer-Verlag, 1990.

[8] L. de Alfaro and T. Henzinger. Concurrent omega-regular games. In *Proc. 15th IEEE Symp. Logic in Comp. Sci.*, pages 141–154, 2000.

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

[10] L. de Alfaro, M. Kwiatkowska, G. Norman, D. Parker, and R. Segala. Symbolic model checking of concurrent probabilistic processes using MTBDDs and the Kronecker representation. In *TACAS: Tools and Algorithms for the Construction and Analysis of Systems*, volume 1785 of *Lect. Notes in Comp. Sci.*, pages 395–410. Springer-Verlag, 2000.

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

[12] H. Everett. Recursive games. In *Contributions to the Theory of Games III*, volume 39 of *Annals of Mathematical Studies*, pages 47–78, 1957.

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

[14] Y. Gurevich and L. Harrington. Trees, automata, and games. In *Proc. 14th ACM Symp. Theory of Comp.*, pages 60–65. ACM Press, 1982.

[15] M. Huth and M. Kwiatkowska. Quantitative analysis and model checking. In *Proc. 12th IEEE Symp. Logic in Comp. Sci.*, pages 111–122, 1997.

[16] D. Kozen. A probabilistic PDL. In *Proc. 15th ACM Symp. Theory of Comp.*, pages 291–297, 1983.

[17] D. Kozen. Results on the propositional μ-calculus. *Theoretical Computer Science*, 27(3):333–354, 1983.

[18] P. Kumar and T. 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.

[19] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer-Verlag, New York, 1991.

[20] D. Martin. The determinacy of Blackwell games. *The Journal of Symbolic Logic*, 63(4):1565–1581, 1998.

[21] A. McIver. Reasoning about efficiency within a probabilitic μ-calculus. In *Proc. of PROBMIV*, pages 45–58, 1998. Technical Report CSR-98-4, University of Birmingham, School of Computer Science.

[22] A. McIver and C. Morgan. Demonic, angelic and unbounded probabilistic choices in sequential programs. *Acta Informatica*, 37(4/5):329–354, 2001.

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

[24] C. Morgan, A. McIver, and K. Seidel. Probabilistic predicate transformers. *ACM Trans. Prog. Lang. Sys.*, 18(3):325–353, 1996.

[25] A. 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.

[26] G. Owen. *Game Theory*. Academic Press, 1995.

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

[28] S. Safra. On the complexity of ω-automata. In *Proc. 29th IEEE Symp. Found. of Comp. Sci.*, pages 319–327. IEEE Computer Society Press, 1988.

[29] S. Safra. Exponential determinization for ω-automata with strong-fairness acceptance condition. In *Proc. ACM Symp. Theory of Comp.*, pages 275–282, 1992.

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

[31] 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.

[32] F. Thuijsman and O. Vrieze. The bad match, a total reward stochastic game. *Operations Research Spektrum*, 9:93–99, 1987.

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

[34] M. Vardi and P. Wolper. Reasoning about infinite computations. *Information and Computation*, 115(1):1–37, 1994.

[35] J. von Neumann and O. Morgenstern. *Theory of games and economic behavior.* Princeton University Press, 1947.

[36] D. Williams. *Probability With Martingales.* Cambridge University Press, 1991.