

Stochastic Games with Lossy Channels

Parosh Aziz Abdulla¹, Noomene Ben Henda¹, Luca de Alfaro²,
Richard Mayr³, and Sven Sandberg¹

¹ Uppsala University, Sweden

² University of California, Santa Cruz, USA

³ NC State University, USA

Abstract. We consider turn-based stochastic games on infinite graphs induced by game probabilistic lossy channel systems (GPLCS), the game version of probabilistic lossy channel systems (PLCS). We study games with Büchi (repeated reachability) objectives and almost-sure winning conditions. These games are pure memoryless determined and, under the assumption that the target set is regular, a symbolic representation of the set of winning states for each player can be effectively constructed. Thus, turn-based stochastic games on GPLCS are decidable. This generalizes the decidability result for PLCS-induced Markov decision processes in [10].

1 Introduction

Background. It is natural to model a reactive system as a 2-player game between the “controller” or player 0, who makes nondeterministic choices of the system, and the “environment” or player 1, who provides malicious inputs to the system. In this model, each state belongs to one of the players, who selects an outgoing transition that determines the next state. Starting in some initial state, the players jointly construct an infinite sequence of states called a *run*. The winning condition is specified as a predicate on runs. Verifying properties of the system corresponds to finding the winner of the game, where the winning condition depends on the property to check.

Systems that have a probabilistic component give rise to stochastic games. These are games where some states belong to “player random”, who selects the next state according to a pre-defined probability distribution. Randomness is useful to model stochastic loss of information such as unreliable communication, as well as randomized algorithms.

Previous work on algorithms for stochastic games has mostly focused on finite-state systems (see, e.g., [26,14,16,12]). However, many systems can only be faithfully modeled using infinitely many states. A lot of recent research has therefore been concerned with probabilistic infinite-state models. Probabilistic versions of lossy channel systems [11,7] and pushdown automata [18,19] use unbounded queues and stacks, respectively. Probabilistic Petri nets [4] model systems with an unbounded number of processes which run in parallel. The recently introduced Noisy Turing machines [8] model computer memories subject to stochastic errors.

We consider infinite-state stochastic games induced by lossy channel systems (LCS) [1,10,24]. LCS consist of finite-state control parts and unbounded channels (queues), i.e., automata where transitions are labeled by send and receive operations. They can model communication protocols such as the sliding window protocol and HDLC [6], where the communication medium is unreliable. In this paper, we introduce *game probabilistic LCS (GPLCS)*. GPLCS are *probabilistic* in the sense that the channels may randomly lose messages; and they are *games* in the sense that the next transition in the control part is selected by one of the players, depending on the current state. We can use player 0 to model nondeterminism in a communication protocol and player 1 to model a malicious cracker trying to break the protocol.

We consider Büchi (repeated reachability) objectives with almost-sure winning conditions. In other words, the goal for player 0 is to guarantee that with probability one, a given set of target states is visited infinitely many times. In the example of the malicious cracker, this corresponds to checking that the system can respond in such a way that it always eventually returns to a “ready state” with probability 1, no matter how the cracker acts.

Related Work. The work closest to ours is [9] where the authors consider the same model. They study GPLCS with simple reachability objectives and different winning conditions; i.e., almost-sure, with positive probability, etc. However, they do not consider GPLCS with Büchi objectives. Previous work on LCS considers several types of nondeterministic [20,6] and probabilistic systems (Markov chains) [22,1,24], as well as Markov decision processes [10] and non-stochastic games [3]. Of these, the work most closely related to ours is [10], which concerns LCS where messages are lost probabilistically and control transitions are taken nondeterministically (i.e., PLCS-induced Markov decision processes). This is a special case of our model in the sense that the game is restricted to only one player. It was shown in [10] that such 1-player Büchi-games are decidable (while coBüchi-games are undecidable). We generalize the decidability result of [10] for PLCS-induced Markov decision processes to 2-player stochastic games. The scheme presented in [10] also differs from ours in the fact that the target set is defined by control-states, while we consider more general *regular sets*. Thus our result is not a direct generalization of [10].

Stochastic games on infinite-state probabilistic recursive systems were studied in [18,19]. However, recursive systems are incomparable to the GPLCS model considered in this paper.

In [3], a model similar to ours is studied. It differs in that the system is not probabilistic, and instead one of the players controls message losses. For this model, [3] proves that safety games are decidable and parity games (which generalize Büchi games) are undecidable.

Two-player *concurrent* (but non-stochastic) games with infinite state spaces are studied in [17]. Concurrency means that the two players independently and simultaneously select actions, and the next state is determined by the combination of the actions and the current state. [17] describes schemes for computing winning sets and strategies for Büchi games (as well as reachability games and

some more general games). The article characterizes classes of games where the schemes terminate, based on properties of certain equivalence relations on states. However, this approach does not work for GPLCS (not even for non-probabilistic LCS), since LCS do not satisfy the necessary preconditions. Unlike the process classes studied in [17], LCS do not have a finite index w.r.t. the equivalences considered in [17].

In [28], a scheme is given to solve non-stochastic *parity* games on infinite state spaces of arbitrary cardinality. The parity condition is more general than the Büchi condition, so the scheme applies to Büchi games too. However, stochastic games are not considered. In fact, if our scheme is instantiated on the special case of non-stochastic Büchi games, it will coincide with the scheme in [28]. Furthermore, [28] does not suggest any class of infinite-state systems for which termination is guaranteed.

Our algorithms are related to the algorithms presented in [16,15] for solving concurrent games with respect to probability-1 ω -regular properties. However, the proofs in [16,15] apply only to finite-state games; we will need to develop entirely new arguments to prove the correctness of our approach for GPLCS.

Contribution. We prove that the almost-sure Büchi-GPLCS problem is decidable: we can compute symbolic representations of the winning sets and winning strategies for both players. The symbolic representations are based on regular expressions, and the result holds under the assumption that the set of target states is also regular. The winning strategies are pure memoryless, i.e., the next state depends only on the current state and is not selected probabilistically. Our result generalizes the decidability result for PLCS-induced Markov decision processes (i.e., 1-player games) in [10].

We now give an overview of our method. First, we give a scheme to compute the winning sets in simple reachability games, where the goal of player 0 is to reach a regular set of target states with a *positive probability*. Next, we give a scheme to construct the winning sets in almost-sure Büchi-games, using the scheme for reachability games as a subroutine. We prove that for GPLCS, both schemes terminate and we show how to instantiate them using *regular state languages* to effectively represent the infinite sets.

Outline. In Section 2, we define stochastic games. In Section 3, we describe GPLCS and show how they induce an infinite-state stochastic game. In Section 4, we show how to construct the winning sets in simple reachability games on GPLCS. In Section 5, we show how to construct the winning sets in Büchi games on GPLCS. Due to space limitations, some proofs are omitted and can be found in [2]; however, the intuitions are given in the main text.

2 Preliminaries

We use \mathbb{R}, \mathbb{N} for the real and natural numbers. If X is a set then X^* and X^ω denote the sets of finite and infinite sequences over X , respectively. The empty word is denoted by ε . For partial functions $f, g : X \rightarrow Y$ which have the same

value when both are defined, we use $f \cup g$ to denote the smallest function that extends both f and g .

A *probability distribution* on a countable set X is a function $f : X \rightarrow [0, 1]$ such that $\sum_{x \in X} f(x) = 1$. We will sometimes need to pick an arbitrary element from a set. To simplify the exposition, we let $\text{select}(X)$ denote an arbitrary but fixed element of the nonempty set X .

Turn-Based Stochastic Games. A *turn-based stochastic game* (or a *game* for short) is a tuple $\mathcal{G} = (S, S^0, S^1, S^R, \longrightarrow, P)$ where:

- S is a countable set of *states*, partitioned into the pairwise disjoint sets of *random states* S^R , states S^0 of player 0, and states S^1 of player 1.
- $\longrightarrow \subseteq S \times S$ is the *transition relation*. We write $s \longrightarrow s'$ to denote that $(s, s') \in \longrightarrow$. Let $\text{Post}(s) := \{s' : s \longrightarrow s'\}$ denote the set of *successors* of s and extend it to sets $Q \subseteq S$ of states by $\text{Post}(Q) := \bigcup_{s \in Q} \text{Post}(s)$. We assume that games are deadlock-free, i.e., each state has at least one successor ($\forall s \in S. \text{Post}(s) \neq \emptyset$).
- The *probability function* $P : S^R \times S \rightarrow [0, 1]$ satisfies both $\forall s \in S^R. \forall s' \in S. (P(s, s') > 0 \iff s \longrightarrow s')$ and $\forall s \in S^R. \sum_{s' \in S} P(s, s') = 1$. Note that for any given state $s \in S^R$, $P(s, \cdot)$ is a probability distribution over $\text{Post}(s)$.

For any set $Q \subseteq S$ of states, we let $\overline{Q} := S - Q$ denote its complement. We define $[Q]^R := Q \cap S^R$, $[Q]^0 := Q \cap S^0$, $[Q]^1 := Q \cap S^1$, and $[Q]^{01} := Q \cap (S^0 \cup S^1)$.

A *run* ρ in a game is an infinite sequence $s_0 s_1 \dots$ of states s.t. $s_i \longrightarrow s_{i+1}$ for all $i \geq 0$. We use $\rho(i)$ to denote s_i . A *path* π is a finite sequence $s_0 \dots s_n$ of states s.t. $s_i \longrightarrow s_{i+1}$ for all $i : 0 \leq i < n$. For any $Q \subseteq S$, we use Π_Q to denote the set of paths that end in some state in Q .

Informally, the two players 0 and 1 construct an infinite run $s_0 s_1 \dots$, starting in some initial state $s_0 \in S$. Player 0 chooses the successor s_{i+1} if $s_i \in S^0$, player 1 chooses s_{i+1} if $s_i \in S^1$, and the successor s_{i+1} is chosen randomly according to the probability distribution $P(s_i, \cdot)$ if $s_i \in S^R$.

Strategies. For $\sigma \in \{0, 1\}$, a *strategy* of player σ is a partial function $f^\sigma : \Pi_{S^\sigma} \rightarrow S$ s.t. $s_n \longrightarrow f^\sigma(s_0 \dots s_n)$ if f^σ is defined. The strategy f^σ prescribes for player σ the next move, given the current prefix of the run. We say that f^σ is total if it is defined for every $\pi \in \Pi_{S^\sigma}$.

A strategy f^σ of player σ is *memoryless* if the next state only depends on the current state and not on the previous history of the game, i.e., for any path $s_0 \dots s_k \in \Pi_{S^\sigma}$, we have $f^\sigma(s_0 \dots s_k) = f^\sigma(s_k)$. A memoryless strategy of player σ can be regarded simply as a function $f^\sigma : S^\sigma \rightarrow S$, such that $s \longrightarrow f^\sigma(s)$ whenever f^σ is defined.

Consider two total strategies f^0 and f^1 of player 0 and 1. A path $\pi = s_0 \dots s_n$ in \mathcal{G} is said to be *consistent* with f^0 and f^1 if the following holds. For all $0 \leq i \leq n-1$, $s_i \in S^0$ implies $f^0(s_0 \dots s_i) = s_{i+1}$ and $s_i \in S^1$ implies $f^1(s_0 \dots s_i) = s_{i+1}$. We define similarly *consistent* runs. In the sequel, whenever the strategies are known from the context, we assume that all mentioned paths and runs are consistent with them.

Probability Measures. We use the standard definition of the probability measure for a set of runs [23]. First, we define the measure for total strategies, and then extend it to general (partial) strategies. We let $\Omega^s = sS^\omega$ denote the set of all infinite sequences of states starting from s . Consider a game $\mathcal{G} = (S, S^0, S^1, S^R, \longrightarrow, P)$, an initial state s , and total strategies f^0 and f^1 of player 0 and 1. For a measurable set $\mathfrak{X} \subseteq \Omega^s$, we define $\mathcal{P}_{f^0, f^1}^s(\mathfrak{X})$ to be the probability measure of \mathfrak{X} under the strategies f^0, f^1 . It is well-known that this measure is well-defined [23]. When the state s is known from context, we drop the superscript and write $\mathcal{P}_{f^0, f^1}(\mathfrak{X})$. For (partial) strategies f^0 and f^1 of player 0 and 1, $\sim \in \{<, \leq, =, \geq, >\}$, and any measurable set $\mathfrak{X} \subseteq \Omega^s$, we define $\mathcal{P}_{f^0, f^1}^s(\mathfrak{X}) \sim x$ iff $\mathcal{P}_{g^0, g^1}^s(\mathfrak{X}) \sim x$ for all total strategies g^0 and g^1 which are extensions of f^0 resp. f^1 . For a single strategy f^σ of player σ , we define $\mathcal{P}_{f^\sigma}^s(\mathfrak{X}) \sim x$ iff $\mathcal{P}_{f^0, f^1}^s(\mathfrak{X}) \sim x$ for all strategies $f^{1-\sigma}$ of player $(1-\sigma)$. If $\mathcal{P}_{f^0, f^1}^s(\mathfrak{X}) = 1$, then we say that \mathfrak{X} happens *almost surely* under the strategies f^0, f^1 .

We assume familiarity with the syntax and semantics of the temporal logic CTL^* (see, e.g., [13]). We use $(s \models \varphi)$ to denote the set of runs starting in s that satisfy the CTL^* path-formula φ . We use $\mathcal{P}_{f^0, f^1}(s \models \varphi)$ to denote the measure of $(s \models \varphi)$ under strategies f^0, f^1 , i.e., we measure the probability of those runs which start in s , are consistent with f^0, f^1 and satisfy the path-formula φ . This set is measurable by [27].

Traps. For a player $\sigma \in \{0, 1\}$ and a set $Q \subseteq S$ of states, we say that Q is a σ -*trap* if player $(1-\sigma)$ has a strategy that forces all runs to stay inside Q . Formally, all successors of states in $[Q]^\sigma \cup [Q]^{1-\sigma}$ are in Q and every state in $[Q]^{1-\sigma}$ has some successor in Q .

Winning Conditions. Our main result considers *Büchi* objectives: player 0 wants to visit a given set $F \subseteq S$ infinitely many times. We consider games with *almost-sure* winning condition. More precisely, given an initial state $s \in S$, we want to check whether player 0 has a strategy f^0 such that for all strategies f^1 of player 1, it is the case that $\mathcal{P}_{f^0, f^1}(s \models \Box \Diamond F) = 1$.

Determinacy and Solvability. A game is said to be *determined* if, from every state, one of the players has a strategy that wins against all strategies of the opponent. Notice that determinacy implies that there is a partitioning W^0, W^1 of S , such that players 0 and 1 have winning strategies from W^0 and W^1 , respectively. A game is *memoryless determined* if it is determined and there are memoryless winning strategies. By *solving* a determined game, we mean giving an algorithm to check, for any state $s \in S$, whether $s \in W^0$ or $s \in W^1$.

3 Game Probabilistic Lossy Channel Systems (GPLCS)

A *lossy channel system (LCS)* [6] is a finite-state automaton equipped with a finite number of unbounded FIFO channels (queues). The system is *lossy* in the sense that, before and after a transition, an arbitrary number of messages may

be lost from the channels. *Probabilistic lossy channel system (PLCS)* [11,7,4] define a probabilistic model for message losses. The standard model assumes that each individual message is lost independently with probability λ in every step, where $\lambda > 0$ is a parameter of the system.

We consider *game probabilistic LCS (GPLCS)*, the 2-player game extension of PLCS. The set of states is partitioned into states belonging to player 0 and 1, and the transitions are controlled by the players. The player who owns the current control-state chooses an enabled outgoing transition. However, message losses occur randomly. While our definition of GPLCS (see below) assumes the same model of independent message loss as in [11,7,4], this is not necessary for our results. We only require the existence of a finite attractor, in the sense described in Section 5. In fact, many other probabilistic message loss models (e.g., burst disturbances, where groups of messages in close proximity are more often affected) satisfy this attractor condition [5].

The players have conflicting goals: player 0 wants to reach a given set of states infinitely often, and player 1 wants to visit it at most finitely many times. This is called a Büchi objective.

Formally, a GPLCS is a tuple $\mathcal{L} = (\mathbf{S}, \mathbf{S}^0, \mathbf{S}^1, \mathbf{C}, \mathbf{M}, \mathbf{T}, \lambda)$ where \mathbf{S} is a finite set of *control-states* partitioned into states $\mathbf{S}^0, \mathbf{S}^1$ of player 0 and 1; \mathbf{C} is a finite set of *channels*, \mathbf{M} is a finite set called the *message alphabet*, \mathbf{T} is a set of *transitions*, and $0 < \lambda < 1$ is the *loss rate*. Each transition $t \in \mathbf{T}$ is of the form $\mathbf{s} \xrightarrow{\text{op}} \mathbf{s}'$, where $\mathbf{s}, \mathbf{s}' \in \mathbf{S}$ and op is one of $c!m$ (send message $m \in \mathbf{M}$ in channel $c \in \mathbf{C}$), $c?m$ (receive message m from channel c), or nop (do not modify the channels).

A GPLCS $\mathcal{L} = (\mathbf{S}, \mathbf{S}^0, \mathbf{S}^1, \mathbf{C}, \mathbf{M}, \mathbf{T}, \lambda)$ induces a game $\mathcal{G} = (S, S^0, S^1, S^R, \longrightarrow, P)$, where $S = \mathbf{S} \times (\mathbf{M}^*)^{\mathbf{C}} \times \{0, 1\}$. That is, each state in the game consists of a control-state, a function that assigns a finite word over the message alphabet to each channel, and one of the symbols 0 or 1. States where the last symbol is 0 are random: $S^R = \mathbf{S} \times (\mathbf{M}^*)^{\mathbf{C}} \times \{0\}$. The other states belong to a player according to the control-state: $S^\sigma = \mathbf{S}^\sigma \times (\mathbf{M}^*)^{\mathbf{C}} \times \{1\}$. Transitions out of states of the form $s = (\mathbf{s}, \mathbf{x}, 1)$ model transitions in \mathbf{T} leaving state \mathbf{s} . On the other hand, transitions leaving states of the form $s = (\mathbf{s}, \mathbf{x}, 0)$ model message losses.

If $s = (\mathbf{s}, \mathbf{x}, 1), s' = (\mathbf{s}', \mathbf{x}', 0) \in S$, then there is a transition $s \longrightarrow s'$ in the game iff one of the following holds:

- $\mathbf{s} \xrightarrow{\text{nop}} \mathbf{s}'$ and $\mathbf{x} = \mathbf{x}'$;
- $\mathbf{s} \xrightarrow{c!m} \mathbf{s}'$, $\mathbf{x}'(c) = \mathbf{x}(c)m$, and for all $c' \in \mathbf{C} - \{c\}$, $\mathbf{x}'(c') = \mathbf{x}(c')$;
- $\mathbf{s} \xrightarrow{c?m} \mathbf{s}'$, $\mathbf{x}(c) = m\mathbf{x}'(c)$, and for all $c' \in \mathbf{C} - \{c\}$, $\mathbf{x}'(c') = \mathbf{x}(c')$.

To model message losses, we introduce the subword ordering \preceq on words: $x \preceq y$ iff x is a word obtained by removing zero or more messages from arbitrary positions of y . This is extended to channel states $\mathbf{x}, \mathbf{x}' : \mathbf{C} \rightarrow \mathbf{M}^*$ by $\mathbf{x} \preceq \mathbf{x}'$ iff $\mathbf{x}(c) \preceq \mathbf{x}'(c)$ for all channels $c \in \mathbf{C}$, and to game states $s = (\mathbf{s}, \mathbf{x}, i), s' = (\mathbf{s}', \mathbf{x}', i') \in S$ by $s \preceq s'$ iff $\mathbf{s} = \mathbf{s}'$, $\mathbf{x} \preceq \mathbf{x}'$, and $i = i'$. For any $s = (\mathbf{s}, \mathbf{x}, 0)$ and any \mathbf{x}' such that $\mathbf{x}' \preceq \mathbf{x}$, there is a transition $s \longrightarrow (\mathbf{s}, \mathbf{x}', 1)$. The probability of random transitions is given by $P((\mathbf{s}, \mathbf{x}, 0), (\mathbf{s}, \mathbf{x}', 1)) = a \cdot \lambda^b \cdot (1 - \lambda)^c$, where a is the number of ways

to obtain \mathbf{x}' by losing messages in \mathbf{x} , b is the total number of messages lost in all channels, and c is the total number of messages in all channels of \mathbf{x}' .

Every state on the form $(\mathbf{s}, \mathbf{x}, 0)$ has at least one successor, namely $(\mathbf{s}, \mathbf{x}, 1)$. If a state $(\mathbf{s}, \mathbf{x}, 1)$ does not have successors according to the rules above, then we add a transition $(\mathbf{s}, \mathbf{x}, 1) \longrightarrow (\mathbf{s}, \mathbf{x}, 0)$, to avoid deadlocks. Intuitively, this means that the run stays in the same control state and only loses messages.

Observe that the game is *bipartite*: every transition goes from a player state to a probabilistic state or the other way around, i.e., $\longrightarrow \subseteq ((S^0 \cup S^1) \times S^R) \cup (S^R \times (S^0 \cup S^1))$.

Problem Statement. We study the problem BÜCHI-GPLCS, defined as follows. The game graph is induced by a GPLCS; and we consider the almost-sure Büchi objective: player 0 wants to ensure that a given target set is visited infinitely often with probability one.

4 Reachability Games on GPLCS

We consider the reachability game where the winning condition is to reach a given target set with positive probability. Reachability games on GPLCS (with this and various other winning conditions) have been studied in [9], where the winning sets are expressed in terms of the target set in a variant of the μ -calculus.

Nevertheless, we give below a more ad-hoc scheme for computing the winning set, in order to keep the article self-contained. Furthermore, many definitions and some more detailed results on the structure of the winning sets and strategies will be needed in the following section on Büchi-games.

We give a scheme for characterizing sets of states from which a player can, with a positive probability, force the game into a given set of target states, while preserving a given invariant. We show that the scheme always terminates for GPLCS, and then give a symbolic representation of the winning sets, based on regular languages. The symbolic representation is valid under the assumption that the set of target states is also regular. Finally, we show correctness of the construction by describing the winning strategies. In fact, we show that if a player can win, then a *memoryless* strategy is sufficient to win.

Scheme. Fix a game $\mathcal{G} = (S, S^0, S^1, S^R, \longrightarrow, P)$ and two sets of states $F, I \subseteq S$, called the *target* and *invariant* sets, respectively. For a player $\sigma \in \{0, 1\}$, we give a scheme for constructing the set $\text{Force}^\sigma(I, F)$ of states where player σ can, with a positive probability, force the run to eventually reach F , while also preserving the property that the run will always remain within I (i.e., states outside I are not visited before F).

The idea of the scheme is to perform backward reachability analysis using the basic operations Pre^σ and $\widetilde{\text{Pre}}^\sigma$, defined as follows. Given $\sigma \in \{0, 1, R\}$ and a set $Q \subseteq S$ of states, let $\text{Pre}^\sigma(Q) := \{s \in S^\sigma : \exists s' \in Q. s \longrightarrow s'\}$ denote the set of states of player σ where it is possible to go to Q in the next step. Define $\widetilde{\text{Pre}}^\sigma(Q) := S^\sigma - \text{Pre}^\sigma(\overline{Q})$ to be the set of states where player σ cannot avoid going to Q in the next step.

The construction is inductive. For $\sigma \in \{0, 1\}$, we define two sequences $\{D_i\}_{i \in \mathbb{N}} : D_0 \subseteq D_1 \subseteq \dots$ and $\{E_i\}_{i \in \mathbb{N}} : E_0 \subseteq E_1 \subseteq \dots$ of sets of states as follows:

$$\begin{aligned} D_0 &:= [F]^R \cap I & E_0 &:= [F]^{01} \cap I \\ D_{i+1} &:= (D_i \cup \text{Pre}^R(E_i)) \cap I & E_{i+1} &:= \left(E_i \cup \text{Pre}^\sigma(D_i) \cup \widetilde{\text{Pre}}^{1-\sigma}(D_i) \right) \cap I. \end{aligned}$$

We let $\text{Force}^\sigma(I, F) := \bigcup_{i \geq 0} D_i \cup E_i$. Intuitively, the set D_i contains those states in S^R from which player σ can force the game to F with positive probability (while remaining in I) within i steps. The set E_i contains the states in $S^0 \cup S^1$ satisfying the same property¹.

Below, we instantiate the above described scheme for GPLCS. In the rest of this section, we consider the game $\mathcal{G} = (S, S^0, S^1, S^R, \longrightarrow, P)$ induced by a GPLCS $\mathcal{L} = (\mathbf{s}, \mathbf{S}^0, \mathbf{S}^1, \mathbf{C}, \mathbf{M}, \mathbf{T}, \lambda)$.

Termination. We recall from [21] that the relation \preceq is a *well quasi-ordering*, i.e., for each infinite sequence w_0, w_1, w_2, \dots of words over \mathbf{M} , there are $j < k$ such that $w_j \preceq w_k$. A set $U \subseteq \mathbf{M}^*$ is said to be *upward closed* if $w \in U$ implies that $w' \in U$ for each $w' \succeq w$. A *channel language* L is a mapping from \mathbf{C} to $2^{\mathbf{M}^*}$. In other words, L maps each channel to a language over \mathbf{M} . We say that L is *upward closed* resp. *regular* if $L(c)$ is upward closed resp. regular for each $c \in \mathbf{C}$. A *state language* L is of the form (\mathbf{s}, L') where $\mathbf{s} \in \mathbf{S}$ and L' is a channel language. We say that L is upward closed (regular) if L' is upward closed (regular). We generalize the definitions above to finite sets M of state languages by M upward closed (regular) if each $L \in M$ is upward closed (regular). The well quasi-ordering property carries directly over from words (mentioned earlier) to our finite sets of state languages (when required to hold for every control-state).

To prove termination of the scheme, we show that sets D_i are “almost” upward closed in the sense that they are closely related to other sets which are upward closed. More precisely, we consider the sequence $D'_0 \subseteq D'_1 \subseteq \dots$ of sets of states where $D'_0 := [F]^R$ and $D'_{i+1} := \text{Pre}^R(E_i)$. Since $\text{Pre}^R(Q)$ is upward closed for any set Q of states, it follows that D'_i is upward closed for each $i > 0$. Upward closedness, together with the well quasi-ordering of \preceq , implies that there is a j such that $D'_j = D'_{j+1} = \dots$. We also observe that $D_i = (D'_0 \cup D'_1 \cup \dots \cup D'_i) \cap I$. This means that $D_{j+1} = D_j$ and consequently $E_{j+2} = E_{j+1}$. Hence, we have the following lemma.

Lemma 1. *For any GPLCS and sets $F, I \subseteq S$ of states, the sequences $\{D_i\}_{i \in \mathbb{N}}$ and $\{E_i\}_{i \in \mathbb{N}}$ converge.*

Forms of Winning Sets. The above termination argument relied on upward closedness of the sets D'_i . In fact, we can derive more information about the structure of the winning sets for games induced by GPLCS. Assuming that the

¹ It is possible to define only one sequence, not separating player states from random states. In later proofs, it will be technically convenient to have the sequence $\{D_i\}_{i \in \mathbb{N}}$ defined, since $\{D_i\}_{i \in \mathbb{N}}$ has properties not shared by $\{E_i\}_{i \in \mathbb{N}}$.

sets F and I are regular state languages, it follows that each set D_i or E_i is also a regular state language. This follows from the fact that regular state languages are closed under the application of Pre^σ and the Boolean operations. Since the scheme terminates (by Lemma 1), the winning set $Q := \text{Force}^\sigma(I, F)$ is also regular. Furthermore, if I and F are upward closed then $[Q]^R$ is also upward closed. This follows from the fact that $\text{Pre}^R(Q)$ is upward closed for any set Q and that the class of upward closed sets is closed under intersection and union. We summarize these properties as properties (1)–(2) of the following lemma. (Properties (2)–(5) are not needed until the next section).

Lemma 2. *Let $Q = \text{Force}^\sigma(I, F)$. Then:*

- (1) *If F and I are regular then Q is regular.*
- (2) *If F and I are upward closed then $[Q]^R$ is upward closed.*
- (3) *Let $s \in I - Q$. If $s \in S^\sigma \cup S^R$, then $\text{Post}(s) \subseteq \overline{Q}$. If $s \in S^{1-\sigma}$, then $\text{Post}(s) \cap \overline{Q} \neq \emptyset$.*
- (4) *$\text{Force}^\sigma(Q, F) = Q$.*
- (5) *$\text{Force}^\sigma(S, F)$ is a σ -trap.*

Correctness. First, we describe a partial memoryless winning strategy $\text{force}^\sigma(I, F)$ for player σ from the states in $[\text{Force}^\sigma(I, F)]^\sigma$. Recall that a memoryless strategy can simply be described as a function that assigns one successor to each state. We define a sequence $e_0 \subseteq e_1 \subseteq e_2 \subseteq \dots$ of strategies for player σ . Let $e_0 := \emptyset$ and define e_{i+1} as follows:

- If $e_i(s)$ is defined then $e_{i+1}(s) := e_i(s)$.
- If $e_i(s)$ is undefined and $s \in [E_{i+1} - E_i]^\sigma$ then $e_{i+1}(s) := \text{select}(\text{Post}(s) \cap D_i)$.

Let $\text{force}^\sigma(I, F) := \bigcup_{i \geq 0} e_i$. From the definitions, we derive the following lemma.

Lemma 3. *In any GPLCS, for any $I, F \subseteq S$, $\sigma \in \{0, 1\}$, and $s \in \text{Force}^\sigma(I, F)$, there exists an $\epsilon_s > 0$ such that $\mathcal{P}_{\text{force}^\sigma(I, F)}(s \models \diamond F) \geq \epsilon_s$.²*

Proof. We recall the construction of the force sets and use induction on i to prove that $\forall i \in \mathbb{N}$ and for any state $s \in (D_i \cup E_i)$ the following holds: There exists an $\epsilon_s > 0$ such that for any extension f^σ of the $\text{force}^\sigma(I, F)$ and any strategy $f^{1-\sigma}$ of the opponent, $\mathcal{P}_{\text{force}^\sigma(I, F), f^{1-\sigma}}(s \models \diamond Q) \geq \epsilon_s$. Observe that $\forall i. D_i \cap E_i = \emptyset$.

The base case $s \in (D_0 \cup E_0) \subseteq F$ holds trivially (take $\epsilon_s := 1$).

Now assume that the claim holds for $i \in \mathbb{N}$. Consider s to be in $D_{i+1} \cup E_{i+1}$. In the case $s \in (D_i \cup E_i)$ the claim already holds by induction hypothesis. The remaining cases are described below.

Case $s \in D_{i+1} - D_i$: This implies that $s \in \text{Pre}^R(E_i)$. Thus there is a state $s' \in E_i$ such that $s \rightarrow s'$ and $\epsilon_{s'} > 0$ by induction hypothesis. We define $\epsilon_s := P(s, s') * \epsilon_{s'} > 0$.

² The weaker statement, i.e., $\mathcal{P}_{\text{force}^\sigma(I, F)}(s \models \diamond F) > 0$, suffices for correctness. However, this stronger version is needed in the sequel.

Case $s \in E_{i+1} - E_i$: This implies one of the following two cases.

- If $s \in [S]^\sigma$ then $s \in \text{Pre}^\sigma(D_i)$. Thus there is a state $s' \in D_i$ which is chosen as successor state to s by the $\text{force}^\sigma(I, F)$ strategy, i.e., $s \longrightarrow s'$ and $s' = \text{force}^\sigma(I, F)(s)$. By induction hypothesis $\epsilon_{s'} > 0$. So we obtain $\epsilon_s := \epsilon_{\text{force}^\sigma(I, Q)(s)} = \epsilon_{s'} > 0$.
- If $s \in [S]^{1-\sigma}$ then $s \in \widetilde{\text{Pre}}^{1-\sigma}(D_i)$. It follows that $\text{Post}(s) \subseteq D_i$. The set $\text{Post}(s)$ is finite, since the system is finitely branching. Furthermore, by induction hypothesis, $\epsilon_{s'} > 0$ for all $s' \in D_i$. Thus we obtain $\epsilon_s := \min_{s' \in \text{Post}(s)}(\epsilon_{s'}) > 0$.

The main result follows since for any $s \in \text{Force}^\sigma(I, Q)$, there exists a finite minimal $i \in \mathbb{N}$ such that $s \in (D_i \cup E_i)$. \square

In the sequel, we use $\text{Force}^\sigma(F)$ to denote $\text{Force}^\sigma(S, F)$, i.e., we do not mention I in case it is equal to S . We define $\text{force}^\sigma(F)$ analogously.

5 Büchi-Games on GPLCS

In this section we consider the BÜCHI-GPLCS problem. We give a scheme for characterizing the winning sets in almost-sure Büchi games, and then instantiate the scheme for GPLCS. In a similar manner to Section 4, we first show that the scheme always terminates for GPLCS, and then describe the winning sets using a symbolic representation based on regular languages. Again, the symbolic representation is valid under the assumption that the set of final states is also regular. We show the correctness of the construction by describing the memoryless winning strategies. Observe that this implies that BÜCHI-GPLCS are memoryless determined and solvable. Throughout this section, we fix a GPLCS $\mathcal{L} = (\mathbf{S}, \mathbf{S}^0, \mathbf{S}^1, \mathbf{C}, \mathbf{M}, \mathbf{T}, \lambda)$ and the induced game $\mathcal{G} = (S, S^0, S^1, S^R, \longrightarrow, P)$. Take $F \subseteq S$; we consider the Büchi goal for player 0 consisting in visiting F infinitely often.

Scheme. We define a sequence $\{X_i\}_{i \in \mathbb{N}} : X_0 \subseteq X_1 \subseteq \dots$ of sets of states which are winning for player 1 with a positive probability. In the definition of $\{X_i\}_{i \in \mathbb{N}}$, we use an auxiliary sequence $\{M_i\}_{i \in \mathbb{N}} : M_0 \supseteq M_1 \supseteq \dots$ of sets of states. The construction is inductive where $X_0 := \emptyset$, $M_0 := S$ and

$$M_{i+1} := \text{Force}^0(\overline{X_i}, F) \quad X_{i+1} := \text{Force}^1(\overline{M_{i+1}})$$

for each $i \geq 0$. Intuitively, the set X_i consists of states “already classified as losing for player 0”. We add states iteratively to these sets. We define M_{i+1} such that $\overline{M_{i+1}}$ is the set of states where player 0 cannot reach F with positive probability while staying always in $\overline{X_i}$. Finally, we claim that the winning states for player 0 are given by $W^0 := \bigcap_{i \geq 0} M_i$, and thus complementarily, the winning states for player 1 are given by $W^1 := \overline{W^0} = \bigcup_{i \geq 0} X_i$.

This property holds by the definitions and will be used later in this section.

Lemma 4. $X_0 \subseteq \overline{M_1} \subseteq X_1 \subseteq \overline{M_2} \subseteq X_2 \subseteq \dots$

The following lemma shows that this construction terminates.

Lemma 5. *The sequence $\{X_i\}_{i \in \mathbb{N}}$ converges for any set $F \subseteq S$ of states.*

Proof. (Sketch; details in [2]) Consider the sequence in Lemma 4. We perform the proof in four steps; namely, we show that (i) there is a K such that $[X_K]^R = [X_{K+1}]^R$; (ii) $X_{K+1} = \overline{M_{K+1}}$; (iii) $M_{K+1} = M_{K+2}$; (iv) $X_{K+1} = X_{K+2}$.

- (i) We show that each $[X_i]^R$ is upward closed, using induction on i . The base case is trivial since $X_0 = \emptyset$. For the induction step we let $Y := \overline{[M_{i+1}]^{01}} \cup [X_i]^R$. Using the definitions of X_i , X_{i+1} , and M_{i+1} , it can be shown that $X_{i+1} = \text{Force}^1(Y)$. Since $[X_i]^R$ is upward closed by the induction hypothesis it follows by Lemma 2(2) that $[X_{i+1}]^R$ is upward closed. From this and well quasi-ordering of \preceq , we get $\exists K. [X_K]^R = [X_{K+1}]^R$. We will use K in the rest of the analysis below.
- (ii) From Lemma 4 and the fact that $[X_K]^R = [X_{K+1}]^R$, we know that $[X_K]^R = \overline{[M_{K+1}]^R} = [X_{K+1}]^R$. This is used to show that $\text{Pre}^R(\overline{M_{K+1}}), \text{Pre}^1(\overline{M_{K+1}}), \widetilde{\text{Pre}}^0(\overline{M_{K+1}}) \subseteq \overline{M_{K+1}}$ which by the definition of X_{K+1} implies $X_{K+1} \subseteq \overline{M_{K+1}}$. Hence, $X_{K+1} = \overline{M_{K+1}}$, by Lemma 4.
- (iii) Since $M_{K+2} = \text{Force}^0(\overline{X_{K+1}}, F)$ and $X_{K+1} = \overline{M_{K+1}}$, we have that $M_{K+2} = \text{Force}^0(M_{K+1}, F)$. From Lemma 2(4) and the fact that $M_{K+1} = \text{Force}^0(\overline{X_K}, F)$, it follows that $M_{K+2} = M_{K+1}$.
- (iv) $X_{K+2} = \text{Force}^1(\overline{M_{K+2}}) = \text{Force}^1(\overline{M_{K+1}}) = X_{K+1}$. \square

Forms of Winning Sets. From Lemma 2(1), it follows that if F is regular then each X_i and M_i is regular. From Lemma 5 we get the following:

Lemma 6. *If F is regular then W^0 and W^1 are regular.*

Winning Strategy for Player 1. We define a sequence $\{x_i\}_{i \in \mathbb{N}}$ of strategies for player 1, such that $x_0 \subseteq x_1 \subseteq \dots$. For each i , the strategy $x_i : [X_i]^1 \rightarrow S$ is memoryless and winning for player 1 from states in X_i . The sequence $\{x_i\}_{i \in \mathbb{N}}$ converges to a memoryless strategy $w^1 := \bigcup_{i \in \mathbb{N}} x_i$ for player 1 which is winning from states in W^1 . We define the sequence using induction on i . We will also motivate why the strategy is winning for player 1. Define $x_0 := \emptyset$. For all $i \geq 0$, we define $x_{i+1}(s)$ by case analysis. By Lemma 4, we know that $X_i \subseteq \overline{M_{i+1}} \subseteq X_{i+1}$. There are three cases, reflecting the membership of s in these three sets:

- (i) If $s \in X_i$ then $x_{i+1}(s) := x_i(s)$. Here, we know by the induction hypothesis that a winning strategy x_i for player 1 has already been defined in s .
- (ii) If $s \in \overline{M_{i+1}} - X_i$ then $x_{i+1}(s) := \text{select}(\text{Post}(s) \cap \overline{M_{i+1}})$. The idea is that player 1 uses a strategy which guarantees that any run either (A) will stay in $\overline{M_{i+1}} - X_i$; or (B) will eventually enter X_i . In (A), player 1 wins since $\overline{M_{i+1}} - X_i$ does not have any states in F by the definition of M_{i+1} . In (B), player 1 wins by the induction hypothesis.

More precisely, we observe that player 1 selects a successor of s which belongs to $\overline{M_{i+1}}$. Such a successor exists by the following argument. First, observe that (by set operations) $\overline{M_{i+1}} - X_i = \overline{X_i} - M_{i+1}$. The result follows

by instantiating Lemma 2(3) with $I = \overline{X_i}$ and $Q = M_{i+1}$. By the same argument, for each $s' \in [\overline{M_{i+1}}]^R \cup [\overline{M_{i+1}}]^0$, all successors of s' belong to $\overline{M_{i+1}}$. This guarantees that either (A) or (B) holds.

- (iii) If $s \in X_{i+1} - \overline{M_{i+1}}$ then $x_{i+1}(s) := \text{force}^1(\overline{M_{i+1}})(s)$. Since, by definition, $X_{i+1} = \text{Force}^1(\overline{M_{i+1}})$, player 1 can use $\text{force}^1(\overline{M_{i+1}})$ to take the game with a positive probability to $\overline{M_{i+1}}$ (Lemma 3). From there, player 1 wins as described above.

Now, consider a state $s \in W^1$. By definition, we know that $s \in X_i$ for some $i \geq 0$. This means that $w^1 = x_i$ is winning for player 1 from s according to the above argument. Hence:

Lemma 7. *For each $s \in W^1$, $\mathcal{P}_{w^1}(s \models \neg \square \diamond F) > 0$.*

Winning Strategy for Player 0. In this paragraph, we define a memoryless strategy w^0 and we prove that it is winning.

To describe how w^0 is defined, we rely on two auxiliary results on games induced by GPLCS. First, we recall the definition of *an attractor*. A set $A \subseteq S$ is called an attractor if $\mathcal{P}(s \models \diamond A) = 1$ for any $s \in S$. In other words, from any state $s \in S$, A is almost surely visited regardless of the strategies of the players. The following result was shown in [11,7,4] for *probabilistic LCS*, where moves in the control graph are taken probabilistically instead of by two competing players. The results straightforwardly generalize to GPLCS.

Lemma 8. *Let $\mathcal{L} = (\mathbf{S}, \mathbf{S}^0, \mathbf{S}^1, \mathbf{C}, \mathbf{M}, \mathbf{T}, \lambda)$ be a GPLCS and let \mathcal{G} be the game induced by \mathcal{L} . The set $A = (\mathbf{S} \times \varepsilon \times \{0, 1\})$ is a finite attractor in \mathcal{G} .*

The second result follows from Lemma 8 and Lemma 3 as described below.

Lemma 9. *Let $\mathcal{G} = (S, S^0, S^1, S^R, \longrightarrow, P)$ be a game induced by a GPLCS. For any $Q, I \subseteq S$ and $\sigma \in \{0, 1\}$, the following holds: For any $s \in \text{Force}^\sigma(I, Q)$, $\mathcal{P}_{\text{force}^\sigma(I, Q)}(s \models \square \text{Force}^\sigma(I, Q) \wedge \neg \square \diamond Q) = 0$.*

Proof. Given $Q, I \subseteq S$, $\sigma \in \{0, 1\}$, and $s \in \text{Force}^\sigma(I, Q)$. We assume that player σ uses an extension f^σ of the $\text{force}^\sigma(I, Q)$ strategy and player $(1 - \sigma)$ uses a strategy $f^{1-\sigma}$. By Lemma 8, the game has a finite attractor A . By definition of the attractor, almost all runs must visit A infinitely often. We define $A' := A \cap \text{Force}^\sigma(I, Q)$.

If $A' = \emptyset$, then $\mathcal{P}_{f^\sigma, f^{1-\sigma}}(s \models \square \text{Force}^\sigma(I, Q) \wedge \neg \square \diamond Q) \leq \mathcal{P}_{f^\sigma, f^{1-\sigma}}(s \models \neg \diamond A) = 0$, where the inequality follows from the assumption and the equality from the definition of an attractor.

Consider now the case where $A' \neq \emptyset$. By Lemma 3 and finiteness of A (and thus A'), we obtain that $\epsilon := \min_{s' \in A'}(\epsilon'_s) > 0$. Almost every run in $(s \models \square \text{Force}^\sigma(I, Q) \wedge \neg \square \diamond Q)$ must visit A' infinitely many times, but Q only finitely many times (and thus have an infinite suffix which never visits Q). Thus,

$$\mathcal{P}_{f^\sigma, f^{1-\sigma}}(s \models \square \text{Force}^\sigma(I, Q) \wedge \neg \square \diamond Q) \leq (1 - \epsilon)^\infty = 0. \quad (1)$$

Therefore $\mathcal{P}_{\text{force}^\sigma(I, Q)}(s \models \square \text{Force}^\sigma(I, Q) \wedge \neg \square \diamond Q) = 0$. \square

Remark 1. Observe that the inequality (1) holds for any strategy $f^{1-\sigma}$ of the opponent and any extension f^σ of the $\text{force}^\sigma(I, Q)$ strategy. In particular, we do **not** require that $f^{1-\sigma}$ is finite-memory. It is possible that $f^{1-\sigma}$ acts quite differently after each of the (possibly infinitely many) visits to the same state in the attractor. The crucial fact is that the quantity $\epsilon_s > 0$ in Lemma 3 is independent of $f^{1-\sigma}$.

Now we are ready to describe the winning strategy w^0 for player 0. The idea of the strategy w^0 is to keep the run within a force set of F with probability 1. This implies that F will be visited infinitely often with probability 1, by Lemma 9. In order to do that, player 0 exploits certain properties of W^0 : By Lemmas 5 and 4, there is an i such that $X_i = \overline{M_{i+1}} = X_{i+1}$. From this and the definition of W^0 it follows that $W^0 = \overline{X_i} = M_{i+1}$. From $W^0 = \overline{X_i}$ and Lemma 2(5) it follows that W^0 is a 1-trap. From $M_{i+1} = \text{Force}^0(\overline{X_i}, F)$, it follows that $W^0 = \text{Force}^0(W^0, F)$. We define w^0 on any state $s \in [W^0]^0$ as follows:

- If $s \in W^0 - F$ then $w^0(s) := \text{force}^0(W^0, F)(s)$. This definition is possible since $W^0 = \text{Force}^0(W^0, F)$.
- If $s \in W^0 \cap F$ then $w^0(s) := \text{select}(\text{Post}(s) \cap W^0)$. This is possible since W^0 is a 1-trap, and therefore s has at least one successor in W^0 .

Consider any run ρ starting from a state inside W^0 , where player 0 follows w^0 . Since W^0 is a 1-trap, ρ will always remain inside W^0 regardless of the strategy of player 1. This implies that $\mathcal{P}_{w^0}(s \models \Box W^0) = 1$. Furthermore, by Lemma 9 and the definitions of w^0 and W^0 , it follows that $\mathcal{P}_{w^0}(s \models \Box W^0 \wedge \neg \Box \Diamond F) = 0$, which gives the following lemma:

Lemma 10. *For any $s \in W^0$, $\mathcal{P}_{w^0}(s \models \Box \Diamond F) = 1$.*

Determinacy and Solvability. Memoryless determinacy of almost-sure Büchi-GPLCS follows from Lemmas 7 and 10. By Lemma 6, for any state $s \in S$, we can check whether $s \in W^0$ or $s \in W^1$. This gives the main result:

Theorem 1. *Büchi-GPLCS are memoryless determined and solvable, for any regular target set F .*

6 Conclusions and Future Work

We have introduced GPLCS and given a terminating algorithm to compute symbolic representations of the winning sets in almost-sure Büchi-GPLCS. The strategies are memoryless, and our construction implies that the games we consider are memoryless determined.

The problem of deciding GPLCS games is not primitive recursive, since it is harder than the control-state reachability problem for LCS, which was shown to be non-primitive recursive by Schnoebelen in [25]. (For a given LCS and control-state q we can construct a GPLCS by defining $S^0 = \emptyset$, $S^1 = S$, making the state q absorbing and defining F as all configurations where the control-state is

not q . Then player 1 has a winning strategy in the GPLCS iff control-state q is reachable in the LCS.)

There are five immediate extensions of our result. (1) Each winning strategy w^0, w^1 wins against any *mixed* strategy of the opponent (i.e. the opponent chooses a probability distribution over the successor states rather than one of them). (2) Our algorithm is easily adapted to almost-sure *reachability*-GPLCS. This is achieved by replacing all outgoing transitions of states in F by self-loops, or, equivalently, replacing the definition of X_{i+1} by $X_{i+1} := \text{Force}^1(\overline{F}, \overline{M_{i+1}})$. (3) Our algorithm can be modified to construct symbolic representations of the winning strategies. A strategy is represented as a finite set $\{L_i, L'_i\}_{i=0}^n$ of pairs of regular state languages, where all L_i are disjoint. Such a finite set represents the strategy f^0 where $f^0(s) = \text{select}(L'_i)$ if $s \in L_i$. (4) We can extend the scheme to *concurrent* games, where the two players move simultaneously, by an appropriate extension of the Pre operator, as in [15]. (5) The algorithm also works when there are *probabilistic* control states in the GPLCS (see, e.g., [1] for definitions), as well as control states owned by the players and probabilistic message losses.

References

1. Abdulla, P.A., Baier, C., Iyer, P., Jonsson, B.: Reasoning about probabilistic lossy channel systems. In: Palamidessi, C. (ed.) CONCUR 2000. LNCS, vol. 1877, pp. 320–333. Springer, Heidelberg (2000)
2. Abdulla, P.A., Ben Henda, N., Mayr, R., Sandberg, S., de Alfaro, L.: Stochastic games with lossy channels. Technical Report 2007-005, Dept. of Information Technology, Uppsala University, Sweden (February 2007)
3. Abdulla, P.A., Bouajjani, A., d’Orso, J.: Deciding monotonic games. In: Baaz, M., Makowsky, J.A. (eds.) CSL 2003. LNCS, vol. 2803, pp. 1–14. Springer, Heidelberg (2003)
4. Abdulla, P.A., Henda, N.B., Mayr, R.: Verifying infinite Markov chains with a finite attractor or the global coarseness property. In: Proc. LICS 2005, 21st IEEE Int. Symp. on Logic in Computer Science, pp. 127–136 (2005)
5. Abdulla, P.A., Henda, N.B., Mayr, R., Sandberg, S.: Eager Markov chains. In: Graf, S., Zhang, W. (eds.) ATVA 2006. LNCS, vol. 4218, pp. 24–38. Springer, Heidelberg (2006)
6. Abdulla, P.A., Jonsson, B.: Verifying programs with unreliable channels. Information and Computation 127(2), 91–101 (1996)
7. Abdulla, P.A., Rabinovich, A.: Verification of probabilistic systems with faulty communication. In: Gordon, A.D. (ed.) ETAPS 2003 and FOSSACS 2003. LNCS, vol. 2620, pp. 39–53. Springer, Heidelberg (2003)
8. Asarin, E., Collins, P.: Noisy Turing machines. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 1031–1042. Springer, Heidelberg (2005)
9. Baier, C., Bertrand, N., Schnoebelen, P.: On computing fixpoints in well-structured regular model checking, with applications to lossy channel systems. In: Hermann, M., Voronkov, A. (eds.) LPAR 2006. LNCS (LNAI), vol. 4246, pp. 347–361. Springer, Heidelberg (2006)
10. Baier, C., Bertrand, N., Schnoebelen, P.: Verifying nondeterministic probabilistic channel systems against ω -regular linear-time properties. ACM Transactions on Comp. Logic (to appear, 2006)

11. Bertrand, N., Schnoebelen, P.: Model checking lossy channels systems is probably decidable. In: Gordon, A.D. (ed.) ETAPS 2003 and FOSSACS 2003. LNCS, vol. 2620, pp. 120–135. Springer, Heidelberg (2003)
12. Chatterjee, K., Jurdziński, M., Henzinger, T.: Simple stochastic parity games. In: Baaz, M., Makowsky, J.A. (eds.) CSL 2003. LNCS, vol. 2803, pp. 100–113. Springer, Heidelberg (2003)
13. Clarke, E., Grumberg, O., Peled, D.: Model Checking. MIT Press, Cambridge (1999)
14. Condon, A.: The complexity of stochastic games. *Information and Computation* 96(2), 203–224 (1992)
15. de Alfaro, L., Henzinger, T.: Concurrent omega-regular games. In: Proc. LICS 2000, 16th IEEE Int. Symp. on Logic in Computer Science, pp. 141–156. IEEE Computer Society Press, Los Alamitos (2000)
16. de Alfaro, L., Henzinger, T., Kupferman, O.: Concurrent reachability games. In: Proc. 39th Annual Symp. Foundations of Computer Science, pp. 564–575. IEEE Computer Society Press, Los Alamitos (1998)
17. de Alfaro, L., Henzinger, T., Majumdar, R.: Symbolic algorithms for infinite-state games. In: Larsen, K.G., Nielsen, M. (eds.) CONCUR 2001. LNCS, vol. 2154, pp. 536–550. Springer, Heidelberg (2001)
18. Esparza, J., Kučera, A., Mayr, R.: Model checking probabilistic pushdown automata. In: Proc. LICS 2004, 20th IEEE Int. Symp. on Logic in Computer Science, pp. 12–21 (2004)
19. Etessami, K., Yannakakis, M.: Recursive Markov decision processes and recursive stochastic games. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 891–903. Springer, Heidelberg (2005)
20. Finkel, A.: Decidability of the termination problem for completely specified protocols. *Distributed Computing* 7(3), 129–135 (1994)
21. Higman, G.: Ordering by divisibility in abstract algebras. *Proc. London Math. Soc.* (3) 2(7), 326–336 (1952)
22. Iyer, P., Narasimha, M.: Probabilistic lossy channel systems. In: Bidoit, M., Dauchet, M. (eds.) CAAP 1997, FASE 1997, and TAPSOFT 1997. LNCS, vol. 1214, pp. 667–681. Springer, Heidelberg (1997)
23. Kemeny, J., Snell, J., Knapp, A.: Denumerable Markov Chains. D Van Nostad Co. (1966)
24. Rabinovich, A.: Quantitative analysis of probabilistic lossy channel systems. In: Baeten, J.C.M., Lenstra, J.K., Parrow, J., Woeginger, G.J. (eds.) ICALP 2003. LNCS, vol. 2719, pp. 1008–1021. Springer, Heidelberg (2003)
25. Schnoebelen, P.: Verifying lossy channel systems has nonprimitive recursive complexity. *Information Processing Letters* 83(5), 251–261 (2002)
26. Shapley, L.S.: Stochastic games. *Proceedings of the National Academy of Sciences* 39(10), 1095–1100 (1953)
27. Vardi, M.: Automatic verification of probabilistic concurrent finite-state programs. In: Proc. FOCS 1985, 26th Annual Symp. Foundations of Computer Science, pp. 327–338 (1985)
28. Zielonka, W.: Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science* 200, 135–183 (1998)