

Probabilistic Modal μ -Calculus with Independent Product

Matteo Mio

LFCS, School of Informatics, University of Edinburgh

Abstract. The probabilistic modal μ -calculus $\text{pL}\mu$ (often called the quantitative μ -calculus) is a generalization of the standard modal μ -calculus designed for expressing properties of probabilistic labeled transition systems. The syntax of $\text{pL}\mu$ formulas coincides with that of the standard modal μ -calculus. Two equivalent semantics have been studied for $\text{pL}\mu$, both assigning to each process-state p a value in $[0, 1]$ representing the probability that the property expressed by the formula will hold in p : a *denotational semantics* and a *game semantics* given by means of two player stochastic games. In this paper we extend the logic $\text{pL}\mu$ with a second conjunction called *product*, whose semantics interprets the two conjuncts as probabilistically independent events. This extension allows one to encode useful operators, such as the modalities *with probability one* and *with non-zero probability*. We provide two semantics for this extended logic: one denotational and one based on a new class of games which we call *tree games*. The main result is the equivalence of the two semantics. The proof is carried out in ZFC set theory extended with Martin's Axiom at the first uncountable cardinal.

1 Introduction

The modal μ -calculus $\text{L}\mu$ [10] is a very expressive logic, for expressing properties of reactive systems (labeled transition systems), obtained by extending classical propositional modal logic with least and greatest fixed point operators. In the last decade, a lot of research has focused on the study of reactive systems that exhibit some kind of probabilistic behavior, and logics for expressing their properties [14,5,9,6,8]. Probabilistic labeled transition systems (PLTS's) [15] are a natural generalization of standard LTS's to the probabilistic setting, as they allow both (countable) non-deterministic and probabilistic choices.

The probabilistic modal μ -calculus $\text{pL}\mu$, introduced in [14,9,5], is a generalization of $\text{L}\mu$ designed for expressing properties of PLTS's. This logic was originally named the *quantitative* μ -calculus, but since other μ -calculus-like logics, designed for expressing properties of non-probabilistic systems, have been given the same name (e.g. [7]), we adopt the *probabilistic* adjective. The syntax of the logic $\text{pL}\mu$ coincides with that of the standard μ -calculus. The denotational semantics of $\text{pL}\mu$ [14,5] generalizes that of $\text{L}\mu$, by interpreting every formula F as a map $\llbracket F \rrbracket : P \rightarrow [0, 1]$, which assigns to each process p a *degree of truth*. In [12], the authors introduce an alternative semantics for the logic $\text{pL}\mu$. This semantics,

given in term of two player stochastic (parity) games, is a natural generalization of the two player (non-stochastic) game semantics for the logic $L\mu$ [16]. As in $L\mu$ games, the two players play a game starting from a configuration $\langle p, F \rangle$, where the objective for Player 1 is to produce a path of configurations along which the outermost fixed point variable X unfolded infinitely often is bound by a greatest fixed point in F . On a configuration of the form $\langle p, G_1 \vee G_2 \rangle$, Player 1 chooses one of the disjuncts G_i , $i \in \{1, 2\}$, by moving to the next configuration $\langle p, G_i \rangle$. On a configuration $\langle p, G_1 \wedge G_2 \rangle$, Player 2 chooses a conjunct G_i and moves to $\langle p, G_i \rangle$. On a configuration $\langle p, \mu X.G \rangle$ or $\langle p, \nu X.G \rangle$ the game evolves to the configuration $\langle p, G \rangle$, after which, from any subsequent configuration $\langle q, X \rangle$ the game again evolves to $\langle q, G \rangle$. On configurations $\langle p, \langle a \rangle G \rangle$ and $\langle p, [a] G \rangle$, Player 1 and 2 respectively choose a transition $p \xrightarrow{a} d$ in the PLTS and move the game to $\langle d, G \rangle$. Here d is a probability distribution (this is the key difference between $pL\mu$ and $L\mu$ games) and the configuration $\langle d, G \rangle$ belongs to Nature, the probabilistic agent of the game, who moves on to the next configuration $\langle q, G \rangle$ with probability $d(q)$. This game semantics allows one to interpret formulae as expressing, for each process p , the (limit) probability of a *property*, specified by the formula, holding at the state p . In [12], the equivalence of the denotational and game semantics for $pL\mu$ on finite models, was proven. The result was recently extended to arbitrary models by the present author [13].

In this paper we consider an extension of the logic $pL\mu$ obtained by adding to the syntax of the logic a second conjunction operator (\cdot) called *product* and its De Morgan dual operator called *coproduct* (\odot) . We call this extension the *probabilistic modal μ -calculus with independent product*, or just $pL\mu^\odot$. The denotational semantics of the product operator is defined as $\llbracket F \cdot G \rrbracket(p) = \llbracket F \rrbracket(p) \cdot \llbracket G \rrbracket(p)$, where the product symbol in the right hand side is multiplication on reals. Such an operator was already considered in [9] as a possible generalization of standard boolean conjunction to the lattice $[0, 1]$. Our logic $pL\mu^\odot$ is novel in containing both ordinary conjunctions and disjunctions (\wedge and \vee) and independent products and coproducts (\cdot and \odot). While giving a denotational semantics to $pL\mu^\odot$ is straightforward, the major task we undertake in this paper is to extend the game semantics of [12] to the new connectives. The game semantics implements the intuition that $H_1 \cdot H_2$ expresses the probability that H_1 and H_2 both hold if verified independently of each other.

To capture formally this intuition we introduce a game semantics for the logic $pL\mu^\odot$ in which independent execution of many instances of the game is allowed. Our games build on those for $pL\mu$ outlined above. Novelty arises in the game interpretation of the game-states $\langle p, H_1 \cdot H_2 \rangle$ and $\langle p, H_1 \odot H_2 \rangle$: when during the execution of the game one of these kinds of nodes is reached, the game is split into two concurrent and independent sub-games continuing their executions from the states $\langle p, H_1 \rangle$ and $\langle p, H_2 \rangle$ respectively. The difference between the game-interpretation of product and coproduct operators is that on a product configuration $\langle p, H_1 \cdot H_2 \rangle$, Player 1 has to win in both generated sub-games, while on a coproduct configuration $\langle p, H_1 \odot H_2 \rangle$ Player 1 needs to win just one of the two generated sub-games.




Fig. 1. Illustrative example

To illustrate the main ideas, let us consider the PLTS of figure 1(a) and the $pL\mu$ formula $F = \langle a \rangle \langle a \rangle \#$ which asserts the possibility of performing two consecutive a -steps. The probability of F being satisfied at p is $\frac{1}{2}$, since after the first a -step, the process $\mathbf{0}$ is reached with probability $\frac{1}{2}$ and no further a -step is possible. Let us consider the $pL\mu^\odot$ formula $H = \mu X.F \odot X$. Figure 1(b) depicts a play in the game starting from the configuration $\langle p, H \rangle$ (fixed-point unfolding steps are omitted). The branching points represent places where coproduct is the main connective, and each T_i represents play in one of the independent subgames for $\langle p, F \rangle$ thereupon generated. We call such a tree, describing play on all independent subgames, a *branching play*. Since all branches are coproducts, and the fixpoint is a least fixpoint, the objective for Player 1 is to win at least one of the games T_i . Since the probability of winning a particular game T_i is $\frac{1}{2}$, and there are infinitely many independent such games, almost surely Player 1 will win one of them. Therefore the game semantics assigns H at p the value 1.

The above example illustrates an interesting application of the new operators, namely the possibility of encoding the *qualitative* probabilistic modalities $\mathbb{P}_{>0}F$ (F holds with *probability greater than zero*) and $\mathbb{P}_{=1}F$ (F holds with *probability one*), which are equivalent to the $pL\mu^\odot$ formulae $\mu X.F \odot X$ and $\nu X.F \cdot X$ respectively. These encodings, which are easily seen to be correct denotationally, provide a novel game interpretation for the qualitative probabilistic modalities, which makes essential use of the new branching features of $pL\mu^\odot$ games (giving a direct game interpretation to the qualitative modalities seems no easier than giving the game semantics for the whole of $pL\mu^\odot$.) Moreover, they show that the interpretation of $pL\mu^\odot$ formulae is, in general, not continuous on the free variables: $\mathbb{P}_{>0}Y$ is an example of $pL\mu^\odot$ formula discontinuous on Y , since $\llbracket \mathbb{P}_{>0}Y \rrbracket_\rho(p) = 1$ if $\rho(Y)(p) > 0$ and $\llbracket \mathbb{P}_{>0}Y \rrbracket_\rho(p) = 0$ otherwise, where ρ interprets every variable Y as a map from process-states to $[0, 1]$. Other useful properties can be expressed by using these probabilistic modalities in the scope of fixed point operators. Some interesting formulae include $\mu X.(\langle a \rangle X \vee (\mathbb{P}_{=1}H))$, $\nu X.(\mathbb{P}_{>0}\langle a \rangle X)$ and $\mathbb{P}_{>0}(\nu X.\langle a \rangle X)$: the first assigns to a process p the probability of eventually reaching, by means of a sequence of a -steps, a state in which

H holds with probability one; the second, interpreted on a process p , has value 1 if there exists an infinite sequence of possible (in the sense of having probability greater than 0) a -steps starting from p , and 0 otherwise; the third formula, express a stronger property, namely it assigns to a process p value 1 if the probability of making (starting from p) an infinite sequence of a -steps is greater than 0, and value 0 otherwise. Moreover, every property expressible in the qualitative fragment of PCTL [8] can be formulated as a $\text{pL}\mu^\odot$ formula.

Formalizing the $\text{pL}\mu^\odot$ games outlined above is a surprisingly technical undertaking. To account for the *branching plays* that arise, we introduce a general notion of *tree game* which is of interest in its own right. Tree games generalize 2-player stochastic games, and are powerful enough to encode certain classes of games of imperfect information such as Blackwell games [2]. Tree games can also be used to formulate other notions that appear in the literature on stochastic games such as *qualitative determinacy* [3,1] and branching-time winning objectives [4]. This, as well as the encoding of qualitative PCTL mentioned above, will appear in the author's forthcoming PhD thesis.

A further level of difficulty arises in expressing when a branching play in a $\text{pL}\mu^\odot$ game is considered an objective for Player 1. This is delicate because branching plays can contain infinitely many interleaved occurrences of product and coproduct operations (so our simple explanation of such nodes above does not suffice). To account for this, branching plays are themselves considered as ordinary 2-player (parity) games with coproduct nodes as Player 1 nodes, and product nodes as Player 2 nodes. Player 1's goal in the *outer* $\text{pL}\mu^\odot$ game is to produce a branching play for which, when itself considered as a game, the *inner* game, he has a winning strategy. To formalize the class of tree games whose objective is specified by means of *inner* games, we introduce the notion of *2-player stochastic meta-game*.

Our main technical result is the equivalence of the denotational semantics and the game semantics for the logic $\text{pL}\mu^\odot$. As in [13] the proof of equivalence of the two semantics is based on the *unfolding method* of [7]. However there are significant complications, notably, the transfinite inductive characterization of the set of winning branching plays in a given $\text{pL}\mu^\odot$ -game (section 6) and the lack of denotational continuity on the free variables taken care by the game-theoretic notion of *robust* Markov branching play (section 7). Moreover, because of the complexity of the objectives described by means of *inner games*, the proof is carried out in ZFC set theory extended with MA_{\aleph_1} (Martin's Axiom at \aleph_1) which is known to be consistent with ZFC. We leave open the question of whether our result is provable in ZFC alone; we do not know if this is possible even restricting the equivalence problem to finite models.

The rest of the paper is organized as follows: in section 2, we fix some terminology and discuss the property MA_{\aleph_1} . In section 3, we define the syntax and the denotational semantics of the logic $\text{pL}\mu^\odot$. In section 4, the class of stochastic tree games, and its sub-class given by two player stochastic meta-parity games, are introduced in detail. In section 5, the game semantics of $\text{pL}\mu^\odot$ is defined in terms of two player stochastic meta-parity games. In section 6, we provide

a transfinite inductive characterization of the objective of the game associated with a formula $\mu X.F$. In section 7, we sketch the main ideas of the proof of the main theorem, which states the equivalence of the two semantics.

2 Background Definitions and Notation

Definition 1 (Probability distributions). A probability distribution d over a set X is a function $d: X \rightarrow [0, 1]$ such that $\sum_{x \in X} d(x) = 1$. The *support* of d , denoted by $\text{supp}(d)$ is defined as the set $\{x \in X \mid d(x) > 0\}$. We denote with $\mathcal{D}(X)$ the set of probability distributions over X .

Definition 2 (PLTS [15]). Given a countable set L of labels, a *Probabilistic Labeled Transition System* is a pair $\langle P, \{\xrightarrow{a}\}_{a \in L} \rangle$, where P is a countable set of states and $\xrightarrow{a} \subseteq P \times \mathcal{D}(P)$ for every $a \in L$. In this paper we restrict our attention to those PLTS such that for every $p \in P$ and every $a \in L$, the set $\{d \mid p \xrightarrow{a} d\}$ is countable. We refer to the countable set $\bigcup_{a \in L} \bigcup_{p \in P} \{d \mid p \xrightarrow{a} d\}$ as the set of probability distributions of the PLTS.

Definition 3 (Lattice operations). Given a set X , we denote with 2^X the set of all subsets $Y \subseteq X$. Given a complete lattice $(X, <)$, we denote with $\sqcup: 2^X \rightarrow X$ and $\sqcap: 2^X \rightarrow X$ the operations of join and meet respectively.

In the following we assume standard notions of basic topology and basic measure theory. The topological spaces we consider will always be 0-dimensional Polish spaces. We specify a probability measure on such a space by assigning compatible values in $[0, 1]$ to basic clopen sets. Such an assignment extends, using standard technology, to a probability measure μ on Borel sets, whence to a complete probability measure, again called μ , on all μ -measurable sets.

Martin’s Axiom (MA), from set theory, states that, for every infinite cardinal $\kappa < 2^\omega$, a certain property MA_κ holds. In this paper we use the property MA_{\aleph_1} as an axiom. This is implied by $\text{MA} + \neg\text{CH}$ (where CH is the *Continuum Hypothesis*), itself implies $\neg\text{CH}$, and is relatively consistent with ZFC set theory. Rather than explaining MA_{\aleph_1} in detail, we instead list the consequences of it that we need. Let μ be a σ -finite Borel measure on a Polish space X , and let Ω be the collection of μ -measurable sets; then every Σ^1_2 subset of X is in Ω and for every $\{X_\alpha\}_{\alpha < \omega_1}$ increasing \subseteq -chain of sets $X_\alpha \in \Omega$ indexed by the ordinals $\alpha < \omega_1$ (where ω_1 is the first uncountable ordinal), the statements $\bigcup_{\alpha < \omega_1} X_\alpha \in \Omega$ (ω_1 -completeness) and $\mu(\bigcup_{\alpha < \omega_1} X_\alpha) = \bigsqcup_{\alpha < \omega_1} \mu(X_\alpha)$ (ω_1 -continuity) hold. We refer [11] for a detailed proof of equivalent properties, in the special case of the Lebesgue measure on reals. As asserted there, the proofs generalize to the measure spaces considered in this paper.

3 The Logic $\text{pL}\mu^\odot$

Given a set Var of propositional variables ranged over by the letters X, Y and Z , and a set of labels L ranged over by the letters a, b and c , the formulae of the logic $\text{pL}\mu^\odot$ are defined by the following grammar:

$$F, G ::= X \mid [a]F \mid \langle a \rangle F \mid F \wedge G \mid F \vee G \mid F \cdot G \mid F \odot G \mid \nu X.F \mid \mu X.F$$

which extends the syntax of the standard μ -calculus with a new kind of conjunction (\cdot) and disjunction (\odot) operators called *product* and *coproduct* respectively. As usual the operators $\nu X.F$ and $\mu X.F$ bind the variable X in F . A formula is *closed* if it has no *free* variables.

Given a PLTS $\langle P, \{\xrightarrow{a}\}_{a \in L} \rangle$ we denote with $(P \rightarrow [0, 1])$ and with $(\mathcal{D}(P) \rightarrow [0, 1])$ the complete lattices of functions from P and from $\mathcal{D}(P)$ respectively, to the real interval $[0, 1]$ with the pointwise order. Given a function $f : P \rightarrow [0, 1]$, we denote with $\bar{f} : \mathcal{D}(P) \rightarrow [0, 1]$ the lifted function defined as follows:

$$\bar{f} \stackrel{\text{def}}{=} \lambda d. \left(\sum_{p \in P} d(p) \cdot f(p) \right).$$

A function $\rho : \text{Var} \rightarrow (P \rightarrow [0, 1])$ is called an *interpretation* of the variables. Given a function $f : P \rightarrow [0, 1]$ we denote with $\rho[f/X]$ the interpretation that assigns f to the variable X , and $\rho(Y)$ to all other variables Y . The denotational semantics $\llbracket F \rrbracket_\rho : P \rightarrow [0, 1]$ of the $\text{pL}\mu^\odot$ formula F , under the interpretation ρ , is defined by structural induction on F as follows:

$$\begin{aligned} \llbracket X \rrbracket_\rho &= \rho(X) \\ \llbracket G \vee H \rrbracket_\rho &= \llbracket G \rrbracket_\rho \sqcup \llbracket H \rrbracket_\rho \\ \llbracket G \wedge H \rrbracket_\rho &= \llbracket G \rrbracket_\rho \sqcap \llbracket H \rrbracket_\rho \\ \llbracket G \odot H \rrbracket_\rho &= \lambda p. (\llbracket G \rrbracket_\rho(p) \odot \llbracket H \rrbracket_\rho(p)) \\ \llbracket G \cdot H \rrbracket_\rho &= \lambda p. (\llbracket G \rrbracket_\rho(p) \cdot \llbracket H \rrbracket_\rho(p)) \\ \llbracket \langle a \rangle G \rrbracket_\rho &= \lambda p. \left(\bigsqcup \{ \llbracket \overline{G} \rrbracket_\rho(d) \mid p \xrightarrow{a} d \} \right) \\ \llbracket [a] G \rrbracket_\rho &= \lambda p. \left(\bigsqcap \{ \llbracket \overline{G} \rrbracket_\rho(d) \mid p \xrightarrow{a} d \} \right) \\ \llbracket \mu X.G \rrbracket_\rho &= \text{least fixed point of } \lambda f. (\llbracket G \rrbracket_{\rho[f/X]}) \\ \llbracket \nu X.G \rrbracket_\rho &= \text{greatest fixed point of } \lambda f. (\llbracket G \rrbracket_{\rho[f/X]}) \end{aligned}$$

where the symbols \cdot and \odot in the definitions of $\llbracket G \cdot H \rrbracket_\rho$ and $\llbracket G \odot H \rrbracket_\rho$ are standard multiplication on reals and the function $x \odot y = x + y - xy$, which is the De Morgan dual of multiplication with respect to the negation $\neg x = 1 - x$. Since the interpretation assigned to every $\text{pL}\mu^\odot$ operator is monotone, the existence of the least and greatest fixed points is guaranteed by the Knaster-Tarski theorem. Moreover the least and the greatest fixed points can be computed inductively as follows: $\llbracket \mu X.G \rrbracket_\rho = \bigsqcup_\alpha \llbracket \mu X.G \rrbracket_\rho^\alpha$ and $\llbracket \nu X.G \rrbracket_\rho = \bigsqcap_\alpha \llbracket \nu X.G \rrbracket_\rho^\alpha$ where α, β are ordinals and $\llbracket \mu X.G \rrbracket_\rho^\alpha$ and $\llbracket \nu X.G \rrbracket_\rho^\alpha$ are defined as follows:

$$\llbracket \mu X.G \rrbracket_\rho^\alpha = \bigsqcup_{\beta < \alpha} \llbracket G \rrbracket_{\rho[\llbracket \mu X.G \rrbracket_\rho^\beta / X]} \quad \text{and} \quad \llbracket \nu X.G \rrbracket_\rho^\alpha = \bigsqcap_{\beta < \alpha} \llbracket G \rrbracket_{\rho[\llbracket \nu X.G \rrbracket_\rho^\beta / X]}$$

4 Stochastic Tree Games

In this (unavoidably long) section we introduce a new class of games which we call *stochastic two player tree games*, or just $2\frac{1}{2}$ -*player tree games*. Stochastic tree games generalizes standard two player stochastic games by allowing a new

class of *branching nodes* on which the execution of the game is split in independent concurrent sub-games. Formally, stochastic tree games are infinite duration games played by Player 1, Player 2 and a third probabilistic agent named *Nature*, on a Arena $\mathcal{A} = \langle (S, E), (S_1, S_2, S_N, B), \pi \rangle$, where (S, E) is a directed graph with countable set of vertices S and transition relation E , (S_1, S_2, S_N, B) is a partition of S and $\pi : S_N \rightarrow \mathcal{D}(S)$. The states in S_1, S_2, S_N and B are called *Player 1 states*, *Player 2 states*, *probabilistic states* and *branching states* respectively. We denote with $E(s)$, for $s \in S$, the set $\{s' \mid (s, s') \in E\}$. As a technical constraint, we require that $\text{supp}(\pi(s)) \subseteq E(s)$, for every $s \in S_N$.

Definition 4 (Paths in \mathcal{A}). We denote with \mathcal{P}^ω and $\mathcal{P}^{<\omega}$ the sets of infinite and finite paths in \mathcal{A} . Given a finite path $\mathbf{s} \in \mathcal{P}^{<\omega}$ we denote with $\text{last}(\mathbf{s})$ the last state $s \in S$ of \mathbf{s} . We denote with $\mathcal{P}_1^{<\omega}$ and $\mathcal{P}_2^{<\omega}$ the sets of finite paths \mathbf{s} such that $\text{last}(\mathbf{s}) \in S_1$ and $\text{last}(\mathbf{s}) \in S_2$ respectively. We write $\mathbf{s} \triangleleft \mathbf{s}'$, with $\mathbf{s}, \mathbf{s}' \in \mathcal{P}^{<\omega}$, if $\mathbf{s}' = \mathbf{s}.s$, for some $s \in S$, where as usual the *dot* symbol denotes the concatenation operator. We denote with \mathcal{P}^t the set of *terminated paths*, i.e. the set of paths \mathbf{s} such that $E(\text{last}(\mathbf{s})) = \emptyset$. We denote with \mathcal{P} the set $\mathcal{P}^\omega \cup \mathcal{P}^t$ and we refer to this set as the set of *completed paths* in \mathcal{A} . Given a finite path $\mathbf{s} \in \mathcal{P}^{<\omega}$, we denote with $O_{\mathbf{s}}$ the set of all completed paths having \mathbf{s} as prefix. We consider the standard topology on \mathcal{P} where the countable basis for the open sets is given by the clopen sets $O_{\mathbf{s}}$, for $\mathbf{s} \in \mathcal{P}^{<\omega}$.

Definition 5 (Tree in \mathcal{A}). A *tree* in the arena \mathcal{A} is a collection $C = \{\mathbf{s}_i\}_{i \in I}$ of finite paths $\mathbf{s}_i \in \mathcal{P}^{<\omega}$, such that

1. C is down-closed: if $\mathbf{s} \in C$ and \mathbf{s}' is a prefix of \mathbf{s} , then $\mathbf{s}' \in C$.
2. C has a root: there exists exactly one finite path $\{s\}$ of length one in C . The path $\{s\}$, denoted by $\text{root}(C)$, is called the root of the tree C .

We consider the nodes \mathbf{s} of C as labeled by the *last* function.

Definition 6 (Uniquely and fully branching nodes of a tree). A node \mathbf{s} in a tree T is said to be *uniquely branching* in T if either $E(\text{last}(\mathbf{s})) = \emptyset$ or \mathbf{s} has a unique successor in T . Similarly, \mathbf{s} is *fully branching* in T if, for every $s \in E(s)$, it holds that $\mathbf{s}.s \in T$.

An outcome of the game in \mathcal{A} , which we call a *branching play*, is a possibly infinite tree T in \mathcal{A} defined as follows:

Definition 7 (Branching play in \mathcal{A}). A *branching play* in the arena \mathcal{A} is a tree T in \mathcal{A} such that, for every node $\mathbf{s} \in T$ the following conditions holds:

1. If $\text{last}(\mathbf{s}) \in S_1 \cup S_2 \cup S_N$ then \mathbf{s} branches uniquely in T .
2. If $\text{last}(\mathbf{s}) \in B$ then \mathbf{s} branches fully in T .

We denote with \mathcal{T} the set of branching plays T in the arena \mathcal{A} .

A branching play T represents a possible execution of the game from the state s labeling the root of T . The nodes of T with more than one child are all labeled with a state $s \in B$ and are the branching points of the game; their children represent the independent instances of play generated at the branching point.

Definition 8 (Topology on \mathcal{T}). Given a finite tree F in \mathcal{A} , we say that F is a *branching-play prefix*, if there exists some $T \in \mathcal{T}$, such that $F \subseteq T$. Given a branching-play prefix F , we denote with $O_F \subseteq \mathcal{T}$ the set of all branching plays T , such that $F \subseteq T$. We fix the topology on \mathcal{T} , where the basis for the open sets is given by the clopen sets O_F , for every branching-play prefix F . It is routine to show that this is a 0-dimensional Polish space.

As usual when working with *stochastic* games, it is useful to look at the possible outcomes of a play up-to the behavior of Nature. In the context of standard two player stochastic games this amounts at considering Markov chains. In our setting the following definition of Markov branching play is natural:

Definition 9 (Markov branching play in \mathcal{A}). A *Markov branching play* in \mathcal{A} is a tree M in \mathcal{A} such that for every node $s \in M$, the following conditions holds:

1. If $last(s) \in S_1 \cup S_2$ then s branches uniquely in T .
2. If $last(s) \in S_N \cup B$ then s branches fully in T .

A Markov branching play, is similar to a branching play except that probabilistic choices of Nature have not been resolved.

Definition 10 (Probability measure \mathcal{M}_M). Every Markov branching play M determines a probability assignment \mathcal{M}_M to every basic clopen set $O_F \subseteq \mathcal{T}$, for F a branching-play prefix, defined as follows:

$$\mathcal{M}_M(O_F) \stackrel{\text{def}}{=} \begin{cases} \prod \{\pi(s, s') \mid s.s.s' \in F \wedge s \in S_N\} & \text{if } F \subseteq M \\ 0 & \text{otherwise} \end{cases}$$

It is the above definition that implements the *probabilistic independence* of the sub-branching plays that follows a branching node s . The assignment \mathcal{M}_M extends to a unique complete probability measure \mathcal{M}_M on the measurable space Ω_M of all \mathcal{M}_M -measurable sets. By MA_{\aleph_1} , the collection Ω_M is ω_1 -complete, and the probability measure \mathcal{M}_M is ω_1 -continuous.

Definition 11 (Measurable space of branching plays in \mathcal{A}). We define the measurable space (\mathcal{T}, Ω) of branching plays in \mathcal{A} taking $\Omega = \bigcap_M \Omega_M$, where M ranges over Markov branching plays in \mathcal{A} . We say that a set $X \subseteq \mathcal{T}$ is *measurable* if $X \in \Omega$. The fact that the σ -algebra Ω is closed under arbitrary ω_1 -unions follows from the remark above that Ω_M , for every M , is ω_1 -complete. Given any Markov branching play M in \mathcal{A} , the (ω_1 -continuous) probability measure \mathcal{M}_M , induced by M on the measurable space Ω_M restricts to a unique (ω_1 -continuous) probability measure on the smaller space Ω , which we denote again with \mathcal{M}_M .

Definition 12 (Two player stochastic tree game). A *two player stochastic tree game* (or a $2\frac{1}{2}$ -player tree game) is given by a pair $\langle \mathcal{A}, \Phi \rangle$, where \mathcal{A} is a stochastic tree game arena as described above, and $\Phi \subseteq \mathcal{T}$, which is the *objective* for Player 1, is a measurable set of branching plays in \mathcal{A} .

Definition 13 (Expected value of a Markov branching play). Let $\langle \mathcal{A}, \Phi \rangle$ be a $2\frac{1}{2}$ -player tree game, and M a Markov branching play in \mathcal{A} . We define the *expected value* of M as: $E(M) = \mathcal{M}(\Phi)$. The value $E(M)$ should be understood as the probability for Player 1 to win the play.

As usual in game theory, players' moves are determined by strategies.

Definition 14 (Deterministic strategies). An (*unbounded memory deterministic*) strategy σ_1 for Player 1 in \mathcal{A} is defined as a function $\sigma_1 : \mathcal{P}_1^{<\omega} \rightarrow S \cup \{\bullet\}$ such that $\sigma_1(\mathbf{s}) \in E(\text{last}(\mathbf{s}))$ if $E(\text{last}(\mathbf{s})) \neq \emptyset$ and $\sigma_1(\mathbf{s}) = \bullet$ otherwise. Similarly a strategy σ_2 for Player 2 is defined as a function $\sigma_2 : \mathcal{P}_2^{<\omega} \rightarrow S \cup \{\bullet\}$. A pair $\langle \sigma_1, \sigma_2 \rangle$ of strategies, one for each player, is called a *strategy profile* and determines the behaviors of both players.

Note that the above definition of strategy captures the intended behavior of the game; both players when acting on a given instance of the game, know all the history of the actions happened on that sub-game, but have no knowledge of the evolution of the other independent parallel sub-games.

Definition 15 ($M_{\sigma_1, \sigma_2}^{s_0}$). Given an initial state $s_0 \in S$ and a strategy profile $\langle \sigma_1, \sigma_2 \rangle$ a unique Markov branching play $M_{\sigma_1, \sigma_2}^{s_0}$ is determined:

1. the root of M is labeled with s_0 ,
2. For every $\mathbf{s} \in M_{\sigma_1, \sigma_2}^{s_0}$, if $\text{last}(\mathbf{s}) = s$ with $s \in S_1$ not a terminal state, then the unique child of \mathbf{s} in $M_{\sigma_1, \sigma_2}^{s_0}$ is $\mathbf{s} \cdot (\sigma_1(\mathbf{s}))$.
3. For every $\mathbf{s} \in M_{\sigma_1, \sigma_2}^{s_0}$, if $\text{last}(\mathbf{s}) = s$ with $s \in S_2$ not a terminal state, then the unique child of \mathbf{s} in $M_{\sigma_1, \sigma_2}^{s_0}$ is $\mathbf{s} \cdot (\sigma_2(\mathbf{s}))$.

Definition 16 (Upper and lower values of a $2\frac{1}{2}$ -player tree game). Let $G = \langle \mathcal{A}, \Phi \rangle$ be a $2\frac{1}{2}$ -player tree game. We define the lower and upper values of G on the state s , denoted by $Val_{\downarrow}^s(G)$ and $Val_{\uparrow}^s(G)$ respectively, as follows:

$$Val_{\downarrow}^s(G) = \bigsqcup_{\sigma_1} \prod_{\sigma_2} E(M_{\sigma_1, \sigma_2}^s) \qquad Val_{\uparrow}^s(G) = \prod_{\sigma_2} \bigsqcup_{\sigma_1} E(M_{\sigma_1, \sigma_2}^s)$$

$Val_{\downarrow}^s(G)$, represents the limit probability of Player 1 winning, when the game begins in s , by choosing his strategy σ_1 first and then letting Player 2 pick an appropriate counter strategy σ_2 . Similarly $Val_{\uparrow}^s(G)$ represents the limit probability of Player 1 winning, when the game begins in s , by first letting Player 2 choose a strategy σ_2 and then picking an appropriate counter strategy σ_1 . In case $Val_{\downarrow}^s(G) = Val_{\uparrow}^s(G)$, we say that the game G at s is *determined*.

Definition 17 (ϵ -optimal strategies). Let $G = \langle \mathcal{A}, \Phi \rangle$ be a $2\frac{1}{2}$ -player tree game. We say that a strategy σ_1 for Player 1 in G is ϵ -optimal, for $\epsilon > 0$, if for every state s , the following inequality holds: $\prod_{\sigma_2} E(M_{\sigma_1, \sigma_2}^s) > Val_{\downarrow}^s(G) - \epsilon$. Similarly we say that a strategy σ_2 for Player 2 in G is ϵ -optimal, for $\epsilon > 0$, if for every state s , the following inequality holds: $\bigsqcup_{\sigma_1} E(M_{\sigma_1, \sigma_2}^s) < Val_{\uparrow}^s(G) + \epsilon$. Clearly ϵ -optimal strategies for Player 1 and Player 2 always exist for every $\epsilon > 0$.

An interesting class of $2\frac{1}{2}$ -player tree games is given by what we call meta-games. A *meta-game* is a $2\frac{1}{2}$ -player tree game G , which we refer to as the *outer* game, in which branching plays are themselves interpreted as (ordinary) two player games and the objective Φ of G is defined as the set of branching plays T in which this *inner game* is winnable for a given player taking part in it. We illustrate this notion by formalizing the class of $2\frac{1}{2}$ -player meta-parity games. A $2\frac{1}{2}$ -player meta-parity game G is specified by a $2\frac{1}{2}$ -player tree game arena $\mathcal{A} = \langle (S, E), \{S_1, S_2, S_N, B\}, \pi \rangle$ and a parity structure \mathbb{P} which is a pair $\langle Pr, Pl \rangle$, where $Pr: S \rightarrow \{0, \dots, n\}$, for some $n \in \mathbb{N}$, and $Pl: B \rightarrow \{1, 2\}$. The function Pr assigns a *priority* to each state $s \in S$. This is needed to define the set $W \subseteq \mathcal{P}$, of completed paths \mathbf{s} such that

- if \mathbf{s} is finite, then $Pr(\text{last}(\mathbf{s}))$ is even,
- if \mathbf{s} is infinite, then the least priority among those that appear (assigned by Pr to the states of \mathbf{s}) infinitely often in \mathbf{s} is even.

The function Pl assigns a *player identifier* to each state $s \in B$. This allows to consider each branching play T in \mathcal{A} as the game G_T played by Player 1 and Player 2 on the tree T : Player 1 moves on a node \mathbf{s} of T , such that $Pl(\text{last}(\mathbf{s})) = 1$, by choosing a successor in the (possibly empty) set of children of \mathbf{s} . Similarly Player 2 moves on the node \mathbf{s} of T such that $Pl(\text{last}(\mathbf{s})) = 2$. The result of a play in the game G_T is a completed path in T . We say that Player 1 wins a play if the resulting path is in W ; Player 2 wins otherwise. Since W is a parity objective we have that G_T is a parity game. The meta-parity game G can therefore be defined formally as a stochastic two player tree game, as follows:

Definition 18 (Two player stochastic meta-parity game). A *two player stochastic meta-parity game* specified by the pair $\langle \mathcal{A}, \mathbb{P} \rangle$, is formally defined as the $2\frac{1}{2}$ -player tree game $\langle \mathcal{A}, \Phi \rangle$ where Φ is defined as follows:

$$\Phi = \{T \mid T \in \mathcal{T} \text{ and Player 1 has a winning strategy in } G_T\}$$

The definition is good because the measurability of Φ follows from $\text{MA}_{\mathbb{N}_1}$ and the following Lemma.

Lemma 1. *The set Φ is a Δ^1_2 set and hence a Σ^1_2 set in \mathcal{T} .*

5 Game Semantics of $\text{pL}\mu^{\odot}$

In this section we describe the $\text{pL}\mu^{\odot}$ game G^F_ρ associated to each triple consisting of a PLTS $\langle P, \{\xrightarrow{a}\}_{a \in L} \rangle$, a (possibly open) $\text{pL}\mu^{\odot}$ formula F , and an interpretation of the variables ρ . For convenience, we assume that F is *normal* [16], i.e., every occurrence of a μ or ν binder binds a distinct variable, and no variable appears both free and bound. The game G^F_ρ is a $2\frac{1}{2}$ -player meta-parity game specified by the arena $\mathcal{A}^F_\rho = \langle (S, E), (S_1, S_2, S_N, B), \pi \rangle$ and parity structure $\mathbb{P} = \langle Pr, Pl \rangle$ defined as follows. The countable set S of vertices of the directed graph (S, E) is given by the set.

$$S = (P \times \text{Sub}(F)) \cup (D \times \text{Sub}(F)) \cup \{\perp, \top\}$$

where P is the set of processes, $Sub(F)$ is the set of sub-formulae of F (defined as usual, e.g. [16]), D is the set of distributions in the PLTS (see definition 2) and $\{\perp, \top\}$ are two special states. The relation E is defined as follows: $E(\langle d, G \rangle) = \{\langle p, G \rangle \mid p \in supp(d)\}$ for every $d \in D$; $E(\langle p, G \rangle)$ is defined by case analysis on the outermost connective of G as follows:

1. if $G = X$, with X free in F , then $E(\langle p, G \rangle) = \{\perp, \top\}$.
2. if $G = X$, with X bound in F by the subformula $\star X.H$, with $\star \in \{\mu, \nu\}$, then $E(\langle p, G \rangle) = \{\langle p, H \rangle\}$.
3. if $G = \star X.H$, with $\star \in \{\mu, \nu\}$, then $E(\langle p, G \rangle) = \{\langle p, H \rangle\}$.
4. if $G = \langle a \rangle H$ then $E(\langle p, G \rangle) = \{\langle d, H \rangle \mid p \xrightarrow{a} d\}$.
5. if $G = [a] H$ then $E(\langle p, G \rangle) = \{\langle d, H \rangle \mid p \xrightarrow{a} d\}$.
6. if $G = H * H'$ with $*$ in $\{\vee, \wedge, \odot, \cdot\}$ then $E(\langle p, G \rangle) = \{\langle p, H \rangle, \langle p, H' \rangle\}$

The relation E is defined on the two special game states \top and \perp as $E(\top) = E(\perp) = \emptyset$. This makes \top and \perp *terminal* states of the game. The partition (S_1, S_2, S_N, B) of S is defined as follows: every state $\langle p, G \rangle$ with G 's main connective in $\{\langle a \rangle, \vee, \mu X\}$ or with $G = X$ where X is a μ -variable, is in S_1 ; dually every state $\langle p, G \rangle$ with G 's main connective in $\{[a], \wedge, \nu X\}$ or with $G = X$ where X is a ν -variable, is in S_2 . Every state of the form $\langle d, G \rangle$ or $\langle p, X \rangle$, with X free in F , is in S_N . Every state $\langle p, G \rangle$ whose G 's main connective is \cdot or \odot is in B . Lastly we define the terminal states \perp and \top to be in S_1 and S_2 respectively. The function $\pi : S_N \rightarrow \mathcal{D}(S)$ assigns a probability distribution to every state under the control of Nature (thus specifying its intended probabilistic behavior) and it is defined as $\pi(\langle d, G \rangle)(\langle p, G \rangle) = d(p)$ on all states of the form $\langle d, G \rangle$; all other states in S_N are of the form $\langle p, X \rangle$, with X free in F ; the function π is defined on these states as follows:

$$\pi(\langle p, X \rangle)(s) \stackrel{\text{def}}{=} \begin{cases} \rho(X)(p) & \text{if } s = \top \\ 1 - \rho(X)(p) & \text{if } s = \perp \\ 0 & \text{otherwise} \end{cases}$$

The priority assignment $Pr : S \rightarrow \{0, \dots, n\}$ is defined, by picking a sufficiently large n , as usual in μ -calculus games: an odd priority is assigned to the states $\langle p, X \rangle$ with X a μ -variable and dually an even priority is assigned to the states $\langle p, X \rangle$ with X a ν -variable, in such a way that if Z subsumes Y in F then $Pr(\langle p, Z \rangle) < Pr(\langle p, Y \rangle)$. Moreover, for every terminal state $s \in S$, we define $Pr(s) = 1$ if $s \in S_1$, and $Pr(s) = 0$ if $s \in S_2$. Lastly, the function $Pl : B \rightarrow \{1, 2\}$ is defined as $Pl(\langle p, G_1 \odot G_2 \rangle) = 1$ and $Pl(\langle p, G_1 \cdot G_2 \rangle) = 2$ for every $p \in P$ and $G_1, G_2 \in Sub(F)$.

Note that if no (co)product operators occur in F , then $B = \emptyset$, and the game is equivalent to the one in [12,13] for the logic $pL\mu$. We are now ready to state our main result.

Theorem 1 (MA_{N1}). *Given a PLTS $\langle P, \{\xrightarrow{a}\}_{a \in L} \rangle$, for every process $p \in P$, interpretation of the variables ρ and $pL\mu^{\odot}$ formula F , the equalities*

$$Val_{\downarrow}^{\langle p, F \rangle}(G_{\rho}^F) = Val_{\uparrow}^{\langle p, F \rangle}(G_{\rho}^F) = \llbracket F \rrbracket_{\rho}(p)$$

hold. In particular the game G_{ρ}^F is determined.

6 Inductive Characterization of the Objective of $G_\rho^{\mu X.F}$

In this section we provide a transfinite inductive characterization of the set $\Phi^{\mu X.G}$ of winning branching plays of the game $G_\rho^{\mu X.F}$ needed in the proof of Theorem 1.

1. Let us consider the game G_ρ^F , where X appears free in F . Note that the two arenas \mathcal{A}_ρ^F and $\mathcal{A}_\rho^{\mu X.F}$ are similar as they differ only in the following aspects:

1. The set of states $S^{\mu X.F}$ of $\mathcal{A}_\rho^{\mu X.F}$ is the same as the set of states S^F of \mathcal{A}_ρ^F , plus the set of states of the form $\langle p, \mu X.F \rangle$, which however play almost no role in the game because these nodes have only one successor ($\langle p, F \rangle$) and are not reachable by any other state.
2. More significantly, the nodes of the form $\langle p, X \rangle$, which are present in both game arenas, are Player 1 states in $G_\rho^{\mu X.F}$ (they have unique successor $\langle p, F \rangle$), and Nature states in G_ρ^F (they have two *terminal* successors \top and \perp reachable with probabilities $\rho(X)(p)$ and $1 - \rho(X)(p)$ respectively).

Moreover observe that the player assignments Pl^F and $Pl^{\mu X.F}$ are identical, and the priority assignments Pr^F and $Pr^{\mu X.F}$ differ only on the game-states s of the form $\langle p, X \rangle$: $Pr^F(s) = 0$ and $Pr^{\mu X.F}(s) = m$ for some odd $m \in \mathbb{N}$. A G_ρ^F branching play T , rooted in $\langle p, F \rangle$ can be depicted as in figure 2(a), where the triangle represents the set of paths in T never reaching a state of the form $\langle q, X \rangle$, for $q \in P$, and the other edges represents the, possibly empty, collection of paths $\{s_i\}_{i \in I \subseteq \mathbb{N}}$ reaching a state of the form $\langle p_i, X \rangle$ which is (necessarily) followed by a terminal state $b_i \in \{\top, \perp\}$. Similarly a branching play T in $G_\rho^{\mu X.F}$, rooted in $\langle p, F \rangle$, can be depicted as in figure 2(b). We extract the common part between the branching plays G_ρ^F and $G_\rho^{\mu X.F}$ by defining the notion of branching pre-play.

Definition 19 (Branching pre-play). Let T be a branching play in G_ρ^F and let I index the (necessarily countable) collection of nodes of the form $\langle p_i, X \rangle$ in T . The *branching pre-play* $T[x_i]_{i \in I}$, which can be depicted as in figure 2(c), is the tree obtained from the branching play T by removing all subtrees rooted in states of the form $\langle p_i, X \rangle$.

Given a I -indexed family $\{b_i\}_{i \in I}$, where $b_i \in \{\top, \perp\}$, we denote with $T[b_i]_{i \in I}$ the branching play in G_ρ^F obtained by adding, for every $i \in I$, the child b_i to the leaf $\langle p_i, X \rangle$ of $T[x_i]_{i \in I}$. Similarly given a family $\{T_i\}_{i \in I}$ of branching plays in $G^{\mu X.F}$, where each T_i is rooted at $\langle p_i, G \rangle$, we denote with $T[T_i]_{i \in I}$ the branching play in $G_\rho^{\mu X.F}$ obtained by adding the subtree T_i after the leaf $\langle p_i, X \rangle$. Clearly every branching play T rooted at $\langle p, F \rangle$ in G_ρ^F is uniquely of the form $T'[b_i]_{i \in I}$ for appropriate $T'[x_i]_{i \in I}$ and $\{b_i\}_{i \in I}$. Similarly every branching play T rooted at $\langle p, F \rangle$ in $G_\rho^{\mu X.F}$ is of the form $T'[T_i]_{i \in I}$ for appropriate $T'[x_i]_{i \in I}$ and $\{T_i\}_{i \in I}$.

Definition 20. The function $m_X : \mathcal{T}^{\mu X.F} \rightarrow \mathcal{T}^F$, from branching plays in $G_\rho^{\mu X.F}$ to branching plays in G_ρ^F , is defined for every subset $X \subseteq \mathcal{T}^{\mu X.F}$ as follows:

$$m_X(T[T_i]_{i \in I}) = T[T_i \in X]_{i \in I} \quad \text{where } T_i \in X \stackrel{\text{def}}{=} \begin{cases} \top & \text{if } T_i \in X \\ \perp & \text{otherwise} \end{cases}$$




Fig. 2. Branching plays and pre-plays

Lemma 2. *If $X \subseteq \mathcal{T}^{\mu X, F}$ is a measurable set, then m_X is a measurable map.*

We now define the operator R , of which $\Phi^{\mu X, G}$ is the least fixed point.

Definition 21. The operator $R: 2^{\mathcal{T}^{\mu X, F}} \rightarrow 2^{\mathcal{T}^{\mu X, F}}$ is defined as follows:

$$R(X) \stackrel{\text{def}}{=} m_X^{-1}(\Phi^F) = \{T[T_i]_{i \in I} \mid T[T_i \in X]_{i \in I} \in \Phi^F\}$$

Theorem 2. *The set $\Phi^{\mu X, F}$ is the least fixed point of the monotone operator R , which is guaranteed to exist by the Knaster-Tarski theorem. Hence the set $\Phi_\rho^{\mu X, F}$ can be defined as $\bigcup_\alpha R^\alpha$ where R^α is defined for every ordinal α as $\bigcup_{\beta < \alpha} R(R^\beta)$.*

The following lemma, which is essential for our proof of Theorem 1, states that the fixed points is reached in at most ω_1 steps.

Lemma 3. $\bigcup_\alpha R^\alpha = \bigcup_{\alpha < \omega_1} R^\alpha$.

7 Proof of Theorem 1

In this section we sketch the main ideas of the proof of Theorem 1. We first introduce a property that is going to be useful in proving the main result:

Definition 22 (Robust Markov branching plays). Fix a $\text{pL}\mu^\odot$ game G_ρ^F , a free variable X in F and an \mathbb{N} -indexed collection $\{\epsilon_n\}_{n \in \mathbb{N}}$ of reals in $(0, 1]$. Let M be a Markov branching play in G_ρ^F . Let $\{x_i\}_{i \in I \subseteq \mathbb{N}}$ be the set of vertices in M labeled with states of the form $\langle p, X \rangle$. Since X is free in F , these vertices are necessarily connected to the two leafs \perp and \top by two edges e_i^\top and e_i^\perp marked with the probabilities λ_i and $1 - \lambda_i$ respectively. Let M^+ be the same Markov branching play where, for each $i \in I$, the probability attached to the edge e_i^\top

is replaced with $\min\{1, \lambda_i + \epsilon_i\}$, and the probability attached to the edge e_i^\perp is replaced by $\max\{0, (1 - \lambda_i) - \epsilon_i\}$. Similarly let M^- be as M where, for all $i \in I$, the probabilities attached to the edge e_i^\top is replaced with $\max\{0, \lambda_i - \epsilon_i\}$, and the probability attached to the edge e_i^\perp is replaced by $\min\{1, (1 - \lambda_i) + \epsilon_i\}$. We say that M is *robust* if and only if both inequalities $E(M^+) \leq E(M) + \sum_{i \in I} \epsilon_i$ and $E(M^-) \geq E(M) - \sum_{i \in I} \epsilon_i$ hold, for every collection $\{\epsilon_n\}_{n \in \mathbb{N}}$.

The notion of robustness can be informally described as follows: in the Markov branching play M^+ we increase the probability associated with the branching plays having paths ending in \top immediately following a configuration $\langle p, X \rangle$; by doing so we increase the value, see Definition 13, of the original M and, in a similar way we decrease the value of M by moving to M^- . A Markov branching play is robust if small changes (either in the direction of M^+ or M^-) in the probabilities (labeling the edges e_i^\top and e_i^\perp , $i \in I$) produce bounded (by $\sum_{i \in I} \epsilon_i$) changes in the overall value of the Markov branching play. Note that altering the probabilities in M uniformly (i.e. taking $\{\epsilon_i\}_{i \in I}$ such that for every $i \neq j$, $\epsilon_i = \epsilon_j$), may produce, in general, unbounded changes in the value of M ; this reflect the discontinuity of the denotational interpretation of $\text{pL}\mu^\odot$ formulae on the free variables.

In order to prove Theorem 1, we prove the following stronger theorem:

Theorem 3 (MA $_{\aleph_1}$). *Given a PLTS $\langle P, \{\xrightarrow{a}\}_{a \in L} \rangle$, for every $\text{pL}\mu^\odot$ formula F and for every interpretation ρ , the following assertions hold for every $p \in P$:*

1. $\llbracket F \rrbracket_\rho(p) \geq \text{Val}_\uparrow^{\langle p, F \rangle}(G_\rho^F)$
2. $\llbracket F \rrbracket_\rho(p) \leq \text{Val}_\downarrow^{\langle p, F \rangle}(G_\rho^F)$
3. *Every Markov branching play M rooted in $\langle p, F \rangle$ in G_ρ^F is robust.*

The proof is by induction on the structure of F , and resembles the *unfold-ing method* of [7,13]. The most difficult case in proving point 1 is when F is of the form $\mu X.H$ (and dually the case $\nu X.H$ is difficult for point 2). This is proven showing that, for every $\epsilon > 0$, there exists a strategy σ_2^ϵ for Player 2 in $G_\rho^{\mu X.H}$, such that for every counter-strategy σ_1 for Player 1, the inequality $E(M_{\sigma_1, \sigma_2^\epsilon}^{\langle p, \mu X.H \rangle}) < \llbracket \mu X.H \rrbracket_\rho(p) + \epsilon$ holds. As in [7,13], the strategy σ_2^ϵ is constructed using δ -optimal strategies for Player 2 in the game $G_{\rho^\gamma}^H$ (where $\rho^\gamma = \llbracket \mu X.H \rrbracket_\rho$) which exist by induction hypothesis. The idea behind the construction of σ_2^ϵ is the following: initially the strategy σ_2^ϵ behaves as some δ_0 -optimal strategy τ_0 for Player 2 in $G_{\rho^\gamma}^H$; if at some point of the play the game reaches a configuration of the form $\langle p, X \rangle$, then Player 2 *improves* his play and, depending on the history of the previously played moves, starts behaving as some δ_1 -optimal strategy τ_1 for Player 2 in $G_{\rho^\gamma}^H$ and so on; the strictly decreasing sequence $\{\delta_i\}_{i \in \mathbb{N}}$ is carefully chosen, so that the desired ϵ -bound follows from the induction hypothesis of robustness for Markov branching plays of $G_{\rho^\gamma}^H$. The desired inequality $E(M_{\sigma_1, \sigma_2^\epsilon}^{\langle p, \mu X.H \rangle}) < \llbracket \mu X.H \rrbracket_\rho(p) + \epsilon$, thanks to Theorem 2, Lemma 3 and the fact that $\mathcal{M}_{\sigma_1, \sigma_2^\epsilon}^{\langle p, \mu X.H \rangle}$ is ω_1 -continuous under MA $_{\aleph_1}$, is equivalent to

$\bigsqcup_{\alpha < \omega_1} \mathcal{M}_{\sigma_1, \sigma_2}^{(p, \mu X, H)}(R^\alpha) < \llbracket \mu X.H \rrbracket_\rho(p) + \epsilon$. The proof that this last inequality holds, is by ordinal induction.

Acknowledgments. I thank my PhD supervisor, Alex Simpson, for helpful suggestions. This research was supported by studentships from LFCS and the IGS at the School of Informatics, Edinburgh, and by EPSRC research grant EP/F042043/1.

References

1. Bertrand, N., Genest, B., Gimbert, H.: Qualitative determinacy and decidability of stochastic games with signals. In: LICS 2009, IEEE Symposium on Logic in Computer Science, Los Angeles, USA (2009)
2. Blackwell, D.: Infinite $G\delta$ games with imperfect information. *Matematyki Applicationes Mathematicae, Hugo Steinhaus Jubilee Volume* (1969)
3. Brázdil, T., Brozek, V., in Kucera, A., Obdržálek, J.: Qualitative reachability in stochastic BPA games. In: STACS, pp. 207–218 (2009)
4. Brázdil, T., Brozek, V., Forejt, V., Kucera, A.: Stochastic games with branching-time winning objectives. In: LICS, pp. 349–358 (2006)
5. de Alfaro, L., Majumdar, R.: Quantitative solution of omega-regular games. *Journal of Computer and System Sciences* 68, 374–397 (2004)
6. Deng, Y., van Glabbeek, R.: Characterising probabilistic processes logically. In: Fermüller, C.G., Voronkov, A. (eds.) LPAR-17. LNCS, vol. 6397, pp. 278–293. Springer, Heidelberg (2010)
7. Fischer, D., Gradel, E., Kaiser, L.: Model checking games for the quantitative μ -calculus. In: *Theory of Computing Systems*. Springer, New York (2009)
8. Hansson, H., Jonsson, B.: A logic for reasoning about time and reliability. *Formal Aspects of Computing* 6 (1994)
9. Huth, M., Kwiatkowska, M.: Quantitative analysis and model checking. In: LICS 1997, Washington, DC, USA, p. 111. IEEE Computer Society, Los Alamitos (1997)
10. Kozen, D.: Results on the propositional mu-calculus. *Theoretical Computer Science*, 333–354 (1983)
11. Martin, D.A., Solovay, R.M.: Internal Cohen extensions. *Ann. Math. Logic* 2, 143–178 (1970)
12. McIver, A., Morgan, C.: Results on the quantitative μ -calculus $qM\mu$. *ACM Trans. Comput. Logic* 8(1), 3 (2007)
13. Mio, M.: The equivalence of denotational and game semantics for the probabilistic μ -calculus. In: 7th Workshop on Fixed Points in Computer Science (2010)
14. Morgan, C., McIver, A.: A probabilistic temporal calculus based on expectations. In: Groves, L., Reeves, S. (eds.) *Proc. Formal Methods*. Springer, Heidelberg (1997)
15. Segala, R.: Modeling and Verification of Randomized Distributed Real-Time Systems. PhD thesis, Laboratory for Computer Science, M.I.T. (1995)
16. Stirling, C.: Modal and temporal logics for processes. Springer, Heidelberg (2001)