

Christel Baier · Tomáš Brázdil ·
Marcus Größer · Antonín Kučera

Stochastic Game Logic

the date of receipt and acceptance should be inserted later

Abstract Stochastic game logic (SGL) is a new temporal logic for multi-agent systems modeled by turn-based multiplayer games with discrete transition probabilities. It combines features of alternating-time temporal logic (ATL), probabilistic computation tree logic (PCTL) and extended temporal logic (ETL). SGL contains an ATL-like modality to specify the individual cooperation and reaction facilities of agents in the multiplayer game to enforce a certain winning objective. While the standard ATL modality states the existence of a strategy for a certain coalition of agents without restricting the range of strategies for the semantics of inner SGL formulas, we deal with a more general modality. It also requires the existence of a strategy for some coalition, but imposes some kind of strategy binding to inner SGL formulas. The PCTL and ETL fragment of SGL serves to formalize qualitative and quantitative, linear or branching time winning objectives.

This paper presents the syntax and semantics of SGL and discusses its model checking problem for different types of strategies. The model checking problem of SGL turns out to be undecidable when dealing with the full class of history-dependent strategies. We show that the SGL model checking problem for memoryless deterministic strategies as well as the model checking problem of the qualitative fragment of SGL for memoryless randomized strategies is PSPACE-complete. Moreover, we establish a reduction from the SGL model checking problem under memoryless randomized strategies into the Tarski algebra which proves the problem to be in EXPSpace.

Christel Baier, Marcus Größer
Technische Universität Dresden, Faculty for Computer Science,
Institute for Theoretical Computer Science, 01062 Dresden, Germany.
E-mail: {baier,groesser}@tcs.inf.tu-dresden.de

Tomáš Brázdil, Antonín Kučera
Faculty of Informatics, Masaryk University, Botanická 68a, 60200 Brno, Czech Republic.
E-mail: {brazdil,kucera}@fi.muni.cz

1 Introduction

Traditional temporal logics, such as linear temporal logic (LTL) or computation tree logic (CTL) are widely used to specify properties of parallel systems [14]. In the classical approach for concurrent systems, the semantics of LTL and CTL relies on an operational model (transition system) where the paths represent the possible interleavings of the processes running in parallel and formulas assert conditions on all paths (LTL) or on the branching structure of states (CTL). This approach is adequate for closed systems where the transition system describes the potential system behaviors from a global perspective. In contrast, the semantics of open systems relies on a *game-based view* where the individual components are considered as *players* (also called *agents*) that interact with each other and their environment. Alur, Henzinger and Kuperman [2] introduced alternating-time temporal logic (ATL) as a variant of CTL with modalities expressing the existence of strategies for coalitions of agents that enforce a certain event. More precisely, ATL extends CTL by formulas of the form $\langle\langle A \rangle\rangle_{\text{ATL}} \varphi$ and $\|A\|_{\text{ATL}} \varphi$ where A is a set of cooperating players (agents) and φ an ATL path formula.¹ Intuitively, the formula $\langle\langle A \rangle\rangle_{\text{ATL}} \varphi$ asserts that there is a strategy for the agents in A such that the event specified by φ holds, no matter how the opponents (i.e., the other agents) behave. The dual operator, denoted by $\|\cdot\|_{\text{ATL}}$, can be understood as universal quantification over strategies. That is, formula $\|A\|_{\text{ATL}} \varphi$ states that φ holds along at least one path under all strategies for A . Stated differently, $\|A\|_{\text{ATL}} \varphi$ asserts the absence of a strategy for the coalition A to avoid φ to hold. The ATL semantics relies on the standard CTL-like approach where all subformulae are interpreted over the “full” structure. I.e., the ATL modalities $\langle\langle \cdot \rangle\rangle_{\text{ATL}}$ and $\|\cdot\|_{\text{ATL}}$ do not restrict the range of strategies for nested $\langle\langle \cdot \rangle\rangle_{\text{ATL}}$ and $\|\cdot\|_{\text{ATL}}$ modalities. For instance, the formula $\langle\langle A \rangle\rangle_{\text{ATL}} \Box \langle\langle B \rangle\rangle_{\text{ATL}} \Diamond p$ asserts the existence of a strategy α for the agents in A such that $\langle\langle B \rangle\rangle_{\text{ATL}} \Diamond p$ holds (in the “full” game) for all states s that can be reached when the agents in A make their decisions according to α , i.e., from these states s the agents in B have a strategy β in the original game (neglecting the strategy α) which ensures that a state where p holds is reached.² Thus in ATL a strategy chosen by the $\langle\langle \cdot \rangle\rangle_{\text{ATL}}$ operator is *not propagated* to the inner ATL state formulae. Therefore, properties stating that a certain agent can react on the choices made by another agent are not expressible in ATL. Several variants of ATL have been proposed that contain modalities for strategy quantification imposing certain bindings of strategies to inner formulas. Examples are game logic [2], ATL with strategy contexts [1, 10] or stochastic logic that contains first-order quantification over strategies [12].

In this paper, we introduce a probabilistic variant of ATL, called *stochastic game logic*, SGL for short. It serves to specify properties of multi-agent systems where actions can have a probabilistic effect. More precisely, we deal with finite-state turn-based games where the game arena is based on a graph structure with several annotations. Each state is either a game configuration where a single agent is declared to choose a successor state (based on some deterministic or random-

¹ Our logic SGL will use other modalities for existential and universal quantification over strategies. To avoid notation overloading, we use the notation $\langle\langle A \rangle\rangle_{\text{ATL}}$ and $\|A\|_{\text{ATL}}$ to denote the standard ATL modalities, while $\langle\langle A \rangle\rangle$ and $\|A\|$ will be used for the SGL-modalities. Later on, we will explain how the ATL-modalities can be derived from $\langle\langle A \rangle\rangle$ and $\|A\|$.

² \Box and \Diamond denote the “always” and “eventually” operator, respectively.

ized strategy) or a probabilistic state where the next state is chosen randomly according to some fixed probabilistic distribution. The logic SGL contains a variant of the ATL-modalities for reasoning about the existence or absence of strategies for coalitions of agents, denoted $\langle\langle A \rangle\rangle$ and $\|A\|$, where – as before – A denotes a coalition of agents, to achieve a certain objective. The difference between the ATL-modalities $\langle\langle A \rangle\rangle_{\text{ATL}}$ and $\|A\|_{\text{ATL}}$ and the SGL modalities $\langle\langle A \rangle\rangle$ and $\|A\|$ becomes apparent when these operators are *nested*. Following the approach of [4, 24, 9], the semantics of the SGL formula $\langle\langle A \rangle\rangle \Phi$ is defined differently. The operator $\langle\langle A \rangle\rangle$ imposes a binding of the strategy α chosen by the agents in A in the same way as first-order quantification $\exists x \phi$ binds the variable x . The scope of the binding is the full formula Φ including its subformulae. However, the nested $\langle\langle A' \rangle\rangle$ operators can revise the binding for the agents in $A \cap A'$.

In SGL, the objectives of coalitions of agents can be linear-time or branching-time properties with qualitative or quantitative probability bounds. For this purpose, the SGL syntax combines features of probabilistic computation tree logic (PCTL) for Markov decision processes [7] with features of extended temporal logic (ECTL) [33, 30, 13]. More precisely, the SGL-operators $\langle\langle \cdot \rangle\rangle$ and $\| \cdot \|$ can be used in combination with PCTL-like properties that might express qualitative or quantitative probability bounds on path-events, or Boolean combinations thereof. We follow here the concept of extended temporal logics especially the concept of ECTL [13] and use deterministic Rabin automata to describe path properties.

With this concept we can formalize typical multi-player game properties such as “*the agents in A have a strategy such that whatever strategy the agents in B choose, the agents in C can react to that strategy so that the winning condition holds*”. This is formalized by SGL formula

$$\langle\langle A \rangle\rangle \|B\| \langle\langle C \rangle\rangle \text{ “the winning condition holds”},$$

where $\|B\| \Phi = \neg \langle\langle B \rangle\rangle \neg \Phi$. This property might or might not be expressible in ATL, depending on the winning condition and whether the game is turn-based or concurrent. In general, the SGL formulae $\langle\langle A \rangle\rangle \|B\| \langle\langle C \rangle\rangle$ “*win.cond.*” and $\langle\langle A \cup C \rangle\rangle$ “*win.cond.*” and *not* equivalent, because C ’s strategies can depend on B ’s decisions.

For an example that illustrates the usefulness of the revision of a strategy chosen for a formula $\langle\langle A \rangle\rangle \Phi$ by another $\langle\langle A \rangle\rangle$ operator inside Φ , we consider the following scenario. A banker or money broker has a certain amount x (say 1 Mio Dollars) to work with. His/her goal is to design a strategy (of buying and selling stock, fixed-term deposit, subscription warrants, etc.) for the upcoming months that guarantees with a given probability (e.g. 90) his/her earnings to become larger than 100.000 Dollars in the next year. On the other hand, if everything goes haywire, the banker wants to be able to have at least 120.000 Dollars at his/her disposal within a day, no matter what happens to the rest of the money.

These are two requirements, that cannot be expressed in a formula of the kind $\mathcal{P}_{\geq \vartheta}(\dots)$. The second one is rather a postulation that allows for a change in the strategy of our banker, which explains the need of the $\langle\langle \cdot \rangle\rangle$ operator in nested form. Thus the appropriate formula looks like this:

$$\langle\langle A \rangle\rangle \left[\left(\mathcal{P}_{\geq 0.9}(\diamond^{\leq 365}(\text{earnings} \geq 100.000)) \wedge \mathcal{P}_{\geq 1}(\Box \langle\langle A \rangle\rangle \mathcal{P}_{\geq 1}(\mathcal{X}(\text{available money} \geq 120.000))) \right) \right]$$

Here, A represents the banker, \mathcal{X} represents the next operator and one step corresponds to one day.

Note that except for the probabilistic operator $\mathcal{P}(\cdot)$ a formula like the one above can also be expressed in ATL.

Our contribution. The novel logic SGL provides a uniform framework for reasoning about qualitative and quantitative linear- and branching-time properties of probabilistic multi-agent systems. We present the syntax and semantics of SGL and study decidability and complexity of the SGL model checking problem for various types of strategies. As a consequence of known results for stochastic games with branching time winning objectives [24,9], we obtain the undecidability of the SGL model checking problem for history-dependent strategies. However, decidability can be established when restricting the range of the strategies for the modalities $\langle\langle A \rangle\rangle$ and $\|A\|$ to memoryless ones. We present a classification of SGL formulas into “types” that yield a perfect match with the polynomial hierarchy when dealing with memoryless strategies only. This yields PSPACE-completeness of the SGL model checking problem under memoryless deterministic strategies and the model checking problem for the qualitative fragment of SGL under memoryless randomized strategies. Using an encoding of the semantics for (arbitrary) SGL formulas in the first-order theory of reals, we obtain an exponentially space-bounded model checking algorithm for full SGL and memoryless randomized strategies.

To the best of our knowledge, this paper (and its preceding conference version [3]) yields the first attempt for defining an ATL-like logic that can express quantitative (PCTL-like) properties. Former approaches with ATL-like modalities for reasoning about concurrent stochastic games have been studied by de Alfaro et al, e.g. [21,20]. However, these papers concentrate on qualitative properties and they do not consider Boolean combination of qualitative properties or the nesting of $\langle\langle \cdot \rangle\rangle$ operators.

The reduction of the SGL model checking problem for memoryless randomized strategies to Tarski algebra reuses ideas that have been presented in [24] for the synthesis of controller from PCTL specifications. However, several non-trivial adaptations to our more general logical framework with potential nestings of ATL-like modalities are required.

Organization. Section 2.1 introduces our model of probabilistic multi-player games (PMG) and related notions. The syntax and semantics of our logic SGL will be introduced in Section 3. The model checking problem for SGL on multi-player games is addressed in Section 4. Section 5 concludes the paper.

2 Preliminaries

We briefly explain our model of probabilistic multi-player games and introduce related notations (Section 2.1) and summarize the relevant features of deterministic Rabin automata (Section 2.2).

2.1 Probabilistic multi-player games

In this paper, we deal with turn-based multi-player games where in each state only one agent makes a move.

Definition 1 (Probabilistic multi-player game (PMG)) A probabilistic multi-player game (PMG) is a tuple $\mathcal{M} = (\text{Agents}, S, \rightarrow, P, \text{Props}, \nu)$ where

- Agents is a finite set of agents,
- S is a set of states, disjointly partitioned into $S = S_{\text{prob}} \cup \bigcup_{a \in \text{Agents}} S_a$,
- $\rightarrow \subseteq S \times S$ is a total transition relation³, i.e., for every $s \in S$ there is $t \in S$ such that $s \rightarrow t$,
- $P : S_{\text{prob}} \times S \rightarrow [0, 1]$ is a probability assignment such that, for all $s \in S_{\text{prob}}$, $\sum_{u \in S} P(s, u) = 1$ and $P(s, t) = 0$ iff $s \not\rightarrow t$,
- Props is a finite set of atomic propositions,
- $\nu : S \rightarrow 2^{\text{Props}}$ is a labeling function that assigns to each state s the set $\nu(s)$ of atomic propositions which hold in s .

A Markov decision process (MDP) is a PMG where the set of agents is a singleton.

We may regard S as a function that assigns to each agent a a set S_a such that $S_a \cap S_b = \emptyset$ if $a \neq b$. The states $s \in S_a$ are called a -states. For an agent set $A \subseteq \text{Agents}$, we write S_A for $\bigcup_{a \in A} S_a$ and refer to the states $s \in S_A$ as A -states. Note that in the a -states, it is agent's a turn to choose a transition $s \rightarrow t$. In the probabilistic states $s \in S_{\text{prob}}$, the successor state is chosen randomly according to P .

So far, no restrictions on \mathcal{M} have been made. When addressing the model checking problem, we consider only finite-state PMG with rational probability assignment (i.e., S is a finite set and $P(s, t)$ is rational for all $(s, t) \in S_{\text{prob}} \times S$).

We write $\text{Paths}(s)$ for the set of all infinite sequences $s_0 s_1 s_2 \dots \in S^\omega$ where $s_0 = s$ and $s_i \rightarrow s_{i+1}$ for all $i \geq 0$. More generally, for a finite sequence $w = s_0 \dots s_k$ of states we use $\text{Paths}(w)$ to denote the set of all $\pi \in \text{Paths}(s_0)$ that start with w (note that $\text{Paths}(w)$ can be empty). We denote by $\text{Succ}(s)$ the set of all successors of s , i.e., $\text{Succ}(s) = \{t \in S \mid s \rightarrow t\}$. For a path $\pi = s_0, s_1, \dots$ and $j \geq 0$, we denote by $\pi(j)$ the state s_j of π .

Given a finite or countably infinite set T , let $\text{Distr}(T)$ be the set of all distributions on T , i.e., functions $\mu : T \rightarrow [0, 1]$ such that $\sum_{t \in T} \mu(t) = 1$. A distribution μ is Dirac if $\mu(t) = 1$ for some $t \in T$.

Definition 2 (Strategy) Let $A \subseteq \text{Agents}$. A history-dependent randomized A -strategy (briefly HR strategy, or simply strategy) is a function $\alpha : S^* S_A \rightarrow \text{Distr}(S)$ such that $\alpha(s_1 \dots s_n s)(t) = 0$ if $s \not\rightarrow t$. An α -path denotes a path $s_0 s_1 s_2 \dots$ which is consistent with α 's decisions, i.e., for all $i \geq 0$, $s_i \in S_A$ implies $\alpha(s_0 \dots s_i)(s_{i+1}) > 0$.

A strategy α is called deterministic (or a HD strategy) if for all $s_1 \dots s_n s \in S^* S_A$, the distribution $\alpha(s_1 \dots s_n s)$ is Dirac. We say that α is memoryless (or an MR strategy) if $\alpha(s_1 \dots s_n s) = \alpha(s)$ for all state-sequences $s_1 \dots s_n$. An MD strategy means a memoryless deterministic strategy. A special type of HR strategies

³ In the rest of this paper, we will write $s \rightarrow t$ instead of $(s, t) \in \rightarrow$.

are finite-memory (FR) strategies, where the decision depends only on the control state entered by some fixed finite-state automaton after reading the sequence $s_1 \dots s_n s \in S^* S_A$. An FD strategy is deterministic finite-memory strategy.

Given a strategy α for all agents and a state $s \in S$, we define the probability space $(Paths(s), \mathcal{F}, Prob^\alpha)$ in the standard way, i.e.,

- \mathcal{F} is the σ -field generated by all $Paths(w)$ where w is a finite sequence of states initiated in s ,
- $Prob^\alpha$ is the unique probability measure such that for all $w = s_0 \dots s_n$ where $s = s_0$ and $n \geq 1$ we have that $Prob^\alpha(Paths(w)) = \prod_{i=0}^{n-1} x_i$. Here x_i is equal to $P(s_i, s_{i+1})$ or $\alpha(s_0 \dots s_i)(s_{i+1})$, depending on whether $s_i \in S_{prob}$ or not, respectively.

Sometimes we also need to consider games induced by A -strategies, where A is just a subset of agents. In particular, this is useful for memoryless strategies. For a given MR A -strategy α , the game \mathcal{M}^α induced by α arises from \mathcal{M} by fixing the decisions for the agents in A according to α (note that the A -states of \mathcal{M} become probabilistic states in \mathcal{M}^α). Formally, we define $\mathcal{M}^\alpha = (Agents \setminus A, S, \rightarrow^\alpha, P^\alpha, Props, \nu)$ where

- $s \rightarrow^\alpha t$ iff $s \in S_A$ and $\alpha(s)(t) > 0$, or $s \notin S_A$ and $s \rightarrow t$;
- $P^\alpha(s, t) = \alpha(s)(t)$ if $s \in S_A$, and $P^\alpha(s, t) = P(s, t)$ if $s \notin S_A$ is a probabilistic state of \mathcal{M} .

2.2 Deterministic Rabin automata

The logic SGL that will be introduced in the next section uses ω -regular languages to specify path properties in the style of the extended computation tree logic ECTL [13]. These languages are expressed by deterministic Rabin automata. We briefly recall here the basic concepts.

Definition 3 (Deterministic Rabin Automata (DRA)) A deterministic Rabin automaton (DRA) \mathcal{A} is a tuple $(Q, \Sigma, q_{init}, \delta, (L_i, R_i)_{i=1}^m)$, where

- Q is a finite set of states,
- Σ is a finite alphabet,
- $q_{init} \in Q$ is the initial state,
- $\delta : Q \times \Sigma \rightarrow Q$ is a transition function, and
- $(L_i, R_i)_{i=1}^m$ is the acceptance condition, where $L_i, R_i \subseteq Q$ for all $1 \leq i \leq m$.

Given an infinite word $\pi = \pi_1 \pi_2 \dots \in \Sigma^\omega$ over the alphabet Σ , we call $r(\pi) = q_1 q_2 q_3 \dots$ where $q_1 = q_{init}$ and $q_{i+1} = \delta(q_i, \pi_i)$ the run of \mathcal{A} for the input word π . By

$$\lim(r(\pi)) = \{q \in Q \mid q_j = q \text{ for infinitely many } j\}$$

we denote the limit of $r(\pi)$, i.e., the set of states that occur infinitely often in $r(\pi)$. We say that a set of states $T \subseteq Q$ is accepting iff there exists an index $j \in \{1, \dots, m\}$ such that $T \cap L_j \neq \emptyset$ and $T \cap R_j = \emptyset$. The language accepted by the Rabin automaton \mathcal{A} is defined as

$$L(\mathcal{A}) = \{\pi \in \Sigma^\omega \mid \lim(r(\pi)) \text{ is accepting}\}.$$

In this paper (particularly in Section 4), we rely on well-known results about optimal values in maximizing and minimizing MDPs with DRA objectives (see [17, 16] or chapter 10 in [5]). To make this paper self-contained, we briefly recall these results (in a form which suits our purposes). Let $\mathcal{M} = (\{a\}, S, \rightarrow, P, \text{Props}, v)$ be an MDP and $\mathcal{A} = (Q, \Sigma, q_{\text{init}}, \delta, (L_i, R_i)_{i=1}^m)$ a DRA where $\Sigma = 2^{\text{Props}}$. For every infinite path $\pi = s_1 s_2 \dots$, let $\hat{\pi} = v(s_1) v(s_2) \dots$ be the corresponding infinite word over Σ . For every $s \in S$, we define its upper and lower \mathcal{A} -value as follows:

$$\begin{aligned} \text{val}^+(\mathcal{A}, s) &= \sup_{\alpha} \text{Prob}^{\alpha}(\{\pi \in \text{Paths}(s) \mid \hat{\pi} \in L(\mathcal{A})\}), \\ \text{val}^-(\mathcal{A}, s) &= \inf_{\alpha} \text{Prob}^{\alpha}(\{\pi \in \text{Paths}(s) \mid \hat{\pi} \in L(\mathcal{A})\}). \end{aligned}$$

Here α ranges over all HR strategies of the only agent a .

The next proposition says that $\text{val}^+(\mathcal{A}, s)$ and $\text{val}^-(\mathcal{A}, s)$ correspond to components of the least solution of efficiently constructible systems of linear (in)equalities.

Proposition 1 (see [7, 17, 16, 5]) *Let $\mathcal{M} = (\{a\}, S, \rightarrow, P, \text{Props}, v)$ be an MDP and $\mathcal{A} = (Q, \Sigma, q_{\text{init}}, \delta, (L_i, R_i)_{i=1}^m)$ a DRA where $\Sigma = 2^{\text{Props}}$. There are systems $\Delta_{\text{acc}}, \Delta_{\text{rej}}$ of linear (in)equalities over the set of variables $\text{Var} = \{x_{s,q} \mid s \in S, q \in Q\}$, constructible in polynomial time, such that for every $s \in S$ we have that*

$$\begin{aligned} - \text{val}^+(\mathcal{A}, s) &= \mu[\Delta_{\text{acc}}](x_{s,q_{\text{init}}}), \\ - \text{val}^-(\mathcal{A}, s) &= 1 - \mu[\Delta_{\text{rej}}](x_{s,q_{\text{init}}}). \end{aligned}$$

Here $\mu[\Delta_{\text{acc}}]$ and $\mu[\Delta_{\text{rej}}]$ denote the least solution of Δ_{acc} and Δ_{rej} in $[0, 1]$, and $\mu[\Delta_{\text{acc}}](x_{s,q})$ and $\mu[\Delta_{\text{rej}}](x_{s,q})$ denote the value of $x_{s,q}$ in $\mu[\Delta_{\text{acc}}]$ and $\mu[\Delta_{\text{rej}}]$, respectively.

Proof (sketch) First, we construct the synchronized product of \mathcal{M} and \mathcal{A} , which is an MDP $\mathcal{M} \times \mathcal{A} = (\{a\}, S \times Q, \rightarrow_{\otimes}, P_{\otimes}, \text{Props}_{\otimes}, v_{\otimes})$, where

- $(S \times Q)_a = S_a \times Q$, $(S \times Q)_{\text{Prob}} = S_{\text{Prob}} \times Q$,
- $(s, q) \rightarrow_{\otimes} (s', q')$ iff $s \rightarrow s'$ and $\delta(q, v(s)) = q'$,
- $P_{\otimes}((s, q), (s', q'))$ is equal either to $P(s, s')$ or 0, depending on whether $(s, q) \rightarrow_{\otimes} (s', q')$ or not, respectively.
- $\text{Props}_{\otimes} = \text{Props}$,
- $v_{\otimes}(s, q) = v(s)$.

An end component of $\mathcal{M} \times \mathcal{A}$ is a set $U \subseteq S \times Q$ such that

- for all $(s, q) \in U$ and $(s', q') \in S \times Q$ such that $s \in S_{\text{Prob}}$ and $(s, q) \rightarrow_{\otimes} (s', q')$ we have that $(s', q') \in U$, i.e., U is closed under successors of probabilistic vertices;
- for every $(s, q) \in U$, where $s \in S_a$, there is $(s', q') \in U$ such that $(s, q) \rightarrow_{\otimes} (s', q')$ and $(s', q') \in U$;
- U is strongly connected, i.e., for each pair of states of U we have that the first state is reachable from the second state by a finite path leading only through the states of U .

Obviously, for each end component U , there is a strategy α for the only agent a such that each state of U is visited infinitely often with probability one, assuming that we start in a state of U . For a given end component U , we use Q_U to denote the set $\{q \in Q \mid (s, q) \in U \text{ for some } s \in S\}$. We say that a given end component U is *accepting* if there is $1 \leq i \leq m$ such that $Q_U \cap L_i \neq \emptyset$ and $Q_U \cap R_i = \emptyset$. Similarly, we say that U is *rejecting* if for all $1 \leq i \leq m$ we have that $Q_U \cap L_i = \emptyset$ or $Q_U \cap R_i \neq \emptyset$. The systems Δ_{acc} and Δ_{rej} are constructed as follows:

- (1) If (s, q) belongs to some accepting/rejecting end component, we add to $\Delta_{acc}/\Delta_{rej}$ the equality $x_{s,q} = 1$;
- (2) otherwise, we add to $\Delta_{acc}/\Delta_{rej}$ either
 - the equality $x_{s,q} = \sum_{(s,q) \rightarrow_{\otimes} (s',q')} P_{\otimes}((s,q), (s',q')) \cdot x_{s',q'}$, if $s \in S_{Prob}$;
 - the inequality $x_{s,q} \geq x_{s',q'}$ for every transition $(s,q) \rightarrow_{\otimes} (s',q')$, if $s \in S_a$.

It follows from the results of [7, 17, 16, 5] that $val^+(\mathcal{A}, s) = \mu[\Delta_{acc}](x_{s,q_{init}})$ and $val^-(\mathcal{A}, s) = 1 - \mu[\Delta_{rej}](x_{s,q_{init}})$. Further, the systems Δ_{acc} and Δ_{rej} are constructible in time polynomial in the size of \mathcal{M} and \mathcal{A} , because the condition of (1) is solvable in polynomial time. \square

3 The logic SGL

For specifying properties of probabilistic multi-player games we introduce a temporal logic, called stochastic game logic (SGL). It borrows ideas from ATL (and the ATL-like formalisms for stochastic games [21, 20]), extended computation tree logic ECTL [13], and the game logic GL of [2]. The probabilistic fragment of SGL contains a PCTL-like probabilistic operator which allows to reason about the probabilities for ω -regular properties, expressed by a deterministic Rabin automaton.

We start with the syntax of SGL formulas (Section 3.1), then presents its formal semantics by interpreting SGL formulas over the states of a probabilistic multi-player game (Section 3.2) and then discuss the expressiveness of SGL (Section 3.3).

3.1 SGL Syntax

Throughout the paper, let AP and AG be countably infinite sets of atomic propositions and agents, respectively.

The abstract syntax of SGL formulae is given by the following equation:

$$\Phi ::= p \mid \neg \Phi \mid \langle\langle A \rangle\rangle \Phi \mid \mathcal{P}_{\bowtie \lambda}(\mathcal{A}; \Phi_1, \dots, \Phi_k)$$

Here p and A range over AP and finite subsets of AG, respectively, $\bowtie \in \{<, \leq, >, \geq\}$ is a comparison operator, $\lambda \in [0, 1]$ is a rational probability bound, and \mathcal{A} is a DRA over the alphabet $2^{\{1, \dots, k\}}$.

Note that SGL does not contain the usual Boolean connectives (\wedge , \vee , etc.) or temporal operators such as \mathcal{X} (next) or \mathcal{U} (until). As we shall see in Example 1,

$s, A, \alpha \models_{XY} p$	iff $p \in v(s)$
$s, A, \alpha \models_{XY} \neg \Phi$	iff $s, A, \alpha \not\models_{XY} \Phi$
$s, A, \alpha \models_{XY} \langle\langle B \rangle\rangle \Phi$	iff there is an XY B -strategy β such that $s, A \cup B, (\alpha \leftarrow \beta) \models_{XY} \Phi$
$s, A, \alpha \models_{XY} \mathcal{P}_{\bowtie \lambda}(\mathcal{A}; \Phi_1, \dots, \Phi_k)$	iff for all HR strategies β for $\text{Agents} \setminus A$ we have that $Prob^{\alpha \leftarrow \beta}(\{\pi \in \text{Paths}(s) \mid \tilde{\pi}_{A, \alpha; XY}^{\Phi_1, \dots, \Phi_k} \in L(\mathcal{A})\}) \bowtie \lambda$

Fig. 1 Semantics of SGL.

these are expressible using appropriate DRA. Actually, negation is also redundant, but we keep it in SGL syntax for technical convenience.

An important syntactic fragment of SGL is *qualitative SGL* where the constant λ in the probabilistic operator $\mathcal{P}_{\bowtie \lambda}$ may only take the value 0 or 1.

3.2 SGL Semantics

The formula $\langle\langle A \rangle\rangle \Phi$ requires the existence of an XY A -strategy⁴ α such that the subformula Φ is satisfied in the game induced by α . However, the decisions already made by α can be changed by another $\langle\langle B \rangle\rangle$ operator in Φ for the agents in $A \cap B$. As decisions once made by a $\langle\langle \cdot \rangle\rangle$ operator might be changed by a nested $\langle\langle \cdot \rangle\rangle$ operator, we need to keep track on the strategy decisions that have already been made.

The $\mathcal{P}_{\bowtie \lambda}(\mathcal{A}; \Phi_1, \dots, \Phi_k)$ operator has the standard PCTL* semantics, meaning that for all HR strategies β of the “remaining” agents, the probability measure of all paths accepted by the automaton \mathcal{A} in the Markov chain induced by combining the “current” strategy α with β matches the probability bound λ . Here, a path is accepted by the automaton \mathcal{A} if its projection to words over $2^{\{1, \dots, k\}}$ indicating which of the formulae Φ_1, \dots, Φ_k are satisfied in each of the states of the path, is in $L(\mathcal{A})$.

Let $\mathcal{M} = (\text{Agents}, S, \rightarrow, P, \text{Props}, v)$ be a PMG where $\text{Agents} \subseteq \text{AG}$, $\text{Props} \subseteq \text{AP}$, and XY a class of strategies (i.e., XY is either MD, MR, HD, or HR). We define a satisfaction relation $s, A, \alpha \models_{XY} \Phi$ where s is a state in \mathcal{M} , $A \subseteq \text{Agents}$, α is an XY A -strategy, and Φ is an SGL formula such that for every $\langle\langle B \rangle\rangle$ operator used in Φ we have that $B \subseteq \text{Agents}$. The intuitive meaning is that s satisfies Φ in the game induced by α . Note that we need to keep track of the strategy decisions already made. The rules for the satisfaction relation are given in Fig. 1. The meaning of the newly employed symbols is the following:

- $(\alpha \leftarrow \beta)$ denotes the strategy for the agents in $A \cup B$ such that the agents in $A \setminus B$ behave according to the strategy α and the agents in B behave according to the strategy β , i.e., decisions already made by an agent in $A \cap B$ are neglected and a new strategy is chosen. Formally, given a path $\pi = s_1, \dots, s_n$, we put

$$(\alpha \leftarrow \beta)(\pi) = \begin{cases} \alpha(\pi) & \text{if } s_n \in S_A \setminus S_B \\ \beta(\pi) & \text{if } s_n \in S_B \end{cases}$$

⁴ Here XY stands for MD, MR, FD, FR, HD, or HR.




Fig. 2 DRA implementing Boolean and modal connectives. The initial state is q_0 .

- $\tilde{\pi}_{A,\alpha;XY}^{\Phi_1,\dots,\Phi_k} \in (2^{\{1,\dots,k\}})^\omega$ is defined as follows.

$$\tilde{\pi}_{A,\alpha;XY}^{\Phi_1,\dots,\Phi_k}(i) = \{j \mid 1 \leq j \leq k \text{ and } \pi(i), A, \alpha \models_{XY} \Phi_j\}.$$

Given a formula Φ , we denote by $Sat_{XY}(\Phi)$ the set of all states of \mathcal{M} that satisfy Φ , i.e.,

$$Sat_{XY}(\Phi) = \{s \in S \mid s, \emptyset, \alpha_\emptyset \models_{XY} \Phi\}.$$

Note that the class of strategies for agents that explicitly appear in some $\langle\langle \cdot \rangle\rangle$ operator is restricted to XY, while the remaining agents can always use unrestricted (i.e., HR) strategies. Intuitively, this is because the remaining agents are usually interpreted as unpredictable intruders, and hence their worst possible behaviour must be taken into account. On the other hand, the strategy for cooperating agents should be as simple as possible. The results in [4] yield that the satisfaction relations \models_{HD} , \models_{HR} , \models_{MD} , and \models_{MR} are pairwise distinct.

Example 1 As we already mentioned, the syntax of SGL does not contain the standard Boolean connectives and temporal operators such as “NextStep” (denoted by \mathcal{X}), “Always” (denoted by \square), or “Until” (denoted by \mathcal{U}). It is perhaps worth noting how to express these operators in SGL. For example, the formula $\mathcal{P}_{\infty\lambda}(\square\Phi_1)$ can be expressed in SGL as $\mathcal{P}_{\infty\lambda}(\mathcal{A}_\square; \Phi_1)$, where the DRA \mathcal{A}_\square is shown in Fig. 2 (together with DRA for some other connectives).

3.3 The relationship between SGL and other logics

In this paragraph we show that formulae of other well-known logics such as CTL, CTL*, PCTL, PCTL*, ATL, etc., can effectively be translated into SGL.

The standard (non-probabilistic) CTL is expressible in SGL. CTL is interpreted over labelled transition systems (Kripke structures) which can be seen as a PMG with no probabilistic states and only one agent. In the CTL semantics each path quantifier \exists, \forall is interpreted over the “full” system. Since in SGL strategies chosen by the $\langle\langle\{1}\rangle\rangle$ operator can be overwritten by another $\langle\langle\{1}\rangle\rangle$ operator, we can embed CTL as follows. Given a labelled transition system T and a CTL formula Φ , let Φ' be the SGL formula obtained from Φ by substituting each occurrence of the existential quantifier $\exists(\cdot)$ by $\langle\langle\{1}\rangle\rangle\mathcal{P}_{\geq 1}(\cdot)$, and each occurrence of the universal quantifier $\forall(\cdot)$ by $\|\{1\}\|\mathcal{P}_{\geq 1}(\cdot)$, where $\|\cdot\| = \neg\langle\langle\cdot\rangle\rangle\neg$. Then,

$$\text{Sat}_{CTL}(\Phi) = \text{Sat}_{MD}^T(\Phi').$$

Note that Φ' is not conform to our SGL syntax as it uses temporal operators like “Always” and “NextStep” instead of Rabin automata to express path properties. But, as indicated in Example 1, the formula Φ' can be transformed into an equivalent SGL formula.

The same transformation embeds CTL* into SGL, but in this case, the SGL formula has to be interpreted over the HD-semantics. That is, given a CTL* formula Φ , it holds that $\text{Sat}_{CTL^*}(\Phi) = \text{Sat}_{HD}^T(\Phi')$. As CTL* uses LTL path formulae we need more complicated automata as the ones introduced in example 1. However, this does not pose any real problems as the languages expressible by LTL formulae are contained in the ω -regular languages, and deterministic Rabin automata are as expressive as ω -regular languages [27, 31, 32].

The standard PCTL (interpreted over Markov decision processes (MDP)) can also be embedded into SGL. Each Markov decision process M can be seen as a PMG with only one agent. In PCTL, there are no path quantifiers like \exists and \forall . The semantics of the PCTL $\mathcal{P}_{\bowtie\lambda}(\cdot)$ operator implicitly quantifies over all strategies in the given MDP M . This is the same as in our SGL semantics. Moreover, given a formula

$$\mathcal{P}_{\bowtie\lambda}(\mathcal{A}; \Phi_1, \dots, \Phi_k),$$

the formulae Φ_1, \dots, Φ_k are interpreted over the same system as the formula $\mathcal{P}_{\bowtie\lambda}(\mathcal{A}; \Phi_1, \dots, \Phi_k)$. Hence, we do not need a transformation from PCTL to SGL as in the CTL case above; we only need the transformation from LTL path formulae to Rabin automata. Given a PCTL formula Φ and an MDP M , it holds that

$$\text{Sat}_{PCTL}(\Phi) = \text{Sat}_{MD}^M(\Phi).$$

Again, the temporal operators have to be substituted by the appropriate automata. Similarly, PCTL* embeds into SGL. Let M be an MDP and Φ be a PCTL* formula. Then,

$$\text{Sat}_{PCTL^*}(\Phi) = \text{Sat}_{HD}^M(\Phi).$$

Remark 1 Let M be an MDP (which can be understood as PMG with one agent), and let Φ be a SGL formula that is obtained from some PCTL* formula in the way indicated above. In particular, note that Φ does not contain the $\langle\langle\cdot\rangle\rangle$ operator. Assume that Φ has nested $\mathcal{P}_{\bowtie\lambda}(\cdot)$ operators, so it might look like this: $\mathcal{P}_{\bowtie\lambda}(\dots \mathcal{P}_{\bowtie\lambda'}(\dots)\dots)$. Let Φ' be the formula obtained from Φ by substituting each occurrence of $\mathcal{P}_{\bowtie\lambda}(\cdot)$ by $\|\{1\}\|\mathcal{P}_{\bowtie\lambda}(\cdot)$. Then,

$$\begin{aligned} Sat_{XY}^M(\Phi) &= Sat_{HD}^M(\Phi') && \text{for each strategy class } XY^5, \text{ whereas} \\ Sat_{HD}^M(\Phi) &\neq Sat_{HD}^M(\|\{1\}\|\Phi) && \text{in general.} \end{aligned}$$

This is because although in Φ' the $\|\{1\}\|$ operator of the outermost $\mathcal{P}_{\bowtie\lambda}(\cdot)$ operator fixes a strategy for the only agent, this strategy can be overwritten by the $\|\{1\}\|$ operator of a nested $\mathcal{P}_{\bowtie\lambda}(\cdot)$ operator. Thus, we get the standard PCTL semantics. On the other hand, the formula $\|\{1\}\|\Phi$ fixes a strategy α by the $\|\{1\}\|$ operator and evaluates the outermost $\mathcal{P}_{\bowtie\lambda}(\cdot)$ operator on the Markov chain M^α . This means that also the nested $\mathcal{P}_{\bowtie\lambda}(\cdot)$ operators are evaluated over M^α which gives the above inequality.

Even ATL is expressible in SGL. In standard ATL, the $\langle\langle\cdot\rangle\rangle_{ATL}$ operator is followed by a path formula. The ATL semantics of the formula $\langle\langle A \rangle\rangle_{ATL} \varphi$ yields the existence of an HD strategy for the A -agents such that for all HD strategies of the agents not in A , the path formula φ holds for the unique path that is determined by the chosen strategies. As already mentioned, the strategy chosen for the A -agents is not propagated to the subformulae. Given a PMG \mathcal{M} without any probabilistic states and an ATL formula Φ , let Φ' be the SGL formula obtained from Φ by substituting each occurrence of $\langle\langle A \rangle\rangle_{ATL} \varphi$ by $\langle\langle A \rangle\rangle_{Agents \setminus A} \mathcal{P}_{\geq 1}(\varphi)$. It holds that

$$Sat_{ATL}(\Phi) = Sat_{HD}^{\mathcal{M}}(\Phi').$$

The corresponding results hold also for ATL* and the SGL satisfaction under the HD-semantics.

In [2], the authors introduce an extension of ATL called game logic (GL). In contrast to ATL, where the operator $\langle\langle\cdot\rangle\rangle_{ATL}$ is followed by a path formula and the semantics implicitly quantifies over all paths, the $\langle\langle\cdot\rangle\rangle$ operator in game logic can also be followed by an existential path quantifier \exists . A formula of the kind $\Phi = \langle\langle A \rangle\rangle(\exists \square \varphi_1 \wedge \exists \square \varphi_2)$ is expressible in GL. Φ assert the existence of a strategy α for the agents in A , such that for some behavior of the remaining agents φ_1 is always true, and for some (possibly different) behavior of the remaining agents φ_2 is always true. Thus, the chosen strategy α is propagated to the inner subformulae. Nevertheless, the semantics of GL does not propagate strategies chosen by $\langle\langle\cdot\rangle\rangle$ operators to nested $\langle\langle\cdot\rangle\rangle$ operators. For example, the GL formula $\langle\langle A \rangle\rangle \langle\langle B \rangle\rangle \Phi$ is equivalent to $\langle\langle B \rangle\rangle \Phi$. Hence, the GL semantics is more alike to the standard CTL* semantics and differs crucially from our SGL semantics. Therefore, GL fails to express typical game properties like “*player B can react to the strategy chosen by player A*”.

ATL-like approaches to reason about stochastic games and qualitative winning objectives have been introduced by de Alfaro et al [21, 20]. They use ATL-like formulae, such as $\langle\langle A \rangle\rangle_{almost} \psi$ or $\langle\langle A \rangle\rangle_{positive} \psi$, to formalize the existence of a strategy for agents in A such that the condition specified by ψ holds almost surely or with positive probability. Our framework generalizes these concepts to the quantitative setting and allows to express, e.g., properties asserting that the agents in A can cooperate so that the probability of the event specified by ψ is within a certain interval, or so that a Boolean combination of such PCTL-like formulae holds, no matter how the other agents behave. The ATL-like formulae $\langle\langle A \rangle\rangle_{almost} \psi$ or $\langle\langle A \rangle\rangle_{positive} \psi$ of [21, 20] are encoded in SGL by the formulae $\langle\langle A \rangle\rangle_{\mathcal{P}=1}(\psi)$ and $\langle\langle A \rangle\rangle_{\mathcal{P}>0}(\psi)$, respectively. However, SGL cannot express the limit operator $\langle\langle A \rangle\rangle_{limit}$ of [21].

4 Model checking SGL

The model checking problem for SGL addresses the question whether for a given finite PMG \mathcal{M} , a state s of \mathcal{M} , and SGL-formula Φ it holds that $s \in \text{Sat}_{XY}(\Phi)$ for a given strategy class XY . We analyze the complexity of the model-checking problem for SGL and its natural fragments with respect to HR, HD, FR, FD, MR, and MD strategy classes.

We start by recalling the standard notion of polynomial hierarchy [25,26]. For every $\ell \geq 0$, the complexity classes Δ_ℓ , Σ_ℓ , and Π_ℓ are defined inductively as follows:

$$\begin{aligned} \Delta_0, \Sigma_0, \Pi_0 &\text{ are equal to P,} \\ \Delta_{i+1} &= \text{P}^{\Sigma_i}, \quad \Sigma_{i+1} = \text{NP}^{\Sigma_i}, \quad \Pi_{i+1} = \text{coNP}^{\Sigma_i}. \end{aligned}$$

A complete problem for Σ_ℓ , where $\ell \geq 1$, is QBF_ℓ . An instance of QBF_ℓ is a quantified Boolean formula (with ℓ quantifier alternations) of the form

$$\exists X_1 \forall X_2 \exists X_3 \forall X_4 \cdots QX_\ell \varphi$$

where the variables which appear in the propositional formula φ are disjointly partitioned into X_1, \dots, X_n , and Q is either \forall or \exists depending on whether ℓ is even or odd, respectively. The question is whether the formula is valid. Without restrictions, we may assume that φ takes the form

$$\begin{aligned} \varphi &= \bigwedge_{i=1}^I \psi_i, & \text{where } I \geq 1 \text{ and} \\ \psi_i &= \bigvee_{j=1}^{J_i} \xi_{i,j}, & \text{where } J_i \geq 1 \text{ and} \\ \xi_{i,j} &= \bigwedge_{k=1}^{K_{i,j}} \rho_{i,j,k}, & \text{where } K_{i,j} \geq 1 \text{ and} \\ \rho_{i,j,k} &= \bigvee_{m=1}^{M_{i,j,k}} \delta_{i,j,k,m}, & \text{where } M_{i,j,k} \geq 1 \text{ and} \\ \delta_{i,j,k,m} &= \bigwedge_{n=1}^{N_{i,j,k,m}} L_{i,j,k,m,n}, & \text{where } N_{i,j,k,m} \geq 1 \text{ and } L_{i,j,k,m,n} \text{ is a literal.} \end{aligned} \tag{1}$$

Here, a literal is a propositional variable or its negation. This assumption is safe because the propositional formula constructed in the proof of Cook's theorem [15] also has the same fixed structure⁶.

⁶ The propositional formula constructed in the proof of Cook's theorem is satisfiable iff a given non-deterministic Turing machine \mathcal{M} running in polynomial time accepts a given input word w . The formula depends on \mathcal{M} and w , but the nesting depth (and structure) of conjunctions and disjunctions is fixed. An explicit construction of the formula can be found in, e.g., [23], and a full justification of our assumption about φ follows from the proof of Σ_ℓ -hardness of QBF_ℓ ; see, e.g., [26], Theorem 17.10.

We show that there is a close correspondence between the polynomial hierarchy and natural syntactic fragments of SGL. For every SGL formula Φ , we define the “type” of Φ , denoted by $\text{Type}(\Phi)$, inductively as follows:

$$\begin{aligned} \text{Type}(p) &= \Delta_0 \\ \text{Type}(\neg\Phi) &= \begin{cases} \Delta_\ell & \text{if } \text{Type}(\Phi) = \Delta_\ell, \\ \Sigma_\ell & \text{if } \text{Type}(\Phi) = \Pi_\ell, \\ \Pi_\ell & \text{if } \text{Type}(\Phi) = \Sigma_\ell. \end{cases} \\ \text{Type}(\langle\langle B \rangle\rangle\Phi) &= \begin{cases} \Sigma_{\ell+1} & \text{if } \text{Type}(\Phi) = \Sigma_\ell \text{ or } \text{Type}(\Phi) = \Pi_\ell, \\ \Sigma_\ell & \text{if } \text{Type}(\Phi) = \Delta_\ell. \end{cases} \\ \text{Type}(\mathcal{P}_{\bowtie\lambda}(\mathcal{A}; \Phi_1, \dots, \Phi_k)) &= \begin{cases} \Delta_{\ell+1} & \text{if } \text{Max}\{\text{Type}(\Phi_1), \dots, \text{Type}(\Phi_k)\} = \Sigma_\ell \\ \Delta_\ell & \text{if } \text{Max}\{\text{Type}(\Phi_1), \dots, \text{Type}(\Phi_k)\} = \Delta_\ell. \end{cases} \end{aligned}$$

Here $\text{Max}\{\text{Type}(\Phi_1), \dots, \text{Type}(\Phi_k)\}$ denotes the “maximal” type in the set. More precisely, let $\ell \geq 0$ be the least number such that for every $1 \leq i \leq k$ we have that $\text{Type}(\Phi_i) = \Lambda_j$ where $j \leq \ell$ and $\Lambda \in \{\Sigma, \Pi, \Delta\}$. If for all $1 \leq i \leq k$ such that $\text{Type}(\Phi_i) = \Lambda_\ell$ we have that $\Lambda = \Delta$, then $\text{Max}\{\text{Type}(\Phi_1), \dots, \text{Type}(\Phi_k)\} = \Delta_\ell$. Otherwise, $\text{Max}\{\text{Type}(\Phi_1), \dots, \text{Type}(\Phi_k)\} = \Sigma_\ell$. Note that we do not distinguish between the types Σ_ℓ and Π_ℓ when defining the maximal type. Intuitively, this is because the maximal type is used as an oracle for a deterministic polynomial-time algorithm which checks the validity of $\mathcal{P}_{\bowtie\lambda}(\mathcal{A}; \Phi_1, \dots, \Phi_k)$, and $P^{\Sigma_\ell} = P^{\Pi_\ell}$ (see Theorem 1).

Now we can define the promised hierarchy of SGL syntactic fragments.

- $\text{SGL}(\Delta_0)$, $\text{SGL}(\Sigma_0)$, and $\text{SGL}(\Pi_0)$ consist of all SGL formulae of type Δ_0 , Σ_0 , and Π_0 , respectively.
- $\text{SGL}(\Delta_{i+1})$, $\text{SGL}(\Sigma_{i+1})$, and $\text{SGL}(\Pi_{i+1})$ consist of all formulae in $\text{SGL}(\Delta_i) \cup \text{SGL}(\Sigma_i) \cup \text{SGL}(\Pi_i)$, and all formulae of type Δ_{i+1} , Σ_{i+1} , and Π_{i+1} , respectively.

4.1 MD strategies

In this section we examine the SGL model-checking problem with respect to MD semantics.

Lemma 1 *Let $\mathcal{M} = (\text{Agents}, S, \rightarrow, P, \text{Props}, v)$ be a PMG where $\text{Agents} \subseteq \text{AG}$, $\text{Props} \subseteq \text{AP}$, and let Φ be an SGL formula. The problem whether $s, A, \alpha \models_{\text{MD}} \Phi$, where $s \in S$, $A \subseteq \text{Agents}$, and α is an MD A -strategy, is solvable in $\text{Type}(\Phi)$.*

Proof We proceed by induction on the structure of Φ . The cases when $\Phi = p$ and $\Phi = \neg\Psi$ are immediate.

If $\Phi = \langle\langle B \rangle\rangle\Psi$, it suffices to guess an appropriate MD B -strategy β and check whether $s, A \cup B, (\alpha \leftarrow \beta) \models_{\text{MD}} \Psi$. Hence, by applying induction hypothesis, the problem whether $s, A, \alpha \models_{\text{MD}} \Phi$ is in $\text{NP}^{\text{Type}(\Psi)}$, which is equal to $\text{Type}(\Phi)$ by the definition of Type . Recall that $\text{NP}^{\Delta_\ell} = \Sigma_\ell$ and $\text{NP}^{\Sigma_\ell} = \text{NP}^{\Pi_\ell} = \Sigma_{\ell+1}$.

If $\Phi = \mathcal{P}_{\bowtie\lambda}(\mathcal{A}; \Psi_1, \dots, \Psi_k)$, we first apply the current MD A -strategy α to \mathcal{M} and construct the PMG \mathcal{M}^α (see Section 2.1). Further, for every state t we compute the set $\sigma(t)$ of all $j \in \{1, \dots, k\}$ such that $t, A, \alpha \models_{MD} \Psi_j$. By induction hypothesis, this is achievable by a deterministic polynomial-time algorithm with $\text{Max}\{\text{Type}(\Psi_1), \dots, \text{Type}(\Psi_k)\}$ oracle. Since the agents in $\text{Agents} \setminus A$ are now considered adversarial, we interpret \mathcal{M}^α as an MDP where the only agent $\{a\}$ controls all $(\text{Agents} \setminus A)$ -states. Further, we interpret $\{1, \dots, k\}$ as the set of atomic propositions of \mathcal{M}^α , and σ as the corresponding valuation. If $\bowtie \in \{>, \geq\}$, we simply check if $\text{val}^-(\mathcal{A}, s) \bowtie \lambda$, where s is considered as a state of \mathcal{M}^α . Note that this is achievable in time polynomial in the size of \mathcal{M}^α and \mathcal{A} by solving the linear program of Proposition 1.

Similarly, if $\bowtie \in \{<, \leq\}$, we check if $\text{val}^+(\mathcal{A}, s) \bowtie \lambda$. Hence, the problem whether $s, A, \alpha \models_{MD} \mathcal{P}_{\bowtie\lambda}(\mathcal{A}; \Psi_1, \dots, \Psi_k)$ is in $P^{\text{Max}\{\text{Type}(\Psi_1), \dots, \text{Type}(\Psi_k)\}}$ which is equal to $\text{Type}(\mathcal{P}_{\bowtie\lambda}(\mathcal{A}; \Psi_1, \dots, \Psi_k))$ by the definition of Type . \square

Now we prove the corresponding lower complexity bound. For the sake of readability, we adopt the following abbreviations, where \mathcal{A}_\vee , \mathcal{A}_\exists , \mathcal{A}_\square , and \mathcal{A}_\forall are the automata of Fig. 2.

$$\begin{aligned} \Phi_1 \Rightarrow \Phi_2 &= \mathcal{P}_{=1}(\mathcal{A}_\vee; \neg\Phi_1, \Phi_2) \\ \mathcal{X}^{\bowtie\lambda} \Phi &= \mathcal{P}_{\bowtie\lambda}(\mathcal{A}_\exists; \Phi) \\ \square^{\bowtie\lambda} \Phi &= \mathcal{P}_{\bowtie\lambda}(\mathcal{A}_\square; \Phi) \\ \Phi_1 \mathcal{U}^{\bowtie\lambda} \Phi_2 &= \mathcal{P}_{\bowtie\lambda}(\mathcal{A}_\forall; \Phi_1, \Phi_2) \\ \|A\| \Phi &= \neg\langle\langle A \rangle\rangle \neg\Phi \end{aligned}$$

We also write $\langle\langle a \rangle\rangle \Phi$ instead of $\langle\langle \{a\} \rangle\rangle \Phi$ for a single agent a .

Lemma 2 *For every $\ell \geq 0$, there is a fixed formula $\Phi \in \text{SGL}(\Sigma_\ell)$ such that the model-checking problem for Φ is Σ_ℓ -hard.*

Proof Let $\exists X_1 \forall X_2 \exists X_3 \forall X_4 \dots QX_\ell \varphi$ be a quantified Boolean formula with ℓ quantifier alternations, where the propositional formula φ takes the special form (1) introduced at the beginning of Section 4. We construct a PMG $\mathcal{M} = (\text{Agents}, S, \rightarrow, P, \text{Props}, \nu)$, a state $s(\varphi) \in S$, and a fixed formula $\Phi \in \text{SGL}(\Sigma_\ell)$ such that $\exists X_1 \forall X_2 \exists X_3 \forall X_4 \dots QX_\ell \varphi$ is valid iff $s(\varphi) \in \text{Sat}_{MD}(\Phi)$.

Let $\{x_1, \dots, x_r\}$ be the set of all propositional variables that appear in φ . The PMG \mathcal{M} is constructed as follows. We put $\text{Agents} = \{a_1, \dots, a_\ell\}$ and $\text{Props} = \{t, f, d, a, e, b\}$. The probability assignment P can be chosen arbitrarily (the precise values of transition probabilities do not influence our arguments; note that P is required to be positive by Definition 1). The states, transitions, and labelling function of \mathcal{M} are defined incrementally. For all $u \in \{1, \dots, r\}$, we put to S a fresh stochastic state B_u which satisfies b and no other proposition. Further, for every $\delta_{i,j,k,m}$ subformula of φ , we add to \mathcal{M} the gadget with initial state δ depicted in Fig. 3. All states of the gadget are fresh, except for B_1, \dots, B_r that are shared by all gadgets (note that the number of outgoing transitions of every B_u is the same and it is equal to the number of $\delta_{i,j,k,m}$ subformulae of φ). The states p_1, \dots, p_r are controlled by the agents, and the other states are stochastic. Intuitively, the agent controlling p_u can set the variable x_u to true or false by selecting the transition

leading to x_u or \bar{x}_u , respectively. The state p_u satisfies either the atomic proposition e or a , depending on whether x_u is quantified existentially (i.e., $x_u \in X_i$ where i is odd) or universally, respectively. For every $i \in \{1, \dots, \ell\}$, the agent a_i sets the variables in X_i , i.e., if $x_u \in X_i$, then p_u is an a_i -state. The propositions t and f are satisfied exactly in x_1, \dots, x_r and $\bar{x}_1, \dots, \bar{x}_r$, respectively, and the proposition d encodes the structure of $\delta_{i,j,k,m}$ in the following way: for every variable x_u , where $1 \leq u \leq r$,

- if $\delta_{i,j,k,m}$ contains both x_u and $\neg x_u$ (as literals), then $d \notin v(x_u)$ and $d \notin v(\bar{x}_u)$;
- if $\delta_{i,j,k,m}$ does not contain x_u and $\neg x_u$, then $d \in v(x_u)$ and $d \in v(\bar{x}_u)$;
- if $\delta_{i,j,k,m}$ contains x_u and does not contain $\neg x_u$, then $d \in v(x_u)$ and $d \notin v(\bar{x}_u)$;
- if $\delta_{i,j,k,m}$ does not contain x_u and contains $\neg x_u$, then $d \notin v(x_u)$ and $d \in v(\bar{x}_u)$.

The other states of the gadget, different from $x_1, \dots, x_r, \bar{x}_1, \dots, \bar{x}_r, B_1, \dots, B_r$, satisfy d as well.

Note that in general, agents can play *inconsistently* by setting the same variable to true and false in different gadgets. However, this can be easily detected. Consider the following formulae:

$$\begin{aligned} \text{Surely}_t &= (\neg t \wedge \neg f) \mathcal{U}^{-1} t \\ \text{Surely}_f &= (\neg t \wedge \neg f) \mathcal{U}^{-1} f \\ \text{Cons}_e &= \square^1 \left((\mathcal{X}^{>0} b \wedge \mathcal{X}^{>0} e) \Rightarrow (\text{Surely}_t \vee \text{Surely}_f) \right) \\ \text{Cons}_a &= \square^1 \left((\mathcal{X}^{>0} b \wedge \mathcal{X}^{>0} a) \Rightarrow (\text{Surely}_t \vee \text{Surely}_f) \right) \end{aligned}$$

Let δ be the initial state of some of the constructed gadgets, and let α be an MD Agents-strategy. We claim that α is consistent iff

$$\delta, \text{Agents}, \alpha \models_{\text{MD}} \text{Cons}_e \wedge \text{Cons}_a.$$

The formula Cons_e encodes the consistency of α with respect to existentially quantified variables, and the formula Cons_a does the same for universally quantified variables (the reason why we treat the existentially/universally quantified variables separately becomes clear later). To see this, consider, e.g., the choice made by α in the states p_u , where x_u is quantified universally. Let q be a predecessor of p_u in some gadget. Then $q, \text{Agents}, \alpha \models_{\text{MD}} \mathcal{X}^{>0} b \wedge \mathcal{X}^{>0} a$. If α behaves consistently in all p_u , i.e., selects either always the transition to x_u or always the transition to \bar{x}_u , then $q, \text{Agents}, \alpha \models_{\text{MD}} \text{Surely}_t$ or $q, \text{Agents}, \alpha \models_{\text{MD}} \text{Surely}_f$, respectively. On the other hand, if α behaves inconsistently in p_u , then $q, \text{Agents}, \alpha \not\models_{\text{MD}} \text{Surely}_t$ and $q, \text{Agents}, \alpha \not\models_{\text{MD}} \text{Surely}_f$. These observations are easy to verify by examining the structure of the gadgets, and also explain the role of B_u states.

Observe that every assignment μ for the propositional variables x_1, \dots, x_r determines a unique consistent MD Agents-strategy α_μ , and vice versa. For every assignment μ and every $\delta_{i,j,k,m}$ subformula we have that $\delta_{i,j,k,m}$ is true in μ iff

$$\delta, \text{Agents}, \alpha_\mu \models_{\text{MD}} \neg b \mathcal{U}^0 \neg d.$$

This follows directly from the construction of the gadgets.

Now we complete the construction of \mathcal{M} by adding the following states and transitions:




Fig. 3 The gadget for $\delta_{i,j,k,m}$ subformula.

- we add a state $s(\varphi)$; further, for every ψ_i subformula we add a state $s(\psi_i)$, for every $\xi_{i,j}$ subformula we add a state $s(\xi_{i,j})$, and for every $\rho_{i,j,k}$ subformula we add a state $s(\rho_{i,j,k})$;
- we add a transition $s(\varphi) \rightarrow s(\psi_i)$ for every $1 \leq i \leq I$;
- for every ψ_i subformula, we add a transition $s(\psi_i) \rightarrow s(\xi_{i,j})$ for all $1 \leq j \leq J_i$;
- for every $\xi_{i,j}$ subformula, we add a transition $s(\xi_{i,j}) \rightarrow s(\rho_{i,j,k})$ for all $1 \leq k \leq K_{i,j}$;
- for every $\rho_{i,j,k}$ subformula and every $1 \leq m \leq M_{i,j,k}$, we add a transition $s(\rho_{i,j,k}) \rightarrow \delta$ where δ is the initial state of the gadget for $\delta_{i,j,k,m}$.

Finally, we define the formula Φ as follows:

$$\langle\langle a_1 \rangle\rangle \|a_2\| \langle\langle a_3 \rangle\rangle \cdots [a_\ell] \mathcal{X}^{=1} \mathcal{X}^{>0} \mathcal{X}^{=1} \mathcal{X}^{>0} (\neg \text{Cons}_a \vee (\text{Cons}_e \wedge \neg b \mathcal{U}^{=0} \neg d))$$

Here $[a_\ell]$ is either $\langle\langle a_\ell \rangle\rangle$ or $\|a_\ell\|$, depending on whether ℓ is odd or even, respectively. Observe that $\Phi \in \text{SGL}(\Sigma_\ell)$ and Φ is a fixed formula for a fixed ℓ . We claim that

$$\exists X_1 \forall X_2 \exists X_3 \forall X_4 \cdots QX_\ell \varphi \text{ is valid iff } s(\varphi) \in \text{Sat}_{MD}(\Phi).$$

The “ \Rightarrow ” direction follows by observing that

- if i is even, then the agent a_i does not gain anything by using an inconsistent strategy, because this inevitably makes the subformula $\neg \text{Cons}_a$ valid;
- if i is odd and all of the previous choices of the agents were consistent (i.e., encode an assignment for the variables in $X_1 \cup \cdots \cup X_{i-1}$), then a_i simply takes an assignment for X_i which makes the formula $\forall X_{i+1} \cdots QX_\ell \varphi$ valid and encodes this assignment in his strategy. If some of the previous choices of the agents was inconsistent, then a_i can play arbitrarily.

The “ \Leftarrow ” direction follows similarly. \square

A direct consequence of Lemma 1 and Lemma 2 is the following:

Theorem 1 *Let $\ell \geq 0$. Then*

- *the model-checking problem for $\text{SGL}(\Sigma_\ell)$ is Σ_ℓ -complete for MD semantics, and the hardness result holds even for a fixed $\text{SGL}(\Sigma_\ell)$ formula;*
- *the model-checking problem for $\text{SGL}(\Pi_\ell)$ is Π_ℓ -complete for MD semantics, and the hardness result holds even for a fixed $\text{SGL}(\Pi_\ell)$ formula;*

- the model-checking problem for $SGL(\Delta_\ell)$ is in Δ_ℓ for MD semantics.

Finally, the model-checking problem for SGL is PSPACE-complete for MD semantics.

Observe that PSPACE-hardness of the model-checking problem for SGL follows immediately from Lemma 2 (we can use the same construction to reduce the QBF problem to SGL model-checking). However, the constructed SGL formula is not fixed anymore because the quantifier alternation depth in QBF instances is not fixed. In fact, it is unlikely that SGL model-checking is PSPACE-hard for some fixed SGL formula Φ , because this would imply collapse of the polynomial hierarchy at $\text{Type}(\Phi)$ level. Let us note that one can also give a simple direct proof for PSPACE-hardness of SGL model-checking which avoids many of the technicalities presented in the proof of Lemma 1.

4.2 MR strategies

We start by observing that if we restrict ourselves to the *qualitative fragment* of SGL (see Section 3.1), then randomization brings only a limited extra power to the agents, and the corresponding results are the same as for MD strategies.

Theorem 2 *Let $\ell \geq 0$. Then*

- the model-checking problem for qualitative $SGL(\Sigma_\ell)$ is Σ_ℓ -complete for MR semantics, and the hardness result holds even for a fixed qualitative $SGL(\Sigma_\ell)$ formula;
- the model-checking problem for qualitative $SGL(\Pi_\ell)$ is Π_ℓ -complete for MR semantics, and the hardness result holds even for a fixed qualitative $SGL(\Pi_\ell)$ formula;
- the model-checking problem for qualitative $SGL(\Delta_\ell)$ is in Δ_ℓ for MR semantics.

Finally, the model-checking problem for qualitative SGL is PSPACE-complete for MR semantics.

Note that the formula Φ constructed in the proof of Lemma 2 is qualitative, and the agents do not gain anything by using randomized strategies. Hence, the proof of Lemma 2 works also for qualitative $SGL(\Sigma_\ell)$ and MR strategies without any change, and thus we obtain the lower bounds of Theorem 2. The upper complexity bounds follow from Lemma 1, which is valid also for MR strategies and qualitative $SGL(\Sigma_\ell)$. However, the subcase $\Phi \equiv \langle\langle B \rangle\rangle \Psi$ in the proof of Lemma 1 requires a slight modification. Realize that for the qualitative fragment of SGL, the exact values of transition probabilities chosen by the agents do not really matter; it is only important which of them are chosen with positive/zero probability. Hence, instead of guessing a single outgoing transition for every B -state, which was enough for MD strategies, we now guess for every B -state a subset of outgoing transitions that are assigned a positive probability (and choose an arbitrary positive distribution over the chosen successors). The rest of the proof of Lemma 1 does not require any modification.

For general SGL formulae, the exact values of transition probabilities are of course relevant, and our decidability proof is based on encoding the SGL model

checking problem into first-order theory of the reals, i.e., $(\mathbb{R}, *, +, \leq)$, which is known to be decidable [29]. Our proof can be seen as an extension of the previous result [24] where it was shown that the MR-controller synthesis for MDPs (viewed as $1\frac{1}{2}$ -player games) and PCTL specifications can effectively be encoded by closed formulae of $(\mathbb{R}, *, +, \leq)$.

In the proof of our next lemma we often construct some finite index set $I = \{i_1, \dots, i_n\}$ and then fix a fresh first-order variable X_i for every $i \in I$. To simplify our notation, we define

$$\tilde{\exists}X_j.\psi = (\exists X_{i_1} \exists X_{i_2} \dots \exists X_{i_n}).\psi.$$

Similarly, if $J \subseteq I$ where $J = \{j_1, \dots, j_m\}$, we write just $(\tilde{\exists}X_j : j \in J).\psi$ instead of $(\exists X_{j_1} \exists X_{j_2} \dots \exists X_{j_m}).\psi$.

Lemma 3 *Let $\mathcal{M} = (\text{Agents}, S, \rightarrow, P, \text{Props}, \nu)$ be a PMG where $\text{Agents} \subseteq \text{AG}$, $\text{Props} \subseteq \text{AP}$, and let Φ be an SGL formula. For every $s \in S$ there is a closed formula $\tau(s, \Phi)$ of $(\mathbb{R}, *, +, \leq)$ such that $s \in \text{Sat}_{\text{MR}}(\Phi)$ iff $\tau(s, \Phi)$ holds.*

Proof For every transition $s \rightarrow t$ of \mathcal{M} , we use $Y_{s,t}$ to denote either the constant whose value is equal to $P(s,t)$ if s is stochastic, or a fresh first-order variable that encodes the probability of $s \rightarrow t$ chosen by the responsible agent. For every subset C of agents, we construct a closed formula $\tau_C(s, \Phi)$ of $(\mathbb{R}, *, +, \leq)$ such that $\tau_C(s, \Phi)$ is valid iff $s, C, \gamma \models_{\text{MR}}^{\mathcal{M}} \Phi$, where the strategy γ is given by the values of the variables $Y_{s,t}$ for $s \in S_C$. Then, we simply put $\tau(s, \Phi) = \tau_{\emptyset}(s, \Phi)$.

The formula $\tau_C(s, \Phi)$ is defined by induction on the structure of Φ . The first three subcases are immediate. We put

$$\begin{aligned} \tau_C(s, p) &= \begin{cases} \text{true} & \text{if } p \in \nu(s), \\ \text{false} & \text{otherwise;} \end{cases} \\ \tau_C(s, \neg\Phi) &= \neg\tau_C(s, \Phi), \\ \tau_C(s, \langle\langle A \rangle\rangle\Phi) &= (\tilde{\exists}Y_{r,t} : r \in S_A, r \rightarrow t) \cdot \left(\tau_{C \cup A}(s, \Phi) \wedge \bigwedge_{r \in S_A} \text{Dist}_r \right), \end{aligned}$$

where

$$\text{Dist}_r = \left(\bigwedge_{r \rightarrow t} 0 \leq Y_{r,t} \leq 1 \right) \wedge \left(\sum_{r \rightarrow t} Y_{r,t} = 1 \right).$$

Now let $\Phi = \mathcal{P}_{\bowtie\lambda}(\mathcal{A}; \Psi_1, \dots, \Psi_k)$, where $\mathcal{A} = (Q, \Sigma, q_{\text{init}}, \delta, (L_i, R_i)_{i=1}^m)$. Let us consider the MDP \mathcal{M}^γ with atomic propositions $\{1, \dots, k\}$ and a labelling function η such that $\eta(t)$ consists of all i where $1 \leq i \leq k$ and $t, C, \gamma \models_{\text{MR}}^{\mathcal{M}} \Psi_i$. We need to encode either the property $\text{val}^+(\mathcal{A}, s) \bowtie \lambda$ or $\text{val}^-(\mathcal{A}, s) \bowtie \lambda$ (where s is interpreted as a state of \mathcal{M}^γ), depending on whether $\bowtie \in \{\leq, <\}$ or $\bowtie \in \{\geq, >\}$, respectively. According to Proposition 1, this means to encode

$$\text{either } \mu[\Delta_{\text{acc}}](x_{s, q_{\text{init}}}) \bowtie \lambda \text{ or } 1 - \mu[\Delta_{\text{rej}}](x_{s, q_{\text{init}}}) \bowtie \lambda,$$

respectively. Note that the first inequality holds iff there exists *some* solution sol of Δ_{acc} in $[0, 1]$ such that $sol(x_{s,q_{init}}) \bowtie \lambda$. Realize that $\mu[\Delta_{acc}] \leq sol$. Similarly, the second inequality holds iff there exists some solution sol of Δ_{rej} in $[0, 1]$ such that $1 - sol(x_{s,q_{init}}) \bowtie \lambda$.

For all $t \in S$ and $q \in Q$, let $Z_{t,q}$ be a fresh first-order variable. The formula $\tau_C(s, \Phi)$ takes the form

$$\tau_C(s, \Phi) = \exists Z_{t,q}. \left(\bigwedge_{t \in S, q \in Q} 0 \leq Z_{t,q} \leq 1 \right) \wedge Sol \wedge Req$$

where

$$Req = \begin{cases} Z_{s,q_{init}} \bowtie \lambda & \text{if } \bowtie \in \{\leq, <\}, \\ 1 - Z_{s,q_{init}} \bowtie \lambda & \text{otherwise.} \end{cases}$$

The subformula Sol says that the variables $Z_{t,q}$ form a solution of Δ_{acc} and Δ_{rej} in $[0, 1]$, respectively.

The construction of Sol is not trivial, because the labeling η depends on the concrete form of Ψ_1, \dots, Ψ_k . Hence, we first need to encode η and the structure of $\mathcal{M}^\gamma \times \mathcal{A}$ symbolically, and then we can proceed with encoding the (in)equalities of Δ_{acc} and Δ_{rej} (cf. the proof of Proposition 1).

Let us fix a fresh variable $X_{t,i}$ for all $t \in S$ and $i \in \{1, \dots, k\}$. We construct a formula Eta which ensures that $X_{t,i}$ is positive iff $t, C, \gamma \models_{MR} \Psi_t$.

$$Eta = \bigwedge_{t \in S, i \in \{1, \dots, k\}} (X_{t,i} > 0 \Leftrightarrow \tau_C(t, \Psi_t))$$

Further, we fix a fresh variable $T_{(t,q),(t',q')}$ for all $(t,q), (t',q') \in S \times Q$ where $(t,q) \neq (t',q')$, and construct a formula $Tran$ which says that $T_{(t,q),(t',q')}$ is either 1 or 0, depending on whether $(t,q) \rightarrow_\otimes (t',q')$ in $\mathcal{M}^\gamma \times A$ or not, respectively.

$$\begin{aligned} Tran = & \bigwedge_{\substack{t,t' \in S \\ q,q' \in Q}} (T_{(t,q),(t',q')} = 1 \vee T_{(t,q),(t',q')} = 0) \\ & \wedge \bigwedge_{\substack{t,t' \in S \\ q,q' \in Q \\ t \neq t'}} T_{(t,q),(t',q')} = 0 \\ & \wedge \bigwedge_{\substack{t,t' \in S \\ q,q' \in Q \\ t \rightarrow t'}} \left(T_{(t,q),(t',q')} = 1 \Leftrightarrow \bigvee_{\substack{A \subseteq \{1, \dots, k\} \\ \delta(q,A) = q'}} \left(\bigwedge_{i \in A} X_{t,i} > 0 \wedge \bigwedge_{i \notin A} X_{t,i} \leq 0 \right) \right) \end{aligned}$$

Now we encode the existence of an accepting/rejecting end component of $\mathcal{M}^\gamma \times \mathcal{A}$. For all $(t,q) \in S \times Q$, we fix a fresh variable $V_{t,q}$. We construct a formula Acc

which says that the set of all (t, q) , where $V_{t,q} > 0$, forms an accepting/rejecting end component.

$$AccRej = \bigwedge_{\substack{t \in S_{Prob} \cup S_C \\ q \in Q}} V_{t,q} > 0 \Rightarrow \left(\bigwedge_{\substack{t' \in S \\ q' \in Q}} T_{(t,q)(t',q')} \Rightarrow V_{t',q'} > 0 \right) \quad (2)$$

$$\wedge \bigwedge_{\substack{t \in S_{Agents \setminus C} \\ q \in Q}} V_{t,q} > 0 \Rightarrow \left(\bigvee_{\substack{t' \in S \\ q' \in Q}} T_{(t,q)(t',q')} \wedge V_{t',q'} > 0 \right) \quad (3)$$

$$\wedge \bigwedge_{\substack{t, t' \in S \\ q, q' \in Q}} (V_{t,q} > 0 \wedge V_{t',q'} > 0) \Rightarrow Reach^{\leq N}(t, q)(t', q') \quad (4)$$

$$\wedge Rabin \quad (5)$$

Formula (2) says that the set of states encoded by $V_{t,q}$ variables is closed under successors of stochastic states, and (3) says that each non-deterministic state in the set has at least one successor in the set. The condition of strong connectedness is encoded by (4), where $N = |S| \cdot |Q|$, and (5) says that the end component is accepting or rejecting, depending on whether $\bowtie \in \{\leq, <\}$ or $\bowtie \in \{\geq, >\}$, respectively. The formula $Reach^{\leq N}(t, q)(t', q')$ says that (t, q) can reach (t', q') by a sequence of at most N transitions, and it is constructed inductively as follows:

$$Reach^{\leq 0}(t, q)(t', q') = \begin{cases} true & \text{if } t = t' \text{ and } q = q', \\ false & \text{otherwise;} \end{cases}$$

$$Reach^{\leq i+1}(t, q)(t', q') = \bigvee_{\substack{t'' \in S \\ q'' \in Q}} (Reach^{\leq i}(t, q)(t'', q'') \wedge T_{(t'', q'')(t', q')} = 1)$$

The formula *Rabin* is easy to construct. We put either

$$Rabin = \bigvee_{1 \leq i \leq m} \left(\bigvee_{\substack{q \in L_i \\ t \in S}} (V_{t,q} > 0) \wedge \bigwedge_{\substack{q \in R_i \\ t \in S}} (V_{t,q} \leq 0) \right)$$

or

$$Rabin = \bigwedge_{1 \leq i \leq m} \left(\bigwedge_{\substack{q \in L_i \\ t \in S}} (V_{t,q} \leq 0) \vee \bigvee_{\substack{q \in R_i \\ t \in S}} (V_{t,q} > 0) \right)$$

depending on whether $\bowtie \in \{\leq, <\}$ or $\bowtie \in \{\geq, >\}$, respectively. Now we can finally express that the $Z_{t,q}$ variables form a solution of Δ_{acc} or Δ_{rej} . Let

$$\begin{aligned} Sat &= \bigwedge_{\substack{r \in S \\ p \in Q}} (\exists V_{t,q}. (AccRej \wedge V_{r,p})) \Rightarrow (Z_{r,p} = 1) \\ &\wedge \bigwedge_{\substack{r \in S_{Prob} \cup S_C \\ p \in Q}} \neg(\exists V_{t,q}. (AccRej \wedge V_{r,p})) \Rightarrow (Z_{r,p} = \sum_{\substack{r' \in S \\ p' \in Q}} T_{(r,p)(r',p')} \cdot Y_{r,r'} \cdot Z_{r',p'}) \\ &\wedge \bigwedge_{\substack{r \in S_{Agents \setminus C} \\ p \in Q}} \neg(\exists V_{t,q}. (AccRej \wedge V_{r,p})) \Rightarrow (\bigwedge_{\substack{r' \in S \\ p' \in Q}} T_{(r,p)(r',p')} = 1 \Rightarrow Z_{r,p} \geq Z_{r',p'}) \end{aligned}$$

Now we put

$$Sol = \exists X_{t,i}. \exists T_{(t,q)(t',q')}. (Eta \wedge Tran \wedge Sat).$$

The correctness of our construction follows by verifying that all subformulae have the intended meaning, which is straightforward. \square

The following theorem is a direct consequence of Theorem 2 and Lemma 3.

Theorem 3 *Let $\ell \geq 0$. Then*

- *the model-checking problem for $SGL(\Sigma_\ell)$ is in EXPTIME and Σ_ℓ -hard for MR semantics; the hardness result holds even for a fixed $SGL(\Sigma_\ell)$ formula;*
- *the model-checking problem for $SGL(\Pi_\ell)$ is in EXPTIME and Π_ℓ -hard for MR semantics; the hardness result holds even for a fixed $SGL(\Pi_\ell)$ formula;*
- *the model-checking problem for $SGL(\Delta_\ell)$ is in EXPTIME for MR semantics.*

Finally, the model-checking problem for SGL is in EXPSPACE and PSPACE-hard for MR semantics.

The lower bounds of Theorem 3 are just inherited from Theorem 2. The upper bounds are obtained by analyzing the size and structure of the formula $\tau(s, \Phi)$ constructed in the proof of Lemma 3, and applying known results about the complexity of $(\mathcal{R}, *, +, \leq)$ and its fragments. Note that the size of $\tau(s, \Phi)$ is *polynomial* in the size of \mathcal{M} and Φ , and the quantifier alternation depth of $\tau(s, \Phi)$ is fixed for every $\ell \geq 0$ (after pushing all negations inside). The general upper bound for $(\mathcal{R}, *, +, \leq)$ is EXPSPACE. The existential fragment of $(\mathcal{R}, *, +, \leq)$ can be decided in polynomial space [11], and every fragment of $(\mathcal{R}, *, +, \leq)$ obtained by restricting the quantifier alternation depth to some fixed level is solvable in exponential time [18]. Thus, we obtain the upper bounds of Theorem 3. Also note that for $\ell = 0$, the upper bounds trivially improve to P (cf. Proposition 1), and for $\ell = 1$, they improve to PSPACE.

4.3 FR, FD, HR, and HD strategies

For history-dependent strategies, the SGL model checking problem becomes undecidable. This follows immediately from the undecidability result for $1\frac{1}{2}$ -player games and PCTL stated in [9]. More precisely, [9] yields the undecidability of the

model checking problem for PMG with a singleton agent set $\{a\}$ and SGL formulae of the form $\langle\langle a \rangle\rangle \Phi$ where Φ is a PCTL formula. This holds for FR, FD, HR, and HD semantics. For HR and HD semantics, the above problem is even shown to be *highly* undecidable, i.e., beyond the arithmetical hierarchy. Thus, we obtain the following:

Theorem 4 *The SGL model checking problem is undecidable for FR, FD, HR, and HD semantics. This result holds even for the $SGL(\Sigma_1)$ and $SGL(\Pi_1)$ fragments.*

The results of [9] do not apply to the qualitative fragment of SGL. In [8], it was shown that the controller synthesis problem for finite-state MDPs and qualitative PECTL* objectives is decidable. Since the underlying argument is quite involved, the question whether this result can be generalized to qualitative SGL is postponed to future work.

5 Conclusion

We introduced a new stochastic game logic (SGL) interpreted over probabilistic multi-player games (PMG). It combines features of alternating time logic (ATL), probabilistic computation tree logic and extended temporal logics. Our logic uses an existential strategy quantifier $\langle\langle \cdot \rangle\rangle$ that, unlike in ATL, propagates the chosen strategies to the subformulae. This enables us to state game properties like “*player B can react to the strategy chosen by player A*”. Whereas the ATL model checking problem is known to be solvable by a polynomially time-bounded algorithm [2], modifying the semantics of the $\langle\langle \cdot \rangle\rangle$ operator so that the strategy decisions are propagated to the subformulae makes the model checking problem PSPACE-hard. The main results of this paper can be summarized as follows.

The model-checking problem for finite state PMG and (full) SGL is

- undecidable for HR and HD strategies,
- PSPACE-complete for MD strategies,
- PSPACE-hard and in EXPSPACE for MR strategies.

The model-checking problem for finite state MPG and the qualitative fragment of SGL is PSPACE-complete for MD and MR strategies.

The decidability of the qualitative fragment of SGL with respect to history dependent strategies remains open.

Acknowledgement

Tomáš Brázdil and Antonín Kučera are supported by the Czech Science Foundation, grant No. P202/12/G061. Christel Baier is supported by the projects PROBPOR and ROCKS funded by the German Research Foundation.

References

1. Thomas Ågotnes, Valentin Goranka and Wojciech Jamroga. Alternating-time temporal logic with irrevocable strategies. *11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK'07)*, pages 15-24, 2007.
2. Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002.
3. Christel Baier, Tomáš Brázdil, Marcus Größer, and Antonín Kučera. Stochastic Game Logic. In *Fourth International Conference on Quantitative Evaluation of Systems (QEST'07)*, pages 227-236, IEEE Computer Society Press, 2007.
4. Christel Baier, Marcus Größer, Martin Leucker, Frank Ciesinski, and Benedikt Bollig. Controller synthesis for probabilistic systems. In *IFIP Worldcongress, Theoretical Computer Science*, 2004.
5. Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. MIT Press, 2008.
6. Michael Ben-Or, Dexter Kozen, and John Reif. The complexity of elementary algebra and geometry. *J. Comput. Syst. Sci.*, 32(2):251–264, 1986.
7. Bianco and de Alfaro. Model checking of probabilistic and nondeterministic systems. In *Proc. FSTTCS*, volume 15 of *Lecture Notes in Computer Science*, pages 499–513, 1995.
8. Tomáš Brázdil, Vojtěch Forejt, and Antonín Kučera. Controller Synthesis and Verification for Markov Decision Processes with Qualitative Branching Time Objectives. In *Proc. ICALP 2008*, volume 5126 of *Lecture Notes in Computer Science*, pages 148–159, 2008.
9. Tomáš Brázdil, Václav Brožek, Vojtěch Forejt, and Antonín Kučera. Stochastic games with branching-time winning objectives. In *Proc. LICS 2006*, pages 349-358, 2006.
10. Thomas Brihaye, Arnaud Da Costa Lopes, François Laroussinie and Nicolas Markey. ATL with Strategy Contexts and Bounded Memory, *International Symposium on Logical Foundations of Computer Science (LFCS)*, volume 5407 of *Lecture Notes in Computer Science*, pages 92-106, 2009.
11. John Canny. Some Algebraic and Geometric Computations in PSPACE. In *Proc. STOC'88*, pages 460–467, 1988.
12. Krishnendu Chatterjee, Thomas A. Henzinger and Nir Piterman. Strategy Logic. *Information and Computation* Volume 208, Issue 6, pages 677-693, 2010.
13. Edmund M. Clarke, Orna Grumberg, and Robert P. Kurshan. A synthesis of two approaches for verifying finite state concurrent systems. *J. Log. Comput.*, 2(5):605–618, 1992.
14. Edmund M. Clarke, Doron Peled and Orna Grumberg. *Model Checking*. MIT Press, 1999.
15. Stephen A. Cook. The Complexity of Theorem-Proving Procedures. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 151–158, 1971.
16. Costas Courcoubetis and Mihalis Yannakakis. The complexity of probabilistic verification. *Journal of the ACM*, 42(4):857–907, July 1995.
17. Luca De Alfaro. *Formal Verification of Probabilistic Systems*. Ph.D.Thesis. Stanford University, Department of Computer Science. 1997.
18. D. Grigoriev. Complexity of Deciding Tarski Algebra. *Journal of Symbolic Computation*, 5(1–2):65–108, 1988.
19. Luca de Alfaro. Formal verification of probabilistic systems. Thesis CS-TR-98-1601, Stanford University, Department of Computer Science, June 1998.
20. Luca de Alfaro and Thomas A. Henzinger. Concurrent omega-regular games. In *Proc. LICS*, pages 141–154, 2000.
21. Luca de Alfaro, Thomas A. Henzinger, and Orna Kupferman. Concurrent reachability games. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 564–575, 1998.
22. Sergiu Hart, Micha Sharir, and Amir Pnueli. Termination of probabilistic concurrent programs. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 5(3):356–380, July 1983.
23. John E. Hopcroft and Jeffrey D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, 1979.
24. Antonín Kučera and Oldřich Stražovský. On the controller synthesis for finite-state Markov decision processes. In *Proc. FSTTCS*, volume 3821 of *Lecture Notes in Computer Science*, pages 541–552, 2005.
25. Albert R. Meyer and Larry Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential space. In *Proceedings of the 13th Annual Symposium on Switching and Automata Theory*, pages 125–129, IEEE, 1972.

26. Christos Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
27. Shmuel Safra. On the complexity of ω -automata. In *29th Annual Symposium on Foundations of Computer Science*, pages 319–327, White Plains, New York, 24–26 October 1988. IEEE.
28. Roberto Segala and Nancy Lynch. Probabilistic simulations for probabilistic processes. *Lecture Notes in Computer Science*, 836:481–496, 1994.
29. Alfred Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5(2):285–309, 1955.
30. Wolfgang Thomas. Computation tree logic and regular omega-languages. *Lecture Notes in Computer Science*, 354:690–713, 1988.
31. Wolfgang Thomas. Automata on infinite objects. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, chapter 4, pages 133–191. Elsevier Science Publishers B. V., 1990.
32. Moshe Y. Vardi. Probabilistic linear-time model checking: An overview of the automata-theoretic approach. *Lecture Notes in Computer Science*, 1601:265–276, 1999.
33. Moshe Y. Vardi and Pierre Wolper. Yet another process logic (preliminary version). *Lecture Notes in Computer Science*, 164:501–512, 1983.