What Agents Can Probably Enforce

Nils Bulling and Wojciech Jamroga

Department of Informatics, Clausthal University of Technology, Germany
bulling,wjamroga@in.tu-clausthal.de

Abstract. Alternating-time Temporal Logic (**ATL**) is probably the most influential logic of strategic ability that has emerged in recent years. The idea of **ATL** is centered around cooperation modalities: $\langle\langle A\rangle\rangle\gamma$ is satisfied if the group A of agents has a collective strategy to enforce temporal property γ against *the worst* possible response from the other agents. So, the semantics of **ATL** shares the "all-or-nothing" attitude of many logical approaches to computation.
Such an assumption seems appropriate in some application areas (life-critical systems, security protocols, expensive ventures like space missions). In many cases, however, one might be satisfied if the goal is achieved with reasonable likelihood. In this paper, we try to soften the rigorous notion of success that underpins **ATL**.

1 Introduction

Alternating-time Temporal Logic (**ATL**) [1] is probably the most influential logic of strategic ability that has emerged in recent years. The idea of **ATL** is centered around *cooperation modalities* $\langle\langle A\rangle\rangle$: $\langle\langle A\rangle\rangle\gamma$ is satisfied if the group of agents A has a collective strategy to enforce temporal property γ. That is, $\langle\langle A\rangle\rangle\gamma$ holds if A has a strategy that succeeds to make γ true against *the worst* possible response from the opponents. So, the semantics of **ATL** shares the "all-or-nothing" attitude of many logical approaches to computation, justified by von Neumann's maximin evaluation of strategies in classical game theory [12]. Such an assumption does seem appropriate in some application areas. For life-critical systems, security protocols, and expensive ventures like space missions it is indeed essential that nothing can go wrong (provided that the assumptions being made are correct). In many cases, however, one might be satisfied if the goal is achieved with reasonable likelihood. Also, it does not seem right to assume that the rest of the agents will behave in the most hostile and destructive way; they may be friendly, indifferent, or simply not powerful enough to do it (for example, due to incomplete knowledge). Thus, to evaluate available strategies, a finer measure of success is needed that takes into account the possibility of a non-adversary response.

A naive (but nevertheless appealing) idea is to evaluate a strategy s by counting against how many opponents' responses it succeeds. If the ratio we get is, say, 50%, we can say that s succeeds in 50% of the cases. Note that this approach is underpinned by the assumption that each response from the other agents is equally likely; that is, we in fact assume that those agents play at random. Putting it in another way: As we do not have any information about the future strategy of the opponents, we assume a uniform distribution over all possible response strategies. On the other hand, assuming the uniform distribution is too strong in many scenarios, where the "proponents" may have a more specific idea of what the opponents will do (obtained e.g. by statistical analysis and/or learning). In order to properly address the issue, we introduce modalities $\langle\langle A\rangle\rangle_{\omega}^{p}\gamma$ that say that *agents A have a collective strategy to enforce γ with probability of at least $p \in [0,1]$, assuming that the expected behavior of the opponents is described by the prediction symbol ω.*

2 Preliminaries

2.1 Alternating-time Temporal Logic

Alternating-time temporal logic (**ATL**) [1] enables reasoning about temporal properties and strategic abilities of agents.

Definition 1 **($\mathcal{L}_{ATL}$).** *Let $\mathbb{A}\mathrm{gt} = \{a_1, \dots, a_k\}$ be a nonempty finite set of all agents, and Π be a set of propositions (we use $\mathsf{p}, \mathsf{q}, \mathsf{r}, \dots$ to denote propositions). $\mathcal{L}_{ATL}(\mathbb{A}\mathrm{gt}, \Pi)$ is defined by the following grammar (where $A \subseteq \mathbb{A}\mathrm{gt}$):*

$$\varphi ::= \mathsf{p} \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle A\rangle\rangle\gamma \quad \textit{where} \quad \gamma ::= \bigcirc\varphi \mid \Box\varphi \mid \varphi\,\mathcal{U}\,\varphi.$$

Formulae φ are called state formulae*, and formulae γ* path formulae*.*

Informally, $\langle\langle A\rangle\rangle\gamma$ expresses that agents A have a *collective strategy to enforce γ*. **ATL** formulae include the usual temporal operators: $\bigcirc$ ("next"), $\Box$ ("always from now on"), and $\mathcal{U}$ (strict "until"). Additionally, $\Diamond$ ("sometime in the future") can be defined as $\Diamond\gamma \equiv \top\,\mathcal{U}\,\gamma$. The semantics of **ATL** is defined by *concurrent game structures*.

Definition 2 (CGS). *A* concurrent game structure (**CGS**) *is a tuple $M = \langle \mathbb{A}\mathrm{gt}, Q, \Pi, \pi, Act, d, o\rangle$, consisting of: a set $\mathbb{A}\mathrm{gt} = \{a_1, \dots, a_k\}$ of* agents*; a set Q of* states*; a set Π of* atomic propositions*; a* valuation of propositions *$\pi : Q \to \mathcal{P}(\Pi)$; and a finite set Act of* actions*. Function $d : \mathbb{A}\mathrm{gt} \times Q \to \mathcal{P}(Act)$ indicates the actions available to agent $a \in \mathbb{A}\mathrm{gt}$ in state $q \in Q$. We often write $d_a(q)$ instead of $d(a,q)$, and use $d(q)$ to denote the set $d_{a_1}(q) \times \dots \times d_{a_k}(q)$ of* action profiles *in state q. Finally, o is a* transition function *which maps each state $q \in Q$ and action profile $\overrightarrow{\alpha} = \langle \alpha_1, \dots, \alpha_k\rangle \in d(q)$ to another state $q' = o(q, \overrightarrow{\alpha})$.*

In this paper, we will only deal with finite models, i.e., we assume that the sets of states and actions in each model are finite.

A *(memoryless) strategy* $s_a : Q \to Act$ is a conditional plan that specifies what $a \in \mathbb{A}\mathrm{gt}$ is going to do for every possible situation.[1] We denote the set of such functions by Σ_a. A *collective strategy* s_A for team $A \subseteq \mathbb{A}\mathrm{gt}$ specifies an individual strategy for each agent $a \in A$; the set of A's collective strategies is given by $\Sigma_A = \prod_{a \in A} \Sigma_a$. A *path* $\lambda = q_0 q_1 \dots$ in model M is an infinite sequence of states that can be effected by subsequent transitions. We use $\lambda[n]$ to denote the nth state in λ; $\lambda[i..j]$ denotes the subpath of λ between positions i and j (also for $j = \infty$). $\Lambda(q)$ denotes the set of all the paths starting in state q. Function $out(q, s_A)$ returns the set of all paths that may result from agents A executing strategy s_A from state q onward.

---

[1] This is a deviation from the original semantics of **ATL** [1], where strategies assign agents' choices to *sequences* of states. We note, however, that both types of strategies yield equivalent semantics for "vanilla" **ATL** [9].

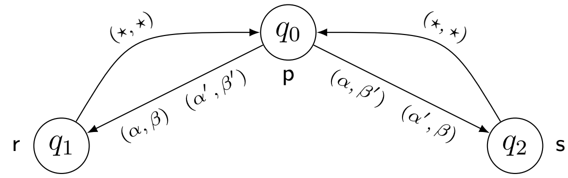

Fig. 1. A simple **CGS** $M_1 = \langle\{1,2\},\{q_0,q_1,q_2\},\{\mathsf{r},\mathsf{s}\},\pi,\{\alpha,\alpha',\beta,\beta'\},d,o\rangle$; π, d, and o can be read off from the figure. By $\star$ we refer to any possible action.

Definition 3 (Semantics of ATL). *Let M be a* ***CGS****, q a state in M, and λ a path in M. The semantics is given by the satisfaction relation $\models$ as follows:*

$M,q \models \mathsf{p}$ *iff* $\mathsf{p} \in \pi(q)$ *(for $\mathsf{p} \in \Pi$);*
$M,q \models \neg\varphi$ *iff* $M,q \not\models \varphi$*;*
$M,q \models \varphi_1 \wedge \varphi_2$ *iff* $M,q \models \varphi_1$ *and* $M,q \models \varphi_2$*;*
$M,q \models \langle\langle A\rangle\rangle\gamma$ *iff there is a collective strategy s_A such that, for every $\lambda \in out(q,s_A)$ we have $M,\lambda \models \gamma$;*
$M,\lambda \models \bigcirc\varphi$ *iff* $M,\lambda[1..\infty] \models \varphi$*;*
$M,\lambda \models \Box\varphi$ *iff* $M,\lambda[i..\infty] \models \varphi$ *for all $i \in \mathbb{N}_0$;*
$M,\lambda \models \varphi_1 \,\mathcal{U}\, \varphi_2$ *iff* $M,\lambda[i..\infty] \models \varphi_2$ *for some $i \geq 0$, and $\lambda[j..\infty] \models \varphi_1$ for all $0 \leq j \leq i$.*

Example 1. Consider a simple two-agent scenario depicted in Figure 1. Agent 1 (resp. 2) can perform actions α and α' (resp. β and β'). For example, strategy profile (α,β), performed in q_0, leads to state q_1 in which r holds. Agent 1 can enforce neither r nor s on its own: $M_1,q_0 \models \neg\langle\langle 1\rangle\rangle\bigcirc\mathsf{r} \wedge \neg\langle\langle 1\rangle\rangle\bigcirc\mathsf{s}$, and neither can agent 2. However, the agents can cooperate to determine the outcome: $M_1,q_0 \models \langle\langle 1,2\rangle\rangle\bigcirc\mathsf{r} \wedge \langle\langle 1,2\rangle\rangle\bigcirc\mathsf{s}$.

2.2 Probability Theory

In this section we recall some basic notions from probability theory. Let X be a non-empty set and let $\mathcal{F} \subseteq \mathcal{P}(X)$ be a set of subsets. $\mathcal{F}$ is called a *(set) algebra over X* iff: (i) $\emptyset \in \mathcal{F}$; (ii) if $A \in \mathcal{F}$ then also $\bar{A} := X \setminus A \in \mathcal{F}$; and (iii) if $A,B \in \mathcal{F}$ then also $A \cup B \in \mathcal{F}$. $\mathcal{F}$ is called a *σ-algebra* if additionally to (i-iii) it also holds (iv) $\bigcup_{i=1}^{\infty} A_i \in \mathcal{F}$ for all $A_1, A_2, \dots \in \mathcal{F}$.

Let $\mathcal{S}$ be a σ-algebra over X. We say that a function $\mu : \mathcal{S} \to \mathbb{R}$ is a *measure (on $\mathcal{S}$)* iff it is non-negative, i.e. $\mu(A) \geq 0$ for all $A \in \mathcal{S}$, and σ-additive, i.e. $\mu(\bigcup_{i=1}^{\infty} A_i) = \sum_{i=1}^{\infty}\mu(A_i)$ whenever each $A_i \in \mathcal{S}$. Finally, we say that the measure μ is a *probability measure* if $\mu(X) = 1$ and call the triple $(X,\mathcal{S},\mu)$ a *probability space*. By $\Xi(\mathcal{S})$ we denote the set of all probability measures over $\mathcal{S}$.

Note that whenever X is finite it is sufficient to define the probabilities of the basic elements $x \in X$. Then, the probability of an event $E \subseteq X$ is given by the sum of the basic probabilities: $\mu(E) = \sum_{x\in E}\mu(\{x\})$, and the corresponding probability measure is uniquely determined over the σ-algebra $\mathcal{P}(X)$. In such cases, we can also write

$\mu(x)$ instead of $\mu(\{x\})$ and $\Xi(X)$ instead of $\Xi(\mathcal{P}(X))$, and also refer to a probability measure over $\mathcal{P}(X)$ as probability measure over X.

3 ATL with Probability

In this section we propose and discuss our new logic **pATL** (***ATL** with probabilistic success*). Firstly, we define the syntax and the semantics on an abstract level. Then, we instantiate the semantics for two different ways of modeling the opponents' behavior: namely, by mixed and behavioral memoryless strategies. Finally, we discuss the relation of **pATL** to "pure" **ATL**.

3.1 Syntax

In **pATL**, cooperation modalities $\langle\langle A\rangle\rangle$ of the original **ATL** are replaced with a richer family of strategic modalities $\langle\langle A\rangle\rangle_{\omega}^{p}$.

Definition 4 **($\mathcal{L}_{pATL}$).** *The basic language $\mathcal{L}_{pATL}(\mathbb{A}\mathrm{gt}, \Pi, \Omega)$ is defined over the nonempty sets Π of propositions, $\mathbb{A}\mathrm{gt} = \{a_1, \dots, a_k\}$ of agents, and Ω of prediction symbols. The language consists of all* state formulae φ *defined as follows:*

$$\varphi ::= \mathsf{p} \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle A\rangle\rangle_{\omega}^{p}\gamma; \quad \textit{where} \quad \gamma ::= \bigcirc\varphi \mid \Box\varphi \mid \varphi\,\mathcal{U}\,\varphi,$$

$\omega \in \Omega$*, and* $p \in [0, 1]$. *Additional temporal operators are defined as before.*

We use p, ω, a, A to refer to a typical proposition, a prediction symbol, an agent, and a group of agents, respectively. The informal reading of formula $\langle\langle A\rangle\rangle_{\omega}^{p}\gamma$ is: *Team A can bring about γ with success level of at least p when the opponents behave according to ω*. The prediction symbols are used to assume some "predicted behavior" of the opponents.

3.2 Semantics: The Abstract Framework

Now we define the semantics of **pATL** in a generic way before considering more concrete settings. Models for **pATL** extend concurrent game structures with *prediction denotation functions* which, given a group of agents, assign prediction symbols to *predicted behaviors*. We use a non-empty set $\mathcal{BH}$ to refer to all possible predicted behaviors. There are several sensible ways how $\mathcal{BH}$ can be instantiated: *Mixed* and *behavioral strategies* provide two well-known possibilities (cf. Sections 3.3 and 3.4).

Definition 5 **(Prediction denotation function).** *Let $\mathcal{BH}$ be a non-empty set representing possible (probabilistic) behaviors of the agents. A* prediction denotation function *is a function $[\![\cdot]\!] : \Omega \times \mathcal{P}(\mathbb{A}\mathrm{gt}) \to \mathcal{BH}$ where $[\![\omega, A]\!]$ denotes a (probabilistic) prediction of A's behavior according to the prediction symbol $\omega \in \Omega$. We write $[\![\omega]\!]_A$ for $[\![\omega, A]\!]$.*

Definition 6 **(Models of pATL).** *A* concurrent game structure with probability (**CGSP**) *is given by a tuple $M = \langle \mathbb{A}\mathrm{gt}, Q, \Pi, \pi, Act, d, o, \Omega, [\![\cdot]\!]\rangle$ where $\langle \mathbb{A}\mathrm{gt}, Q, \Pi, \pi, Act, d, o\rangle$ is a **CGS**, Ω is a set of prediction symbols, and $[\![\cdot]\!]$ is a prediction denotation function.*

Our semantics of $\langle\langle A\rangle\rangle_{\omega}^{p}$ is based on the generic notion of a *success measure*. The measure indicates "how successful" a group of agents is wrt property γ if the opponents behave according to their predicted behavior. The semantics of **pATL**, parameterized by a success measure, updates the **ATL** semantics from the previous section by replacing the rule for the cooperation modalities.

Definition 7 (Success measure). *A* success measure $success$ *is a function that takes a strategy of the proponents* s_A*, a probabilistic prediction* $[\![\omega]\!]_{\mathbb{A}\mathrm{gt}\setminus A}$ *of the opponents' behavior, the current state of the system* q*, and a* ***pATL*** *path formula* γ *and returns a score* $success(s_A, [\![\omega]\!]_{\mathbb{A}\mathrm{gt}\setminus A}, q, \gamma)$ *from* $[0, 1]$.

Definition 8 (Semantics of pATL). *Let* M *be a* ***CGSP****. The semantics of* ***pATL*** *updates the clauses from Definition 3 by replacing the clause for* $\langle\langle A\rangle\rangle$ *with the following:*

$M, q \models \langle\langle A\rangle\rangle_{\omega}^{p}\gamma$ *iff there is* $s_A \in \Sigma_A$ *such that* $success(s_A, [\![\omega]\!]_{\mathbb{A}\mathrm{gt}\setminus A}, q, \gamma) \geq p$.

Various success measures may prove appropriate for different purposes; they inherently depend on the type of the prediction denotation functions and therewith on the possible predicted behaviors represented by $\mathcal{BH}$.

3.3 Opponents' Play: Mixed Strategies

As the first instantiation of the generic framework, we consider *mixed memoryless strategies* which are probability distributions over pure memoryless strategies of the opponents. This notion of behavior fits well our initial intuition of *counting the favorable opponents' responses* in order to determine the success level of a strategy.

Definition 9 (Mixed memoryless strategy). *A* mixed memoryless strategy *(mms)* σ_A *for* $A \subseteq \mathbb{A}\mathrm{gt}$ *is a probability measure over* $\mathcal{P}(\Sigma_A)$.

Definition 10 (mms denotation function). *A* mms denotation function *is a prediction denotation function with* $\mathcal{BH} = \bigcup_{A\subseteq\mathbb{A}\mathrm{gt}} \Xi(\Sigma_A)$*, such that* $[\![\omega]\!]_A \in \Xi(\Sigma_A)$*.* $[\![\omega]\!]_A(s)$ *denotes the probability that* s *will be played by* A *according to the prediction symbol* ω.

In this paper, we take the success measure of a mms wrt property γ to be the *expected probability* of making γ true. For this purpose, we first define the *outcome of a strategy*.

Definition 11 (Outcome of a strategy against a mms). *The* outcome *of strategy* s_A *against a mixed memoryless strategy* $\sigma_{\mathbb{A}\mathrm{gt}\setminus A}$ *at state* q *is the probability measure over* $\Lambda(q)$ *given by:*

$$\mathcal{O}(s_A, \sigma_{\mathbb{A}\mathrm{gt}\setminus A}, q)(\lambda) := \sum_{t\in Resp(s_A,\lambda)} \sigma_{\mathbb{A}\mathrm{gt}\setminus A}(t)$$

where $Resp(s_A, \lambda) = \{t \in \Sigma_{\mathbb{A}\mathrm{gt}\setminus A} \mid \lambda \in out(q, \langle s_A, t\rangle)\}$ *is the set of all response strategies* t *of the opponents, that, together with* A*'s strategy* s_A*, result in path* λ.[2]

[2] Note that for a deterministic strategy profile $\langle s_A, t_{\mathbb{A}\mathrm{gt}\setminus A}\rangle$ the outcome set contains exactly one path.

Thus, $\mathcal{O}(s_A, \sigma_{\mathbb{A}\mathrm{gt}\setminus A}, q)(\lambda)$ sums up the probabilities of all responses in $Resp(s_A, \lambda)$, for each path λ. In consequence, $\mathcal{O}(s_A, \sigma_{\mathbb{A}\mathrm{gt}\setminus A}, q)(\lambda)$ denotes the probability that the opponents will play a strategy resulting in λ. Note also that, when memoryless strategies are played, the same action vector is performed every time a particular state is revisited, which restricts the set of paths than can occur.

Definition 12 (Minimal periodic path, $\Lambda^{mp}(q)$). *We say that a path $\lambda \in \Lambda(q)$ is* minimal periodic *if, and only if, the path can be written as $\lambda = \lambda[0, j]\lambda[j+1, i] \ldots \lambda[j+1, i]$ where $i \in \mathbb{N}_0$ is the* minimal *natural number such that there is some $j < i$ and $\lambda[i] = \lambda[j]$. The set of all minimal periodic paths starting in q is denoted by $\Lambda^{mp}(q)$. We note that, for a finite model, the set $\Lambda^{mp}(q)$ consists of only finitely many paths.*

Proposition 1. *$\mathcal{O}(s_A, \sigma_{\mathbb{A}\mathrm{gt}\setminus A}, q)$ is a probability measure over $\Lambda(q)$ and over $\Lambda^{mp}(q)$.*

Proof. That $\mathcal{O}(s_A, \cdot, q)$ is non-negative follows from the fact that $\sigma_{\mathbb{A}\mathrm{gt}\setminus A}(t) \geq 0$ for all response strategies t. It is easy to see that all non minimal periodic paths have probability zero since we consider memoryless strategies only. This implies that there are only finitely many paths with non-zero probability. Thus, $\mathcal{O}(s_A, \sigma_{\mathbb{A}\mathrm{gt}\setminus A}, q)$ is σ-additive, and the following holds: $\mathcal{O}(s_A, \sigma_A, q)(\Lambda(q)) = \mathcal{O}(s_A, \sigma_A, q)(\Lambda^{mp}(q)) = \sum_{\lambda \in \Lambda^{mp}(q)} \sum_{t \in Resp(s_A, \lambda)} \sigma_B(t) = \sum_{t \in \hat{Resp}(s_A)} \sigma_B(t)$ where $\hat{Resp}(s_A)$ consists of all strategies $t \in \Sigma_B$ such that there is a path $\lambda \in \Lambda^{mp}(q)$ with $\lambda \in out(q, \langle s_A, t \rangle)$. But then $\hat{Resp}(s_A) = \Sigma_B$ and thus the sum is equal to 1. $\square$

Definition 13 (Success measure with mms). *The success measure against mixed memoryless strategies is defined as below:*

$$success(s_A, \sigma_{\mathbb{A}\mathrm{gt}\setminus A}, q, \gamma) = \sum_{\lambda \in \Lambda(q)} holds_\gamma(\lambda) \cdot \mathcal{O}(s_A, \sigma_{\mathbb{A}\mathrm{gt}\setminus A}, q)(\lambda),$$

$$\text{where } holds_\gamma(\lambda) = \begin{cases} 1 & \text{if } M, \lambda \models \gamma \\ 0 & \text{else.} \end{cases}$$

Function $holds_\gamma : \Lambda \to \{0, 1\}$ can be seen as a characteristic function of the path formula γ: It indicates, for each path λ, whether γ holds on λ or not.

By Proposition 1, $success(s_A, \sigma_{\mathbb{A}\mathrm{gt}\setminus A}, q, \gamma)$ is indeed an expected value, and it is actually defined by a *finite* sum. Moreover, measuring the success of strategy s_A by counting the favorable vs. all responses of the opponents is a special case, obtained by setting $[\![\omega]\!]_{\mathbb{A}\mathrm{gt}\setminus A}$ to the uniform probability distribution over $\Sigma_{\mathbb{A}\mathrm{gt}\setminus A}$.

Example 2. Consider the system from Example 1. We have discussed in Section 2.1 that 1 is able to enforce neither r nor s. However, it might be the case that additional information about 2's behavior is available, namely that 2 plays action β' more often than β (say, seven out of every ten times). This kind of observation can be formalized by a probability measure σ over $\{\beta, \beta'\}$ with $\sigma(\beta) = 0.3$ and $\sigma(\beta') = 0.7$.

Using **ATL**, it was not possible to state any "positive" fact about 1's power. **pATL** allows a finer-grained analysis. We can now state that 1 can enforce any outcome (r or s) with probability at least 0.7. Formally, let $[\![\omega]\!]_2 = \sigma$. We have that $M, q_0 \models \langle\langle 1 \rangle\rangle_\omega^{0.7} \bigcirc \mathsf{r} \wedge \langle\langle 1 \rangle\rangle_\omega^{0.7} \bigcirc \mathsf{s}$. If 1 desires r, he should play α' since $\langle \alpha', \beta' \rangle$ leads to r; otherwise the agent should select action α in q_0.

3.4 Opponents' Play: Behavioral Strategies

In this section we present an alternative instantiation of the semantics, where the prediction of opponents' play is based on the notion of *behavioral strategies* (which follows the Markovian assumption that the probability of taking an action depends only on the state where it is executed). We show that the semantics is well defined for **pATL**.

Definition 14 (Behavioral strategy). *A* behavioral strategy *for $A \subseteq \mathbb{A}\mathrm{gt}$ is a function $\beta_A : Q \to \bigcup_{q\in Q} \Xi(d_A(q))$ such that $\beta_A(q)$ is a probability measure over $d_A(q)$, i.e., $\beta_A(q) \in \Xi(d_A(q))$. We use $\mathcal{B}_A$ to denote the set of behavioral strategies of A.*

Definition 15 (Behavioral strategy denotation function). *A* behavioral strategy denotation function *is a prediction denotation function with $\mathcal{BH} = \bigcup_{A\subseteq\mathbb{A}\mathrm{gt}} \mathcal{B}_A$, such that $[\![\omega]\!]_A \in \mathcal{B}_A$. Thus, $[\![\omega]\!]_A(q)(\overrightarrow{\alpha})$ denotes the probability that the collective action $\overrightarrow{\alpha}$ will be played by agents A in state q according to the prediction symbol ω.*

As in the case of mixed memoryless strategies (cf. Definition 11), the outcome of a strategy against behavioral predictions is a probability measure over paths. However, the setting is more complicated now. For mixed predictions it suffices to consider a probability distribution over the finite set of pure strategies which induces a probability measure over the set of paths. Indeed, only finite prefixes of paths, namely the non-looping parts, are relevant for the outcome. For behavioral strategies, actions (rather than strategies) are probabilistically determined, which makes it possible for different actions to be executed when the system returns to a previously visited state. Thus, the probability of a specific set of paths depends on the *whole* paths that belong to the set.

To define the outcome of a behavioral strategy we first need to define the probability space induced by the probabilities of one-step transitions; to this end, we follow the construction from [8]. Recall that $\Lambda(q)$ denotes the set of all infinite paths starting in q. The probability of a set of paths is defined inductively by consistently assigning probabilities to all finite initial segments (prefixes) of a path. The intuition is that prefix h can be used to represent the set of infinite paths that extend h. By imposing closure wrt complement and (countable) union, we obtain a probability measure for some sets of paths. Of course, not *every* set of paths can be constructed this way, but we prove (in Proposition 2) that all the relevant sets can.

We use $\Lambda^n(q)$ to denote the set of finite prefixes (histories) of length n of the paths from $\Lambda(q)$; note that $\Lambda^n(q)$ is always finite for finite models. Now, we define $\mathcal{F}^n(q)$ and $\mathcal{F}(q)$ to be the following sets of subsets of $\Lambda(q)$:

$$\mathcal{F}^n(q) := \big\{\{\lambda \mid \lambda[0, n-1] \in T\} \quad \big| \quad T \subseteq \Lambda^n(q)\big\} \quad \text{and} \quad \mathcal{F}(q) := \bigcup_{n=0}^{\infty} \mathcal{F}^n(q).$$

That is, for each set of prefixes $T \subseteq \Lambda^n(q)$, the set $\mathcal{F}^n(q)$ includes the set of all their infinite extensions. Note that every $\mathcal{F}^n(q)$ is a σ-algebra. Each element S of $\mathcal{F}^n(q)$ (often called *cylinder set*) can be written as a finite union of *basic cylinder sets* $[h_i] := \{\lambda \in \Lambda(q) \mid h_i \leq \lambda\}$ where $h_i \in \Lambda^n(q)$ is a history of length n and $h_i \leq \lambda$ denotes that h_i is an initial prefix of λ; so, $S = \bigcup_i [h_i]$ for appropriate $h_i \in \Lambda^n(q)$. We use these basic cylinder sets to define an appropriate probability measure.

A basic cylinder set $[h_i]$ consists of all extensions of h_i; hence, the probability that one of h_i's extensions $\lambda \in [h_i]$ will occur is equal to the probability that h_i will take place. Given a strategy s_A and a behavioral response $\beta_{\mathbb{A}\mathrm{gt}\setminus A}$, the probability for $[h_i], h_i = q_0 \dots q_n$, is defined as the product of subsequent transition probabilities:

$$\nu^{s_A}_{\beta_{\mathbb{A}\mathrm{gt}\setminus A}}([h_i]) := \prod_{i=0}^{n-1} \sum_{\overrightarrow{\alpha} \in Act(s_A, q_i, q_{i+1})} \beta_{\mathbb{A}\mathrm{gt}\setminus A}(q_i)(\overrightarrow{\alpha})$$

where $Act(s_A, q_i, q_{i+1}) = \{\overrightarrow{\alpha} \in d_{\mathbb{A}\mathrm{gt}\setminus A}(q_i) \mid q_{i+1} = o(q_i, \langle s_A(q_i), \overrightarrow{\alpha}\rangle)\}$ consists of all action profiles which can be performed in q_i and which lead to q_{i+1} given the choices s_A of agents A. According to [8], function $\nu^{s_A}_{\beta_{\mathbb{A}\mathrm{gt}\setminus A}}$ is uniquely defined on $\mathcal{F}(q)$ and the restriction of $\nu^{s_A}_{\beta_{\mathbb{A}\mathrm{gt}\setminus A}}$ to $\mathcal{F}^n(q)$ is a measure on $\mathcal{F}^n(q)$ for each n. Still, $\mathcal{F}(q)$ is not a σ-algebra.

Therefore, we take $\mathcal{S}(q)$ to be the smallest σ-algebra containing $\mathcal{F}(q)$ and extend $\nu^{s_A}_{\beta_{\mathbb{A}\mathrm{gt}\setminus A}}$ to a measure on $\mathcal{S}(q)$ as follows:

$$\mu^{s_A}_{\beta_{\mathbb{A}\mathrm{gt}\setminus A}}(S) := \inf_{C \in \mathcal{H}(S)} \left\{ \nu^{s_A}_{\beta_{\mathbb{A}\mathrm{gt}\setminus A}} \left(\bigcup C \right) \right\}$$

where $S \in \mathcal{S}(q)$ and $\mathcal{H}(S)$ denotes the denumerable set of coverings of S by basic cylinder sets. That is, $\mathcal{H}(S)$ consists of sets $\{[h_1], [h_2], \dots\}$ such that $S \subseteq \bigcup_i [h_i]$. According to [8], we have that $(\Lambda(q), \mathcal{S}(q), \mu^{s_A}_{\beta_{\mathbb{A}\mathrm{gt}\setminus A}})$ is a probability space. Actually, $\mu^{s_A}_{\beta_{\mathbb{A}\mathrm{gt}\setminus A}}$ is the *unique* extension of $\nu^{s_A}_{\beta_{\mathbb{A}\mathrm{gt}\setminus A}}$ on $\mathcal{F}^n(q)$ to the σ-algebra $\mathcal{S}(q)$ [8, Theorem 1.19]; in particular, this means that both measures coincide on all sets from $\mathcal{F}(q)$. We refer to $\mu^{s_A}_{\beta_{\mathbb{A}\mathrm{gt}\setminus A}}$ as the *probability measure on $\mathcal{S}(q)$ induced by the pure strategy s_A and the behavioral strategy $\beta_{\mathbb{A}\mathrm{gt}\setminus A}$*.

Definition 16 (Success measure with behavioral memoryless strategies). *Like in the previous section, the success measure of strategy s_A wrt the formula γ is defined as the expected value of the characteristic function of γ (i.e., $holds_\gamma$) over $(\Lambda(q), \mathcal{S}(q), \mu^{s_A}_{\beta_{\mathbb{A}\mathrm{gt}\setminus A}})$.*

$$success(s_A, \beta_{\mathbb{A}\mathrm{gt}\setminus A}, q, \gamma) := E[holds_\gamma] = \int_{\Lambda(q)} holds_\gamma d\mu^{s_A}_{\beta_{\mathbb{A}\mathrm{gt}\setminus A}}.$$

Note that the formulation uses a *Lebesgue integral* over the σ-algebra $\mathcal{S}(q)$. Now we can show that the semantics of **pATL** with behavioral strategies is well-defined. We first prove that $holds_\gamma$ is $\mathcal{S}(q)$-measurable (i.e., every preimage of $holds_\gamma$ is an element of $\mathcal{S}(q)$ and thus can be assigned a measure); then, we show that $holds_\gamma$ is integrable.

Proposition 2. *Function $holds_\gamma$ is $\mathcal{S}(q)$-measurable and $\mu^{s_A}_{\beta_{\mathbb{A}\mathrm{gt}\setminus A}}$-integrable for any **pATL**-path formula γ.*

Proof. In particular, we have to show that $holds_\gamma^{-1}(A) := \{\lambda \in \Lambda(q) \mid holds_\gamma(\lambda) \in A\}$ is measurable for every $A \subseteq \{0, 1\}$ (i.e., $holds_\gamma^{-1}(A) \in \mathcal{S}(q)$). The cases $\emptyset$ and $\{0, 1\}$ are trivial. The case for $\{0\}$ is clear if we have shown it for $A = \{1\}$ (cf. property

(ii) of σ-algebras, Section 2.2). Therefore, let $f_\gamma := holds_\gamma^{-1}(\{1\})$. The proof proceeds by structural induction on γ.

I. Case "$\Box$": (i) Let $\gamma = \Box p$ where p is a propositional logic formula (e.g. $p = \mathsf{r} \wedge \neg \mathsf{s}$). We define $L_n^{\Box p} := \{\lambda \in \Lambda(q) \mid \forall i \in \mathbb{N}_0 (0 \leq i < n \rightarrow M, \lambda[i] \models p\}$. We have that each $L_n^{\Box p} \in \mathcal{F}^n(q) \subseteq \mathcal{S}(q)$ and that $\bigcap_{n \in \mathbb{N}} L_n^{\Box p} = f_\gamma$; hence, also that $f_\gamma \in \mathcal{S}(q)$ because of property (ii) and (iv) of σ-algebras (cf. Section 2.2). That f_γ is integrable follows from Lebesgue's Dominated Convergence Theorem: f_γ is measurable and $|f_\gamma|$ is bounded by the $\mu^{s_A}_{\beta_{\mathbb{A}\mathrm{gt} \setminus A}}$-integrable (constant) function 1. **(ii)** Let $\gamma = \Box \langle\langle B \rangle\rangle^p_{\omega'} \gamma'$ and suppose $f_{\gamma'}$ is already proven to be integrable. Then, $L_n^{\Box \langle\langle B \rangle\rangle^p_{\omega'} \gamma'}$ can be defined in the same way as above. **(iii)** Suppose that for each sub path formula γ' contained in φ_1 and φ_2 we have proven that $f_{\gamma'}$ is integrable, then L_n^γ can be defined in the same way as above for $\gamma = \Box \neg \varphi_1$ and $\gamma = \Box(\varphi_1 \wedge \varphi_2)$.

II. Case "$\bigcirc$": Similar to I(i) we define $L_n^{\bigcirc p} := \{\lambda \in \Lambda(q) \mid n > 1 \text{ and } M, \lambda[1] \models p\}$. Then, we have that $\bigcup_{n \in \mathbb{N}} L_n^{\bigcirc p} = f_{\bigcirc p} \in \mathcal{S}(q)$. The rest of the proof is done analogously to I.

III. Case "$\mathcal{U}$": Here, we also just consider the part corresponding to I(i). We set $L_n^{p \mathcal{U} q} := \{\lambda \in \Lambda(q) \mid \exists j (0 \leq j < n \rightarrow (M, \lambda[j] \models q \wedge \forall i \in \mathbb{N}_0 (0 \leq i < j \rightarrow M, \lambda[i] \models p))\}$; then, we have that $\bigcup_{n \in \mathbb{N}_0} L_n^{p \mathcal{U} q} = f_{p \mathcal{U} q} \in \mathcal{S}(q)$. $\Box$

Note that **pATL** with behavioral strategies can be seen as a special case of the multi-agent Markov Temporal Logic **MTL** from [7], since $\langle\langle A \rangle\rangle^p_\omega \gamma$ can be rewritten as the **MTL** formula $p \preccurlyeq (\mathbf{str}_{\mathbb{A}\mathrm{gt} \setminus A} \omega) \langle\langle A \rangle\rangle \gamma$.

3.5 Relationship to ATL

Firstly, we observe that an analogous success measure can be constructed for **ATL**:

$$success_{\text{ATL}}(s_A, q, \gamma) = \min_{\lambda \in out(s_A, q)} \{holds_\gamma(\lambda)\}.$$

Then, $M, q \models_{\text{ATL}} \langle\langle A \rangle\rangle \gamma$ iff there is a $s_A \in \Sigma_A$ such that $success_{\text{ATL}}(s_A, q, \gamma) = 1$. Thus, the abstract framework can be instantiated in a way that embraces the original semantics of **ATL**. Alternatively, we can try to embed **ATL** in **pATL** using the probabilistic success measures we have already defined.

Embedding ATL in pATL with Mixed Strategies We consider **pATL** with mixed memoryless strategies. The idea is to require that every response strategy has a non-zero probability. Note that a given **CGSP** M *induces* a **CGS** M' in a straightforward way: Only the set of prediction symbols and the prediction denotation function must be left out. In the following we will also use **CGSP**'s together with **ATL** formulae (without probability) by implicitly considering the induced **CGS**'s.

Theorem 1. *Let γ be an **ATL** path formula with no cooperation modalities, and let ω be a prediction symbol describing a mixed memoryless strategy such that $[\![\omega]\!]_{\mathbb{A}\mathrm{gt} \setminus A}(t) > 0$ for every $t \in \Sigma_{\mathbb{A}\mathrm{gt} \setminus A}$. Then, for all models M and states q in M it holds that:*

$$M, q \models_{\textit{ATL}} \langle\langle A \rangle\rangle \gamma \textit{ iff } M, q \models_{\textit{pATL}} \langle\langle A \rangle\rangle^1_\omega \gamma.$$

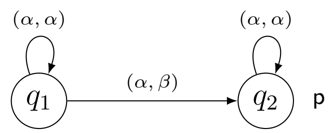

Fig. 2. CGS M_2 with actions α and β. The $\star \in \{\alpha, \beta\}$ refers to any of the two actions.

Proof. Let $\bar{A} := \mathbb{A}\mathrm{gt} \setminus A$ for $A \subseteq \mathbb{A}\mathrm{gt}$. "$\Rightarrow$": Assume that $s_A \in \Sigma_A$ and that for all $\lambda \in out(q, s_A)$ it holds that $M, \lambda \models \gamma$. Now suppose that $M, q \not\models_{\mathbf{pATL}} \langle\!\langle A \rangle\!\rangle^1_\omega \gamma$. In particular that would mean that $success(s_A, \sigma_A, q, \gamma) = \sum_{\lambda \in \Lambda(q)} holds_\gamma(\lambda) \cdot \sum_{t \in Resp(s_A, \lambda)} \sigma_{\bar{A}}(t) < 1$. This can only be caused by two cases: **(1)** There is a path $\lambda \in \Lambda(q)$ a strategy $t \in Resp(s_A, \lambda)$ with $\sigma_{\bar{A}}(t) > 0$ and $holds_\gamma(\lambda) = 0$. But then $\lambda \in out(q, s_A)$ contradicts the assumption that s_A is successfully.
(2) There is a strategy $t \in \Sigma_{\bar{A}}$ with $\sigma_{\bar{A}}(t) > 0$ and for all $\lambda \in \Lambda(q)$ it holds that $t \notin Resp(s_A, \lambda)$ (*). But there must be a path λ with $\{\lambda\} = out(q, (s_A, t))$ and thus $t \in Resp(s_A, \lambda)$, which contradicts (*).

"$\Leftarrow$": Assume that $s_A \in \Sigma_A$ and $success(s_A, \sigma_A, q, \gamma) = 1$. Suppose that there is a path $\lambda \in out(q, s_A)$ with $M, \lambda \not\models \gamma$. This means that strategy t with $out(q, (s_A, t)) = \{\lambda\}$ is in $Resp(s_A, \lambda)$ but plays no role in the calculation of the success value since $holds_\gamma(\lambda) = 0$. This contradicts the assumption that $success(s_A, \sigma_A, q, \gamma) = 1$. □

Condition $[\![\omega]\!]_{\mathbb{A}\mathrm{gt} \setminus A}(t) > 0$ ensures that no "bad response" of the opponents is neglected because of zero probability. Since we only deal with finite models, the uniform distribution over Σ_A is always well defined.

Corollary 1. *Let* $\mathsf{u_A}$ *be a prediction symbol that denotes the uniform distribution over strategies of the agents in A, and let* $tr(\varphi)$ *replace all occurrences of* $\langle\!\langle A \rangle\!\rangle$ *by* $\langle\!\langle A \rangle\!\rangle^1_{\mathsf{u}_{\mathbb{A}\mathrm{gt}\setminus \mathsf{A}}}$ *in* φ*. Then,* $M, q \models_{\mathbf{ATL}} \varphi$ *iff* $M, q \models_{\mathbf{pATL}} tr(\varphi)$*.*

ATL vs. pATL with Behavioral Strategies In Theorem 1 we have shown that, under the semantics based on mixed response strategies, the **ATL** operator $\langle\!\langle A \rangle\!\rangle$ can be replaced by $\langle\!\langle A \rangle\!\rangle^1_\omega$ if all response strategies have non-zero probability according to ω. One could expect the same for behavioral strategies if it is assumed that each "response action" is left possible; however, an analogous result does not hold. That is because we consider probabilities over all infinite paths in the system, which makes for a continuous probability space. Thus, the probability that a particular path will occur is zero, while it still *can* occur, cf. Example 3. Proposition 3 is an immediate corollary: **pATL** with behavioral predictions cannot simulate plain **ATL** operators in a straightforward way. Still, as Proposition 4 shows, that can be done in the subclass of *acyclic* ***CGS*** (the result will become important for the model checking analysis in Section 4.2).

Example 3. Let M_2' be the **CGSP** based on **CGS** M_2 shown in Figure 2. Note that $M, q_1 \models \neg\langle\!\langle a_1 \rangle\!\rangle \Diamond \mathsf{p}$. What happens if agent a_2 behaves according to a behavioral strategy? Let β_{a_2} be the behavioral strategy specified as follows: $\beta_{a_2}(q_1)(\alpha) = \epsilon$,

$\beta_{a_2}(q_1)(\beta) = 1-\epsilon$, and $\beta_{a_2}(q_2)(\alpha) = 1$ where $0 < \epsilon < 1$. This behavioral strategy assigns non-zero probability to all actions of a_2. Then, for a symbol ω with $[\![\omega]\!]_{a_2} = \beta_{a_2}$ we have that $M, q_1 \models \langle\!\langle a_1 \rangle\!\rangle^1_\omega \Diamond \mathsf{p}$. Thus, a_1 has a strategy which guarantees $\Diamond \mathsf{p}$ *with expected probability 1*. The reason for that is due to the fact that the only possible path which can prevent $\Diamond \mathsf{p}$ is $q_1 q_1 q_1 \ldots$. But the probability that this is going to happen is $\lim_{n\to\infty} \prod_{i=1}^{n} \epsilon = 0$.

Proposition 3. *There is an **ATL** path formula γ, a model M and a state q such that $M, q \models_{ATL} \neg\langle\!\langle A \rangle\!\rangle\gamma$ but $M, q \models_{pATL} \langle\!\langle A \rangle\!\rangle^1_\omega \gamma$ for every behavioral strategy.*

Let us define a *sink state* as a state with a loop to itself being the only outgoing transition. A **CGS** (resp. **CGSP**) is *acyclic* iff it contains no cycles except for the loops at sink states. Such a model includes only a finite number of paths, so the following proposition can be proven analogously to Theorem 1.

Proposition 4. *Let M be an acyclic **CGS** and ω denote a behavioral prediction for $\mathbb{A}\mathrm{gt} \setminus A$ in which every action is possible (i.e., $[\![\omega]\!]_{\mathbb{A}\mathrm{gt}\setminus A}(q)(\overrightarrow{\alpha}) > 0$ for every $q \in Q, \overrightarrow{\alpha} \in d_{\mathbb{A}\mathrm{gt}\setminus A}(q)$). Then, $M, q \models_{ATL} \langle\!\langle A \rangle\!\rangle\gamma$ iff $M, q \models_{pATL} \langle\!\langle A \rangle\!\rangle^1_\omega \gamma$.*

4 Model Checking

In this section, we discuss the complexity of model checking formulae of our "**ATL** with probabilistic success". We have presented two alternative semantics for the logic, underpinned by two different ways of assuming the opponents' behavior. The semantics based on mixed strategies seems to be the simpler of the two, as the success measure is based on a finite probability distribution, and hence can be computed as a finite sum of elements. In contrast, the semantics based on behavioral strategies refers to an integral of a continuous probability distribution – so one might expect that checking formulae of **pATL** in the latter case is much harder. Surprisingly, it turns out to be the opposite.

4.1 Model Checking pATL with Mixed Opponents' Strategies

We study the model checking problem with respect to the number of transitions in the model (m) and the length of the formula (l). As the number of memoryless strategies is usually exponential in the number of transitions, we need a compact way of representing mixed strategies (representing them explicitly as arrays of probability values would yield structures of exponential size). For the rest of this section, we assume that a mixed strategy is represented as a sequence of pairs $[\langle \mathcal{C}_1, p_1\rangle, \ldots, \langle \mathcal{C}_n, p_n\rangle]$, where the length of the sequence is polynomial in m, l, every $\mathcal{C}_i$ is a condition on strategies that can be checked in polynomial time wrt m, l, and every $p_i \in [0, 1]$ is a probability value with a polynomial representation wrt m, l. For simplicity, we assume that conditions $\mathcal{C}_i$ are mutually exclusive. The idea is that the probability of strategy s is determined as $p(s) = p_i$ by the condition $\mathcal{C}_i$ which holds for s; if no $\mathcal{C}_i$ holds for s then the probability of s is 0. We also assume that the distribution is normalized, i.e., $\sum_{s\in\Sigma} p(s) = 1$ where $p(s)$ denotes the probability of s determined by the representation given above.

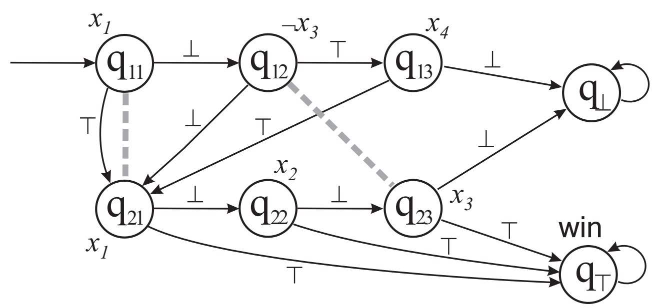

Fig. 3. The concurrent epistemic game structure for formula $F \equiv (x_1 \vee \neg x_3 \vee x_4) \wedge (x_1 \vee x_2 \vee x_3)$. States q_{11}, q_{21} and q_{12}, q_{23} are indistinguishable for the agent: the same action (valuation) must be specified in both within a uniform strategy.

In this setting, model checking **pATL** with mixed memoryless strategies turns out to be at least $\mathbf{PP}$-hard, where $\mathbf{PP}$ is the class of decision problems solvable by a probabilistic Turing machine in polynomial time, with an error probability of less than 1/2 for all instances [5]. We prove it by a polynomial-time reduction of "Majority SAT", a typical $\mathbf{PP}$-complete problem. Since $\mathbf{PP}$ contains both $\mathbf{NP}$ and co-$\mathbf{NP}$, we obtain $\mathbf{NP}$-hardness and co-$\mathbf{NP}$-hardness as an immediate corollary.

Definition 17 (MAJSAT). *The problem MAJSAT is formulated as follows: Given a Boolean formula F in conjunctive normal form with propositional variables $x_1, \ldots, x_n$, answer YES if more than half of all assignments of $x_1, \ldots, x_n$ make F true, and NO otherwise.*

Proposition 5. *Model checking **pATL** with mixed memoryless strategies is $\mathbf{PP}$-hard.*

Proof (sketch). We prove the hardness by a reduction of MAJSAT. First, we take the formula F and construct a single agent concurrent epistemic game structure M in a way similar to [9]. The model includes 2 special states: $q_\top$ (the winning state) and $q_\bot$ (the losing state), plus one state for each literal instance in F. The "literal" states are organized in levels, according to the clause they appear in: q_{ij} refers to the jth literal of clause i. At each "literal" state, the agent can declare the underlying proposition true or false. If the declaration validates the literal, then the system proceeds to the next clause; otherwise it proceeds to the next literal in the same clause. For example, if q_{12} refers to literal $\neg x_3$, then action "true" makes the system proceed to q_{13} (in search of another literal that would validate clause 1), while action "false" changes the state to q_{21} (to validate the next clause). In case the last literal in a clause has been invalidated, the system proceeds to $q_\bot$; when a literal in the last clause is validated, a transition to $q_\top$ follows. There is a single atomic proposition win in the model, which holds only in state $q_\top$. An example of the construction is shown in Figure 3.

Every two nodes with the same underlying proposition are connected by an indistinguishability link to ensure that strategies consistently assign variables $x_1, \ldots, x_n$ with Boolean values. To achieve this, it is enough to require that only *uniform* strategies are used by the agent; a strategy is uniform iff it specifies the same choices in indistinguishable states. Now we observe the following facts:

- There is a 1–1 correspondence between assignments of $x_1, \ldots, x_n$ and uniform strategies of the validating agent. Also, each uniform strategy s determines exactly one path $\lambda(s)$ starting from q_{11};
- By the above, the number of uniform strategies is equal to the number of different assignments of $x_1, \ldots, x_n$. Thus, there are $D = 2^n$ uniform strategies in total;
- A uniform strategy successfully validates F iff it enforces path $\lambda(s)$ that achieves $q_\top$, i.e., one for which $\lambda(s) \models \Diamond \mathsf{win}$;
- Uniformity of a strategy can be checked in time polynomial wrt m (the number of transitions in the model). Let $\mathcal{C}$ be an encoding of the uniformity condition; then, mixed strategy $[\langle \mathcal{C}, \frac{1}{D} \rangle]$ assigns the same importance to every uniform strategy and discards all non-uniform ones. We define symbol ω to denote that strategy;
- MAJSAT(F)=YES iff $\frac{\#\text{ assignments } V \text{ of } x_1,\ldots,x_n \text{ such that } V \models F}{\#\text{ all assignments of } x_1,\ldots,x_n} \geq 0.5$ iff $\frac{\#\text{ uniform strategies } s \text{ such that } \lambda(s) \models \Diamond \mathsf{win}}{\#\text{ all uniform strategies}} \geq 0.5$ iff $M, q_{11} \models \langle\langle \emptyset \rangle\rangle^{0.5}_{\omega} \Diamond \mathsf{win}$, which concludes the reduction. □

Corollary 2. *Model checking **pATL** with mms's is* $\mathbf{NP}$*-hard and co-*$\mathbf{NP}$*-hard.*

For the upper bound, we present a $\mathbf{PSPACE}$ algorithm for model checking **pATL** with mms's. The algorithm uses an $\mathbf{NP}^{\#\mathbf{P}}$ procedure, i.e., one which runs in nondeterministic polynomial time with calls to an oracle that counts the number of accepting paths of a nondeterministic polynomial time Turing machine [11]. The class $\mathbf{NP}^{\#\mathbf{P}}$ is known to lie between $\mathbf{PH}$ and $\mathbf{PSPACE}$ [10].

Theorem 2. *Model checking **pATL** with mixed memoryless strategies is in* $\mathbf{PSPACE}$.

Proof (Sketch). Let γ be a path formula that does not include cooperation modalities. The following procedure checks if $M, q \models \langle\langle A \rangle\rangle^{p}_{\omega} \gamma$:

1. Nondeterministically choose a strategy s_A of agents A; /requires at most m steps/
2. For each $\langle \mathcal{C}_i, p_i \rangle \in [\![\omega]\!]$, execute $T_i := oracle(s_A, \mathcal{C}_i)$; /polynomially many calls/
3. Answer YES if $\sum_i p_i T_i \geq p$ and NO otherwise. /computation polynomial in the representation of p_i and T_i/

The oracle computes the number of $\mathbb{A}\mathrm{gt} \setminus A$'s strategies $t_{\mathbb{A}\mathrm{gt}\setminus A}$ such that $t_{\mathbb{A}\mathrm{gt}\setminus A}$ obeys $\mathcal{C}_i$ and $\langle s_A, t_{\mathbb{A}\mathrm{gt}\setminus A} \rangle$ generate a path that satisfies γ. That is, the oracle counts the accepting paths of the following nondeterministic Turing machine:

1. Nondet. choose a strategy $t_{\mathbb{A}\mathrm{gt}\setminus A}$ of agents $\mathbb{A}\mathrm{gt} \setminus A$; /requires at most m steps/
2. Check whether $t_{\mathbb{A}\mathrm{gt}\setminus A}$ satisfies $\mathcal{C}_i$; /polynomially many steps/
3. If so, "trim" model M by removing choices that are not in $\langle s_A, t_{\mathbb{A}\mathrm{gt}\setminus A} \rangle$, then model-check the **CTL** formula $\mathsf{A}\gamma$ in the resulting model and return the answer of that algorithm; otherwise return NO. /m steps + **CTL** model checking which is polynomial in m, l [3]/

The main procedure runs in time $\mathbf{NP}^{\#\mathbf{P}}$, and hence the task can be done in polynomial space [10]. For the case when γ includes nested strategic modalities, the procedure is applied recursively (bottom-up). That is, we get a deterministic Turing machine with adaptive calls to the $\mathbf{PSPACE}$ procedure. Since $\mathbf{P}^{\mathbf{PSPACE}} = \mathbf{PSPACE}$, we obtain the upper bound. □

4.2 Model Checking pATL with Behavioral Opponents' Strategies

The semantics of **pATL** with opponents' behavior modeled by behavioral strategies is mathematically more advanced than for mixed strategies. So, one may expect the corresponding model checking problem to be even harder than the one we studied in Section 4.1. Surprisingly, it turns out that checking **pATL** with behavioral strategies can be done in polynomial time wrt the number of transitions in the model (m) and the length of the formula (l). Below, we sketch the procedure $mcheck(M, q, \varphi)$ that checks whether $M, q \models \varphi$:

- $\varphi \equiv p$, $\neg\psi$, or $\psi_1 \wedge \psi_2$: proceed as usual;
- $\varphi \equiv \langle\!\langle A \rangle\!\rangle^{p}_{\omega} \Box \psi$: (for $\varphi \equiv \langle\!\langle A \rangle\!\rangle^{p}_{\omega} \bigcirc \psi$ and $\varphi \equiv \langle\!\langle A \rangle\!\rangle^{p}_{\omega} \psi_1 \,\mathcal{U}\, \psi_2$ analogously)
  1. Model check ψ in M recursively. Replace ψ with a new proposition yes holding in exactly those states $st \in Q$ for which $mcheck(M, st, \psi) = $ YES;
  2. Reconstruct M as a 2-player **CGSP** M' with agent 1 representing team A and 2 representing $\mathbb{A}\mathrm{gt} \setminus A$. That is, $d'_1(st) = \prod_{a \in A} d_a(st)$, $d'_2(st) = \prod_{a \in \mathbb{A}\mathrm{gt} \setminus A} d_a(st)$ for each $st \in Q$, and the transition function o' is updated accordingly.
  3. Fix the behavior of agent 2 in M' according to $[\![\omega]\!]_{\mathbb{A}\mathrm{gt} \setminus A}$. That is, construct the probabilistic transition function o'' so that, for each $st, st' \in Q, \alpha_1 \in d'_1(st)$: $o''(st, \alpha_1, st') = \sum_{\{\alpha_2 \in d'_2(st) \mid o'(st, \alpha_1, \alpha_2) = st'\}} [\![\omega]\!]_{\mathbb{A}\mathrm{gt} \setminus A}(st, \alpha_2)$. Also, reconstruct proposition yes as a reward function that assigns 1 at state st if $\mathsf{yes} \in \pi'(st)$ and 0 otherwise. Note that the resulting structure M'' is a Markov Decision Process [2];
  4. Model check the formula $\exists\Box\mathsf{yes}$ of "Discounted **CTL**" [4] in M'', q and return the answer. This can be done in time polynomial in the number of transitions in M'' and exponential in the length of the formula [4]. Note, however, that the length of $\exists\Box\mathsf{yes}$ is constant.

Since part 2-4 requires $O(m)$ steps, and it is repeated at most l times (once per subformula of φ), we get that the procedure runs in time $O(ml)$.

For the lower bound, we observe that reachability in And-Or-Graphs [6] can be reduced (in constant time) to model checking of the fixed **ATL** formula $\langle\!\langle a \rangle\!\rangle \Diamond p$ over acyclic **CGS** (cf. [1]). By Proposition 4, this reduces (again in constant time) to model checking of **pATL** with behavioral predictions. In consequence, we get the following.

Theorem 3. *Model checking **pATL** with the opponents' behavior modeled by behavioral memoryless strategies is* $\mathbf{P}$*-complete with respect to the number of transitions in the model and the length of the formula.*

Thus, it turns out that the model checking problem associated with the more sophisticated semantics can be done in linear time wrt the input size, while model checking the seemingly simpler semantics is much harder ($\mathbf{NP}$- and co-$\mathbf{NP}$-hard).

5 Conclusions

In this paper, we combine the rigorous approach to success of **ATL** with a quantitative analysis of the possible outcome of strategies. The resulting logic goes well beyond the usual "all-or-nothing" reasoning: Instead of always looking at the opponents' most dangerous response, we assume them to select strategies according to some probability measure. To this end, we define new cooperation modalities $\langle\langle A \rangle\rangle_\omega^p \gamma$ with the intuitive reading that group A has a strategy to enforce γ with probability p assuming that the opponents behave according to the predicted behavior denoted by ω. Although we introduce two specific notions of success (one based on mixed response strategies, the other on behavioral predictions), the idea of the success measure is generic and can be implemented according to the designer's needs. This enables the framework to be used in a very flexible way and in various scenarios.

We show that the semantics of **pATL** based on mixed responses embeds **ATL**, while the semantics of **pATL** based on behavioral responses does not (or, at least, not in a straightforward way). Furthermore, we prove that model checking **pATL** with mixed responses is located between **PP** and **PSPACE**, while the same problem for behavioral predictions can be done in linear time wrt the input size (i.e., no worse than for original **ATL**). Thus, we obtain the surprising result that the first semantics (which looked more intuitive and less mathematically advanced at the first glance) turns out to be considerably handicapped in terms of complexity when compared to the other semantics.

We thank Valentin Goranko and the anonymous referees for their comments and discussion, and Hendrik Baumann for his help. The research was partially conducted within the Polish development project no. O R000024 04.

References

1. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *Journal of the ACM*, 49:672–713, 2002.
2. R. Bellman. A Markovian decision process. *Journal of Mathematics and Mechanics*, 6:679–684, 1957.
3. E. Clarke and E. Emerson. Design and synthesis of synchronization skeletons using branching time temporal logic. In *Proceedings of Logics of Programs Workshop*, volume 131 of *Lecture Notes in Computer Science*, pages 52–71, 1981.
4. L. de Alfaro, M. Faella, T. Henzinger, R. Majumdar, and M. Stoelinga. Model checking discounted temporal properties. In *Proceedings of TACAS'04*, volume 2988 of *LNCS*, pages 57–68, 2004.
5. J. Gill. Computational complexity of probabilistic Turing machines. *SIAM Journal on Computing*, 6(4), 1977.
6. N. Immerman. Number of quantifiers is better than number of tape cells. *Journal of Computer and System Sciences*, 22(3):384–406, 1981.
7. W. Jamroga. A temporal logic for multi-agent MDP's. In *Proceedings of the AAMAS Workshop on Formal Models and Methods for Multi-Robot Systems*, pages 29–34, 2008.
8. J. G. Kemeny, L. J. Snell, and A. W. Knapp. *Denumerable Markov Chains*. 1966.
9. P. Y. Schobbens. Alternating-time logic with imperfect recall. *Electronic Notes in Theoretical Computer Science*, 85(2), 2004.
10. S. Toda. On the computational power of PP and P. In *Proceedings of IEEE FOCS'89*, pages 514–519, 1989.
11. L. G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8:189–201, 1979.
12. J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behaviour*. Princeton University Press: Princeton, NJ, 1944.