Quantified Epistemic and Probabilistic ATL

Henning Schnoor
Institut für Informatik, Christian-Albrechts-Universität zu Kiel, 24098 Kiel, Germany

Keywords: ATL, Multi-agent systems, Epistemic Logic.

Abstract: We introduce QAPI (quantified ATL with probabilism and incomplete information), which extends epistemic and probabilistic ATL with a flexible mechanism to reason about strategies in the object language, allowing very flexible treatment of the behavior of the "counter-coalition". QAPI can express complex strategic properties such as equilibria. We show how related logics can be expressed in QAPI, provide bisimulation relations, and study the issues arising from the interplay between quantifiers and both epistemic and temporal operators.

1 INTRODUCTION

ATL (Alternating-time temporal logic) (Alur et al., 2002) is a logic to reason about strategic properties of games. Its strategy operator $\langle\langle A\rangle\rangle\,\varphi$ expresses "there is a strategy for coalition A to achieve φ." We introduce QAPI (quantified ATL with probabilism and incomplete information), a powerful epistemic and probabilistic extension of ATL with quantification of and explicit reasoning about strategies. QAPI's key features are:

- *Strategy Variables* allow explicit reasoning about strategies in the object language,
- A *generalized Strategy Operator* flexibly binds the behavior of some coalitions to strategies, while the remaining players exhibit standard ATL "worst-case" behavior,
- *Quantification* of strategy variables expresses dependence between strategies.

Existential quantification of strategies already appears as part of the $\langle\langle.\rangle\rangle$-operator of ATL, however QAPI makes this more explicit and allows separating the *quantification* of a strategy and the *reasoning* about it in the formulas. To this end, the logic can reason directly about the effect of a coalition following a strategy and express statements as "if coalition A follows strategy s, then φ is true."

QAPI properly includes e.g., ATL*, strategy logic (Chatterjee et al., 2007), ATLES (Walther et al., 2007), (M)IATL (Ågotnes et al., 2007), ATEL-R* and ATOL (Jamroga and van der Hoek, 2004). QAPI can reason about equilibria and express that a coalition *knows* a strategy to be successful. This requirement is often useful, and is e.g., hard-coded into the strategy definition in (Schobbens, 2004). In addition, QAPI features probabilistic reasoning, i.e., can express that events occur with a certain probability bound.

We illustrate QAPI's advantages with an important example. When evaluating $\langle\langle A\rangle\rangle\,\varphi$ in ATL, the behavior of players not in A (we denote this "counter-coalition" with $\overline{A}$) is universally quantified: A must succeed for every possible behavior of $\overline{A}$. Hence A has a strategy for φ only if such a strategy works even in the worst-case setting where

- $\overline{A}$'s only goal is to stop A from reaching the goal,
- the players in $\overline{A}$ know A's goal,
- $\overline{A}$'s actions may depend on unknown information.

These issues are particularly relevant when players have incomplete information about the game. Variants of ATL for this case were suggested in e.g., (Jamroga, 2004; Schobbens, 2004; Jamroga and van der Hoek, 2004; Herzig and Troquard, 2006; Schnoor, 2010b). These logics restrict agents to strategies that can be implemented with the available information, but still require them to be successful for every possible behavior of the counter-coalition. Hence the above limitations still apply—for example, "A can achieve φ against every strategy of $\overline{A}$ that uses only information available to $\overline{A}$" cannot be expressed.

QAPI's direct reasoning about strategies provides a flexible way to specify the behavior of all players, and in particular addresses the above-mentioned shortcomings with a fine-grained specification of the behavior of the "counter-coalition" $\overline{A}$. For example, the following behaviors of $\overline{A}$ can be specified:

- $\overline{A}$ continues a strategy for their own goal—i.e., $\overline{A}$

Schnoor H..
Quantified Epistemic and Probabilistic ATL.
DOI: 10.5220/0004189300140023
In *Proceedings of the 5th International Conference on Agents and Artificial Intelligence* (ICAART-2013), pages 14-23
ISBN: 978-989-8565-39-6

is unaware of (or not interested in) what A does,

- $\overline{A}$ follows a strategy tailor-made to counteract the goal φ, but that can be implemented with information available to $\overline{A}$—here $\overline{A}$ reacts to A with "realistic" capabilities, i.e., strategies based on information actually available to $\overline{A}$,
- $\overline{A}$ plays an arbitrary sequence of actions, which does not have to correspond to an implementable strategy—this is the pessimistic view of the logics mentioned above: A must be successful against every possible behavior of the players in $\overline{A}$.

As we will show, detailed reasoning about the counter-coalition is only one advantage of QAPI. Our results are as follows:

1. We prove that QAPI has a natural notion of bisimulation which is more widely applicable than the one in (Schnoor, 2010b), even though QAPI is considerably more expressive. In particular, our definition can establish strategic and epistemic equivalence between finite and infinite structures.
2. We discuss the effects of combining quantification, epistemic, and temporal operators in detail. The combination of these operators can lead to unnatural situations, which motivate the restriction of QAPI to infix quantification.
3. We prove complexity and decidability results for model checking QAPI. In the memoryless case QAPI's added expressiveness compared to ATL* comes without significant cost: The complexity ranges from PSPACE to 3EXPTIME for games that are deterministic or probabilistic. Hence the deterministic case matches the known PSPACE-completeness for ATL* with memoryless strategies (Schobbens, 2004). As expected, the problem is undecidable in the perfect-recall case.

Related Work. We only mention the most closely related work (in addition to the papers mentioned above) from the rich literature. QAPI is an extension of the ATL*-semantics introduced in (Schnoor, 2010b), and utilizes the notion of a strategy choice introduced there. In this paper, we extend the semantics and the results of (Schnoor, 2010b) by the use of strategy variables, quantification, and explicit strategy assignment, which lead to a much richer language. In particular, the semantics in (Schnoor, 2010b) does not handle negation of the strategy operator in a satisfactory way in the incomplete-information setting. Further, our notion of a bisimulation is much more general than the one suggested in (Schnoor, 2010b). QAPI's approach of allowing first-order like quantification of strategies is very similar to the treatment of strategies in strategy logic (Chatterjee et al., 2007). However, the combination of epistemic aspects and quantification reveals some surprising subtleties, which we discuss in Section 4, and to the best of our knowledge, there are no results on bisimulations for strategy logic.

Relaxations of ATL's universal quantification over the counter-coalition's behavior were studied in (Ågotnes et al., 2007; Walther et al., 2007) for the complete-information case. In (Schnoor, 2012), QAPI is used to specify strategic and epistemic properties of cryptographic protocols, the bisimulation results from the present paper are used to obtain a protocol verification algorithm.

All proofs can be found in the technical report (Schnoor, 2010a).

2 Syntax and Semantics of QAPI

2.1 Concurrent Game Structures

We use the definition of concurrent game structures from (Schnoor, 2010b), which extends the one from (Alur et al., 2002) with probabilistic (see also (Chen and Lu, 2007)) and epistemic aspects (see also (Jamroga and van der Hoek, 2004)):

Definition 1. *A* concurrent game structure (CGS) *is a tuple* $\mathcal{C} = (\Sigma, Q, \mathbb{P}, \pi, \Delta, \delta, \mathtt{eq})$*, where*

- Σ *and* $\mathbb{P}$ *are finite sets of* players *and* propositional variables, Q *is a (finite or infinite) set of* states,
- $\pi \colon \mathbb{P} \to 2^Q$ *is a* propositional assignment,
- Δ *is a* move function *such that* $\Delta(q,a)$ *is the set of* moves *available at state* $q \in Q$ *to player* $a \in \Sigma$*. For* $A \subseteq \Sigma$ *and* $q \in Q$*, an* (A,q)-move *is a function* c *such that* $c(a) \in \Delta(q,a)$ *for all* $a \in A$.
- δ *is a* probabilistic transition function *which for each state* q *and* (Σ,q)*-move* c*, returns a discrete*[1] *probability distribution* $\delta(q,c)$ *on* Q *(the state obtained when in* q*, all players perform their move as specified by* c*),*
- $\mathtt{eq}$ *is an* information function $\mathtt{eq}\colon \{1,\dots,n\} \times \Sigma \to \mathcal{P}(Q \times Q)$*, where* $n \in \mathbb{N}$ *and for each* $i \in \{1,\dots,n\}$ *and* $a \in \Sigma$*,* $\mathtt{eq}(i,a)$ *is an equivalence relation on* Q*. We also call each* $i \in \{1,\dots,n\}$ *a* degree of information.

Moves are merely "names for actions" and only have meaning in combination with the transition function δ. A subset $A \subseteq \Sigma$ is a *coalition of* $\mathcal{C}$. We leave out "of $\mathcal{C}$" when $\mathcal{C}$ is clear from the context, omit set brackets for singletons, etc. The coalition $\Sigma \setminus A$ is denoted with $\overline{A}$. We write $\Pr(\delta(q,c) = q')$ for

[1] A probability distribution Pr on Q is discrete, if there is a countable set $Q' \subseteq Q$ such that $\sum_{q\in Q'} \Pr(q) = 1$.

$(\delta(q,c))(q')$, i.e., consider $\delta(q,c)$ as a random variable on Q. The function eq expresses incomplete information: It specifies pairs of states that a player cannot distinguish. By specifying several relations $\mathsf{eq}(1,a),\ldots,\mathsf{eq}(n,a)$ for each player, we can specify how much information a player may use to reach a certain goal. This is useful e.g., in security definitions (Cortier et al., 2007; Schnoor, 2012).

$\mathcal{C}$ is *deterministic* if all distributions $\delta(q,c)$ assign 1 to one state and 0 to all others, $\mathcal{C}$ has *complete information* if $\mathsf{eq}(i,a)$ is always the equality relation.

2.2 Strategies, Strategy Choices, and Formulas

The core operator of QAPI is the *strategy operator*: $\langle\langle A:\mathrm{S}_1,\ B:\mathrm{S}_2\rangle\rangle_i^{\geq\alpha}\varphi$ expresses "if coalition A follows S_1 and B follows S_2, where both coalitions base their decisions only on information available to them in information degree i, the run of the game satisfies φ with probability $\geq\alpha$, no matter what players from $\overline{A\cup B}$ do." Here, S_1 and S_2 are variables for *strategy choices* which generalize strategies (see below). While similar to the ATL-operator $\langle\langle.\rangle\rangle$, the strategy operator is much more powerful: It allows to flexibly find a strategy to a coalition. This allows, for example, to model that a coalition *commits* to a strategy (in ATL*, a strategy is revoked when the $\langle\langle.\rangle\rangle$-operator is nested) and much more (see examples below).

Definition 2. *Let $\mathcal{C}$ be a CGS with n degrees of information. Then the set of* strategy formulas for $\mathcal{C}$ *is defined as follows:*

- *A propositional variable of $\mathcal{C}$ is a state formula,*
- *conjunctions and negations of state (path) formulas are state (path) formulas,*
- *every state formula is a path formula,*
- *if $A_1,\ldots,A_m$ are coalitions, $1\leq i\leq n$, $0\leq\alpha\leq 1$, and $\blacktriangleleft$ is one of $\leq,<,>,\geq$, and ψ is a path formula, and S_i is an A_i-strategy choice variable for each i, then $\langle\langle A_1:\mathrm{S}_1,\ldots,A_m:\mathrm{S}_m\rangle\rangle_i^{\blacktriangleleft\alpha}\psi$ is a state formula,*
- *if A is a coalition, $1\leq i\leq n$, ψ is a state formula, and $k\in\{D,E,C\}$ then $\mathcal{K}_{A,i}^k\psi$ is a state formula,*
- *If φ_1 and φ_2 are path formulas, then $\mathsf{X}\varphi_1$, $\mathsf{P}\varphi_1$, $\mathsf{X}^{-1}\varphi_1$, and $\varphi_1\mathsf{U}\varphi_2$ are path formulas.*

The values D, E, and C indicate different notions of knowledge, namely *distributed knowledge*, *everybody knows*, and *common knowledge*. We use standard abbreviations like $\varphi\vee\psi=\neg(\neg\varphi\wedge\neg\psi)$, $\Diamond\varphi=\mathsf{true}\,\mathsf{U}\varphi$, and $\Box\varphi=\neg\Diamond\neg\varphi$. A $\langle\langle.\rangle\rangle$-formula is one whose outmost operator is the strategy operator. In a CGS with only one degree of information, we omit the i subscript of the strategy operator; in a deterministic CGS we omit the probability bound $\blacktriangleleft\alpha$ (and understand it to be read as ≥ 1). Quantified strategy formulas are strategy formulas in which the appearing strategy choice variables are quantified:

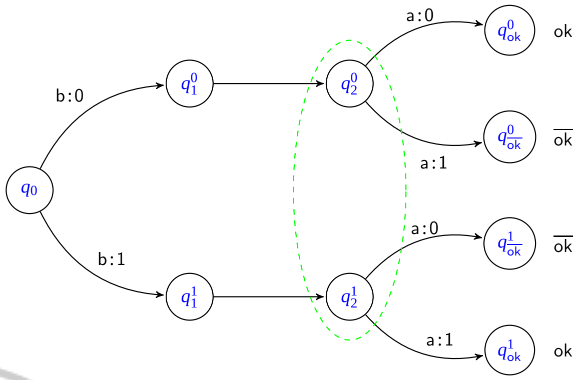

Figure 1: Strategy choices.

Definition 3. *Let $\mathcal{C}$ be a CGS, let φ be a strategy formula for $\mathcal{C}$ such that every strategy choice variable appearing in φ is one of $\mathrm{S}_1,\ldots,\mathrm{S}_n$. Then*

$$\forall\mathrm{S}_1\exists\mathrm{S}_2\forall\mathrm{S}_3\ldots\exists\mathrm{S}_n\varphi$$

is a quantified strategy formula for $\mathcal{C}$.

Requiring a strict $\forall\exists\ldots$-alternation is without loss of generality and can be obtained via dummy variables. On the other hand, allowing quantification only in the prefix is a deliberate restriction of QAPI, the reasons for which we discuss in detail in Section 4.

Definition 4. *For a player a, an a-*strategy *in a CGS $\mathcal{C}=(\Sigma,Q,\mathbb{P},\pi,\Delta,\delta,\mathsf{eq})$ is a function s_a with $s_a(q)\in\Delta(q,a)$ for each $q\in Q$. For an information degree i, s_a is i-*uniform *if $q_1\sim_{\mathsf{eq}_i(a)}q_2$ implies $s_a(q_1)=s_a(q_2)$. For $A\subseteq\Sigma$, an A-strategy is a family $(s_a)_{a\in A}$, where each s_a is an a-strategy.*

Our strategies are *memoryless*: A move only depends on the current state, not on the history of the game. With incomplete information, the question how players can *identify* suitable strategies is relevant. Consider the CGS in Figure 1. The players are a and b, the game starts in q_0. The first move by b controls whether the next state is q_1^0 or q_1^1. For $x\in\{0,1\}$, q_1^x is always followed by q_2^x. In q_2^x, the move 0 leads to a state satisfying ok iff $x=0$; move 1 is successful iff $x=1$. Player a cannot distinguish q_2^0 and q_2^1. We ask whether he has a strategy leading to ok that is successful started in both q_1^0 and q_1^1. If a can only use strategies, he must play the same move in q_2^0 and in q_2^1, and thus fails in one of them. However, if a can *decide* on a strategy and remember this decision, the player can choose in q_1^0 (q_1^1) a strategy playing 0 (1) in every state, and be successful.

Strategy choices (Schnoor, 2010b) formalize how a player chooses a strategy, and distinguish between states where a strategy is *identified* and where it is *executed*: In state q_1^0 or q_1^1, player a uses his information to choose the strategy that he follows from then on. When using only strategies, the knowledge has to be present at the time of *performing* a move. Hence strategy choices give players additional capabilities over the pure memoryless setting, by allowing to remember *decisions*. In contrast to the *perfect recall* case, where players remember the entire run of a game, there is no significant computational price, whereas perfect recall makes the model checking problem undecidable (cp. Section 6).

Definition 5. *A* strategy choice *for a coalition A in a CGS* $\mathcal{C} = (\Sigma, Q, \mathbb{P}, \pi, \Delta, \delta, \mathsf{eq})$ *is a function* S *such that for each* $a \in A$, $q \in Q$, *each* $\langle\langle . \rangle\rangle_i$*-formula* φ, $\mathsf{S}(a,q,\varphi)$ *is an i-uniform a-strategy in* $\mathcal{C}$*, and if* $q_1 \sim_{\mathsf{eq}_i(a)} q_2$*, then* $\mathsf{S}(a,q_1,\varphi) = \mathsf{S}(a,q_2,\varphi)$.

In the definition of a strategy choice, syntax and semantics meet, since one input to a strategy choice is the goal a coalition is supposed to achieve—such a goal is best specified with a formula. The formula also specifies the coalition working together to achieve the goal. For a coalition A, and a strategy choice S for A, the strategy chosen for A by S in a state q to reach the goal φ is the A-strategy $(s_a)_{a\in A}$ with $s_a = \mathsf{S}(a,q,\varphi)$ for each a. We denote this strategy with $\mathsf{S}(A,q,\varphi)$. Strategy choices model the *decision* of a single player to use a certain strategy. For coalitions, they model strategies agreed upon before the game for possible goals. This allows their members to predict the each other's behavior without in-game communication. As mentioned above, the crucial point is that strategy choices distinguish between states where a strategy is *identified* and where it is *executed*: In state q_1^0 or q_1^1 of the above example, player a uses his information to choose a strategy which he then follows. When using only strategies, the knowledge has to be present at the time of *playing* a move. A strategy choice hence allows players to "remember" previous decisions. For coalitions, it models prior agreement helpful in e.g., coordination games.

The *strategy operator* binds the behavior of the players in the appearing coalitions to the strategies specified by the assigned strategy choices (see below). The remaining players (the "counter-coalition") are treated as "free agents" in QAPI: Every possible behavior of these players is taken into account. Such a behavior may not even follow any strategy, for example they may perform different moves when encountering the same state twice during the game. This is formalized as a *response* (cp. (Schnoor, 2010b)) to a coalition A, which is a function r such that $r(t,q)$ is a $(\overline{A},q)$-move for each $t \in \mathbb{N}$ and each $q \in Q$. This models an arbitrary reaction to the outcomes of an A-strategy: In the i-th step of a game, $\overline{A}$ performs the move $r(i,q)$, if the current state is q.

When a coalition A follows the strategy s_A, and the behavior of $\overline{A}$ is defined by the response r, the moves of all players are fixed; the game is a Markov process. This leads to the following definition of "success probability." A *path* in a CGS $\mathcal{C}$ is a sequence $\lambda = \lambda[0]\lambda[1]\ldots$ of states of $\mathcal{C}$.

Definition 6. *Let* $\mathcal{C}$ *be a CGS, let* s_A *be an A-strategy, let r be a response to A. For a set M of paths over* $\mathcal{C}$*, and a state* $q \in Q$, $Pr(q \to M \mid s_A + r)$ *is the probability that in the Markov process resulting from* $\mathcal{C}$*,* s_A*, and r with initial state q, the resulting path is in M.*

A key feature of QAPI is the flexible binding of strategies to coalitions, which is done using the strategy operator. As a technical tool to resolve possible ambiguities, we introduce a "join" operation on strategy choices: If the coalitions $A_1, \ldots, A_n$ follow strategy choices $\mathsf{S}_1, \ldots, \mathsf{S}_n$, the resulting "joint strategy choice" for $A_1 \cup \cdots \cup A_n$ is $\mathsf{S}_1 \circ \cdots \circ \mathsf{S}_n$. This is a "union" of the S_i with a tie-breaking rule for players appearing in several of the coalitions: These always follow the "left-most" applicable strategy choice. We define the (associative) operator $\circ$ as follows:

$$\mathsf{S}_1 \circ \mathsf{S}_2(a,q,\varphi) = \begin{cases} \mathsf{S}_1(a,q,\varphi), & \text{if } a \in A_1, \\ \mathsf{S}_2(a,q,\varphi), & \text{if } a \in A_2 \setminus A_1. \end{cases}$$

This definition ensures that if a coalition $A_1 \cup \cdots \cup A_n$ is instructed to follow the strategy choice $\mathsf{S}_1 \circ \cdots \circ \mathsf{S}_n$, then even if $A_i \cap A_j \neq \emptyset$, for each agent the strategy choice to follow is well-defined.

2.3 Evaluating Formulas

In the same manner as the syntax, we also define QAPI's semantics in two stages: We first handle strategy formulas, where instantiations for the appearing strategy choice variables are given. This naturally leads to the semantics definition for quantified formulas. Our semantics is very natural: Propositional variables and operators are handled as usual, temporal operators behave as in linear-time temporal logic, and $\langle\langle A_1 : \mathsf{S}_1, \ldots, A_n : \mathsf{S}_n \rangle\rangle_i^{\geq\alpha} \psi$ expresses that when coalitions $A_1, \ldots, A_n$ follow the strategy choices $\mathsf{S}_1, \ldots, \mathsf{S}_n$ with information degree i available, the formula ψ is satisfied with probability $\geq \alpha$. The knowledge operator $\mathcal{K}$ models group knowledge, see below.

Definition 7. *Let* $\mathcal{C} = (\Sigma, Q, \mathbb{P}, \pi, \Delta, \delta, \mathsf{eq})$ *be a CGS, let* $\overrightarrow{\mathsf{S}} = (\mathsf{S}_1, \ldots, \mathsf{S}_n)$ *be a sequence of strategy choices* *instantiating*[2] *the strategy choice variables* $\mathsf{S}_1, \ldots, \mathsf{S}_n$*. Let* φ *be a state formula, let* ψ_1*,* ψ_2 *be path formulas, let* λ *be a path over Q, let* $t \in \mathbb{N}$*. We define*

- $\mathcal{C}, \vec{\mathsf{S}}, q \models p$ *iff* $q \in \pi(p)$ *for* $p \in \mathbb{P}$,
- *conjunction and negation are handled as usual,*
- $(\lambda,t), \vec{\mathsf{S}} \models \varphi$ *iff* $\mathcal{C}, \vec{\mathsf{S}}, \lambda[t] \models \varphi$,
- $(\lambda,t), \vec{\mathsf{S}} \models \mathsf{X}\psi_1$ *iff* $(\lambda, t+1), \vec{\mathsf{S}} \models \psi_1$,
- $(\lambda,t), \vec{\mathsf{S}} \models \mathsf{P}\psi_1$ *iff there is some* $t' \leq t$ *and* $(\lambda,t'), \vec{\mathsf{S}} \models \psi_1$,
- $(\lambda,t), \vec{\mathsf{S}} \models \mathsf{X}^{-1}\psi_1$ *iff* $t \geq 1$ *and* $(\lambda, t-1), \vec{\mathsf{S}} \models \psi_1$,
- $(\lambda,t), \vec{\mathsf{S}} \models \psi_1 \mathsf{U} \psi_2$ *iff there is some* $i \geq t$ *such that* $(\lambda,i), \vec{\mathsf{S}} \models \psi_2$ *and* $(\lambda,j), \vec{\mathsf{S}} \models \psi_1$ *for all* $t \leq j < i$,
- *If* $k \in \{D,E,C\}$*, then* $\mathcal{C}, \vec{\mathsf{S}}, q \models \mathcal{K}^k_{A,i}\varphi$ *iff* $\mathcal{C}, \vec{\mathsf{S}}, q' \models \varphi$ *for all* $q' \in Q$ *with* $q \sim^k_{A,i} q'$ *(see below),*
- $\mathcal{C}, \vec{\mathsf{S}}, q \models \underbrace{\langle\langle A_{i_1} : \mathsf{S}_{i_1}, \ldots, A_{i_m} : \mathsf{S}_{i_m} \rangle\rangle_i^{\blacktriangleleft \alpha} \psi_1}_{=:\varphi_1}$ *iff for every response r to* $A_{i_1} \cup \cdots \cup A_{i_m}$*, we have* $Pr\Big(q \to \Big\{\lambda \mid (\lambda,0), \vec{\mathsf{S}} \models \psi_1\Big\} \mid \mathsf{S}_{i_1} \circ \cdots \circ \mathsf{S}_{i_m}(A_{i_1} \cup \cdots \cup A_{i_m}, q, \varphi_1) + r\Big) \blacktriangleleft \alpha$.

The relations $\sim^D_{A,i}$, $\sim^E_{A,i}$, and $\sim^C_{A,i}$ referenced in Definition 7 represent different possibilities to model group knowledge. For a coalition A and an information degree i, they are defined as follows:

- $\sim^D_{A,i} = \cap_{a \in A} \mathsf{eq}(i,a)$ expresses *distributed knowledge*: $\mathcal{K}^D_{A,i}\varphi$ is true if φ can be deduced from the combined knowledge of every member of A (with respect to information degree i),
- $\sim^E_{A,i} = \cup_{a \in A} \mathsf{eq}(i,a)$ models *everybody knows*: $\mathcal{K}^E_{A,i}\varphi$ is true if every agent in A on his own has enough information to deduce that φ holds (with respect to information degree i),
- $\sim^C_{A,i}$ is the reflexive, transitive closure of $\sim^E_{A,i}$. This models *common knowledge*: $\mathcal{K}^C_{A,i}\varphi$ expresses that (in A, with information degree i), everybody knows that φ is true, and everybody knows that everybody knows that φ is true, …, etc.

These concepts have proven useful to express the knowledge of a group. See (Halpern and Moses, 1990) for detailed discussion.

For quantified formulas, we define:

Definition 8. *Let* $\mathcal{C}$ *be a CGS, let* $\psi = \forall \mathsf{S}_1 \exists \mathsf{S}_2 \forall \mathsf{S}_3 \ldots \exists \mathsf{S}_n \varphi$ *be a quantified strategy formula for* $\mathcal{C}$*, let q be a state of* $\mathcal{C}$*. Then* ψ *is* satisfied *in* $\mathcal{C}$ *at q, written* $\mathcal{C}, q \models \psi$*, if for each* $i \in \{2,4,\ldots n\}$*, there is a function* s_i *such that for all strategy choices* $\mathsf{S}_1, \mathsf{S}_3, \ldots, \mathsf{S}_{n-1}$*, if* S_i *is defined as* $s_i(\mathsf{S}_1, \ldots, \mathsf{S}_{i-1})$ *for even i, then* $\mathcal{C}, (\mathsf{S}_1, \ldots, \mathsf{S}_n), q \models \varphi$*.*

Constant strategy choices (which only depend on the player, not on the state or the formula) are essentially strategies. We introduce quantifiers $\exists_c$ and $\forall_c$ quantifying over constant strategy choices.

[2] I.e., if S_i is an A-strategy choice variable for some coalition A, then S_i is a strategy choice for A.

2.4 MQAPI

MQAPI (*M*emory-enabled QAPI), is QAPI with *perfect recall*. The semantics can be defined in the straight-forward way by encoding history in the states of a system, see, e.g.,(Schnoor, 2010b).

3 Examples

3.1 Restricted Adversaries

The following expresses "A can achieve φ against every uniform strategy of $\overline{A}$:"

$$\exists \mathsf{S}_1 \forall \mathsf{S}_2 \left\langle\left\langle A : \mathsf{S}_1, \overline{A} : \mathsf{S}_2 \right\rangle\right\rangle_1 \varphi.$$

This is weaker than $\exists \mathsf{S}_1 \langle\langle A : \mathsf{S}_1 \rangle\rangle_1 \varphi$: In the latter, $\overline{A}$ is not restricted to any strategy at all, while in the former, $\overline{A}$ has to follow a uniform strategy.

3.2 Sub-coalitions Changing Strategy

Often, when a coalition $A' \subsetneq A$ changes the strategy, they rely on $A \setminus A'$ to continue the current one. Assume that A works together to reach a state where $A' \subsetneq A$ has strategies for φ_1 and φ_2, *if players in* $A \setminus A'$ *continue their earlier strategy*. We express this as

$$\exists_c \mathsf{S}_A \exists \mathsf{S}_{A'} \langle\langle A : \mathsf{S}_A \rangle\rangle_1 \Diamond \big(\langle\langle A' : \mathsf{S}_{A'}, A : \mathsf{S}_A \rangle\rangle_1 \Diamond \varphi_1 \wedge \langle\langle A' : \mathsf{S}_{A'}, A : \mathsf{S}_A \rangle\rangle_1 \Diamond \varphi_2 \big).$$

This expresses that A uses a *fixed strategy* and does not change behavior depending on whether A' attempts to achieve φ_1 or φ_2. In particular, $A \setminus A'$ does not need to know which of these goals A' attempts to achieve. We use the same strategy choice for φ_1 and φ_2 to require A' to identify the correct strategy with the available information.

3.3 Knowing whether a Strategy is Successful

The following expresses "there is an A-strategy such that there is no B-strategy such that the coalition C can know that its application successfully achieves φ:"

$$\exists_c \mathsf{S}_A \forall_c \mathsf{S}_B \neg \mathcal{K}^C \langle\langle A : \mathsf{S}_A, B : \mathsf{S}_B \rangle\rangle_1 \varphi.$$

This is very different from expressing that A has a strategy preventing φ, i.e., $\exists \mathsf{S}_A \langle\langle A : \mathsf{S}_A \rangle\rangle_1 \neg\varphi$, since (i) There may be a successful strategy for B, but not enough information for C to determine that it is successful, (ii) the goal φ may still be reachable if B does not follow a (uniform) strategy.

3.4 Winning Secure Equilibria (WSE)

If player a (b) has goal φ_a (φ_b), a WSE (Chatterjee et al., 2006) is a pair of strategies (s_a, s_b) such that both goals are achieved when a and b play s_a and s_b, and if b plays such that φ_a is not reached anymore, but a still follows s_a, then b's goal φ_b is also not satisfied anymore (same for player a). QAPI can express this as follows: Both goals are reached if (s_a, s_b) is played, and neither player can reach his goal without reaching that of the other player as well, if the latter follows the WSE strategy.

$$\begin{aligned} \exists_c \mathsf{S}_a \exists_c \mathsf{S}_b \quad & \langle\langle a : \mathsf{S}_A,\ b : \mathsf{S}_B \rangle\rangle_1 (\varphi_a \wedge \varphi_b) \\ \wedge \quad & \langle\langle a : \mathsf{S}_A \rangle\rangle_1 (\varphi_b \rightarrow \varphi_a) \\ \wedge \quad & \langle\langle b : \mathsf{S}_B \rangle\rangle_1 (\varphi_a \rightarrow \varphi_b)\,. \end{aligned}$$

3.5 Expressing ATEL-R* and ATOL

ATOL (Jamroga and van der Hoek, 2004) requires *identifying* strategies with the agent's knowledge. ATOL's key operator is defined as follows (right-hand side in our notation)—in the following, A is the coalition *playing*, and Γ the one *identifying* the strategy:

> $\mathcal{C}, q \models \langle\langle A \rangle\rangle_{\mathcal{K}(\Gamma)} \varphi$ iff there is a constant strategy choice S_A such that for all $q' \in \mathcal{C}$ with $q' \sim_\Gamma q$, we have that $\mathcal{C}, q' \models \langle\langle A : \mathsf{S}_A \rangle\rangle_1 \varphi$.

The above can be translated into QAPI by writing

$$\mathcal{C}, q' \models \mathcal{K}^\Gamma \langle\langle A : \mathsf{S}_A \rangle\rangle_1 \varphi,$$

where S_A's quantification depends on the parity of negation and is restricted to constant strategy choices.[3] In (Jamroga and van der Hoek, 2004), it is stated that requiring "Γ knows that A has a strategy to achieve φ" is insufficient to express $\langle\langle A \rangle\rangle_{\mathcal{K}(\Gamma)} \varphi$. It suffices in QAPI since we quantify S_A *before* the $\mathcal{K}$-operator, hence Γ knows that the *fixed* A-strategy is successful. ATEL-R* would quantify the strategy *after* the $\mathcal{K}$-operator in a formula like $\mathcal{K}_\Gamma \langle\langle A \rangle\rangle \varphi$: A could choose a *different* strategy in each state. ATEL-R* (ATOL with recall) can be expressed in MQAPI analogously. The above highlights the usefulness of QAPI's ability to directly reason about strategy choices. Strategy logic (Chatterjee et al., 2007), ATLES (Walther et al., 2007), and (M)IATL (Ågotnes et al., 2007) can be expressed similarly.

[3] It is not sufficient to rely on the uniformity of strategy choices (the same strategy must be chosen in A-indistinguishable states), since there must be a single strategy that is successful in all Γ-indistinguishable states, and Γ might have less information than A.

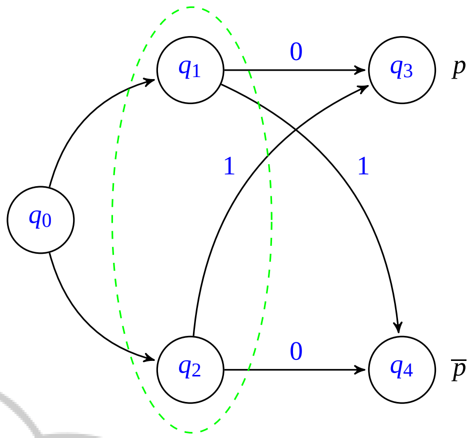

Figure 2: Infix quantification example.

4 QUANTIFICATION AND EPISTEMIC/TEMPORAL OPERATORS

We now study the interplay between quantifiers and temporal or epistemic operators: Applying quantifiers in the scope of epistemic or temporal operators often leads to highly counter-intuitive behavior. This behavior is the reason why QAPI only allows quantification in a quantifier block prefixing the formula. The issues we demonstrate here are not specific to QAPI or the concept of strategy choices, but are general effects that arise in any formalism combining the operators we discuss here with some mechanism of forcing agents to "know" which strategy to apply. The core issue is that an unrestricted $\exists$-quantifier adds a high degree of non-uniformity to the agent's choices, which is incompatible with the epistemic setting.

To demonstrate these issues, in this section, we consider QAPI^{infix}, which is QAPI with arbitrary nesting of quantifiers and other operators. The semantics is defined by applying quantification in every state in the obvious way. Clearly, quantification can always be pulled outside of the scope of propositional and $\Diamond$-operators. The remaining temporal and epistemic operators cannot be handled so easily.

4.1 Quantification in the Scope of Temporal Operators

Consider the following QAPI^{infix}-formula:

$$\mathbb{A}\Box\exists \mathsf{S}_A \langle\langle A : \mathsf{S}_A\rangle\rangle_1^{\geq 1} \psi.$$

The quantifier $\mathbb{A}$ abbreviates $\exists \mathsf{S}_\emptyset \langle\langle \emptyset : \mathsf{S}_\emptyset\rangle\rangle_1^{\geq 1}$ and expresses quantification over all reachable paths (essentially $\mathbb{A}$ is CTL's **A**-operator). The formula expresses that in all reachable states, there is a strategy choice for A that accomplishes ψ. There are no uniformity or epistemic constraints on the $\exists$-quantifier: Even in states that look identical for all members of A, completely different strategy choices can be applied. This is problematic in an epistemic setting: Consider the CGS with two players a and b in Figure 2. We only indicate the moves of player a. The game is turn-based, where it is b's turn in the state q_0 and a's turn in the remaining states. The first action of b chooses whether the next state is q_1 or q_2, these two states are indistinguishable for a. In q_1, player a must play 0 to reach a state where p holds, in state q_2, a must play 1 to achieve this. Now consider the following formula (we consider the coalition $A = \{a\}$):

$$\mathbb{A}\mathsf{X}\exists \mathsf{S}_A \langle\langle A : \mathsf{S}_A\rangle\rangle_1^{\geq 1} p.$$

This formula is true in q_0: In both possible follow-up states, there is a strategy choice that allows player a to enforce that p is true in the next state: In q_1 (q_2), we choose a strategy choice S_1 that for every possible goal and in every state always plays the move 0 (1). *Individually*, these strategy choices satisfy every imaginable uniformity condition, since they fix one move forever. However, intuitively in q_1, player a cannot achieve $\mathsf{X}p$, since a cannot identify the correct strategy choice to apply. This shows that having an existential quantifier in the scope of a temporal operator yields counter-intuitive results.

A natural way to address this problem is to restrict quantification to be "uniform" and demand that the quantifier chooses the same strategy choice in the states indistinguishable for A. We can express this in QAPI^{infix} by requiring that the strategy choice "returned" by the quantifier is successful in all indistinguishable states—in other words, requiring A to know that the strategy choice is successful. In this case, the same strategy choice can be used in all indistinguishable states as intended. In the above example, we therefore would consider the following formula (for singleton-coalitions, all notions of knowledge coincide, we use *common knowledge* in the example):

$$\mathbb{A}\mathsf{X}\exists \mathsf{S}_A \mathcal{K}_{A,1}^{C} \langle\langle A : \mathsf{S}_A\rangle\rangle_1^{\geq 1} p.$$

If we follow the above suggestion and always combine existential quantification with requiring the knowledge that the introduced strategy choice accomplishes its goal, the behavior is much more natural—however, as we now demonstrate, these are exactly the cases which already can be expressed in QAPI.

To do this, we need to decide on a suitable notion of group knowledge to apply in formulas of the above structure. If we use *distributed* knowledge, we essentially allow coordination inside the coalition A as part of the existential quantifier. This is similar to the behavior of ATL/ATL*, where the $\langle\langle . \rangle\rangle$-operator also allows coordination. Hence *distributed knowledge* does not achieve the desired effect. However, *everyone knows* and *common knowledge* do not suffer from these issues: In both cases, each agent on his own can determine whether the current strategy "works." We now show that this intuition is supported by formal arguments: In the case of *everyone knows* or *common knowledge*, the existential quantifier can indeed be exchanged with the $\Box$ operator, the same does not hold for *distributed knowledge*.

Proposition 9. *Let φ be a formula in which the variable S_A does not appear, and which does not use past-time operators, and let $k \in \{E, C\}$. Then* $\Box\exists \mathsf{S}_A \mathcal{K}_{A,i}^{k} \langle\langle A : \mathsf{S}_A\rangle\rangle_i^{\geq\alpha} \varphi \equiv \exists \mathsf{S}_A \Box \mathcal{K}_{A,i}^{k} \langle\langle A : \mathsf{S}_A\rangle\rangle_i^{\geq\alpha} \varphi.$

We require that φ does not contain S_A, since the idea of the above discussion is the direct coupling of the existential quantification of S_A and the group knowledge about the effects of its application. Requiring that φ does not have past-time operators is clearly crucial for memoryless strategies: If φ, e.g., requires to play a specific move if and only if that move has been played previously, then the strategy choice clearly must depend on the history and the above equivalence does not hold. Proposition 9 does not hold for distributed knowledge instead:

Example 10. Consider a CGS $\mathcal{C}$ with players a and b and two Boolean variables x and y, where player a (b) only sees the value of variable x (y) and the values of the variables change randomly in every transition. Each player always has the moves 0 and 1 available. Consider the coalition $A = \{a, b\}$ and the formula φ expressing "a moves according to y and b moves according to x"[4] Since the distributed knowledge of A allows to identify the values of both x and y, in each state there is a strategy choice achieving φ, however clearly there is no single strategy choice which works in all states. Hence, the formula $\Box\exists \mathsf{S}_A \mathcal{K}_{A,1}^{D} \langle\langle A : \mathsf{S}_A\rangle\rangle_1^{\geq 1} \varphi$ is always true in $\mathcal{C}$, while $\exists \mathsf{S}_A \Box \mathcal{K}_{A,1}^{D} \langle\langle A : \mathsf{S}_A\rangle\rangle_1^{\geq 1} \varphi$ is always false.

[4] To express this as a variable, the CGS needs to record the last move of each player in the state in the obvious way.

Proposition 9 can be generalized in several directions. For ease of presentation we only present the above simple form of Proposition 9 which supports the main argument of this section: "Intuitively sensible" applications of quantifications inside $\Box$-operators can be eliminated.

4.2 Quantification in the Scope of Epistemic Operators

We now show that quantification in the scope of epistemic operators leads to similar issues as the case of temporal operators considered above. We again consider the CGS in Figure 2. In q_0, the formula

$$\mathbb{A}\mathsf{X}\mathcal{K}^{d}_{A,1}\exists \mathsf{S}_A \langle\langle A : \mathsf{S}_A\rangle\rangle^{\geq 1}_{1} \mathsf{X}p$$

is true: Agent a (who alone forms the coalition A) knows that there is a successful strategy choice, since there is one in both q_1 and in q_2. However, as seen above, he does not know this strategy choice.

We now present a similar result to Proposition 9, for quantification in the scope of epistemic operators, and identify cases in which these operators commute. For this, we exhibit a "maximal" class of formulas for which knowledge and quantification can always be exchanged. When discussing whether quantification of a variable S_i commutes with an operator (epistemic or otherwise), clearly we are only interested in formulas in which the variable S_i actually plays a non-trivial role. To formalize this, we extend the concept of a "relevant" variable which is well-known in propositional logic, to the class of strategy variables:

Definition 11. *Let φ be a formula with free strategy variables among $\{\mathsf{S}_1,\ldots,\mathsf{S}_n\}$. We say that the variable S_i is* relevant *for φ if there exists a CGS $\mathcal{C}$, a state q of $\mathcal{C}$, and strategy choices $\mathsf{S}_1,\ldots,\mathsf{S}_n,\mathsf{S}'_i$ such that $\mathcal{C},(\mathsf{S}_1,\ldots,\mathsf{S}_n),q \models \varphi$ and $\mathcal{C},(\mathsf{S}_1,\ldots,\mathsf{S}_{i-1},\mathsf{S}'_i,\mathsf{S}_{i+1},\ldots,\mathsf{S}_n),q \not\models \varphi$.*

This means that there exists a situation where it matters which strategy choice is used to instantiate the variable S_i. Examples for an irrelevant variable S_A are $\langle\langle A : \mathsf{S}_A\rangle\rangle^{\geq 1}_{i} (\Diamond x \vee \Box\neg x)$ or $\langle\langle A : \mathsf{S}_A\rangle\rangle^{\geq 0}_{i} \Diamond x$.

Definition 12. *For a coalition A and a degree of information i, $k \in \{D,E,C\}$, a formula φ is k-i-*simple *in S_A, if one of the following conditions is true:*

- *S_A is an irrelevant variable of φ, or*
- *φ is equivalent to a formula of the form $\mathcal{K}^{k}_{A,i}\psi$.*

Formulas that are k-i-simple give a "natural" semantics when prefixed with an existential quantifier, since in the same way as there, the non-uniformity of the existential quantifier is reduced using the epistemic operator. We now show that in these cases, infix quantification again is not necessary, as here, the existential and the epistemic operators commute:

Lemma 13. *If φ is k-i-simple and has a single free strategy variable, then for all CGS $\mathcal{C}$ and states q,*

$$\mathcal{C},q \models \mathcal{K}^{k}_{A,i}\exists s_A\varphi \text{ if and only if } \mathcal{C},q \models \exists s_A \mathcal{K}^{k}_{A,i}\varphi.$$

This class of formulas is maximal—as soon as we have a formula that depends on the variables S_A and of which A's knowledge does not suffice to determine the truth, we cannot swap the above operators.

Proposition 14. *Let φ be a formula such that φ is not k-i-simple in S_A and the coalition A is bound to S_A in the entire formula, then $\exists \mathsf{S}_A \mathcal{K}^{k}_{A,i}\varphi \not\equiv \mathcal{K}^{k}_{A,i}\exists \mathsf{S}_A\varphi$.*

The prerequisite that A is bound to S_A in the entire formula is necessary to e.g., preclude cases where S_A is only used in a non-meaningful way. It is not a strong requirement, as (with infix quantification) usually the subformula directly succeeding the existential quantifier will be the one "talking about" the quantified strategy choice. It is possible to strengthen Proposition 14, however again the simple form here is suffices to show that in the cases where quantification in the scope of an epistemic operator gives a satisfactory semantics, the quantifier can be moved out of scope of that operator, and hence QAPI suffices.

4.3 Discussion

Nesting of quantification and epistemic or temporal operators leads to counter-intuitive behavior, since quantification introduces a degree of non-uniformity, whereas a core issue in the epistemic setting is to enforce sufficient uniformity to ensure that agents have enough knowledge to decide on the "correct" move to play in every situation. Although we did not give a complete characterization of the cases in which temporal/epistemic operators and quantifiers commute and it is notoriously difficult to give a good definition of a "natural" semantics, our results give strong evidence for our claim: In the cases where infix quantification leads to a natural semantics, the quantifiers can be swapped with the temporal/epistemic operators, hence infix quantification is unneeded.

Another reason why QAPI only allows quantifiers in the prefix of a formula is that in the presence of strategy choices, infix quantification does not seem to be particularly useful: Quantification of *strategies* that may be different in any state can be handled by strategy choices in a way that is compatible with the epistemic setting, since strategy choices may return different strategies in states that are distinguishable for an agent. On the other hand, infix quantification of *strategy choices* is very unnatural: Strategy choices express “global behavior” of coalitions allowing prior agreement, but during the game only rely on communication that is part of the game itself. Quantification inside formulas would express “prior agreement” *during the game*, which defeats its purpose.

There may be interesting properties that can only be expressed using QAPI^{infix}, but usuallyz QAPI is sufficient and avoids the above problems.

5 SIMULATIONS

Bisimulations relate structures in a truth-preserving way. They allow to obtain decidability results for game structures with infinite state spaces (if a bisimilar finite structure exists), or can reduce the state space of a finite system. In (Schnoor, 2012), our bisimulation results are used to obtain a model-checking algorithm on an infinite structure by utilizing a bisimulation between this structure and a finite one. We give the following definition, which is significantly less strict than the one in (Schnoor, 2010b): For example, our definition can establish bisimulations between structures with different numbers of states (see example below). This is not possible in the definition from (Schnoor, 2010b), since there a bisimulation is essentially a relation Z which is a simulation in *both directions simultaneously*. Since a simulation in the sense of (Schnoor, 2010b) is a function between state spaces, this implies that Z must contain, for every state in one CGS, *exactly* one related state in the other. Hence such a Z induces a bijection between state spaces, and is essentially an isomorphism. The following definition is somewhat simplified to increase readability, it only treats game structures that have a single degree of information, which is therefore omitted here.

Definition 15. *Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be CGSs with state sets Q_1 and Q_2, the same set of players, and the same set of propositional variables. A* probabilistic bisimulation *between $\mathcal{C}_1$ and $\mathcal{C}_2$ is a pair of functions (Z_1, Z_2) where $Z_1\colon Q_1 \to Q_2$ and $Z_2\colon Q_2 \to Q_1$ such that there are move transfer functions Δ_1 and Δ_2 such that for $\{i,\bar{i}\} = \{1,2\}$ and all $q_i \in Q_i$, $q_{\bar{i}} = Z_i(q_i)$, and all coalitions A:*

- *q_i and $q_{\bar{i}}$ satisfy the same propositional variables,*
- *if c_i is a (A,q_i) move, the $(A,q_{\bar{i}})$-move $c_{\bar{i}}(a) = \Delta_i(a,q_i,c_{\underline{i}}(a))$ for all $a \in A$ satisfies that for $\{j,\bar{j}\} = \{1,2\}$ and all $(\overline{A},q_j)$-moves $c_j^{\overline{A}}$, there is a $(\overline{A},q_{\bar{j}})$-move $c_{\bar{j}}^{\overline{A}}$ such that for all $q_i' \in Q_i$, $Pr\left(Z_{\bar{i}}(\delta(q_{\bar{i}},c_{\bar{i}} \cup c_{\bar{i}}^{\overline{A}'})) = q_i'\right) = Pr\left(\delta(q_i,c_i \cup c_i^{\overline{A}}) = q_i'\right)$.*
- *if $q_i \sim_a q_i'$, then $\Delta_i(a,q_i,c) = \Delta_i(a,q_i',c)$ for all c*
- *if $q_i \sim_a q_i'$, then $Z_i(q_i) \sim_a Z_i(q_i')$*
- *if $q_{\bar{i}} \sim_A q_{\bar{i}}'$, there is q_i' with $Z_i(q_i') = q_{\bar{i}}'$ and $q_i \sim_A q_i'$.*
- *$Z_1 \circ Z_2$ and $Z_2 \circ Z_1$ are idempotent.*

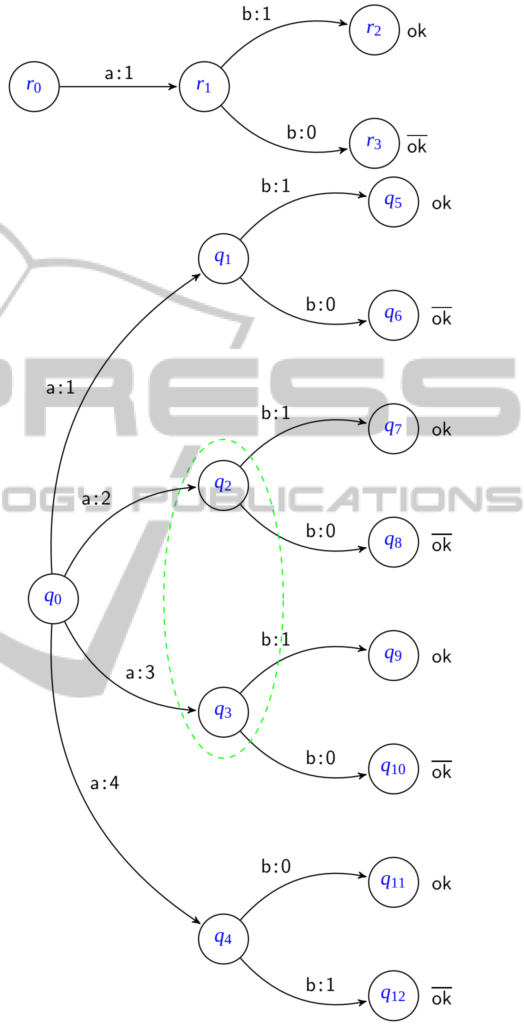

Figure 3: Game Structures $\mathcal{C}_1$ and $\mathcal{C}_2$

Theorem 16. *Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be concurrent game structures, let (Z_1, Z_2) be a probabilistic bisimulation between $\mathcal{C}_1$ and $\mathcal{C}_2$, let q_1 and q_2 be states of $\mathcal{C}_1$ and $\mathcal{C}_2$ with $Z_1(q_2) = q_1$ and $Z_2(q_1) = q_2$. Let φ be a quantified strategy state formula. Then $\mathcal{C}_1, q_1 \models \varphi$ if and only if $\mathcal{C}_2, q_2 \models \varphi$.*

Consider the games $\mathcal{C}_1$ and $\mathcal{C}_2$ in Figure 3. In both, player a starts, he has a single choice in $\mathcal{C}_1$ and 4 choices in $\mathcal{C}_2$. The move by b then determines whether ok holds in the final state. In states r_1 of $\mathcal{C}_1$ and q_1, q_2, and q_3 of $\mathcal{C}_2$, a must play 1 to make ok true, in state q_4 of $\mathcal{C}_2$, he must play 0. States q_2 and q_3 are indistinguishable for a in $\mathcal{C}_2$. CGSs $\mathcal{C}_1$ and $\mathcal{C}_2$

with state sets Q_1 and Q_2 are bisimilar via (Z_1, Z_2), where $Z_2 \colon Q_2 \to Q_1$ is defined as follows:

- $Z_2(q_0) = r_0$,
- $Z_2(q_1) = Z_2(q_2) = Z_2(q_3) = Z_2(q_4) = r_1$,
- $Z_2(q_5) = Z_2(q_7) = Z_2(q_9) = Z_2(q_{11}) = r_2$,
- $Z_2(q_6) = Z_2(q_8) = Z_2(q_{10}) = Z_2(q_{12}) = r_3$.

The move transfer function swaps moves 0 and 1 when transferring from r_1 to q_4. $Z_1 \colon Q_1 \to Q_2$ maps r_0 to q_0, r_1 to q_1, r_2 to q_5 and r_3 to q_6, the move transfer functions map all of a's possible moves in q_0 to the move 1, the moves of b are mapped to themselves (note that q_4 is not used in this direction). It is easy to check that (Z_1, Z_2) is a bisimulation.

Theorem 16 states that state related via both Z_2 and Z_1 satisfy the same formulas. This applies to (r_0, q_0), (r_1, q_1), (r_2, q_5), and (r_3, q_6). The example shows a bisimulation between structures with complete and incomplete information, and with different cardinalities.

6 MODEL CHECKING COMPLEXITY

Model checking is the problem to determine, for a CGS $\mathcal{C}$, a quantified strategy formula φ, and a state q, whether $\mathcal{C}, q \models \varphi$. We state the following results for completeness, the proofs are straight-forward using results and techniques from the literature (Alur et al., 2002; Brázdil et al., 2006; Chatterjee et al., 2007; Schnoor, 2010b). We note that the model-checking problem for MQAPI is undecidable except for restrictions that reduce QAPI to strategy logic.

Theorem 17. *The QAPI model-checking problem is*

1. PSPACE-*complete for deterministic CGSs,*
2. *solvable in* 3EXPTIME *and* 2EXPTIME-*hard for probabilistic structures.*

REFERENCES

Ågotnes, T., Goranko, V., and Jamroga, W. (2007). Alternating-time temporal logics with irrevocable strategies. In (Samet, 2007), pages 15–24.

Alur, R., Henzinger, T. A., and Kupferman, O. (2002). Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713.

Brázdil, T., Brozek, V., Forejt, V., and Kucera, A. (2006). Stochastic games with branching-time winning objectives. In *LICS*, pages 349–358. IEEE Computer Society.

Chatterjee, K., Henzinger, T. A., and Jurdzinski, M. (2006). Games with secure equilibria. *Theor. Comput. Sci.*, 365(1-2):67–82.

Chatterjee, K., Henzinger, T. A., and Piterman, N. (2007). Strategy logic. In Caires, L. and Vasconcelos, V. T., editors, *CONCUR*, volume 4703 of *Lecture Notes in Computer Science*, pages 59–73. Springer.

Chen, T. and Lu, J. (2007). Probabilistic alternating-time temporal logic and model checking algorithm. In Lei, J., editor, *FSKD (2)*, pages 35–39. IEEE Computer Society.

Cortier, V., Küsters, R., and Warinschi, B. (2007). A cryptographic model for branching time security properties - the case of contract signing protocols. In Biskup, J. and Lopez, J., editors, *ESORICS*, volume 4734 of *Lecture Notes in Computer Science*, pages 422–437. Springer.

Halpern, J. Y. and Moses, Y. (1990). Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37:549–587.

Herzig, A. and Troquard, N. (2006). Knowing how to play: uniform choices in logics of agency. In Nakashima, H., Wellman, M. P., Weiss, G., and Stone, P., editors, *AAMAS*, pages 209–216. ACM.

Jamroga, W. (2004). Some remarks on alternating temporal epistemic logic. In *Proceedings of Formal Approaches to Multi-Agent Systems (FAMAS 2003*, pages 133–140.

Jamroga, W. and van der Hoek, W. (2004). Agents that know how to play. *Fundamenta Informaticae*, 63(2-3):185–219.

Samet, D., editor (2007). *Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2007), Brussels, Belgium, June 25-27, 2007*.

Schnoor, H. (2010a). Explicit strategies and quantification for ATL with incomplete information and probabilistic games. Technical Report 1008, Institut für Informatik, Christian-Albrechts-Universität zu Kiel.

Schnoor, H. (2010b). Strategic planning for probabilistic games with incomplete information. In van der Hoek, W., Kaminka, G. A., Lespérance, Y., Luck, M., and Sen, S., editors, *AAMAS*, pages 1057–1064. IFAAMAS.

Schnoor, H. (2012). Deciding epistemic and strategic properties of cryptographic protocols. In Foresti, S., Yung, M., and Martinelli, F., editors, *ESORICS*, volume 7459 of *Lecture Notes in Computer Science*, pages 91–108. Springer.

Schobbens, P.-Y. (2004). Alternating-time logic with imperfect recall. *Electronis Notes in Theoretical Computer Science*, 85(2):82–93.

Walther, D., van der Hoek, W., and Wooldridge, M. (2007). Alternating-time temporal logic with explicit strategies. In (Samet, 2007), pages 269–278.