INSTITUT FÜR INFORMATIK

Explicit Strategies and Quantification for ATL with Incomplete Information and Probabilistic Games

Henning Schnoor

Bericht Nr. 1008
August 2010

CHRISTIAN-ALBRECHTS-UNIVERSITÄT

ZU KIEL

Institut für Informatik der
Christian-Albrechts-Universität zu Kiel
Olshausenstr. 40
D – 24098 Kiel

Explicit Strategies and Quantification for ATL with Incomplete Information and Probabilistic Games

Henning Schnoor

Bericht Nr. 1008
August 2010

e-mail: schnoor@ti.informatik.uni-kiel.de

Explicit Strategies and Quantification for ATL with Incomplete Information and Probabilistic Games

Henning Schnoor

Abstract

We introduce QAPI (quantified ATL with probabilism and incomplete information), an extension of ATL [AHK02] that provides a flexible mechanism to reason about strategies that can be identified and followed by agents that do not have complete information about the state of the system. QAPI allows reasoning about strategies directly in the object language, which allows to express complex strategic properties as equilibria. We show how several other logics can be expressed in QAPI, and provide suitable bisimulation relations, as well as complexity and decidability results for the model checking problem.

Introduction

ATL [AHK02] is the standard logic to reason about strategic properties of temporal systems, naturally modelled as games. The main strategic feature of the language is the operator $\langle\langle . \rangle\rangle$, where $\langle\langle A \rangle\rangle\,\varphi$ expresses "the coalition A has a strategy to achieve the goal φ." In the standard semantics, it is assumed that the players in A have complete information about the current state of the system, hence they are able to *follow* any strategy (a strategy is a function from the state set into the sets of possible moves for a player). Moreover, the existential quantifier is unrestricted and allows strategies that cannot possibly be *identified* by the coalition A.

There are several extensions of ATL that address these issues, see, e.g., [Jam04, JvdH04, HT06, Sch04, Sch10]. There are also extensions of ATL that allow reasoning about strategies directly in the object language [CHP07], which allows expressing complex properties of strategies directly in the formulas (as for example, equilibria conditions). Additionally, there are extensions of ATL to treat probabilistic games [CL07, BJ09].

We introduce the logic QAPI (quantified ATL with probabilism and incomplete information[1]), which combines all of the above advantages and more:

1. QAPI allows restriction to strategies that can be *identified* and *followed* with the knowledge available to the players,
2. we allow reasoning about strategies in the object language, and quantification over these,
3. we cover probabilistic game structures as well as probabilistic statements.

To treat the subtle interplay between identifying and following a strategy, we use strategy choices as introduced in [Sch10]. We show that even in our more general setting, this leads to strictly stronger expressiveness than reasoning about strategies alone.

In addition to combining the above-listed advantages, a key feature of QAPI is the flexibility in the treatment of the "counter-coalition:" Usually in ATL, $\langle\langle A\rangle\rangle\,\varphi$ "frees" the coalition $\overline{A}$ from any commitment, i.e., they abandon any strategies they may currently be following. This is undesirable in several natural scenarios. To address this, several approaches have been suggested which express binding "commitment" of a coalition to a strategy, even when another coalition changes its strategy [WvdHW07, ÅGJ07]. In QAPI, the behavior of $\overline{A}$ is treated very flexibly: It allows formulas to specify whether the coalition $\overline{A}$

- Continues to follow a strategy they are currently implementing—this models that $\overline{A}$ are unaware of what A is attempting to do, or,
- may follow a strategy which is tailor-made to counteract A's efforts to achieve the goal φ—this models $\overline{A}$ being able to react to A's strategy, while still bound to a strategy (possibly with restricted information) themselves, or,
- performs an arbitrary sequence of actions, which may be completely non-uniform and does not need to correspond to any implementable strategy—this expresses the traditional "pessimistic" view of ATl by requiring that the players in A must be successful against every possible behavior of the players in $\overline{A}$.

Integration of all the above features leads to a very powerful logic. In particular, QAPI properly includes previously-suggested logics as strategy logic, ATLES, (M)IATL, ATEL-R*, and ATOL as examples (see Section 2 for details on these logics). Additionally, the flexible approach of QAPI allows to reason about complete-information strategies and incomplete-information strategies in the same formula, and, using concepts from epistemic logic, can express statements as a coalition

[1] In this paper we adopt the notation that *incomplete* information refers to player's uncertainty about the current state of the game, in contrast to *imperfect* information which is often used to express uncertainty about the past.

knowing that their strategy will be successful. This is an often-useful requirement (and is, for example, hard-coded into the definition of admissible strategies in [Sch04]).

Our epistemic operators can be used to treat three notions of knowlegde, which express concepts as "everybody knows," "distributed knowledge," and "common knowledge."

We show that QAPI has a natural notion of a bisimulation, which is a significant relaxation of the one suggested in [Sch10]. In particular, our new definition can establish equivalence of finite and infinite structures, which is impossible with the definition from [Sch10]. We also provide decidability results for the model checking problem of QAPI: The problem is decidable for memoryless strategies, but undecidable in the history-aware case, even for very strong restrictions of both the syntax of QAPI and the games under consideration.

We defer the discussion of related literature to Section 2, where we explain how to express previous extensions of ATL in QAPI.

Organization of the paper. In Section 1, we introduce syntax and semantics of QAPI, and discuss some of its aspects. Section 2 contains examples for the expressiveness of QAPI, including comparison to other extensions of ATL. In Section 3, we introduce bisimulations and prove their relevant properties. Section 4 contains our results on complexity and decidability of the model checking problem, we conclude in Section 4.

1 Semantics of QAPI

1.1 Concurrent game structures

Our definition of a concurrent game structure extends the one from [AHK02] with probabilism (see also [CL07]) and treats incomplete information (see also [JvdH04]):

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

- Σ is a finite set of *players*,
- Q is a set of *states*,
- $\mathbb{P}$ is a finite set of *propositional variables*,
- $\pi\colon \mathbb{P} \to 2^Q$ 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[2] probability distribution $\delta(q, c)$ on Q (the state obtained when in q, all players perform their move as specified by c),
- eq is an *information function* $\mathsf{eq}\colon \{1, \ldots, n\} \times \Sigma \rightarrow \mathcal{P}(Q \times Q)$, where $n \in \mathbb{N}$ and for each $i \in \{1, \ldots, n\}$ and $a \in \Sigma$, $\mathsf{eq}(i, a)$ is an equivalence relation on Q. We also call each $i \in \{1, \ldots, n\}$ a *degree of information*.

A subset $A \subseteq \Sigma$ is a *coalition of $\mathcal{C}$*. We omit "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 often write $\Pr\left(\delta(q, c) = q'\right)$ for $\left(\delta(q, c)\right)(q')$, i.e., we consider $\delta(q, c)$ as a random variable on Q. The function eq expresses incomplete information: For each player it defines an equivalence relation specifying states that a cannot distinguish. By specifying several relations $\mathsf{eq}(1, a), \ldots, \mathsf{eq}(n, a)$ for each player, we can reason over different degrees of information required for goals. We write $q_1 \sim_{\mathsf{eq}_i(A)} q_2$ for $(q_1, q_2) \in \cap_{a \in A} \mathsf{eq}(i, a)$ (no member of A can distinguish q_1 and q_2). $\mathcal{C}$ is *deterministic* if the distributions returned by δ assign 1 to a single state and 0 to all others. A CGS has *complete information* if $\mathsf{eq}(i, a)$ is the equality relation for all i and a.

1.2 Strategies, strategy choices, and formulas

We now introduce the syntax and semantics for QAPI, which is based on ATL*. Informally, $\langle\langle A : \mathrm{S}_1,\ B : \mathrm{S}_2\rangle\rangle_i^{\geq\alpha}\,\varphi$ expresses "if coalition A follows S_1 and coalition B follows S_2, where both coalitions base their decisions only on information available to them in information degree i, then the run of the game satisfies φ with probability $\geq \alpha$." This operator is called *strategy operator*, S_1 and S_2 are variables for so-called *strategy choices* that generalize strategies, see below. We say that an A-strategy choice variable is a symbol S representing strategy choices for A.

Definition 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 formulas are state formulas,
- if A and B 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_1 (S_2) is an A- (B-) strategy choice variable, then $\langle\langle A : \mathsf{S}_1,\ B : \mathsf{S}_2\rangle\rangle_i^{\blacktriangleleft\alpha}\,\psi$ is a state formula,

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

- if A is a coalition, $1 \leq i \leq n$, and ψ is a state formula, then $\mathcal{K}_i^A \psi$ is a state formula,
- every state formula is a path formula,
- conjunctions and negations of path formulas are path formulas,
- 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.

Strategy formulas contain placeholders (strategy choice variables) to be instantiated with strategy choices. 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 often omit the i subscript of the strategy operator, similarly in a deterministic CGS we usually omit the probability bound $\blacktriangleleft \alpha$ (and understand it to be read as ≥ 1 in deterministic structures). When quantifying the appearing strategy choice variables in strategy formulas, one obtains quantified strategy formulas:

Definition 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 $\mathrm{S}_1, \ldots, \mathrm{S}_n$, for an even 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* $\mathcal{C}$. We call φ the *kernel* of the formula.

Note that the above definition does not require that every one of $\mathrm{S}_1, \ldots, \mathrm{S}_n$ actually appears in φ. In the following, we will omit quantification for symbols that do not appear in φ, hence the above also allows formulas that start with existentially quantification, have an odd number of quantifiers, no strict alternation, etc. For the definition of the semantics of quantified formulas, we follow a "constructive" approach, where the values for existentially-quantified variables are given by a function which may depend on the values for the universally quantified formulas. This is technically convenient in our setting, also note that using the Axiom of Choice, this is equivalent to a non-constructive definition.

We allow quantification only in a quantifier prefix of the formula, and not inside the kernel. See Section 1.5.2 for comments on this restriction. We introduce some notation for our semantics definition: A *path* in a CGS $\mathcal{C}$ is a sequence λ of states of $\mathcal{C}$. With $\lambda[i]$ we denote the ith state in λ, with $\lambda[i,k]$ the sequence $\lambda[i], \ldots, \lambda[k]$, and with $\lambda[i,\infty]$ the (possibly infinite) sequence $\lambda[i], \lambda[i+1], \ldots$. A *strategy* fixes a move for a player in each state, it is *uniform* if the moves in indistinguishable states are identical (i.e., the player can follow the strategy given his knowledge).

Definition Let $\mathcal{C}$ be a CGS with state set Q, move function Δ. For a player a, an *a-strategy* in $\mathcal{C}$ is a function s_a such that $s_a(q) \in \Delta(q, a)$ for each $q \in Q$. For a degree of information 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 coalition A, 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 may only depend on the current state and not on the history of the game. History-aware strategies can be handled in the canonical way embedding histories directly into the states, see Section 1.5.1. For technical reasons, we assume that there is some move which every player can play in every situation—since moves can be "renamed" without any consequence, this is not a strong restriction, note that of course the consequences of this move can be very different in every state. However, existence of such a move ensures that i-uniform strategies always exists.

In addition to *follow* a strategy, a player needs to *identify* the correct one. To do this, a player usually only knows the equivalence class of the current state of the game, the goal, and the coalition he is working with, as well as a potential adversial coalition trying to achieve a contrasting goal. In QAPI, this information is directly encoded into formulas, hence the choice of a strategy by a player is based on the current state and the formula representing the goal. Similar uniformity conditions as in the *application* of a strategy of course also apply when *choosing* one: In states that a player cannot distinguish, he has to choose the same strategy. The following notion of a *strategy choice*, introduced in [Sch10], is the natural formalization of this idea:

Definition Let $\mathcal{C}$ be a CGS with state set Q, and let A be a coalition. A *strategy choice for A in $\mathcal{C}$* is a function S such that for each $a \in A$, $q \in Q$, each $\langle\langle . \rangle\rangle_i$-formula φ for $\mathcal{C}$, $\mathsf{S}(a, q, \varphi)$ is an i-uniform a-strategy in $\mathcal{C}$, and if $q_1 \sim_{\mathsf{eq}_i(a)} q_2$, then $\mathsf{S}(a, q_1, \varphi) = \mathsf{S}(a, q_2, \varphi)$.

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 $\mathsf{S}(A, q, \varphi)$. Specifying the information-degree i in the formula allows to specify different amounts of knowledge available to different coalition, or for different goals. Strategy choices are useful for several reasons: In Section 1.7.1 we show they are essential already to correctly handle situations with a single player. For coalitions A containing several players, they have an additional application: Here they model a set of strategies that A agreed upon before the game for a set of possible goals. This allows to express situations where players can rely on shared information to predict the behavior of other members of the coalition without requiring in-game communication.

In our semantics, we are concerned with the case where two (not necessarily disjoint) coalitions A and B are each following their own strategy choice, S_A and S_B, respectively. The "joint strategy choice" they are following is denoted by $\mathsf{S}_A \circ \mathsf{S}_B$, which is a strategy choice for $A \cup B$ and defined as

$$\mathsf{S}_A \circ \mathsf{S}_B(a, q, \varphi) = \begin{cases} \mathsf{S}_A(a, q, \varphi), & \text{if } a \in A, \\ \mathsf{S}_B(a, q, \varphi), & \text{otherwise, i.e., if } a \in B \setminus A. \end{cases}$$

While players in the scope of a strategy operator follow the corresponding strategy, the remaining players (often called the "counter-coalition"), have two possibilities: They can follow a (possibly uniform) strategy on their own, or can behave in an arbitrary way that cannot be defined using a strategy. For example, they may perform different moves when encountering the same state twice in the run of a game. For the second case, we use the following notion: A *response* to a coalition A is a function r such that $r(t, q)$ is a $(\overline{A}, q)$-move for each $t \in \mathbb{N}$ and each $q \in Q$. Hence a response is an arbitrary reaction to the (probabilistic) outcomes of a possible strategy chosen by A. It models that in the i-th step of a run, the coalition $\overline{A}$ performs the move $r(i, q)$, if the current state is q. A third possibility is to require that not only does the counter-coalition have to follow a (uniform) strategy, but more strongly require that they keep following a strategy that they are currently implementing due to a previous application of a strategy operator, see Section 2.2.

When 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 allows us to define the "success probability" of strategies in a natural way:

Definition 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 $q \in Q$,

$$\Pr\left(q \to M \mid s_A + r\right)$$

is the probability that in the Markov process resulting from $\mathcal{C}$, s_A, and r with initial state q, the resulting path is an element of M.

1.3 Evaluating Formulas

We now define the semantics for QAPI. We proceed in two stages: We first define truth for strategy formulas, where interpretations for the appearing strategy choice variables are given. In a second step, we then define truth for formulas where variables are quantified. The definition is the natural one: Propositional variables and operators are handled as usual, temporal operators behave as in linear-time temporal logic, and $\langle\langle A_1 : \mathsf{S}_1,\ B : \mathsf{S}_2\rangle\rangle_i^{\geq\alpha}\,\psi$ expresses that when coalitions A_1 and A_2 follow the strategy choices S_1 and S_2, with information degree i available, the probability that the run of the game satisfies ψ is at least α. The semantics of the knowledge operator $\mathcal{K}$ is standard. In the following, the strategy choice variables $\mathrm{S}_1, \ldots, \mathrm{S}_n$ are instantiated with strategy choices $\mathsf{S}_1, \ldots, \mathsf{S}_n$.

Definition Let $\mathcal{C} = (\Sigma, Q, \mathbb{P}, \pi, \Delta, \delta, \mathsf{eq})$ be a CGS, let $A_1, \ldots, A_n$ be coalitions, and for each i let S_i be an A_i-strategy choice variable, and let S_i be a strategy choice for A_i. Let φ_1, φ_2 be state formulas, let ψ_1, ψ_2 be path formulas, let λ be a path over Q, let $t \in \mathbb{N}$. We use $\overrightarrow{\mathsf{S}}$ as abbreviation for $(\mathsf{S}_1, \ldots, \mathsf{S}_n)$. We define

- $\mathcal{C}, \overrightarrow{\mathsf{S}}, q \models p$ iff $q \in \pi(p)$ for $p \in \mathbb{P}$,
- $\mathcal{C}, \overrightarrow{\mathsf{S}}, q \models \neg\varphi_1$ iff $C, \overrightarrow{\mathsf{S}}, q \not\models \varphi_1$,
- $\mathcal{C}, \overrightarrow{\mathsf{S}}, q \models \varphi_1 \wedge \varphi_2$ iff $C, \overrightarrow{\mathsf{S}}, q \models \varphi_1$ and $C, \overrightarrow{\mathsf{S}}, q \models \varphi_2$,
- $(\lambda, t), \overrightarrow{\mathsf{S}} \models \varphi_1$ iff $\mathcal{C}, \overrightarrow{\mathsf{S}}, \lambda[t] \models \varphi_1$,
- $(\lambda, t), \overrightarrow{\mathsf{S}} \models \neg\psi_1$ iff $(\lambda, t), \overrightarrow{\mathsf{S}} \not\models \psi_1$,
- $(\lambda, t), \overrightarrow{\mathsf{S}} \models \psi_1 \wedge \psi_2$ iff $(\lambda, t), \overrightarrow{\mathsf{S}} \models \psi_1$ and $(\lambda, t), \overrightarrow{\mathsf{S}} \models \psi_2$,
- $(\lambda, t), \overrightarrow{\mathsf{S}} \models \mathsf{X}\psi_1$ iff $\lambda[t+1, \infty], \overrightarrow{\mathsf{S}} \models \psi_1$,
- $(\lambda, t), \overrightarrow{\mathsf{S}} \models \mathsf{P}\psi_1$ iff there is some $t' <= t$ and $(\lambda, t'), \overrightarrow{\mathsf{S}} \models \psi_1$,
- $(\lambda, t), \overrightarrow{\mathsf{S}} \models \mathsf{X}^{-1}\psi_1$ iff $t >= 1$ and $(\lambda, t-1), \overrightarrow{\mathsf{S}} \models \psi_1$,
- $(\lambda, t), \overrightarrow{\mathsf{S}} \models \psi_1 \mathsf{U} \psi_2$ iff there is some $i \geq t$ such that $\lambda[i, \infty], \overrightarrow{\mathsf{S}} \models \psi_2$ and $\lambda[j, \infty], \overrightarrow{\mathsf{S}} \models \psi_1$ for all $t \leq j < i$,
- $\mathcal{C}, \overrightarrow{\mathsf{S}}, q \models \mathcal{K}_i^A \varphi_1$ if $\mathcal{C}, \overrightarrow{\mathsf{S}}, q' \models \varphi_1$ for all $q' \in Q$ with $q' \sim_{\mathsf{eq}_i(A)} q$,
- $\mathcal{C}, \overrightarrow{\mathsf{S}}, q \models \underbrace{\langle\langle A_i : \mathsf{S}_1,\ A_j : \mathsf{S}_2\rangle\rangle_i^{\blacktriangleleft\alpha}\,\psi}_{=:\varphi_1}$ iff for every response r to $A_i \cup A_j$, we have

$$\Pr\left(q \rightarrow \left\{\lambda \mid (\lambda, 0), \overrightarrow{\mathsf{S}} \models \psi_1\right\} \mid \mathsf{S}_i \circ \mathsf{S}_j(A_i \cup A_j, q, \varphi_1) + r\right) \blacktriangleleft \alpha,$$

The availabe knowledge is relevant in the above definition in three ways: The semantics of the $\mathcal{K}$-operator uses the indistinguishability relations directly, and the application of the strategy choices results in the mentioned requirements that strategies can be identified and followed with the knowledge available to the players. The above can be generalized to mention more than 2 coalitions in the natural way, i.e., a formula $\langle\langle A_1 : \mathsf{S}_1,\ A_2 : \mathsf{S}_2,\ A_3 : \mathsf{S}_3\rangle\rangle_i^{\alpha}\,\psi$ can be used to let coalitions A_1, A_2, and A_3 play the strategies for ψ chosen by strategy choices S_1, S_2, and S_3 for information degree, which is defined using the strategy choice $\mathsf{S}_1 \circ \mathsf{S}_2 \circ \mathsf{S}_3$ defined in the obvious way as $\mathsf{S}_1 \circ (\mathsf{S}_2 \circ \mathsf{S}_3)$. Similarly, we write $\langle\langle A : \mathsf{S}\rangle\rangle_i^{\blacktriangleleft\alpha}\,\varphi$ for $\langle\langle A : \mathsf{S},\ \emptyset : \emptyset\rangle\rangle_i^{\blacktriangleleft\alpha}\,\varphi$.

Given the above definition covering the case where "instantiations" of all appearing variables are known, the semantics definition for quantified strategy formulas is the natural one:

Definition Let $\mathcal{C}$ be a CGS, let $\psi = \forall \mathrm{S}_1 \exists \mathrm{S}_2 \forall \mathrm{S}_3 \ldots \exists \mathrm{S}_n \varphi$ 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, $\mathcal{C}, q \models \psi$, if for each $i \in \{2, 4, \ldots n\}$, there is a function s_i such that for all strategy choices 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 $\mathcal{C}, \mathsf{S}_1, \ldots, \mathsf{S}_n, q \models \varphi$.

It is often convenient to be able to restrict quantification to a smaller set of strategy choices. In particular, *constant* strategy choices (which may 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, with the obvious semantics. This allows us to express equality of probabilities in formulas: $\langle\langle A \rangle\rangle_i^{=\alpha} \psi$ can be expressed as $\langle\langle A \rangle\rangle_i^{\geq\alpha} \psi \wedge \langle\langle A \rangle\rangle_i^{\leq\alpha} \psi$ and requiring the strategy choice to be constant in the quantifier block. Additionally, with $\exists_m$ and $\forall_m$ we denote quantification restricted to so-called "maximal" strategy choices that proved useful in [Sch10].

See Section 2.2 for an application of restricted quantification.

1.4 Different notions of knowledge

In the above definition of the semantics of the knowledge operator $\mathcal{K}$, we have adopted the notion of *distributed knowledge*—a fact is known by a group of agents if it is derivable from their combined knowledge. There are several other ways to define knowledge of a coalition, see e.g., [JvdH04]: One can define "Coalition A knows φ in q" to be true if $\mathcal{C}, \mathsf{S}, q' \models \varphi$ for all q' such that

"Everybody knows" $q' \sim_A^E q$, where $\sim_A^E$ is the union of all $\sim_a$ for $a \in A$ (i.e., every agent has to know φ individually),

"Common knowledge" $q' \sim_A^C$, where $\sim_A^C$ is the reflexive, transitive closure of $\sim_A^E$ (i.e., every agent has to know that every agent knows, etc.),

"Distributed knowledge" $q' \sim_A^D$, where $\sim_A^D$ is the intersection of all $\sim_a$ for $a \in A$ (i.e., as in the definition of our $\mathcal{K}$-operator).

We briefly explain how "Everybody knows" and "Common knowledge" can be expressed in our formalism:

- *Everybody knows φ (wrt. information degree i)* can be expressed as $\bigwedge_{a \in A} \mathcal{K}_i^a \varphi$,
- *φ is common knowledge (wrt. information degree i)* can be expressed as: For every $n \in N$, for every $a_1, a_2, \ldots, a_n \in A$, it holds that $\mathcal{K}_i^{a_1} \mathcal{K}_i^{a_2} \ldots \mathcal{K}_i^{a_n} \varphi$.

Note that the translation for common knowledge into our semantics leads to an infinite conjunction of formulas, and hence enriching the language with a special operator for common knowlegde is sensible to keep formulas finite. However, we use the above translation to demonstrate that all three notions of knowledge can be expressed in our semantics, which allows us to consider only a single case when proving the correctness results for simulations in Section 3.

It is thus possible to introduce operators $\mathcal{K}^E$ and $\mathcal{K}^C$ to denote the above versions of knowledge (renaming the standard operator to $\mathcal{K}^D$), and all the results in the paper are also true for the richer language—it is clear that the algorithms in Section 4 can easily be modified to handle these operators.

1.5 Discussion of Semantics

1.5.1 History-Dependence

History-dependent strategies allow a player to determine his current move based on the entire *history* of the game, not merely the current state. Formally, QAPI only treats memoryless strategies, however history-dependence can be treated canonically: For a CGS $\mathcal{C}$, $\mathcal{C}^{hst}$ is obtained from $\mathcal{C}$ by encoding histories into states in the canonical way (see also [ÅGJ07, Sch10]).

The following formal definition uses the obvious treatment of incomplete information, where a player a can distinguish histories if they have different length of if he can distinguish between individual points in time.

Definition Let $\mathcal{C} = (\Sigma, Q, \mathbb{P}, \Pi, \Delta, \delta, \mathsf{eq})$ be a $\mathcal{C}$ with n degrees of information. Then $\mathcal{C}^{hst}$, the *history-dependent version of* $\mathcal{C}$, is defined as $\mathcal{C}^{hst} = (\Sigma, Q^+, \mathbb{P}, \Pi', \Delta', \delta', \mathsf{eq}')$, where

- Q^+ is the set of non-empty, finite sequences over Q,
- for $p \in \mathbb{P}$, $\Pi'(p) = \{q_1 \dots q_n q \mid q \in \Pi(p), q_1, \dots, q_n \in Q\}$,
- $\Delta'(q_1 \dots q_n, a) = \Delta(q_n, a)$,
- For a state $q_1 \dots q_n \in Q^+$ and a (Σ, q)-move c, $\delta'(q_1 \dots q_n, c)$ is defined as $q_1 \dots q_n \delta(q_n, c)$,
- $\mathsf{eq}'(i, a)$ is defined as follows: For sequences $q^1 = q_1 \dots q_n$ and $q^2 = q'_1 \dots q'_m$ we have $q^1 \sim_{\mathsf{eq}'_i(a)} q^2$ if and only if $n = m$ and $q_j \sim_{\mathsf{eq}_i(a)} q'_j$ for all $j \leq n$.

1.5.2 In-line quantification

Our definition restricts quantification over strategy choices to a quantifier block at the beginning of a formula. This is natural since the purpose of strategy choices as introduced in [Sch10] is to express "global behavior" of coalitions, allowing a maximal amount of pre-agreement among a coalition, but only using available information during the game. In particular, quantification over strategy choices inside formulas represents performing "prior agreement" *during the run of the game*, which defeats its purpose.

1.5.3 Mixing Uniformity Requirements

In QAPI, players have the same information when identifying and when following a strategy. This is consistent with the intuitive notion of basing decisions on available information. Nonetheless, one could also consider strategy choices where this is not the case, and allow a certain non-uniformity in the choice of strategies. As an example this allows to move quantifiers over *strategies* appearing inside formulas to the quantifier block at the beginning of the formula: Consider the formula

$$\Box \exists_c \mathrm{S}_a \, \langle\langle a : \mathrm{S}_a \rangle\rangle_2 \, \varphi,$$

expressing that at every future state, player a has a strategy (note that S_a must reference a constant strategy choice, i.e., a strategy) to achieve φ. This formula cannot be equivalently rewritten by moving the quantifier to the front: In this case, the strategy may not depend on the state anymore, but must choose its actions based on the information that a has about the state. Using strategy choices allowing to use complete knowledge to *identify* a strategy, but still requiring strategies to be 2-uniform, we can rewrite the formula as

$$\exists^* \mathrm{S}_a \, \langle\langle a : \mathrm{S}_a \rangle\rangle_2 \, \varphi,$$

where the quantification $\exists^* \mathrm{S}_a$ is interpreted as quantifying over strategy choices as explained above (note that the dual case, where the application of a strategy has more information available than its choice is not interesting—the "missing" information for choosing the strategy can be hard-coded into the strategy itself. However, obviously the case of incomparable degrees of information is relevant).

This construction is very unnatural: Allowing, *at the same state of the game, for the same player*, different information for identifying and for following a strategy is unlikely to capture realistic situations. In particular, one loses the intended purpose of strategy choices to identify correct strategies with the available information. One ostensible advantage of the above construction is to embed the semantics of *strategy logic* [CHP07] into QAPI (strategy logic uses quantification over strategies inside the formulas). However, strategy logic is concerned only with the complete-information, history-dependent setting, hence the above-mentioned issues to not appear, and rewriting can be performed with the standard QAPI semantics.

1.6 Possible Generalizations

There are several ways in which QAPI can be extended to gain additional expressivity. An obvious possibility is to generalize the temporal operators allowed to construct path formulas. In stochastic game logic [BBGK07], Rabin automata are used to describe path properties. QAPI can be generalized in this direction in the straight-forward way. Our results on simulation (see Section 3) remain true for this generalization, since simulations ensure that corresponding paths have the exact same atomic and strategic properties. Such a generalization of QAPI properly contains stochastic game logic.

Another possible generalization is to combine history-dependent and memoryless strategies in a single model, and allow formulas to specify (via special quantification) whether strategies may use history. A way to express this feature directly in the semantics of QAPI is to introduce degrees of information in $\mathcal{C}^{hst}$ that model bounded memory (see also Section 2.9). Again, our simulation result still holds for this case, as the question whether a strategy uses history to determine a move is invariant under our construction. However, Theorem 4.2 implies that the model checking problem remains undecidable even when only a single quantifier ranging over history-dependent strategy choices may appear in the formula, even when the game structure is deterministic or has complete information—in the case where both constraints are met simultaneously, we are in the case of strategy logic [CHP07], and thus decidable.

Additionally, to further increase the ability of formulas to reason about strategies, one could allow atomic formulas of the form $\mathit{strategychoice\ variable}_1(a_1, \varphi_1) = \mathit{strategychoice\ variable}_2(a_2, \varphi_2)$, which evaluate to true if the strategy choices instantiating the strategy choice variables $\mathit{strategychoice\ variable}_1$ and $\mathit{strategychoice\ variable}_2$ return the same strategies for players a_1 and a_2 for the goals φ_1 and φ_2 in the current state. Such formulas would be "atomic" even though φ_1 and φ_2 may be arbitrarily complicated, since φ_1 and φ_2 are not evaluated, but only used as input to the strategy choices. This would allow to let formulas "force" players to use the same strategies for different formulas, etc. Such a generalization is straight-forward, and clearly both our simulation and complexity results are unchanged for this addition.

Finally, we introduced restricted quantification over *constant* strategy choices. In other settings, quantification over strategy choices that return the same strategy in the same state, but not necessarily for different formulas may be convenient. However, such a notation essentially is "syntactic sugar," since one can rewrite

QAPI-formulas such that different strategy choices are used for a different subformula.

As for notation, there are many standard constructs for which "syntactic sugar" may be introduced: For example one may use a notation that expresses a binding commitment of a coalition to a strategy choice which is also propagated to subformulas (which can be expressed in QAPI by repeating the relevant coalition and strategy choice in every appearing strategy operator of a subformula), etc.

1.7 Strategies vs. Strategy Choices

In many situations, considering strategies instead of strategy choices is sufficient in the *quantified* setting: Using quantification one can force coalitions to use the same strategy in different states, for example in $\exists_c \mathrm{S} \Box \langle\langle A : \mathrm{S} \rangle\rangle_2 \varphi$. This avoids the non-uniformity in the choice of strategies in classical ATL* as in the formula $\Box \langle\langle A \rangle\rangle \varphi$. In particular, our translations from Section 2 only uses quantification over constant strategy choices to express previously-introduced extensions of ATL* into QAPI. However, it turns out that it is not possible to restrict QAPI to strategies only without losing expressivity. In this section, we show that there are natural situations that cannot be expressed when using only strategies, but require the more general setting of strategy choices. Since, however, restriction to strategies is sufficient in many settings, we then introduce simplified notation for this case.

1.7.1 Strategy choices are more expressive than strategies

Consider the CGS in Figure 1. The players are a and b, the game starts in q_0. In every state, a has two moves m_1 and m_2; however their only influence is in the final step. The first move is by b, and controls whether the next state is q_1^A or q_1^B. For every possible move, q_1^X is followed by q_2^X for $X \in \{A, B\}$. Played in q_2^X, the move m_1 leads to a state satisfying the propositional variable ok if and only if $X = A$, conversely m_2 is successful only if $X = B$. Player a has complete information except that he cannot distinguish between q_2^A and q_2^B. This expresses a situation where a has insufficient memory to remember all of the information learned during the run of the game.

We now ask whether a has a strategy to achieve the goal $\Diamond\mathsf{ok}$ starting from q_1^A or q_1^B. Using only strategies, the situation is expressed with the formula

$$\exists_c \mathrm{S}\ \mathsf{X} \langle\langle a : \mathrm{S} \rangle\rangle\ \mathsf{X} \Diamond \mathsf{ok},$$

using quantification over constant strategy choices to simulate quantification over strategies. Clearly, there is no suitable uniform strategy for a: Such a strategy needs to play the same move in q_2^A and in q_2^B, and thus is unsuccessful in one of

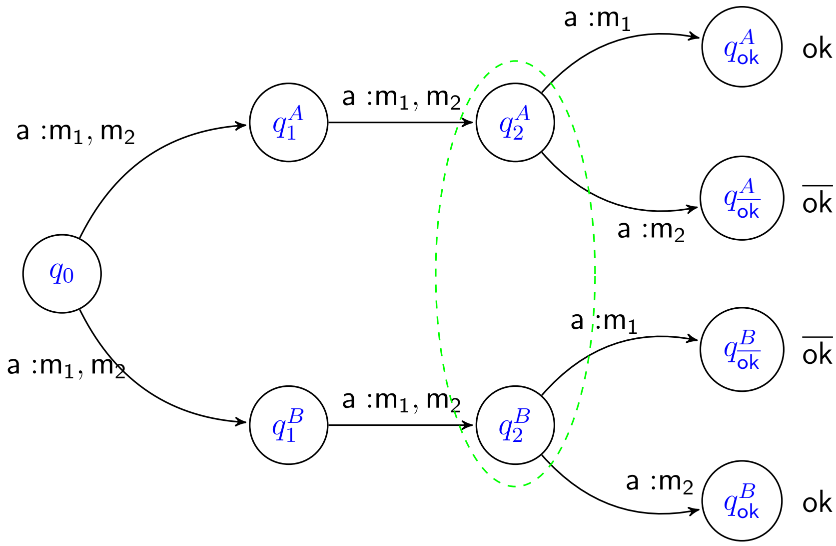

Figure 1: Strategy choices are necessary

these—the above formula is not satisfied in q_0. However, if we allow using a *strategy choice*, the corresponding formula

$$\exists \mathsf{S}\ \mathsf{X}\ \langle\langle a : \mathrm{S} \rangle\rangle\ \mathsf{X} \Diamond \mathsf{ok}$$

is true in q_0: For instantiating S, we use a strategy choice that in q_1^A returns a strategy playing m_1 in every state, and in q_1^B returns one always playing m_2. Both strategies are trivially uniform and clearly, this strategy choice achieves the goal.

The crucial point here is the distinction between states where a strategy is *identified* and where it is *executed*: In state q_1^A or q_1^B, player a uses the available information to choose the strategy to play from then on. When using only strategies, the knowledge has to be present at the time of *performing* a move. The use of strategy choices allows a player to "remember" decisions made in previous states.

If "information loss" does not appear, strategies are sufficient: The decision which strategy to use made by the strategy choice can be encoded in the strategies themselves, as the relevant information is still available later. Informally, the strategy can be defined as in state q, to "look back" to the state q' where a strategy choice S would be applied, evaluate S at q', and return the move intended for the state q. In particular, in structures $\mathcal{C}^{hst}$, strategies can be used instead of strategy choices, as long as different strategies are used for different goals (another feature of a strategy choice is that the strategy may depend on the goal; however this easily can be ensured by using different strategy-symbols for subformulas describing different goals).

The example above also shows that strategy choices are not only relevant when considering coalitions of size at least 2—in which case they can be used to express prior agreement—but are also essential for reasoning about the strategic possibilities of a single player.

1.7.2 Simplifying QAPI: Considering Strategies Only

Despite the above-mentioned shortcomings, a restriction of QAPI to only strategies is useful in some applications. We define simplified notation for this case. We define formulas etc. as before, the only difference is that quantifiers explicitly mention the information degree of the strategy: $\exists_i \mathrm{S}$ means "there exists an i-uniform strategy." Since the information degree is handled in the quantification, it is not repeated in the strategy operator; this operator (for coalitions A_1, …, A_n and variables S_1, …, S_n representing strategies) is now

$$\langle\langle A_1 : \mathrm{S}_1, \ldots, \ A_n : \mathrm{S}_n \rangle\rangle^{\blacktriangleleft \alpha} .$$

The semantics are the obvious ones inherited from the more general setting. Note that here it is possible to mix strategies of different degrees of information in a single coalition operator. Hence a translation into the standard setting needs to add new information degrees constructed from Cartesian products of the original information degrees. Note, however, that the application of our simulation results (see Section 3) is still possible without any further requirements: Since mixing of different degrees only appears in the choice (and thus translation) of strategies, and not in the knowledge operator, the knowledge transfer property does not need to hold for the new information degrees obtained by Cartesian products—for translating strategies, only move uniformity is required, and this property is only concerned with the knowledge of a single player.

2 Examples

We now give applications of QAPI. We first note that trivially, QAPI is a generalization[3] of both ATL* and the semantics introduced in [Sch10]. We will see below that other interesting extensions of ATL/ATL* can be expressed in QAPI as well. The logics that we "simulate" only talk about single strategies. In our translation, we will still use strategy choices, since in many cases, the strategy in question may depend on the state—for example, in standard ATL, the formula $\langle\langle A \rangle\rangle \Box \langle\langle B \rangle\rangle \varphi$

[3] Note that when "translating" ATL* into QAPI, the implicit existential quantification over strategies in ATL* is moved to the quantifier block at the beginning of the formula; depending on the parity of nested negations, the quantifier remains existential or becomes universal. We also note that the same construction addresses the problematic behaviour of negation in [Sch10], which in that paper was addressed with the introduction of so-called maximal strategy choices.

expresses that the coalition A has a way to ensure that in every future state, the coalition B has a strategy to enforce φ. The strategy for B may be different for every state. Hence the most natural way to express this situation is to use a strategy choice for B, which may choose a different strategy in every state.[4] Also, some of the logics we mention below do not treat incomplete information, hence whenever in the following we do not mention the information degree, we assume that the involved CGSs only have a single degree of information, which is complete information for all players. Similarly, previous logics do not treat probabilism, hence we omit probabilities in the following discussion, and understand them to all be ≥ 1.

2.1 An adversary with incomplete information

An obvious application is modeling situations where a coalition A plays against $\overline{A}$, and the latter also is assumed to follow a strategy that can be implemented using incomplete information. To express "A can achieve φ against every possible strategy of B," we use the formula

$$\exists \mathrm{S}_1 \forall \mathrm{S}_2 \, \langle\langle A : \mathrm{S}_1,\ B : \mathrm{S}_2 \rangle\rangle_i^{\geq \alpha} \, \varphi,$$

which expresses the informal statement. Note that this is strictly weaker than saying $\exists \mathrm{S}_1 \, \langle\langle A : \mathrm{S}_1 \rangle\rangle_{\geq \alpha}^{i} \, \varphi$, since in the latter formula, the coalition $\overline{A}$ is not restricted to following any (uniform or otherwise) strategy, whether in the first one, $\overline{A}$'s behavior has to follow a uniform strategy.

2.2 Sub-coalitions changing their strategy

In many natural scenarios, when a sub-coalition A' of a coalition A changes their strategy, they can rely on the players in the remainder of the coalition to continue following their previously-agreed strategies. In standard ATL*, as well as in the semantics from [Sch10], this cannot be expressed: The operator $\langle\langle A' \rangle\rangle$ considers every possible behavior of players not in A'. In other words, the binding of the operator $\langle\langle . \rangle\rangle$ does not propagate to $\langle\langle . \rangle\rangle$-subformulas.

QAPI allows to express this situation. Consider a coalition A working together to reach a state from which on $A' \subseteq A$ has a strategy to achieve a goal φ_1 and a strategy to achieve goal φ_2, *if they can rely on players in $A \setminus A'$ to continue their earlier strategy.* We can express this as

[4]Since a strategy itself may also depend on the state, it may be possible to work with a single strategy—i.e., a constant strategy choice— instead, however we feel that the existential quantification in each state is more naturally expressed using existential quantification of strategy choices. See Section 1.7.1 for a formal example why it is not always possible to only consider strategies instead of strategy choices.

$$\exists_c \mathrm{S}_A \exists \mathrm{S}_{A'} \langle\langle A : \mathrm{S}_A \rangle\rangle \Diamond \left(\langle\langle A' : \mathrm{S}_{A'},\ A : \mathrm{S}_A \rangle\rangle \Diamond\varphi_1 \wedge \langle\langle A' : \mathrm{S}_{A'},\ A : \mathrm{S}_A \rangle\rangle \Diamond\varphi_2 \right),$$

recall that $\exists_c$ quantifies over constant strategy choices. This expresses that A sticks to a fixed strategy and does not change their behavior depending on whether the subcoalition A' attempts to achieve the goal φ_1 or φ_2. Recall that in $\langle\langle A' : \mathrm{S}_{A'},\ A : \mathrm{S}_A \rangle\rangle$, the members of A' are bound to $\mathrm{S}_{A'}$, even if $A \cap A' \neq \emptyset$.

2.3 Knowing whether a strategy is successful

In QAPI, it is possible to reason about a coalition's knowledge about the success probability of their own strategies. For example, the following formula expresses "there is a strategy for A such that there is no strategy for B such that the coalition C can know that its application successfully achieves φ:"

$$\exists \mathrm{S}_A \forall \mathrm{S}_B \neg \mathcal{K}_2^C \langle\langle A : \mathrm{S}_A,\ B : \mathrm{S}_B \rangle\rangle_2^{\geq 1} \varphi.$$

Note that this is very different from expressing that A has a strategy preventing φ, i.e., $\exists \mathrm{S}_A \langle\langle A : \mathrm{S}_A \rangle\rangle_2^{\geq 1} \neg\varphi$, 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.

2.4 Winning Secure Equilibria

A *winning secure equilibrium*, defined in [CHJ06], is a special case of a Nash equilibrium: In a game with two players a and b, where a (b) wants to achieve the goal φ_a (φ_b), a winning secure equilibrium is a pair of strategies (s_a, s_b) such that both goals are achieved when both a and b stick to the strategies, and as soon as b plays a strategy s'_b such that the joint strategy obtained from s_a and s'_b does not satisfy φ_a anymore, then b's goal φ_b is not satisfied anymore either, symmetrically for player a. In QAPI, this can easily be expressed with

$$\exists_c \mathrm{S}_a \exists_c \mathrm{S}_b \langle\langle a : \mathrm{S}_A,\ b : \mathrm{S}_B \rangle\rangle (\varphi_a \wedge \varphi_b) \wedge \langle\langle a : \mathrm{S}_A \rangle\rangle (\varphi_b \to \varphi_a) \wedge \langle\langle b : \mathrm{S}_B \rangle\rangle (\varphi_a \to \varphi_b).$$

The formula expresses that both goals are reached if both players stick to their strategies, and (from player a's point of view), b cannot achieve φ_b without a achieving φ_a, as long as player a follows S_A. Note that we again use quantification over constant strategy choices here, since none of the players is likely to know whether the other one continues playing the strategy from the equilibrium.

2.5 Expressing Strategy logic

Strategy logic, introduced in [CHP07], is very similar to our semantics. Essentially, strategy logic extends ATL* with explicit quantification over strategies, and then express statements as "if player a plays strategy x, player b plays strategy y, the resulting game satisfies the formula Ψ." Since strategy logic covers games with complete information and history-awareness only, quantification can be moved to the start of the formula (see Section 1.5). Therefore, our semantics strictly contains strategy logic (containment is strict since we also model incomplete information and probabilistic games). Note that since strategy logic contains many other logics as for example game logic, this holds true for our semantics as well.

2.6 Expressing ATLES

ATLES, ATL with explicit strategies, was introduced by Walther, van der Hoek, and Wooldridge in [WvdHW07]. The focus of their extension of ATL is to model that players can *commit* to a given strategy. They use a *commitment function* ρ, fixing strategies for players in its domain. The formula $\langle\langle A\rangle\rangle_\rho\,\varphi$ is interpreted as (for the natural formal definition, see [WvdHW07]):

> "Given the commitments of the agents $b \in \mathsf{dom}(\rho)$ to use strategy $\rho(b)$, the agents $a \in A \setminus \mathsf{dom}(\rho)$ have a strategy such that, no matter what the agents $c \in \Sigma \setminus (\mathsf{dom}(\rho) \cup A)$ will do, φ will result." [WvdHW07]

In QAPI, the role of ρ is played by a constant strategy choice for the coalition $B = \mathsf{dom}(\rho)$, which we denote with S_ρ. S_ρ is obtained by returning, for each player $a \in B$, the strategy fixed for a by ρ. The formula $\langle\langle A\rangle\rangle_\rho\,\varphi$ can be written as $\langle\langle B : \mathsf{S}_\rho,\ A : \mathsf{S}_A\rangle\rangle\,\varphi$, where S_A is (existentially or universally, depending on even or odd number of negations) quantified to a constant strategy choice, and appears only once in the kernel. This mirrors the semantics of [WvdHW07] exactly: The players in B follow the strategy fixed by ρ, the players in A follow an existentially quantified strategy, and the remaining players act arbitrarily.

The above treatment leaves out some details from [WvdHW07], they especially address the question of how to represent commitment functions in formulas using *strategy terms*. It is clear that an analogous representation can be used in our setting. Additionally, their paper contains a complete axiomatization of ATLES, and complexity results which give more efficient algorithms than the ones obtained by our translation. Finally, they have decidability results for the history-dependent case, which is undecidable in the more general QAPI.

2.7 Expressing (M)IATL

These variants of ATL were introduced by Ågotnes, Goranko, and Jamroga in [ÅGJ07]. Their key feature here is to treat strategies as *irrevocable*: Once a coalition A follows a strategy after an application of the $\langle\langle A\rangle\rangle$-operator, it does not change this strategy when the nested formula is evaluated. Formally, when applying $\langle\langle A\rangle\rangle\,\varphi$, a strategy for A is chosen (existentially quantified as usual) and the CGS is modified to essentially hard-code this chosen strategy into the structure.

To translate IATL into our semantics, replace every formula $\langle\langle A\rangle\rangle\,\varphi$ with $\langle\langle A : \mathsf{S}\rangle\rangle\,\varphi$, and in φ, recursively replace every $\langle\langle A_1 : \mathsf{S}_1,\ A_2 : \mathsf{S}_2,\ \dots\rangle\rangle\,\psi$ with $\langle\langle A : \mathsf{S},\ A_1 : \mathsf{S}_1,\ A_2 : \mathsf{S}_2,\ \dots\rangle\rangle\,\psi$, where S is existentially/universally (again depending on the parity of negation) quantified at the beginning of the formula.

The variant MIATL extends IATL with history-aware strategies (IATL with **m**emory). To model hard-coding of strategies into the CGS, the authors use an "unfolding" of the CGS $\mathcal{C}$ in the same way as in the definition of $\mathcal{C}^{hst}$. The authors also observe that IATL and MIATL are not invariant under a standard notion of bisimulation, showing that IATL can express things which ATL cannot. More generally they prove that the expressiveness of ATL and IATL (MIATL) is not comparable. Additionally, they provide complexity results for model checking. Note that our results in Section 3 imply that IATL and MIATL are invariant under our notion of bisimulations, as they are contained in QAPI.

2.8 Expressing ATEL-R*

The logic ATEL-R* was introduced by Jamroga and van der Hoek in [JvdH04], and is (among other things) concerned with the same issues that lead to our definition of strategy choices, i.e., the observation that it is not sufficient to require strategies to be uniform, but that the act of *identifying* a strategy also should be consistent with the knowledge of the agent(s) in question. The semantics of the key operator of ATEL-R* are defined as follows (right-hand side is in our notation)—in the following, both A and Γ are coalitions, where A is the coalition *playing*, and Γ is the coalition *identifying* the strategy:

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

The above can be translated into QAPI by writing

$$\mathcal{C}^{hst}, q' \models \mathcal{K}_1^\Gamma\,\langle\langle A : \mathsf{S}_A\rangle\rangle_1\,\varphi,$$

where again the quantification of S depends on the parity of negation and is restricted to constant strategy choices.[5] We note that in [JvdH04], it is explicitly mentioned that requiring that Γ knows that A has a strategy to achieve φ is not sufficient to express their definition of $\langle\langle A\rangle\rangle_{\mathcal{K}(\Gamma)}\varphi$. It works in our case due to the difference in quantification: In QAPI, S_A is quantified *before* the application of the knowledge-operator for Γ, and thus Γ knows that a *fixed* A-strategy will be successful. In their semantics, quantification of the A-strategy would happen *after* the application of the knowledge-operator in a formula like $\mathcal{K}_\Gamma\langle\langle A\rangle\rangle\varphi$, hence the coalition A would be allowed to choose a *different* strategy in each relevant state. We note that the authors of [JvdH04] also considered explicit quantification over strategies in their setting, but did not pursue the idea, since they are concerned with obtaining a logic that is still propositional—a feature that our notion of quantified strategy choices gives up. Also, they obtain decidability results where our translation of course leads to an undecidable logic.

2.9 Expressing ATOL

ATOL (where O stands for observation) was also introduced by Jamroga and van der Hoek in [JvdH04], and is essentially "ATEL-R* with bounded recall," i.e., the semantics are the same as for ATEL-R*, but (in our notation) evaluated over $\mathcal{C}$ and not over $\mathcal{C}^{hst}$. Hence the translation is obtained in an analogous way as for ATEL-R* in Section 2.8 above.

Further, the authors of [JvdH04] discuss "mixing" of ATEL-R* and ATOL, i.e., to talk about memoryless and history-aware strategies simultaneously. This feature can be simulated in our logic by considering models of the form $\mathcal{C}^{hst}$, and defining, for every agent, two degrees of information, one which may take the history into account, and another one that does not (this can be formalized in a straight-forward way, given the construction of $\mathcal{C}^{hst}$).

3 Simulations

In [Sch10], simulation relations were introduces that allow to transfer strategy choices from one game structure to another. Since the translation is unidirectional, it does not preserve truth of quantified formulas. A notion of bisimulation was hinted at as well, which was defined as a relation Z which is a simulation in both directions—essentially, an isomorphism. In particular, bisimulations in this sense only exists between structures where the state set has the same cardinality.

[5] It is not sufficient to rely on the uniformity of strategy choices here (which require the same strategy to 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.

We show that the goal of a bisimulation—preserving truth of formulas in both directions—can be obtained with a much more relaxed definition that in particular allows bisimulations between finite and infinite structures. Our definition extends the one given in [Sch10] in a natural way. We first recall the definition of a simulation from [Sch10].

In the following, when Z is a binary relation on state sets, then for a state q, we write $Z(q)$ to denote the set $\{q' \mid (q, q') \in Z\}$.

Definition Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be CGSs with state sets Q_1 and Q_2, the same set of players, the same set of propositional variables, and n degrees of information. Then a relation $Z \subseteq Q_1 \times Q_2$ is a *probabilistic uniform strong alternating simulation for a coalition A from $\mathcal{C}_1$ to $\mathcal{C}_2$* if for all $(q_1, q_2) \in Z$, all $i \in \{1, \ldots, n\}$, and all players $a \in A$, there is a function $\Delta^{1\to2}_{(i,a,q_1,q_2)}$ such that for all $A' \subseteq A$ we have

- *propositional equivalence:* q_1 and q_2 satisfy the same propositional variables,
- for all (A', q_1)-moves c_1, the (A', q_2)-move c_2 with $c_2(a) = \Delta^{1\to2}_{(i,a,q_1,q_2)}(c_1(a))$ has the
  1. *Forward Move Property:* for each $(\overline{A'}, q_1)$-move $c_1^{\overline{A'}}$, there is a $(\overline{A'}, q_2)$-move $c_2^{\overline{A'}}$ such that for all $q_1' \in Q_1$, we have
  $$\Pr\left(\delta(q_2, c_2 \cup c_2^{\overline{A'}}) \in Z(q_1')\right) = \Pr\left(\delta(q_1, c_1 \cup c_1^{\overline{A'}}) = q_1'\right).$$
  2. *Backward Move Property:* for each $(\overline{A'}, q_2)$-move $c_2^{\overline{A'}}$, there is a $(\overline{A'}, q_1)$-move $c_1^{\overline{A'}}$ such that for all $q_1' \in Q_1$, we have
  $$\Pr\left(\delta(q_2, c_2 \cup c_{q_2}^{\overline{A'}}) \in Z(q_1')\right) = \Pr\left(\delta(q_1, c_1 \cup c_1^{\overline{A'}}) = q_1'\right).$$
- *Move Uniformity:* If $(q_1, q_2), (q_1', q_2') \in Z$ with $q_1 \sim_{\mathsf{eq}_1^i(a)} q_1'$ and $q_2 \sim_{\mathsf{eq}_1^i(a)} q_2'$, then $\Delta^{1\to2}_{(i,a,q_1,q_2)} = \Delta^{1\to2}_{(i,a,q_1',q_2')}$,
- *Uniformity:* for all $a \in A$, and all $(q_1', q_2') \in Z$, if $q_2 \sim_{\mathsf{eq}_2^i(a)} q_2'$, then $q_1 \sim_{\mathsf{eq}_1^i(a)} q_1'$.
- *Knowledge Transfer:* if $q_1' \sim_{\mathsf{eq}_i^1(A')} q_1$, then there is some $q_2' \in Q_2$ such that $q_2' \sim_{\mathsf{eq}_i^2(A')} q_2$ and $(q_1', q_2') \in Z$.
- *Uniqueness:* For all $q_2 \in Q_2$, there is exactly one $q_1 \in Q_1$ with $(q_1, q_2) \in Z$ (i.e., $Z^{-1} \colon Q_2 \to Q_1$ is a function).

Note that even though if Z is a probabilistic uniform strong alternating simulation then Z^{-1} is a function, we still write Z as a relation as is usually done with bisimulation-like concepts. The conditions in the above definition are natural, and (in addition to standard definitions of bisimulations) ensure that moves for players can be transferred from one CGS to the other in a way that leads to the same "effect" and can be "computed" by the players with the information available to them. Uniformity and Knowledge Transfer ensure that the degrees of information in the different structures are "compatible."

The results on simulation proven in [Sch10], about simulations can be stated in our notation as follows: If there is a probabilistic uniform strong alternating simulation from $\mathcal{C}_1$ to $\mathcal{C}_2$, then truth of a quantified strategy formula with a single existential quantifier is transferred from $\mathcal{C}_1$ to $\mathcal{C}_2$. As mentioned above, this result does not extend to formulas that also use universal quantification, or to transfer truth of formulas from $\mathcal{C}_2$ to $\mathcal{C}_1$. In particular, as mentioned in [Sch10], a probabilistic uniform strong alternating simulation does *not* imply that the structures are "strategically invariant." The statement in [Sch10] can be generalized as follows: A probabilistic uniform strong alternating simulation allows to transfer the truth value of a formula in which only existentially quantified variables appear, from $\mathcal{C}_1$ to $\mathcal{C}_2$. Analogously, if only universally quantified variables appear than the truth of formulas is transferred in the other direction.

3.1 Bisimulation Definition

To handle formulas in which both types of quantification appear, we need simulations in both directions. Obviously, these simulations cannot be completely unrelated to each other, they need to "agree" on the relationship between the states in a certain way.

Definition Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be concurrent game structures. Then a *probabilistic bisimulation* for a coalition A between $\mathcal{C}_1$ and $\mathcal{C}_2$ is a pair of relations (Z_1, Z_2) such that

- Z_1 is a probabilistic strategy simulation for A from $\mathcal{C}_1$ to $\mathcal{C}_2$,
- Z_2 is a probabilistic strategy simulation for A from $\mathcal{C}_2$ to $\mathcal{C}_1$,
- $Z_1^{-1} \circ Z_2^{-1}$ and $Z_2^{-1} \circ Z_1^{-1}$ are idempotent.

Idempotence ensures that if we apply both simulations simultaneously (one for "transferring" existentially quantified strategy choices, and the other one to transfer universally quantified strategy choices), then the transferred strategies "talk about the same states." This definition of a bisimulation is significantly less strict than the one from [Sch10]. The following theorem shows that our notion of a bisimulation allows to transfer truth of quantified formulas in the obvious way:

Theorem 3.1 *Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be concurrent game structures, let $\mathbb{A}$ be a set of coalitions such that (Z_1, Z_2) is a probabilistic bisimulation for every $A \in \mathbb{A}$ between $\mathcal{C}_1$ and $\mathcal{C}_2$, let q_1 be a state of $\mathcal{C}_1$, let q_2 be a state of $\mathcal{C}_2$ such that $(q_1, q_2) \in Z_1$ and $(q_2, q_1) \in Z_2$. Let φ be a quantified strategy formula for $\mathcal{C}_1$ (and thus for $\mathcal{C}_2$) such that every coalition appearing in φ is an element of $\mathbb{A}$. Then $\mathcal{C}_1, q_1 \models \varphi$ if and only if $\mathcal{C}_2, q_2 \models \varphi$.*

Note that in particular, if (Z_1, Z_2) is a probabilistic bisimulation for the set of all players Σ, then truth of every formula is maintained between the two structures. The introduction of the set $\mathbb{A}$ in Theorem 3.1 allows to use relations Z that satisfy some of the conditions (usually, the knowledge-related ones) only for a subset of Σ, to still transfer truth of formulas in which only a subset of the possible coalitions appears.

Also note that results from [Sch10] carry over in a natural way, for example we mention:

Proposition 3.2 *Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be concurrent game structures such that there is a probabilistic bisimulation between $\mathcal{C}_1$ and $\mathcal{C}_2$. Then there also is a probabilistic bisimulation between ${\mathcal{C}_1}^{hst}$ and ${\mathcal{C}_2}^{hst}$.*

Proof. The result follows directly from the proof of [Sch10, Proposition 4.3] and the observation that the construction in that proof also implies that the "lifted" simulation relations satisfy idempotence. □

Note, however, that clearly [Sch10, Proposition 4.2] does not apply to bisimulations, i.e., usually the structures $\mathcal{C}$ and $\mathcal{C}^{hst}$ are not bisimilar. As mentioned in [Sch10], this already holds for (unidirectional) probabilistic strategy simulations.

3.2 An example

We now give a very simple example for a bisimulation, which shows that a "large" CGS can be represented using a small "core." Consider the CGSs $\mathcal{C}_1$ and $\mathcal{C}_2$in Figure 3.2. The games presented here are turn-based, i.e., only the moves of one of the two players a or b influence the successor state. In both games, player a makes the first move, where he has 4 choices, leading to 4 different follow-up states in $\mathcal{C}_1$, and only a single choice in $\mathcal{C}_2$. The next (and final) move is by player b, and determines whether in the final state, the variable x is true or false. In the states q_1, q_2, and q_3 of $\mathcal{C}_1$, as well as in state r_1 of $\mathcal{C}_2$, a needs to play move 1 to ensure that x is true in the next state, in state q_4 of $\mathcal{C}_1$, the move 2 leads to x being true. Additionally, the states q_2 and q_3 are indistinguishable for b in $\mathcal{C}_1$, whereas in $\mathcal{C}_2$, both players have complete information (which is also true for player a in $\mathcal{C}_1$).

The structures $\mathcal{C}_1$ with state set Q_1 and $\mathcal{C}_2$ with state set Q_2 are bisimilar, via (Z_1, Z_2), where the function $Z_1^{-1}\colon Q_1 \to Q_2$ is defined as follows:

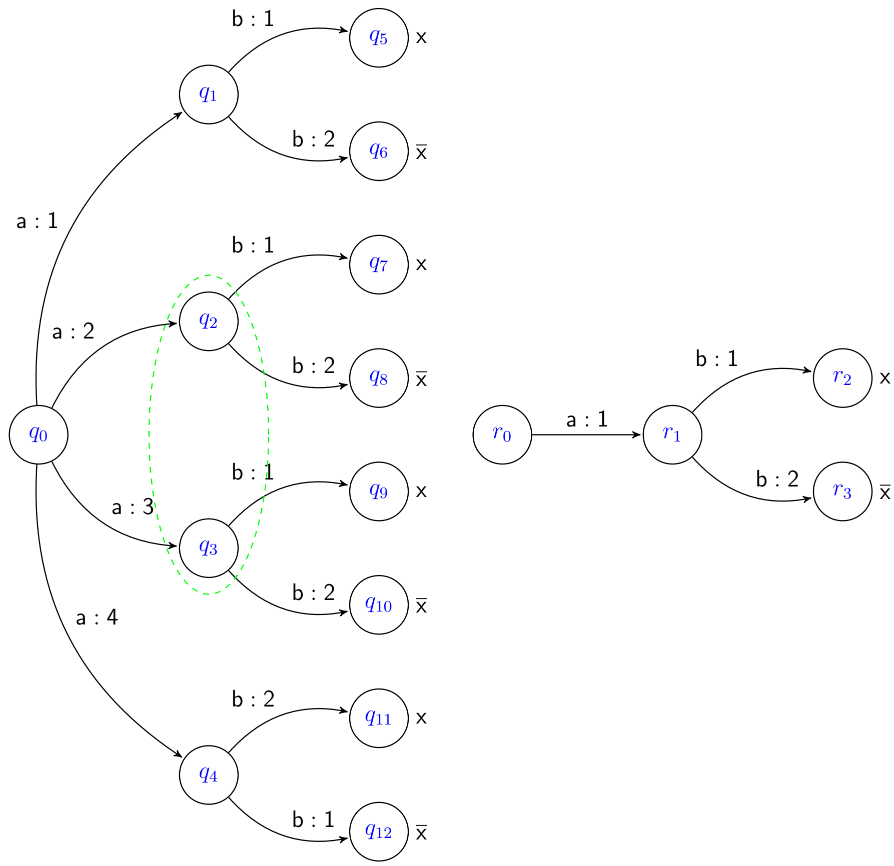

Figure 2: Structures $\mathcal{C}_1$ and $\mathcal{C}_2$

- $Z_1^{-1}(q_0) = r_0$,
- $Z_1^{-1}(q_1) = Z_1^{-1}(q_2) = Z_1^{-1}(q_3) = Z_1^{-1}(q_4) = r_1$,
- $Z_1^{-1}(q_5) = Z_1^{-1}(q_7) = Z_1^{-1}(q_9) = Z_1^{-1}(q_{11}) = r_2$,
- $Z_1^{-1}(q_6) = Z_1^{-1}(q_8) = Z_1^{-1}(q_{10}) = Z_1^{-1}(q_{12}) = r_3$.

The move transfer function for player A maps all of A's possible moves in q_0 to the move 1, the moves of B are mapped to themselves in the state q_1, q_2, and q_3, while in the state q_4, B's move 2 is mapped to the move 1 and vice versa. It is obvious that this translation respects incomplete information (i.e., move uniformity), since the move transfer function for B is the same in the states q_2 and q_3 that B cannot distinguish.

The converse embedding, $Z_2^{-1}\colon Q_2 \to Q_1$ 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 are the identity (alternatively, one could also define $Z_2^{-1}(r_1) = q_4$ and let the move transfer function reverse B's moves).

It is straight-forward to check that (Z_1, Z_2) is indeed a bisimulation, note that the forward- and backward move properties are very simple in these turn-based games.

Theorem 3.1 states that truth of any quantified strategy formula is transferred between $\mathcal{C}_1$ and $\mathcal{C}_2$, in states that are related via both Z_1 and Z_2. In the example, this applies to the pairs (q_0, r_0), (q_1, r_1), (q_5, r_2), and (q_6, r_3). It is clear that, for example, the states q_2 and r_1 are "strategically equivalent" as well. This also can be formally shown using bisimulations: Note that the choice of target states of Z_2^{-1} in the above definition was somewhat arbitrary, instead of choosing the "q_1-branch of $\mathcal{C}_1$," it is also possible to target the other branches, which then gives strategic equivalence of the remaining state pairs.

While strategic equivalence in this small example is not surprising, it is worth noting that the example shows that bisimulations can exist between structures with state sets of different cardinality, and also between structures with complete and with incomplete information. The above example is deterministic for easier presentation, a probabilistic one can be obtained by, instead of letting the moves reach a state where x is true or false, assign probabilities like 0.8/0.2 to these events.

3.3 Proof of Theorem 3.1

The remainder of this section proves Theorem 3.1. In the following, we let $\mathcal{C}_1$, $\mathcal{C}_2$, Z_1, and Z_2 refer to the structures and relations from the statement of Theorem 3.1. We first introduce some notation.

Definition Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be concurrent game structures, let Z be a probabilistic strategy simulation from $\mathcal{C}_1$ to $\mathcal{C}_2$. For a strategy choice S_1 for $\mathcal{C}_1$, we denote with $Z(\mathsf{S}_1)$ the strategy choice obtained by the construction of [Sch10, Theorem 4.1].

Note that for the probabilistic case, the construction of the strategy choice in the proof of [Sch10, Theorem 4.1] is deterministic, hence the above definition is well-defined. We introduce some abbreviations that will prove useful in the remainder of the section. For a strategy choice S_1 (S_2) of $\mathcal{C}_1$ ($\mathcal{C}_2$), the application of the simulation relation gives a strategy choice $Z_1(\mathsf{S}_1)$ ($Z_2(\mathsf{S}_2)$) for $\mathcal{C}_1$ ($\mathcal{C}_2$). From now on, we often regard Z_1 and Z_2 as functions on strategy choices. Since the proof of Theorem 3.1 transfers strategy choices back and forth multiple times, we introduce the abbreviation $Z^{\alpha_1\alpha_2\dots\alpha_n}$ for $Z_{\alpha_1} \circ Z_{\alpha_2} \circ \cdots \circ Z_{\alpha_n}$ for $\alpha_i \in \{1, 2\}$ (note that this is only meaningful if $\alpha_i \neq \alpha_{i+1}$ for all applicable i). For example, with $Z^{1212}(\mathsf{S}_2)$ we denote the strategy choice $Z_1(Z_2(Z_1(Z_2(\mathsf{S}_2))))$.

In the following proposition, we list some simple properties of bisimulations:

Proposition 3.3 *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 S_1 and S_2 be strategy choices for*

$\mathcal{C}_1$. *Then the following hold:*

- $Z_1(\mathsf{S}_1 \circ \mathsf{S}_2) = Z_1(\mathsf{S}_1) \circ Z_1(\mathsf{S}_2)$,
- $Z^{2121}(\mathsf{S}_1) = Z^{21}(\mathsf{S}_1)$.

Proof. The first part is obvious, the second immediately follows from the idempotence of the composition simulation relations and the fact that the construction of the strategy choice in the proof of [Sch10, Theorem 4.1] is obtained by essentially applying the simulation relations to transfer states from one CGS to the other. □

We now prove Theorem 3.1.

Proof. It is obviously sufficient to show one direction. Hence assume that $\mathcal{C}_1, q_1 \models \varphi$, where $\varphi = \forall \mathrm{S}_1 \exists \mathrm{S}_2 \forall \mathrm{S}_3 \dots \exists \mathrm{S}_n \psi$, where for each applicable i, S_i is a symbol for a strategy choice for a coalition A_i. We show that $\mathcal{C}_2, q_2 \models \varphi$. To make the instantiation of strategy choice variables with strategy choices more explicit, in the following we write $\mathcal{C}, q \models \psi[\mathrm{S}_1/\mathsf{S}_1, \dots, \mathrm{S}_n/\mathsf{S}_n]$ instead of $\mathcal{C}, (\mathsf{S}_1, \dots, \mathsf{S}_n), q \models \psi$ Since $\mathcal{C}_1, q_1 \models \varphi$, there are functions s^1_2, s^1_4, $\dots$, s^1_n such that for all strategy choices S^1_1, S^1_3, $\dots$, S^1_{n-1}, the strategy choices S^1_2, S^1_4, $\dots$, S^1_n defined as $\mathsf{S}^1_i = s^1_i(\mathsf{S}^1_1, \mathsf{S}^1_3, \dots, \mathsf{S}^1_{i-1})$ have the property that $\mathcal{C}_1, q_1 \models \psi[\mathrm{S}_1/\mathsf{S}^1_1, \dots, \mathrm{S}_n/\mathsf{S}^1_n]$.

To prove that $\mathcal{C}_1, q_1 \models \varphi$, we construct functions $s^2_2, s^2_4, \dots, s^2_n$ with the required properties. For each relevant i, and for strategy choices S^2_1, S^2_3, $\dots$, S^2_{n-1} for $\mathcal{C}_2$, let

$$s^2_i(\mathsf{S}^2_1, \mathsf{S}^2_3, \dots, \mathsf{S}^2_{i-1}) = Z^1(s^1_i(Z^2(\mathsf{S}^2_1), Z^2(\mathsf{S}^2_3), \dots, Z^2(\mathsf{S}^2_{i-1}))).$$

Note that this is well-defined: If S^2_1, S^2_3, $\dots$, S^2_{i-1} are strategy choices for $\mathcal{C}_2$, then $Z^2(\mathsf{S}^2_1)$, $Z^2(\mathsf{S}^2_3)$, $\dots$, $Z^2(\mathsf{S}^2_{i-1})$ are strategy choices for $\mathcal{C}_1$, thus we can apply the $s^1_{\dots}$-functions, and use Z^1 to transfer the resulting strategy choices to $\mathcal{C}_2$.

We prove that the functions $s^2_{\dots}$ defined above indeed satisfy the requirements. Hence let S^2_1, S^2_3, $\dots$, S^2_{n-1} be strategy choices for $\mathcal{C}_2$ (for coalitions A_1, A_3, $\dots$, A_{i-1}), and for $i \in \{2, 4, \dots, n\}$, define $\mathsf{S}^2_i = s^2_i(\mathsf{S}^2_1, \mathsf{S}^2_3, \dots, \mathsf{S}^2_{i-1})$. We prove that $\mathcal{C}_2, q_2 \models \psi[\mathrm{S}_1/\mathsf{S}^2_1, \dots, \mathrm{S}_n/\mathsf{S}^2_n]$.

For $i \in \{1, 3, \dots, n-1\}$, define $\mathsf{S}^1_i = Z^2(\mathsf{S}^2_i)$, and for $i \in \{2, 4, \dots, n\}$, define $\mathsf{S}^1_i = s^1_i(\mathsf{S}^1_1, \mathsf{S}^1_3, \dots, \mathsf{S}^1_{i-1})$. Then for all $i \in \{2, 4, \dots, n\}$, it follows that $\mathsf{S}^2_i = Z_1(\mathsf{S}^1_i)$.

Due to the choice of the $s^1_{\dots}$-functions, we know that

$$\mathcal{C}_1, q_1 \models \psi[\mathrm{S}_1/\mathsf{S}^1_1, \dots, \mathrm{S}_n/\mathsf{S}^1_n].$$

Since Z_1 is a probabilistic strategy simulation, we can apply[6] [Sch10, Theorem 4.1] and conclude that

[6]Since [Sch10, Theorem 4.1] only considers the case of a single strategy choice, application of the theorem here formally requires (straight-forward) induction.

$$\mathcal{C}_2, q_2 \models \psi[\mathrm{S}_1/Z_1(\mathsf{S}^1_1), \ldots, \mathrm{S}_n/Z_1(\mathsf{S}^1_n)].$$

Since Z_2 is a probabilistic strategy simulation as well, we can also apply Z_2, and (using the abbreviations introduced above), conclude that

$$\mathcal{C}_1, q_1 \models \psi[\mathrm{S}_1/Z^{21}(\mathsf{S}^1_1), \ldots, \mathrm{S}_n/Z^{21}(\mathsf{S}^1_n)].$$

Since for odd values of i, we know that $\mathsf{S}^1_i = Z_2(\mathsf{S}^2_i)$, this implies that

$$\mathcal{C}_1, q_1 \models \psi[\mathrm{S}_1/Z^{212}(\mathsf{S}^2_1), \mathrm{S}_2/Z^{21}(\mathsf{S}^1_2), \ldots, \mathrm{S}_{n-1}/Z^{212}(\mathsf{S}^1_{n-1}), \mathrm{S}_n/Z^{21}(\mathsf{S}^1_n)].$$

Applying the simulation Z_1 once more implies

$$\mathcal{C}_2, q_2 \models \psi[\mathrm{S}_1/Z^{1212}(\mathsf{S}^2_1), \mathrm{S}_2/Z^{121}(\mathsf{S}^1_2), \ldots, \mathrm{S}_{n-1}/Z^{1212}(\mathsf{S}^1_{n-1}), \mathrm{S}_n/Z^{121}(\mathsf{S}^1_n)].$$

Now indirectly assume that the theorem does not hold, i.e., that $\mathcal{C}_2, q_2 \not\models \psi[\mathrm{S}_1/\mathsf{S}^2_1, \ldots, \mathrm{S}_n/\mathsf{S}^2_n]$. It then follows that

$$\mathcal{C}_2, q_2 \models \neg\psi[\mathrm{S}_1/\mathsf{S}^2_1, \ldots, \mathrm{S}_n/\mathsf{S}^2_n].$$

Applying simulations Z_1, Z_2, Z_1, and Z_2, we obtain (again using [Sch10, Theorem 4.1]) that

$$\mathcal{C}_2, q_2 \models \neg\psi[\mathrm{S}_1/Z^{1212}(\mathsf{S}^2_1), \ldots, \mathrm{S}_n/Z^{1212}(\mathsf{S}^2_n)].$$

Since for even values of i, we know that $\mathsf{S}^2_i = Z_1(\mathsf{S}^1_i)$, this implies that

$$\mathcal{C}_2, q_2 \models \neg\psi[\mathrm{S}_1/Z^{1212}(\mathsf{S}^2_1), \mathrm{S}_2/Z^{12121}(\mathsf{S}^1_2), \ldots, \mathrm{S}_{n-1}/Z^{1212}(\mathsf{S}^2_{n-1}), \mathrm{S}_n/Z^{12121}(\mathsf{S}^1_n)].$$

An application of Proposition 3.3 to the even values of i implies that

$$\mathcal{C}_2, q_2 \models \neg\psi[\mathrm{S}_1/Z^{1212}(\mathsf{S}^2_1), \mathrm{S}_2/Z^{121}(\mathsf{S}^1_2), \ldots, \mathrm{S}_{n-1}/Z^{1212}(\mathsf{S}^2_{n-1}), \mathrm{S}_n/Z^{121}(\mathsf{S}^1_n)].$$

This is a contradiction to the above, hence the functions $s^2_{\ldots}$ indeed satisfy the requirements.

It remains to show that the truth of formulas using restricted quantification (i.e., over constant or φ-maximal strategies) is preserved as well. For the constant case this is trivial, as it follows from the definition of $Z_1(\mathsf{S}_1)$ that if S_1 is constant, then $Z_1(\mathsf{S}_1)$ is constant as well. For the φ-maximal case, we show that we can express φ-maximality using quantified formulas that only use standard quantification—hence the restricted quantifier is essentially "syntactic sugar," and the invariance of truth for formulas using this quantification follows from the above. For the definition of maximality as well as the appearing notation, see [Sch10].

Hence let φ be a fixed formula for which we want to express φ-maximality. First note that $\mathsf{sat}_\varphi(\mathsf{S}_1, i) \subseteq \mathsf{sat}_\varphi(\mathsf{S}_2, i)$ can be expressed by the formula

$$\bigwedge_{\langle\langle A\rangle\rangle_i^{\blacktriangleleft\alpha}\psi \in \mathsf{sf}_i(\varphi)} \langle\langle \emptyset : \emptyset\rangle\rangle^{\geq 1} \Box \left(\langle\langle A : S_1\rangle\rangle_i^{\blacktriangleleft\alpha} \psi \rightarrow \langle\langle A : S_2\rangle\rangle_i^{\blacktriangleleft\alpha} \psi\right)$$

where $\mathsf{sf}_i(\varphi)$ denotes the $\langle\langle . \rangle\rangle$-subformulas of φ with strategic depth i. Note that in the outmost strategy operator we omitted the information degree as this is not relevant for the empty coalition. In the sequel, we use $\mathsf{S}_1 \leq_\varphi^i \mathsf{S}_2$ for the above formula, and use $\mathsf{S}_1 \equiv_\varphi^i \mathsf{S}_2$ as shorthand for $(\mathsf{S}_1 \leq_\varphi^i \mathsf{S}_1) \wedge (\mathsf{S}_2 \leq_\varphi^i \mathsf{S}_1)$. Now $\mathsf{S}_1 \leq_\varphi \mathsf{S}_2$ can be expressed as

$$\bigwedge_{i \leq \mathsf{sd}(\varphi)} (\mathrm{S}_1 \equiv_\varphi^i \mathrm{S}_2) \vee \left(\bigvee_{i \leq \mathsf{sd}(\varphi)} \left(\left(\bigwedge_{j<i} \mathrm{S}_1 \equiv_\varphi^i \mathrm{S}_2\right) \wedge (\mathrm{S}_1 \leq_\varphi^i \mathrm{S}_2) \wedge \neg(\mathrm{S}_2 \leq_\varphi^i \mathrm{S}_1)\right)\right).$$

Now a strategy choice is φ-maximal if and only if it satisfies $\mathsf{max}_\varphi(\mathrm{S})$ defined as $\forall \mathrm{S}' (\mathrm{S} \leq_\varphi \mathrm{S}' \rightarrow \mathrm{S}' \leq_\varphi \mathrm{S})$. A formula ψ using the restricted $\exists_m \mathrm{S}_i$ or $\forall_m \mathrm{S}_j$ quantifiers can now be rewritten by consecutively replacing the kernel χ of ψ with

- $\mathsf{max}_\varphi(\mathsf{S}_i) \wedge \chi$ if S_i is existentially quantified, or with
- $\mathsf{max}_\varphi(\mathsf{S}_j) \rightarrow \chi$ if S_j is universally quantified.

Note that this only checks maximality on the reachable parts of the CGS. If the entire formula is evaluated at, let's say, a singleton and there is an unconnected remainder of the graph that does not have a maximal strategy choice, then this is not checked. However, by applying the bisimulation in other parts of the CGS, it follows that in that case a maximal strategy choice does not exist in either of the structures. Therefore, truth of formulas containing restricted quantification is invariant under bisimulation as well, which proves the theorem. □

3.4 Bisimulations for Deterministic Game Structures

In [Sch10], a relaxed notion of simulation was introduced for deterministic structures and it was proved that for these structures, the simpler notion is sufficient to transfer strategy choices. This notion, there called a uniform strong alternating simulation, can be used as the base of defining bisimulations in the same way as probabilistic uniform strong alternating simulation: The proof of Theorem 3.1 uses the simulation result from [Sch10] essentially in a black-box fashion, and hence since the result from [Sch10] is true for the relaxed, deterministic version, this is true for our Theorem 3.1 as well. There are two details of the proof in [Sch10] that we used in our proof:

1. we references some details of the construction of the transfer of strategy choices between the game structures, note that these are the same for deterministic and probabilistic structures,
2. we stated that for probabilistic structures, the construction of $Z(\mathsf{S})$ is deterministic—this is since the definition of a probabilistic uniform strong alternating simulation establishes a unique function $Z^{-1}\colon \mathcal{C}_1 \to \mathcal{C}_2$. This is not true for the deterministic case, since the definition of a uniform strong alternating simulation does not have such a requirement. However, using the Axiom of Choice, such a function can be constructed. It therefore should be noted that for deterministic structures that are not countable, the analog of Theorem 3.1, relaxed to uniform strong alternating simulation, requires the Axion of Choice.

4 Model Checking: Complexity and Decidability

We now consider the model checking problem for QAPI. Model checking is the question, given a game structure $\mathcal{C}$ and a quantified strategy formula φ, whether φ is satisfied in a given state of $\mathcal{C}$. Our algorithms for the decidable cases are constructive: They can be used to produce, for a given instantiation of universally-quantified strategy choices, suitable instantiations for the existentially quantified ones. As expected, the question whether strategies are allowed to base their decisions on history or are required to be memoryless has a huge impact on the questions of decidability. Formally, we define the problems $Q\mathsf{MC}$ and $Q\mathsf{MC}^{hst}$. In both cases, the input consists of a CGS $\mathcal{C}$, a state q of $\mathcal{C}$, and a quantified strategy formula φ. For $Q\mathsf{MC}$, the question is whether $\mathcal{C}, q \models \varphi$, in the case of $Q\mathsf{MC}^{hst}$, we ask whether $\mathcal{C}^{hst}, q \models \varphi$.

It turns out that there is no significant computational price to pay for the additional expressiveness provided by quantifiers, compared to the semantics from [Sch10]. In addition to showing that the (upper and lower) complexity bounds from that paper still hold, we also show that the lower bounds can be made more tight: We prove that in the history-aware case, even if we allow all players to have complete information, the model checking problem is undecidable.

Theorem 4.1 *The problem* $Q\mathsf{MC}$ *is*

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

Proof. The algorithm for both cases are very similar, and consist of two components:

- Perform a complete search over all possible instantiations of the symbols for strategy choices appearing in the given formula φ,
- for each of these, check whether their combinations satisfies the kernel of φ.

The first component can easily be done in polynomial space, since the representation of a strategy choice is quartic in the size of the input: For each player, and each subformula of φ, and each state q of the given CGS $\mathcal{C}$, it specifies a valid move in that state. Since the number of strategy choices to be considered simultaneously is exactly the number of quantifiers in the formula φ. only polynomial space is required for going through all relevant combinations of strategy choices. Note that the uniformity conditions, checking whether a given strategy choice is constant or maximal can be performed in polynomial time (for the maximality condition, see [Sch10]). Clearly, the necessary uniformity conditions for strategy choices can be performed in polynomial time, given their representation and that of the game structure. We can now use the algorithms of [Sch10, Theorem 5.1] to check, for each of these combinations, whether the kernel of φ is satisfied. Since PSPACE with a PSPACE-oracle still gives polynomial space complexity, this establishes the PSPACE-result. To see that the 3EXPTIME upper bound still holds, note that the game structures used in the calls of the 3EXPTIME model checking algorithm that the proof of [Sch10, Theorem 5.1] uses as oracle is always given an input that is a substructure of $\mathcal{C}$.

The lower bounds follow directly from [Sch10, Theorem 5.1], which is contained in QAPI's model checking problem as a special case. □

For memory-dependent strategies, it is known that ATL* remains decidable [AHK02]. Adding quantification to ATL* essentially leads to strategy logic [CHP07], which is still decidable. However, when adding the capabilities to reason about incomplete information, or probabilistic games, then the model-checking problem of the resulting logic becomes undecidable:

Theorem 4.2 *The problem* QMC^{hst} *is undecidable, even when restricted to a single existential quantifier and*

1. *game structures with complete information, or*
2. *deterministic game structures.*

Restricted to game structures that have complete information and are deterministic simultaneously, QMC^{hst} *is decidable even with no restrictions on the quantifiers.*

Proof. 1. The case with complete information, but a probabilistic transition function follows directly from Theorem 3.4 in [BBFK06]: Their result proves that the question whether in a randomized game with a single player, a deterministic strategy to achieve a goal expressed by a PCTL-formula exists. Since equalities of a success probability and some value α can easily be expressed in our semantics with requiring the same strategy has success probability $\leq \alpha$ and $\geq \alpha$, it follows that PCTL can be expressed with probabilistic ATL; undecidability of QMC^{hst} for the complete-information case follows. Since their result is only concerned with finding a single strategy, the problem remains undecidable when restricted to formulas with a single existential quantifier.

2. For deterministic structures, but incomplete information, the result follows from [Sch10, Theorem 5.2].

Now consider game structures that are both deterministic and have complete information. In this case, it is obvious that strategy choices can always assumed to be constant: The fact that strategy choices may depend on the formula can be handled by (in the case of non-constant strategy choices) using different strategy choices for different subformulas. The fact that strategy choices allow the choice of strategy to depend on possibly more information than the choice of moves performed by the strategy later is obviously irrelevant in complete-information games, formally the function choosing the strategy in the strategy choice and the choices performed by the selected strategy can be combined into a (complete-information) strategy. Hence in this case, QAPI is equivalent to strategy logic, which is decidable due to [CHP07]. □

Undecidability of model checking for QAPI in the history-aware case does not come as a surprise, given that less expressive logics already have undecidable model checking problems. Even though the obvious restrictions of QAPI covered by Theorem 4.2 still have undecidable model checking problems, it is an interesting question to identify sufficiently large subsets of QAPI that are still decidable even with history-awareness. Since QAPI contains logics with decidable model checking problems as special cases (see Section 2), it is likely that relevant decidable fragments do in fact exist.

Conclusion

We have introduced QAPI, which compared to the semantics of ATL* introduced in [Sch10] leads to a significant increase in expressivity:

1. Formulas can reason about strategies explicitly,
2. it is possible to force the "counter-coalition" to use restricted knowledge as well,

3. in re-using symbols for strategy choices, one can express knowledge that a coalition has over another coalitions strategies.

The logic QAPI is expressive enough to include several previous extensions of ATL*.

An interesting extension of QAPI would be to reason about *mixed* strategies, where the players themselves may randomize over their next current move. It is likely that our simulation results still hold for a thus-generalized setting, however the complexity of the model checking problem in this setting is an open question.

References

[ÅGJ07] Thomas Ågotnes, Valentin Goranko, and Wojciech Jamroga. Alternating-time temporal logics with irrevocable strategies. In Samet [Sam07], pages 15–24.

[AHK02] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.

[BBFK06] Tomás Brázdil, Václav Brozek, Vojtech Forejt, and Antonín Kucera. Stochastic games with branching-time winning objectives. In *LICS*, pages 349–358. IEEE Computer Society, 2006.

[BBGK07] Christel Baier, Tomás Brázdil, Marcus Größer, and Antonín Kucera. Stochastic game logic. In *QEST*, pages 227–236. IEEE Computer Society, 2007.

[BJ09] Nils Bulling and Wojciech Jamroga. What agents can probably enforce. *Fundamenta Informaticae*, 93(1-3):81–96, 2009.

[CHJ06] Krishnendu Chatterjee, Thomas A. Henzinger, and Marcin Jurdzinski. Games with secure equilibria. *Theor. Comput. Sci.*, 365(1-2):67–82, 2006.

[CHP07] Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. Strategy logic. In Luís Caires and Vasco Thudichum Vasconcelos, editors, *CONCUR*, volume 4703 of *Lecture Notes in Computer Science*, pages 59–73. Springer, 2007.

[CL07] Taolue Chen and Jian Lu. Probabilistic alternating-time temporal logic and model checking algorithm. In J. Lei, editor, *FSKD (2)*, pages 35–39. IEEE Computer Society, 2007.

[HT06] Andreas Herzig and Nicolas Troquard. Knowing how to play: uniform choices in logics of agency. In Hideyuki Nakashima, Michael P. Wellman, Gerhard Weiss, and Peter Stone, editors, *AAMAS*, pages 209–216. ACM, 2006.

[Jam04] Wojciech Jamroga. Some remarks on alternating temporal epistemic logic. In *Proceedings of Formal Approaches to Multi-Agent Systems (FAMAS 2003*, pages 133–140, 2004.

[JvdH04] Wojciech Jamroga and Wiebe van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63(2-3):185–219, 2004.

[Sam07] Dov Samet, editor. *Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2007), Brussels, Belgium, June 25-27, 2007*, 2007.

[Sch04] Pierre-Yves Schobbens. Alternating-time logic with imperfect recall. *Electronis Notes in Theoretical Computer Science*, 85(2):82–93, 2004.

[Sch10] Henning Schnoor. Strategic planning for probabilistic games with incomplete information. In *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, pages 1057–1064, 2010.

[WvdHW07] Dirk Walther, Wiebe van der Hoek, and Michael Wooldridge. Alternating-time temporal logic with explicit strategies. In Samet [Sam07], pages 269–278.