

An Epistemic Strategy Logic (Extended Abstract)

Xiaowei Huang

The University of New South Wales

Ron van der Meyden

The University of New South Wales

Introduction

There are many subtle issues concerning agent knowledge in settings where multiple agents act strategically. In the process of understanding these issues, there has been a proliferation of modal logics dealing with epistemic reasoning in strategic settings, e.g., [14, 12, 9]. The trend has been for these logics to contain large numbers of operators, each of which combines several different concerns, such as the existence of strategies, and knowledge that groups of agents may have about these strategies. We argued in a previous work [8] that epistemic temporal logic already has the expressiveness required for many applications of epistemic strategy logics, provided that one works in a semantic framework in which strategies are explicitly rather than (as in most alternating temporal epistemic logics) implicitly represented, and makes the minor innovation of including new agents whose local states correspond to the strategies being used by other agents. This gives a more compositional basis for epistemic strategic logic. In the case of imperfect recall strategies and knowledge operators, and a CTL* temporal basis, this leads to a temporal epistemic strategy logic with a PSPACE complete model checking problem.

However, some of our results in [8] required a restriction to cases not involving a common knowledge operator, because certain notions could not be expressed. In the present paper, we develop a remedy for this weakness. We propose an epistemic strategy logic which, like [3, 11], supports explicit naming and quantification over strategies. However we achieve this in a slightly more general way: we first generalize temporal epistemic logic to include operators for quantification over global states and reference to their components, and then apply this generalization to a system that includes strategies encoded in the global states and references these using the “strategic” agents of [8]. The resulting framework can express many of the subtly different notions that have been the subject of proposals for alternating temporal epistemic logics. In particular, it generalizes the expressiveness of the logic in [8] but is able to also deal with the common knowledge issues that restricted the scope of that work. The new logic retains the pleasant compositional capabilities of the prior proposal. Moreover, the extra expressiveness comes without extra cost: model checking remains PSPACE-complete, at least for a fragment based on CTL temporal operators.

An extended temporal epistemic logic

We extend temporal epistemic logic with a set of variables Var , an operator $\exists x.$ and a construct $e_i(x)$, where x is a variable and $\exists x.\phi$ says, intuitively, that there exists in the system a global state x such that ϕ holds at the current point, and $e_i(x)$ says that agent i has the same local state at the current point and at the global state x . Let $Prop$ be a set of atomic propositions and let Ags be a set of agents. Formally, the language ETLK($Ags, Prop, Var$) has syntax given by the grammar:

$$\phi \equiv p \mid \neg\phi \mid \phi_1 \vee \phi_2 \mid A\phi \mid \bigcirc\phi \mid \phi_1 U \phi_2 \mid \exists x.\phi \mid e_i(x) \mid D_G\phi \mid C_G\phi$$

where $p \in Prop$, $x \in Var$, $i \in Ags$, and $G \subseteq Ags$. The construct $D_G\phi$ expresses that agents in G have distributed knowledge of ϕ , i.e., could deduce ϕ if they pooled their information, and $C_G\phi$ says that ϕ is common knowledge to group G . The temporal formulas $\bigcirc\phi$, $\phi_1 U \phi_2$, $A\phi$ have the standard meanings from CTL^* , i.e., $\bigcirc\phi$ says that ϕ holds at the next moment of time, $\phi_1 U \phi_2$ says that ϕ_1 holds until ϕ_2 does, and $A\phi$ says that ϕ holds in all possible evolutions from the present situation. Other operators can be obtained in the usual way, e.g., $\phi_1 \wedge \phi_2 = \neg(\neg\phi_1 \vee \neg\phi_2)$, $\Diamond\phi = (true U \phi)$, $\Box\phi = \neg\Diamond\neg\phi$, etc. The universal form $\forall x.\phi = \neg\exists x.\neg\phi$ expresses that ϕ holds for all global states x . For an agent $i \in Ags$, we write $K_i\phi$ for $D_{\{i\}}\phi$; this expresses that agent i knows the fact ϕ . The notion of everyone in group G knowing ϕ can then be expressed as $E_G\phi = \bigwedge_{i \in G} K_i\phi$. We write $e_G(x)$ for $\bigwedge_{i \in G} e_i(x)$. This says that at the current point, the agents in G have the same local state as they do at the global state named by variable x .

The semantics of $ETLK(Ags, Prop, Var)$ builds straightforwardly on the following definitions used in the standard semantics for temporal epistemic logic [4]. Consider a system for a set of agents Ags . A *global state* is an element of the set $\mathcal{G} = L_e \times \prod_{i \in Ags} L_i$, where L_e is a set of states of the environment and each L_i is a set of *local states* for agent i . A *run* is a mapping $r : \mathbb{N} \rightarrow \mathcal{G}$ giving a global state at each moment of time. A *point* is a pair (r, m) consisting of a run r and a time m . An *interpreted system* is a pair $\mathcal{I} = (\mathcal{R}, \pi)$, where \mathcal{R} is a set of runs and π is an *interpretation*, mapping each point (r, m) with $r \in \mathcal{R}$ to a subset of $Prop$. For $n \leq m$, write $r[n \dots m]$ for the sequence $r(n)r(n+1) \dots r(m)$. Elements of $\mathcal{R} \times \mathbb{N}$ are called the *points* of \mathcal{I} . For each agent $i \in Ags \cup \{e\}$, we write $r_i(m)$ for the component of $r(m)$ in L_i , and then define an equivalence relation on points by $(r, m) \sim_i (r', m')$ if $r_i(m) = r'_i(m')$. We also define $\sim_G^D \equiv \bigcap_{i \in G} \sim_i$, and $\sim_G^E \equiv \bigcup_{i \in G} \sim_i$, and $\sim_G^C \equiv (\bigcup_{i \in G} \sim_i)^*$ for $G \subseteq Ags$. We take \sim_\emptyset^D to be the universal relation on points, and \sim_\emptyset^E and \sim_\emptyset^C to be the identity relation.

To extend this semantic basis for temporal epistemic logic to a semantics for $ETLK(Ags, Prop, Var)$, we just need to add a construct that interprets variables as global states. A *context* for an interpreted system \mathcal{I} is a mapping Γ from Var to global states occurring in \mathcal{I} . We write $\Gamma[g/x]$ for the context Γ' with $\Gamma'(x) = g$ and $\Gamma'(y) = \Gamma(y)$ for all variables $y \neq x$. The semantics of the language $ETLK$ is given by a relation $\Gamma, \mathcal{I}, (r, m) \models \phi$, representing that formula ϕ holds at point (r, m) of the interpreted system \mathcal{I} , relative to context Γ . This is defined inductively on the structure of the formula ϕ , as follows:

- $\Gamma, \mathcal{I}, (r, m) \models p$ if $p \in \pi(r, m)$;
- $\Gamma, \mathcal{I}, (r, m) \models \neg\phi$ if not $\Gamma, \mathcal{I}, (r, m) \models \phi$;
- $\Gamma, \mathcal{I}, (r, m) \models \phi \wedge \psi$ if $\Gamma, \mathcal{I}, (r, m) \models \phi$ and $\Gamma, \mathcal{I}, (r, m) \models \psi$;
- $\Gamma, \mathcal{I}, (r, m) \models A\phi$ if $\Gamma, \mathcal{I}, (r', m) \models \phi$ for all $r' \in \mathcal{R}$ with $r[0 \dots m] = r'[0 \dots m]$;
- $\Gamma, \mathcal{I}, (r, m) \models \bigcirc\phi$ if $\Gamma, \mathcal{I}, (r, m+1) \models \phi$;
- $\Gamma, \mathcal{I}, (r, m) \models \phi U \psi$ if there exists $m' \geq m$ such that $\Gamma, \mathcal{I}, (r, m') \models \psi$ and $\Gamma, \mathcal{I}, (r, k) \models \phi$ for all k with $m \leq k < m'$;
- $\Gamma, \mathcal{I}, (r, m) \models \exists x.\phi$ if $\Gamma[r'(m')/x], \mathcal{I}, (r, m) \models \phi$ for some point (r', m') of \mathcal{I} ;
- $\Gamma, \mathcal{I}, (r, m) \models e_i(x)$ if $r_i(m) = \Gamma(x)_i$;
- $\Gamma, \mathcal{I}, (r, m) \models D_G\phi$ if $\Gamma, \mathcal{I}, (r', m') \models \phi$ for all (r', m') such that $(r', m') \sim_G^D (r, m)$;
- $\Gamma, \mathcal{I}, (r, m) \models C_G\phi$ if $\Gamma, \mathcal{I}, (r', m') \models \phi$ for all (r', m') such that $(r', m') \sim_G^C (r, m)$.

The definition is standard, except for the constructs $\exists x.\phi$ and $e_i(x)$. The clause for the former says that $\exists x.\phi$ holds at a point (r, m) if there exists a global state $g = r'(m')$ such that ϕ holds at the point (r, m) , provided, we interpret x as referring to g . Note that it is required that g is attained at some point (r', m') , so actually occurs in the system \mathcal{I} . The clause for $e_i(x)$ says that this holds at a point (r, m) if the local

state of agent i , i.e., $r_i(m)$, is the same as the local state $\Gamma(x)_i$ of agent i at the global state $\Gamma(x)$ that interprets the variable x according to Γ .

We remark that these novel constructs introduce some redundancy, in that the set of epistemic operators D_G could be reduced to the “universal” operator D_\emptyset , since

$$D_G\phi \equiv \exists x.(\mathbf{e}_G(x) \wedge D_\emptyset(\mathbf{e}_G(x) \Rightarrow \phi)) .$$

Evidently, given the complexity of this formulation, D_G remains a useful notation.

Strategic Environments

In order to deal with agents that operate in an environment by strategically choosing their actions, we introduce a richer type of transition system that models the available actions and their effects on the state. An *environment* for agents Ags is a tuple $E = \langle S, I, Acts, \rightarrow, \{O_i\}_{i \in Ags}, \pi \rangle$, where S is a set of states, I is a subset of S , representing the initial states, $Acts = \prod_{i \in Ags} Acts_i$ is a set of joint actions, where each $Acts_i$ is a nonempty set of actions that may be performed by agent i , component $\rightarrow \subseteq S \times Acts \times S$ is a transition relation, $O_i : S \rightarrow L_i$ is an observation function, and $\pi : S \rightarrow \mathcal{P}(Prop)$ is a propositional assignment. An environment is said to be finite if all its components, i.e., $S, Ags, Acts_i, L_i$ and $Prop$ are finite. Intuitively, a joint action $a \in Acts$ represents a choice of action $a_i \in Acts_i$ for each agent $i \in Ags$, performed simultaneously, and the transition relation resolves this into an effect on the state. We assume that \rightarrow is serial in the sense that for all $s \in S$ and $a \in Acts$ there exists $t \in S$ such that $(s, a, t) \in \rightarrow$.

A *strategy* for agent $i \in Ags$ in such an environment E is a function $\alpha : S \rightarrow \mathcal{P}(Acts_i) \setminus \{\emptyset\}$, selecting a set of actions of the agent at each state.¹ We call these the actions *enabled* at the state. A *group strategy*, or *strategy profile*, for a group G is a tuple $\alpha_G = \langle \alpha_i \rangle_{i \in G}$ where each α_i is a strategy for agent i . A strategy α_i is *deterministic* if $\alpha_i(s)$ is a singleton for all s . A strategy α_i for agent i is *uniform* if for all states s, t , if $O_i(s) = O_i(t)$, then $\alpha_i(s) = \alpha_i(t)$. A strategy $\alpha_G = \langle \alpha_i \rangle_{i \in G}$ for a group G is *locally uniform (deterministic)* if α_i is uniform (respectively, deterministic) for each agent $i \in G$. Given an environment E , we write $\Sigma^{det}(E)$ for the set of deterministic strategies, $\Sigma^{unif}(E)$ for the set of all locally uniform joint strategies, and $\Sigma^{unif, det}(E)$ for the set of all deterministic locally uniform joint strategies.

We now define an interpreted system that contains all the possible runs generated when agents Ags behave by choosing a strategy from some set Σ of joint strategies in the context of an environment E . One innovation, introduced in [8], is that the construction introduces new agents $\sigma(i)$, for each $i \in Ags$. The observation of $\sigma(i)$ is the strategy currently being used by agent i . Agent $\sigma(i)$ is not associated with any actions, and is primarily for use in epistemic operators, to allow reference to what can be deduced were agents to reason using information about each other’s strategies. For $G \subseteq Ags$, we write $\sigma(G)$ for the set $\{\sigma(i) \mid i \in G\}$. Additionally, we include an agent e for representing the state of the environment. (This agent, also, is not associated with any actions.)

Formally, given an environment $E = \langle S, I, Acts, \rightarrow, \{O_i\}_{i \in Ags}, \pi \rangle$ for agents Ags , where $O_i : S \rightarrow L_i$ for each $i \in Ags$, and a set $\Sigma \subseteq \prod_{i \in Ags} \Sigma_i$ of joint strategies for the group Ags , we define the *strategy space* interpreted system $I(E, \Sigma) = (\mathcal{R}, \pi')$. The system $I(E, \Sigma)$ has global states $\mathcal{G} = S \times \prod_{i \in Ags} L_i \times \prod_{i \in Ags} \Sigma_i$. Intuitively, each global state consists of a state of the environment E , a local state for each agent i in E , and a strategy for each agent i . We index the components of this cartesian product by e , the elements of Ags and the elements of $\sigma(Ags)$, respectively. We take the set of runs \mathcal{R} of $I(E, \Sigma)$ to be the set of all runs $r : \mathbb{N} \rightarrow \mathcal{G}$ satisfying the following constraints, for all $m \in \mathbb{N}$ and $i \in Ags$

¹More generally, a strategy could be a function of the history, but we focus here on strategies that depend only on the final state.

1. $r_e(0) \in I$ and $\langle r_{\sigma(i)}(0) \rangle_{i \in \text{Ags}} \in \Sigma$,
2. $r_i(m) = O_i(r_e(m))$,
3. $(r_e(m), a, r_e(m+1)) \in \rightarrow$ for some joint action $a \in \text{Acts}$ such that for all $j \in \text{Ags}$ we have $a_j \in \alpha_j(r_j(m))$, where $\alpha_j = r_{\sigma(j)}(m)$
4. $r_{\sigma(i)}(m+1) = r_{\sigma(i)}(m)$.

The first constraint, intuitively, says that runs start at an initial state of E , and the initial strategy profile at time 0 is one of the profiles in Σ . The second constraint states that the agent i 's local state at time m is the observation that agent i makes of the state of the environment at time m . The third constraint says that evolution of the state of the environment is determined at each moment of time by agents choosing an action by applying their strategy at that time to their local state at that time. The joint action resulting from these individual choices is then resolved into a transition on the state of the environment using the transition relation from E . The final constraint says that agents' strategies are fixed during the course of a run. Intuitively, each agent picks a strategy, and then sticks to it. The interpretation π' of $\mathcal{I}(E, \Sigma)$ is determined from the interpretation π' of E by taking $\pi'(r, m) = \pi(r_e(m))$ for all points (r, m) .

Our epistemic strategy logic is now just an instantiation of the extended temporal epistemic logic in the strategy space generated by an environment. That is, we start with an environment E and an associated set of strategies Σ , and then work with the language $\text{ETLK}(\text{Ags} \cup \sigma(\text{Ags}) \cup \{e\}, \text{Prop}, \text{Var})$ in the interpreted system $\mathcal{I}(E, \Sigma)$. We call this instance of the language $\text{ESL}(\text{Ags}, \text{Prop}, \text{Var})$, or just ESL when the parameters are implicit.

Applications

In [8], we proposed a logic $\text{CTL}^*\text{K}(\text{Ags} \cup \sigma(\text{Ags}), \text{Prop})$ extending temporal epistemic logic with strategy agents to allow the reasoning about knowledge and strategy by standard epistemic operators. The language introduced above is a generalization of the definitions in [8], to which we have added the constructs $\exists x.\phi$ and $e_i(x)$. For formulas without these constructs, the semantics of ESL ignores the context Γ , so this component of the triple $\Gamma, \mathcal{I}(E, \Sigma), (r, m)$ can be removed from the definition, and it collapses to the definitions for $\text{CTL}^*\text{K}(\text{Ags} \cup \sigma(\text{Ags}), \text{Prop})$ in [8].

In the system $\mathcal{I}(E, \Sigma)$ we may refer, using distributed knowledge operators D_G where G contains the new strategic agents $\sigma(i)$, to what agents would know, should they take into account not just their own observations, but also information about other agent's strategies. For example, the distributed knowledge operator $D_{\{i, \sigma(i), \sigma(j)\}}$ captures what agent i would know, taking into account its own strategy and the strategy being used by agent j . Various applications of the usefulness of these distributed knowledge operators containing strategic agents are given in [8]. For example, we describe an application to *erasure policies* in computer security in which we write formulas such as

$$\neg D_{\emptyset} \neg (\text{done} \wedge \neg \text{exploited} \wedge EF \bigvee_{x \in \text{CCN}} D_{\{A, \sigma(A), \sigma(M)\}} (\text{cc} \neq x))$$

to state that it is possible for an attacker A on an e-commerce payment gateway to obtain information about a credit card number cc even after the transaction is done, provided that the attacker reasons using knowledge about their own observations, their own strategy, but also knowledge of the strategy being used by the merchant M . Here $\text{done} \wedge \neg \text{exploited}$ captures a situation where the transaction is done but the attacker has not run any exploit, and $D_{\{A, \sigma(A), \sigma(M)\}} (\text{cc} \neq x)$ says that the attacker is able to rule out the specific credit card number x from the range of possible values CCN for the actual credit card

number cc (so the attacker has at least one bit of information about the actual credit card number). The modality $\neg D_0 \neg$ is used to state that there is a point of the system where the formula holds. In particular, since the system builds in all possible strategies for the players, this modality captures a quantification over strategies.

In further applications given in [8], we showed that $CTL^*K(Ags \cup \sigma(Ags), Prop)$ can be used to express game theoretic equilibria, to reason about knowledge-based programs [4], and that many variants of alternating temporal epistemic logics that have been proposed in the literature can be expressed using $CTL^*K(Ags \cup \sigma(Ags), Prop)$. We refer the reader to [8] for details.

However, we had to make a restriction for some of these expressiveness results to formulas that do not contain uses of a common knowledge operator. We now show how the extended language of the present paper can remove this restriction.

Jamroga and van der Hoek [10] formulate a construct $\langle\langle H \rangle\rangle_{\mathcal{K}(G)}^\bullet \phi$ that says, effectively, that there is a strategy for a group H that another group G knows (for notion of group knowledge \mathcal{K} , which could be E for everyone knows, D for distributed knowledge, or C for common knowledge) to achieve goal ϕ . The semantics of this construct is given with respect to an environment E and a state s , and (in outline) is given by $E, s \models \langle\langle H \rangle\rangle_{\mathcal{K}(G)}^\bullet \phi$ if there exists a uniform strategy α for group H such that for all states t with $s \sim_G^\mathcal{K} t$, we have that all paths ρ from t that are consistent with α satisfy ϕ . Here $\sim_G^\mathcal{K}$ is the appropriate epistemic indistinguishability relation on states of E . We show in [8] how $\langle\langle H \rangle\rangle_{\mathcal{K}(G)}^\bullet \phi$ can be expressed in $CTL^*K(Ags \cup \sigma(Ags), Prop)$ for the cases where \mathcal{K} is either E or D .

In the case of the operator $\langle\langle H \rangle\rangle_{C(G)}^\bullet \phi$, the definition involves the common knowledge that a group G of agents would have if they were to reason taking into consideration the strategy being used by another group H . This does not appear to be expressible using $CTL^*K(Ags \cup \sigma(Ags), Prop)$. In particular, the formula $C_{G \cup \sigma(H)} \phi$ does not give the intended meaning. Instead, what needs to be expressed is the greatest fixpoint X of the equation $X \equiv \bigwedge_{i \in G} D_{\{i\} \cup \sigma(H)}(X \wedge \phi)$. The language $CTL^*K(Ags \cup \sigma(Ags), Prop)$ does not include fixpoint operators and it does not seem that the intended meaning is expressible. On the other hand, it can be expressed with $ESL(Ags, Var, Prop)$ in a natural way by a formula

$$C_G(e_{\sigma(H)}(x) \Rightarrow \phi)$$

which says that it is common knowledge to the group G that ϕ holds if the group H is running the strategy profile capture by the variable x . Using this idea, the construct $\langle\langle H \rangle\rangle_{C(G)}^\bullet \phi$ can be represented with ESL as

$$\exists x. C_G(e_{\sigma(H)}(x) \Rightarrow \phi) .$$

(We remark that a carefully stated equivalence result requires an appropriate treatment of initial states in the environment E . We refer to [8] for details.) Applying similar ideas, ESL can also be used to eliminate, from the results on reasoning about knowledge-based programs presented in [8], the restriction to knowledge-based programs not containing common knowledge operators.

Model Checking

Since interpreted systems are always infinite objects, we use environments to give a finite input for the model checking problem. For an environment E , a set of strategies Σ for E , and a context Γ for $I(E, \Sigma)$, we write $\Gamma, E, \Sigma \models \phi$ if $\Gamma, I(E, \Sigma), (r, 0) \models \phi$ for all runs r of I of $I(E, \Sigma)$. (Often, the formula ϕ will be a sentence, i.e., will have all variables x in the scope of an operator $\exists x$. In this case the statement $\Gamma, E, \Sigma \models \phi$ is independent of Γ and we may write simply $E, \Sigma \models \phi$) The model checking problem is to

determine whether $\Gamma, E, \Sigma \models \phi$ for a finite state environment E , a set Σ of strategies and a context Γ , where ϕ is an $\text{ESL}(\text{Ags}, \text{Var}, \text{Prop})$ formula.

For generality, we abstract Σ to a parameterized class such that for each environment E , the set $\Sigma(E)$ is a set of strategies for E . We say that the parameterized class $\Sigma(E)$ is *PTIME-presented*, if it is presented by means of an algorithm that runs in time polynomial in the size of E and verifies if a given strategy α is in $\Sigma(E)$. For example, the class $\Sigma(E)$ of all strategies for E can be PTIME-presented, as can $\Sigma^{\text{unif}}(E)$, $\Sigma^{\text{det}}(E)$ and $\Sigma^{\text{unif}, \text{det}}(E)$.

A naive model checking algorithm would construct a transition system over the set of states $S \times \Sigma(E)$, where S is the set of states of E , then apply model checking techniques on it. Note that a joint strategy for an environment E can be represented in space $|S| \times |\text{Acts}|$. Thus, the size of $S \times \Sigma(E)$ is exponential as a function of the size of E . This means that the naive procedure would require not less than exponential time. In fact, it is possible to do better than this (assuming that PSPACE is better than EXPTIME), provided we restrict to the CTL-based fragment of the language. This is the fragment in which the temporal operators occur only in the forms $A \bigcirc \phi$, $\neg A \neg \bigcirc \phi$, $A \phi_1 U \phi_2$, and $\neg A \neg \phi_1 U \phi_2$.

Theorem 1 *Let $\Sigma(E)$ be a PTIME presented class of strategies for environments E . The complexity of deciding, given an environment E , an ESL formula ϕ in the CTL-based fragment, and a context Γ for the free variables in an ESL formula ϕ relative to E and $\Sigma(E)$, whether $\Gamma, E, \Sigma(E) \models \phi$, is PSPACE-complete.*

That is, the additional constructs $\exists x.\phi$ and $e_i(x)$ that we have added to the logic of [8] to obtain ESL do not increase the complexity of model checking, provided we restrict to the CTL-based fragment. It remains to be resolved whether this holds for the full language.

Conclusions

Hybrid logic [1] is an approach to the extension of modal logics that uses “nominals”, i.e., propositions p that hold at a single world. These can be used in combination with operators such as $\exists p$, which marks an arbitrary world as the unique world at which nominal p holds. Our construct $\exists x$ is closely related to the hybrid construct $\exists p$, but we work in a setting that is richer in both syntax and semantics than previous works. There have been a few works using hybrid logic ideas in the context of epistemic logic [7, 13] but none are concerned with temporal logic. Hybrid temporal logic has seen a larger amount of study [2, 6, 5, 15], with variances in the semantics used for the model checking problem.

We note that if we were to view the variable x in our logic as a propositional constant, it would be true at a set of points in the system $I(E, \Sigma)$, hence not a nominal in that system. Results in [2], where a hybrid linear time temporal logic formula is checked in all paths in a given model, suggest that a variant of ESL in which x is treated as a nominal in $I(E, \Sigma)$ would have a complexity of model checking at least non-elementary, compared to our PSPACE complexity result.

Our model checking result seems to be more closely related to the a result in [5] that model checking a logic $\text{HL}(\exists, @, F, A)$ is PSPACE-complete. Here F is essentially a branching time future operator and A is a universal operator (similar to our D_0), the construct $@_p \phi$ says that ϕ holds at the world marked by the nominal p , and $\exists p(\phi)$ says that ϕ holds after marking some world by p . The semantics in this case does not unfold the model into either a tree or a set of linear structures before checking the formula, so the semantics of the hybrid existential \exists is close to our idea of quantifying over global states. Our language, however, has a richer set of operators, even in the temporal dimension, and introduces the strategic dimension in the semantics. It would be an interesting question for future work to consider fragments of our language to obtain more precise statement of the relationship with hybrid temporal logics.

Strategy Logic [3] is a (non-epistemic) generalization of ATL for perfect information strategies in which strategies may be explicitly named and quantified. Work on identification of more efficient variants of quantified strategy logic includes [11], which formulates a variant with a 2-EXPTIME-complete model checking problem. In both cases, strategies are perfect recall strategies, rather than the imperfect recall strategies that form the basis for our PSPACE-completeness result for model checking. The exploration of our logic over such a richer space of strategies is an interesting topic for future research.

References

- [1] P. Blackburn & J. Seligman (1998): *What are hybrid languages?* In M. de Rijke, H. Wansing & M. Zakharyashev, editors: *Advances in Modal Logic*, 1, CSLI Publications, pp. 41–62.
- [2] Laura Bozzelli & Ruggero Lanotte (2010): *Complexity and succinctness issues for linear-time hybrid logics*. *Theor. Comput. Sci.* 411(2), pp. 454–469. Available at <http://dx.doi.org/10.1016/j.tcs.2009.08.009>.
- [3] Krishnendu Chatterjee, Thomas A. Henzinger & Nir Piterman (2010): *Strategy logic*. *Inf. Comput.* 208(6), pp. 677–693.
- [4] R. Fagin, J.Y. Halpern, Y. Moses & M.Y. Vardi (1995): *Reasoning About Knowledge*. The MIT Press.
- [5] Massimo Franceschet & Maarten de Rijke (2006): *Model checking hybrid logics (with an application to semistructured data)*. *J. Applied Logic* 4(3), pp. 279–304. Available at <http://dx.doi.org/10.1016/j.jal.2005.06.010>.
- [6] Massimo Franceschet, Maarten de Rijke & Bernd-Holger Schlingloff (2003): *Hybrid Logics on Linear Structures: Expressivity and Complexity*. In: *10th International Symposium on Temporal Representation and Reasoning / 4th International Conference on Temporal Logic (TIME-ICTL 2003)*, pp. 166–173. Available at <http://doi.ieeecomputersociety.org/10.1109/TIME.2003.1214893>.
- [7] Jens Ulrik Hansen (2011): *A Hybrid Public Announcement Logic with Distributed Knowledge*. *Electr. Notes Theor. Comput. Sci.* 273, pp. 33–50. Available at <http://dx.doi.org/10.1016/j.entcs.2011.06.011>.
- [8] Xiaowei Huang & Ron van der Meyden (2014): *A Temporal Logic of Strategic Knowledge*. To appear KR’14, extended version available at http://www.cse.unsw.edu.au/~meyden/research/atl_obs.pdf.
- [9] Wojciech Jamroga & Thomas Ågotnes (2007): *Constructive knowledge: what agents can achieve under imperfect information*. *Journal of Applied Non-Classical Logics* 17(4), pp. 423–475.
- [10] Wojciech Jamroga & Wiebe van der Hoek (2004): *Agents that Know How to Play*. *Fundamenta Informaticae* 62, pp. 1–35.
- [11] Fabio Mogavero, Aniello Murano & Moshe Y. Vardi (2010): *Reasoning About Strategies*. In: *FSTTCS*, pp. 133–144.
- [12] Sieuwert van Otterloo & Geert Jonker (2005): *On Epistemic Temporal Strategic Logic*. *ENTCS* 126, pp. 77–92.
- [13] O. Roy (2009): *A dynamic-epistemic hybrid logic for intentions and information changes in strategic games*. *Synthese* 171(2), pp. 291–320. Available at <http://dx.doi.org/10.1007/s11229-009-9644-3>.
- [14] Pierre-Yves Schobbens (2004): *Alternating-time logic with imperfect recall*. *ENTCS* 85(2), pp. 82–93.
- [15] Thomas Schwentick & Volker Weber (2007): *Bounded-Variable Fragments of Hybrid Logics*. In: *Proc. STACS 2007, 24th Annual Symposium on Theoretical Aspects of Computer Science*, Springer LNCS 4393, pp. 561–572. Available at http://dx.doi.org/10.1007/978-3-540-70918-3_48.