Proceedings of the Sixteenth International Conference on
Principles of Knowledge Representation and Reasoning (KR 2018)

Changing Observations in Epistemic Temporal Logic

Aurèle Barrière
ENS Rennes

Bastien Maubert, Aniello Murano, Sasha Rubin
Università degli Studi di Napoli "Federico II"

Abstract

We study dynamic changes of agents' observational power in logics of knowledge and time. We consider CTL*K, the extension of CTL* with knowledge operators, and enrich it with a new operator that models a change in an agent's way of observing the system. We extend the classic semantics of knowledge for agents with perfect recall to account for changes of observation, and we show that this new operator strictly increases the expressivity of CTL*K. We also provide a model-checking procedure for the logic we introduce, which has the same complexity as the best known model-checking procedure for the less expressive logic CTL*K.

Introduction

In epistemic temporal logics such as LTLK and CTLK, an agent's view of a particular state of the system is given by an observation of that state, and an agent's observation of a given state does not change over time. In other words, these frameworks have no primitive for reasoning about agents whose observation power can change. Because this phenomenon occurs in real scenarios, for instance when a user of a system is granted a higher security level giving access to more data, we propose here to tackle this problem.

We extend CTL*K with a new unary operator, Δ^o, that represents changes of observation power, and is read "the agent changes her observation power to o". Formula $\Delta^{o_1} AF(\Delta^{o_2}(Kp \vee K\neg p))$ for instance expresses that "For an agent with initial observation power o_1, in all possible futures there exists a point where, if the agent updates her observation power to o_2, she learns whether or not the proposition p holds". If in this example o_1 and o_2 represent different "security levels" and p is sensitive information, then the formula expresses a possible avenue for attack. Another motivation for studying such logics comes from the recently introduced Strategy Logic with Imperfect Information (Berthon et al. 2017), an extension of Strategy Logic (Mogavero et al. 2014) in which agents can dynamically change observation power when changing strategies.

For memoryless agents, adapting the semantics of CTL*K to include the observation-change operator is straightforward, and model-checking algorithms also can be easily adapted at no cost in complexity: the model-checking problem remains PSPACE-complete as for CTL*K (Raimondi and Lomuscio 2005; Kong and Lomuscio 2017).

The case of agents with perfect recall, which we study in this work, is more delicate. The model-checking problem for LTLK and CTL*K is nonelementary decidable (van der Meyden and Shilov 1999; Bozzelli, Maubert, and Pinchinat 2015), with k-EXPTIME upper-bound for formulas with at most k nested knowledge operators. The same upper-bounds are known for CTLK (Dima 2009). In this work we show that the observation-change operator increases expressivity but, as for the memoryless semantics, it does not increase the known complexity of the model-checking problem.

CTL*KΔ

Let $\mathcal{AP}$ be a countably infinite set of atomic propositions, $Ag = \{a_1, \ldots, a_m\}$ a finite set of agents, and $\mathcal{O}$ a finite set of *observations*, that represent possible observational powers of agents. We extend the syntax of CTL*K, with an observation change operator Δ^o_a for each agent a.

Definition 1 (Syntax)**.** *The sets of history formulas φ and path formulas ψ of* CTL*KΔ *are defined by the following grammar:*

$$\begin{array}{lcl} \varphi & ::= & p \mid \neg\varphi \mid \varphi \wedge \varphi \mid A\psi \mid K_a\varphi \mid \Delta^o_a\varphi \\ \psi & ::= & \varphi \mid \neg\psi \mid \psi \wedge \psi \mid X\psi \mid \psi U \psi, \end{array}$$

where $p \in \mathcal{AP}$, $a \in Ag$ and $o \in \mathcal{O}$.

Operators X and U are the classic *next* and *until* operators, and A is the universal path quantifier from branching-time temporal logics. K_a is the knowledge operator from epistemic logics, and $K_a\varphi$ reads as "agent a knows that φ is true". Our new *observation change* operator, Δ^o_a, reads as "agent a now observes the system with observation o".

Models of CTL*KΔ are Kripke structures equipped with one relation $\sim_o$ on states for each observation o.

Definition 2 (Models)**.** *A* Kripke structure with observations *is a structure $M = (\text{AP}, S, T, V, \{\sim_o\}_{o\in\mathcal{O}}, s^\iota, \boldsymbol{o}^\iota)$, where*

- AP $\subset \mathcal{AP}$ *is a finite subset of atomic propositions,*
- *S is a set of states,*
- *$T \subseteq S \times S$ is a left-total*[1] *transition relation,*

[1] i.e., for every $s \in S$ there exists $s' \in S$ such that sTs'. This cosmetic restriction is made to avoid having to deal with finite runs ending in deadlocks.

- $V : S \to 2^{\mathrm{AP}}$ *is a valuation function,*
- $\sim_o \subseteq S \times S$ *is an equivalence relation, for each* $o \in \mathcal{O}$*,*
- $s^\iota \subseteq S$ *is an initial state, and*
- $\boldsymbol{o}^\iota \in \mathcal{O}^{Ag}$ *is an initial observation for each agent.*

A *path* is an infinite sequence of states $\pi = s_0 s_1 \dots$ such that for all $i \geq 0$, $s_i T s_{i+1}$, and a *history* h is a finite prefix of a path.

Observation records. To define which histories the agent cannot distinguish, we need to keep track of how she observed the system at each point in time. To do so, we record each observation change as a pair (o, n), where o is the new observation and n is the time when this change occurs.

Definition 3. *An* observation record r *is a finite word over* $\mathcal{O} \times \mathbb{N}$*, i.e.,* $r \in (\mathcal{O} \times \mathbb{N})^*$.

We write $r \cdot (o, n)$ for the observation record obtained by appending (o, n) to the observation record r, and $r[n]$ for the subrecord consisting of all pairs (o, m) in r such that $m = n$. An observation record r *stops at* n if $r[m]$ is empty for all $m > n$, and r *stops at history* h if it stops at $|h| - 1$.

We shall write $\boldsymbol{r}$ for a *record tuple* $\{r_a\}_{a \in Ag}$. Given a record tuple $\boldsymbol{r} = \{r_a\}_{a \in Ag}$ and $a \in Ag$ we write $\boldsymbol{r}_a$ for r_a, and for an observation o and time n we let $\boldsymbol{r} \cdot (o, n)_a$ be the record tuple $\boldsymbol{r}$ where $\boldsymbol{r}_a$ is replaced with $\boldsymbol{r}_a \cdot (o, n)$.

Observations at time n**.** We let $ol_a(\boldsymbol{r}, n)$ be the list of observations used by agent a at time n:

$$ol_a(\boldsymbol{r}, 0) = \boldsymbol{o}^\iota_a \cdot o_1 \cdot \ldots \cdot o_k, \quad \text{if } \boldsymbol{r}_a[0] = (o_1, 0) \cdot \ldots \cdot (o_k, 0), \text{ and}$$
$$ol_a(\boldsymbol{r}, n+1) = \mathrm{last}(ol_a(\boldsymbol{r}, n)) \cdot o_1 \cdot \ldots \cdot o_k, \quad \text{if } \boldsymbol{r}_a[n+1] = (o_1, n+1) \cdot \ldots \cdot (o_k, n+1).$$

Observe that $ol_a(\boldsymbol{r}, n)$ is never empty: if no observation change occurs at time n, $ol(r, n)$ only contains the last observation taken by the agent. If r is empty, the latter is the initial observation o_ι defined by the model.

Synchronous perfect recall. The usual definition of synchronous perfect recall states that for an agent with observation o, histories h and h' are indistinguishable if they have the same length and are point-wise indistinguishable, i.e., $|h| = |h'|$ and for each $i < |h|$, $h_i \sim_o h'_i$. We adapt this definition to changing observations: two histories are indistinguishable if, at each point in time, the states are indistinguishable for all observations used at that time.

Definition 4 (Dynamic synchronous perfect recall)**.** *Given a record tuple* $\boldsymbol{r}$*, two histories* h *and* h' *are equivalent for agent* a*, written* $h \approx^{\boldsymbol{r}}_a h'$*, if*

$$|h| = |h'| \text{ and } \forall i < |h|,\ \forall o \in ol_a(\boldsymbol{r}, i),\ h_i \sim_o h'_i.$$

Definition 5 (Semantics)**.** *Fix a model* M*. A history formula* φ *is evaluated in a history* h *and a record tuple* $\boldsymbol{r}$*. A path formula* ψ *is interpreted on a run* π*, a point in time* $n \in \mathbb{N}$ *and a record ruple* $\boldsymbol{r}$*. The semantics is defined by induction:*

$h, \boldsymbol{r} \models p$	*if*	$p \in V(last(h))$
$h, \boldsymbol{r} \models \neg\varphi$	*if*	$h, \boldsymbol{r} \not\models \varphi$
$h, \boldsymbol{r} \models \varphi_1 \wedge \varphi_2$	*if*	$h, \boldsymbol{r} \models \varphi_1$ *and* $h, \boldsymbol{r} \models \varphi_2$
$h, \boldsymbol{r} \models A\psi$	*if*	$\forall \pi$ *s.t.* $h \preccurlyeq \pi,\ \pi, \lvert h\rvert - 1, \boldsymbol{r} \models \psi$
$h, \boldsymbol{r} \models K_a\varphi$	*if*	$\forall h'$ *s.t.* $h' \approx^{\boldsymbol{r}}_a h,\ h', \boldsymbol{r} \models \varphi$
$h, \boldsymbol{r} \models \Delta^o_a\varphi$	*if*	$h, \boldsymbol{r} \cdot (o, \lvert h\rvert - 1)_a \models \varphi$
$\pi, n, \boldsymbol{r} \models \varphi$	*if*	$\pi_{\leq n}, \boldsymbol{r} \models \varphi$
$\pi, n, \boldsymbol{r} \models \neg\psi$	*if*	$\pi, n, \boldsymbol{r} \not\models \psi$
$\pi, n, \boldsymbol{r} \models \psi_1 \wedge \psi_2$	*if*	$\pi, n, \boldsymbol{r} \models \psi_1$ *and* $\pi, n, \boldsymbol{r} \models \psi_2$
$\pi, n, \boldsymbol{r} \models X\psi$	*if*	$\pi, (n+1), \boldsymbol{r} \models \psi$
$\pi, n, \boldsymbol{r} \models \psi_1 U \psi_2$	*if*	$\exists m \geq n$ *s.t.* $\pi, m, \boldsymbol{r} \models \psi_2$ *and* $\forall k$ *s.t.* $n \leq k < m,\ \pi, k, \boldsymbol{r} \models \psi_1$

If there is only one possible observation power, then the observation-change operator is useless, and in that case $\mathsf{CTL}^*\mathsf{K}\Delta$ and $\mathsf{CTL}^*\mathsf{K}$ are equi-expressive. However, if there are at least two distinct possible observation powers, then we can prove that the observation-change operator Δ^o_a adds expressivity to $\mathsf{CTL}^*\mathsf{K}$:

Theorem 1. *If* $|\mathcal{O}| > 1$ *then* $\mathsf{CTL}^*\mathsf{K}\Delta$ *is strictly more expressive than* $\mathsf{CTL}^*\mathsf{K}$*.*

We can also prove that this added expressivity comes for free in terms of complexity of model-checking, as we can establish the same upper-bounds as those known for $\mathsf{CTL}^*\mathsf{K}$:

Theorem 2. *The model-checking problem for* $\mathsf{CTL}^*\mathsf{K}\Delta$ *is in* k-EXPTIME *for formulas of knowledge depth at most* k*.*

We are currently working on establishing matching lower-bounds for $\mathsf{CTL}^*\mathsf{K}$, as they are not known yet. We would inherit them also for $\mathsf{CTL}^*\mathsf{K}\Delta$, which would prove that the model-checking procedure we established is essentially optimal, and that the additional expressivity provided by the observation change operator really comes for free.

References

Berthon, R.; Maubert, B.; Murano, A.; Rubin, S.; and Vardi, M. Y. 2017. Strategy logic with imperfect information. In *LICS*, 1–12.

Bozzelli, L.; Maubert, B.; and Pinchinat, S. 2015. Uniform strategies, rational relations and jumping automata. *Information and Computation* 242:80 – 107.

Dima, C. 2009. Revisiting satisfiability and model-checking for ctlk with synchrony and perfect recall. In *CLIMA IX-2008*, 117–131.

Kong, J., and Lomuscio, A. 2017. Symbolic model checking multi-agent systems against CTL*K specifications. In *AAMAS*, 114–122.

Mogavero, F.; Murano, A.; Perelli, G.; and Vardi, M. Y. 2014. Reasoning about strategies: On the model-checking problem. *ACM Trans. Comput. Log.* 15(4):34:1–34:47.

Raimondi, F., and Lomuscio, A. 2005. The complexity of symbolic model checking temporal-epistemic logics. In *CS&P*, 421–432.

van der Meyden, R., and Shilov, N. V. 1999. Model checking knowledge and time in systems with perfect recall (extended abstract). In *FSTTCS*, 432–445.