TECHNISCHE UNIVERSITÄT DRESDEN

Dresden University of Technology
Institute for Theoretical Computer Science
Chair for Automata Theory

LTCS–Report

Complexity and Succinctness of Public Announcement Logic

Carsten Lutz

LTCS-Report 05-09

Lehrstuhl für Automatentheorie
Institut für Theoretische Informatik
TU Dresden
http://lat.inf.tu-dresden.de

Hans-Grundig-Str. 25
01062 Dresden
Germany

Complexity and Succinctness of Public Announcement Logic

Carsten Lutz
Institute for Theoretical Computer Science
TU Dresden, Germany
lutz@tcs.inf.tu-dresden.de

ABSTRACT

There is a recent trend of extending epistemic logic (EL) with dynamic operators that allow to express the evolution of knowledge and belief induced by knowledge-changing actions. The most basic such extension is public announcement logic (PAL), which is obtained from EL by adding an operator for truthful public announcements. In this paper, we consider the computational complexity of PAL and show that it coincides with that of EL. This holds in the single- and multi-agent case, and also in the presence of common knowledge operators. We also prove that there are properties that can be expressed exponentially more succinct in PAL than in EL. This shows that, despite the known fact that PAL and EL have the same expressive power, there *is* a benefit in adding the public announcement operator to EL: it exponentially increases the succinctness of formulas without having negative effects on computational complexity.

1. INTRODUCTION

One of the most prominent applications of logic in agent-based systems is reasoning about the knowledge and belief of agents. Although traditionally, epistemic logic (EL) is the basic logical tool for this purpose [14], it has always been clear that EL is too simple for many relevant applications in this area. Most strikingly, basic EL does not include any syntactic or semantic means for representing dynamic and evolutionary aspects of knowledge. Since it is crucial for almost all agent-based systems that the knowledge and belief of agents are subject to change [1], such expressive means are often indispensible. In the literature, there are two dominant approaches to adding dynamics to EL: first, EL can be extended with a temporal component that allows to reason about the evolution of knowledge over time [9]. And second, EL can be extended with dynamic operators that allow to describe the ramifications of knowledge-changing actions. The latter approach is a relatively recent development, and the resulting extensions of EL are often called *dynamic epistemic logics (DELs)* [8].

By now, a large number of DELs has been proposed, and the various proposals differ considerably in expressive power [2, 11, 16, 3, 4, 6, 7]. However, there is a dynamic operator that is included in almost all proposed logics: the *public announcement* operator that has first been introduced in [16]. This operator allows to state that, after some announcement that is publicly made by an outsider to all agents simultaneously, some property holds true. Both the announcement and the property may include epistemic statements such as "agent a knows fact F" or "agent a believes that agent b knows fact F". The announcement is assumed to be truthful, i.e., the person making the announcement does not lie. The effect of the announcement being public is that everybody knows the announced fact, everybody knows that everybody knows it, and so forth. It is interesting to note that the announced fact is not necessarily true anymore *after* the announcement. For example, this is the case if the announced fact is "agent a knows fact F, but agent b doesn't know that" (because, after the announcement, agent b knows that agent a knows F).

The appropriateness of the public announcement operator for incorporating dynamics into EL has been demonstrated by elegantly modelling a number of standard problems involving public announcements such as the muddy children puzzle [3, 8]. It is also witnessed by the inclusion of this operator in almost all proposed DELs—see the papers cited above. Existing research about the operator has mainly concentrated on the expressiveness and axiomatics of the obtained extensions of EL. For example, it is known that EL with the public announcement operator has the same expressive power as EL without it.[1] However, performing *actual reasoning* is of great importance when applying DELs in agent-based systems, and thus, the computational properties of such logics need to be analyzed. While the computational complexity of model checking DELs is easily pinpointed, not much is known about deciding satisfiabilityand validity (henceforth called *reasoning*). Therefore, the purpose of the current paper is to *analyze the computational complexity of reasoning in epistemic logics extended with the public announcement operator.* Since computational complexity, expressiveness, and succinctness issues turn out to be intimately and subtly related to computational complexity in the considered logics, we include the latter in our analysis.

We start our investigation with *public announcement logic (PAL)*, the extension of basic EL with the public announcement operator. As noted above, the expressive power of PAL is known to be identical to the expressive power of EL: there exists an equivalence-preserving translation from the former to the latter [16, 8]. Computationally, this translation is only moderately useful: it yields decidability of reasoning in PAL, but it does not produce tight complexity bounds due to an exponential blowup in formula size. More precisely, the known translation yields upper complexity bounds for PAL that are identical to those for EL, but raised by one

9[1]This should not be taken as an indication that the operator is not worth studying; see below.

exponential. In contrast, the best known lower bounds are the ones from EL.

In this paper, we show that the existing upper bounds can be improved by one exponential, and thus reasoning in PAL is of the same complexity as reasoning in EL. To this end, we first propose a novel, equivalence-preserving translation from PAL to EL. Like the existing one, this translation induces an exponential blowup in formula size. The advantage of the new translation is that it can be modified such that it is only satisfiability-preserving, but avoids the exponential blowup in formula size. The modified translation takes formulas of single-agent PAL to formulas of single-agent EL, and formulas of multi-agent PAL to formulas of multi-agent EL extended with an "everybody knows" operator. Thus, it can be used to prove that (i) single-agent PAL is NP-complete, and (ii) multi-agent PAL is PSPACE-complete. We then extend our equivalence-preserving translation and its satisfiability-preserving modification to PAL extended with (two variants of) common knowledge. In this case, the target language of the translation is propositional dynamic logic (PDL), and we obtain EXPTIME-completeness results.

Due to the facts that PAL and EL are equally expressive and of the same computational complexity, one may be tempted to think that the addition of the public announcement operator to EL is only syntactic sugar and not of much interest. However, it has been convincingly argued in [3, 8] that PAL is a much more intuitive and natural formalism for talking about the dynamics of knowledge than EL. In this paper, we add another, more concrete advantage of PAL: we prove that there are properties that can be expressed exponentially more succinct in PAL than in EL. Thus, the public announcement operator contributes to the succinctness of the logic, and this exponential increase in succinctness is not even penalised by an increase in computational complexity. Of course, our succinctness result also implies that one cannot hope to find an *equivalence-preserving* translation from PAL to EL that avoids an exponential blowup. A limitation of our (current) succinctness result is that it applies only to the class of all Kripke structures, and not to the class of epistemic structures.

2. PUBLIC ANNOUNCEMENT LOGIC

Let PL and N be countable sets of *propositional letters* and *agents*. The formulas of *public announcement logic with common knowledge (PALC)* are built according to the following syntax rule:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_a\varphi \mid [\varphi]\psi \mid C_A\varphi$$

where p ranges over PL, a ranges over N, and A ranges over 2^N. The operator $[\varphi]\psi$ is the *public announcement operator*: $[\varphi]\psi$ states that, after φ is publically announced, ψ holds true. A detailed introduction to PALC and how it can be used to model the dynamics of knowledge can be found in [3, 8]. Here, we confine ourselves to a simple example: the formula

$$K_a(K_b\mathsf{secret} \wedge [K_b\mathsf{secret}]K_c\mathsf{leaked}_{a,b}),$$

with secret and $\mathsf{leaked}_{a,b}$ propositional letters, states the following: agent a knows that agent b knows the secret, and if this is publically announced, then agent c will know that agent a has leaked the secret to b.

By dropping the *common knowledge operator* $C_A\varphi$ from PALC, we get formulas of *public announcement logic (PAL)*. By dropping the public announcement operator from PAL, we get formulas of *epistemic logic (EL)*. As usual, we use $\varphi \vee \psi$ as an abbreviation for $\neg(\neg\varphi \wedge \neg\psi)$, $\varphi \to \psi$ for $\neg\varphi \vee \psi$, $\widehat{K}_a\varphi$ for $\neg K_a\neg\varphi$, $\langle\varphi\rangle\psi$ for $\neg[\varphi]\neg\psi$, and $\top$ for $p \vee \neg p$, with p an arbitrary (but fixed) propositional letter. When talking of *single-agent PAL or EL*, we assume that N is a singleton. In this case, we drop the index $\cdot_a$ from $K_a\varphi$ and $\widehat{K}_a\varphi$.

$[\varphi]p$	$\rightsquigarrow$	$\varphi \to p$
$[\varphi](\psi \wedge \chi)$	$\rightsquigarrow$	$[\varphi]\psi \wedge [\varphi]\chi$
$[\varphi]\neg\psi$	$\rightsquigarrow$	$\varphi \to \neg[\varphi]\psi$
$[\varphi]K_a\psi$	$\rightsquigarrow$	$\varphi \to K_a[\varphi]\psi$

Figure 1: The standard translation from PAL to EL.

The semantics of PALC is defined by means of Kripke structures. An *epistemic model* (or *model* for short) is a triple $\mathcal{M} = (S, \sim, V)$ with S a non-empty set of states, $\sim = (\sim^a)_{a\in N}$ a family of equivalence relations on S, and V a function assigning a set of states $V(p) \subseteq S$ with each $p \in \mathsf{PL}$. Given a model $\mathcal{M} = (S, \sim, V)$ and an $s \in S$, we define:

$$\begin{array}{rll}
\mathcal{M}, s \models p & \text{iff} & s \in V(p) \\
\mathcal{M}, s \models \neg\varphi & \text{iff} & \mathcal{M}, s \not\models \varphi \\
\mathcal{M}, s \models \varphi \wedge \psi & \text{iff} & \mathcal{M}, s \models \varphi \text{ and } \mathcal{M}, s \models \psi \\
\mathcal{M}, s \models K_a\varphi & \text{iff} & s \sim_a t \text{ implies } \mathcal{M}, t \models \varphi, \text{ for all } t \in S \\
\mathcal{M}, s \models [\varphi]\psi & \text{iff} & \mathcal{M}, s \models \varphi \text{ implies } \mathcal{M}|\varphi, s \models \psi \\
\mathcal{M}, s \models C_A\varphi & \text{iff} & s\, (\bigcup_{a\in A} \sim_a)^*\, t \text{ implies } \mathcal{M}, t \models \varphi, \\
 & & \qquad\qquad \text{for all } t \in S
\end{array}$$

where $\cdot^*$ denotes the reflexive-transitive closure operator on binary relations, and the model $\mathcal{M}|\psi := (T, \approx^i, U)$ is defined as follows:

$$\begin{array}{rcl}
T & := & \{t \in S \mid \mathcal{M}, t \models \psi\} \\
\approx_a & := & \sim_a \cap (T \times T) \\
U(p) & := & V(p) \cap T.
\end{array}$$

It is easily checked that the semantics of the derived modalities is as follows:

$$\begin{array}{rll}
\mathcal{M}, s \models \widehat{K}_a\varphi & \text{iff} & s \sim_a t \text{ and } \mathcal{M}, t \models \varphi, \text{ for some } t \in S \\
\mathcal{M}, s \models \langle\varphi\rangle\psi & \text{iff} & \mathcal{M}, s \models \varphi \text{ and } \mathcal{M}|\varphi, s \models \psi.
\end{array}$$

Given a model $\mathcal{M}$, we use $S(\mathcal{M})$ to denote the set of states of $\mathcal{M}$. It is not too difficult to see that, via the second argument, the public announcement operator can be conceived as a normal modal operator (in the sense of, e.g., [5]) that behaves in many ways like a modal operator for a functional modality. For example, $\langle\varphi\rangle\psi$ implies $[\varphi]\psi$.

3. EXPRESSIVITY AND SUCCINCTNESS

As first observed by Plaza [16], the basic epistemic logic with public announcements PAL has exactly the same expressive power as EL. Indeed, by exhaustively applying the rewrite rules in Figure 1, one can convert every PAL formula into an equivalent EL formula [8]. For the purposes of this paper, however, it is convenient to work with a translation that is different from this standard one. To introduce the new translation, we need a bit of notation: for φ a PAL formula and $\sigma = \varphi_1 \cdots \varphi_k$ a finite sequence of PAL formulas, we use

- $|\varphi|$ to denote the length of φ, i.e. the number of symbols needed to write down φ, including symbols such as "[" and "]"; likewise, $|\sigma|$ denotes $|\varphi_1| + \cdots + |\varphi_k|$;
- $\mathcal{M}|\sigma$ as an abbreviation for $(((\mathcal{M}|\varphi_1)|\varphi_2)\cdots|\varphi_k)$ with $\mathcal{M}|\varepsilon = \mathcal{M}$;
- $\mathsf{pre}(\sigma)$ to denote the set of all true prefixes of σ including the empty sequence ε;
- for each $\nu \in \mathsf{pre}(\sigma)$, ν/σ to denote the leftmost symbol of σ that is not in ν.

We now define, for every PAL formula φ and every finite sequence of PAL formulas σ, an EL formula φ^σ. The definition of the formulas φ^σ proceeds by induction on $|\varphi| + |\sigma|$ as shown in Figure 2. In the $K_a\varphi$ case, the conjunction collapses to true if $\sigma = \varepsilon$. To see that we really do induction on $|\varphi| + |\sigma|$, note that the symbols "[" and "]" contribute to the size of the left-hand side of the last line.

The most important property of the formulas φ^σ is given by the following lemma. It can be proved by induction on $|\varphi| + |\sigma|$, for details consult the Appendix A.

Lemma 1. *For all models $\mathcal{M} = (S, \sim, V)$, PAL formulas φ, finite sequences of PAL formulas σ, and states $s \in S(\mathcal{M}|\sigma)$, we have $\mathcal{M}, s \models \varphi^\sigma$ iff $\mathcal{M}|\sigma, s \models \varphi$.*

Lemma 1 clearly implies that each PAL formula φ is equivalent to the EL formula φ^ε, which gives us the new translation. It can easily be seen that this translation usually produces different formulas than the standard translation given in Figure 1. For example, the result of translating $[p_1][p_2]\cdots[p_k]q$ with our translation results in the formula

$$p_1 \to (p_2 \to (\cdots \to (p_k \to q))),$$

which is of length $\mathcal{O}(k)$. In contrast, the standard translation produces a highly redundant formula of length $2^{\mathcal{O}(k)}$. Thus, there are formulas on which our translation is exponentially more succinct. In general, however, the new translation does not avoid an exponential blowup in formula size. See Theorem 2 below for example formulas.

Given the exponential blowup induced by both translations, it is a natural question whether the exponential blowup in translating PAL to EL can be avoided at all. We answer this question to the negative: PAL is exponentially more succinct than epistemic logic, at least on *unrestricted models*, i.e., on models whose relations are not required to be equivalence relations. Our aim is to prove the following:

Theorem 2. *For $i \geq 0$, define*

- $\varphi_0 := \top$;
- $\varphi_{i+1} := \langle\langle\varphi_i\rangle\widehat{K}_a\top\rangle\widehat{K}_b\top$.

On unrestricted models, every EL formula ψ equivalent to φ_i is of length at least 2^i, for all $i \geq 0$.

Define a sequence of EL formulas $\psi_0, \psi_0', \psi_1, \psi_1', \psi_2, \ldots$ as follows:

- $\psi_0 := \top$;
- $\psi_i' := \psi_i \wedge \widehat{K}_a\psi_i$;
- $\psi_{i+1} := \psi_{i-1}' \wedge \widehat{K}_b\psi_{i-1}'$.

$$\begin{array}{rcl} p^\sigma & := & p \\ (\neg\varphi)^\sigma & := & \neg\varphi^\sigma \\ (\varphi \wedge \psi)^\sigma & := & \varphi^\sigma \wedge \psi^\sigma \\ (K_a\varphi)^\sigma & := & K_a(\bigwedge_{\nu \in \mathsf{pre}(\sigma)} (\nu/\sigma)^\nu \to \varphi^\sigma) \\ ([\varphi]\psi)^\sigma & := & \varphi^\sigma \to \psi^{\sigma\cdot\varphi} \end{array}$$

Figure 2: The new translation.

Using the translation from PAL to EL given in Section 3, which also applies to the case of unrestricted models, it is straightforward to prove by induction on i that, for $i \geq 0$, φ_i is equivalent to ψ_i. Thus, for proving Theorem 2 it suffices to prove that every EL formula χ that is equivalent to ψ_i on unrestricted models is of length at least 2^i, for all $i \geq 0$.

Let $N = \{a, b\}$ be the set of relevant agents. In what follows, a *path set* is a subset of N^*. For φ an EL formula φ over the set of agents N, we define the *path-set* P_φ of φ by structural induction as follows:

- $P_p := \{\varepsilon\}$;
- $P_{\neg\varphi} := P_\varphi$;
- $P_{\varphi\wedge\psi} := P_\varphi \cup P_\psi$;
- $P_{K_a\varphi} := \{\varepsilon\} \cup \{aw \mid w \in P_\varphi\}$.

Now, let χ be an EL formula that is equivalent to ψ_i, for some $i > 0$. We show the following:

$$P_{\psi_i} \subseteq P_\chi. \qquad (*)$$

To prove $(*)$, assume to the contrary that there is a $\widehat{w} \in P_{\psi_i} \setminus P_\chi$. Define a model $\mathcal{M} = (S, \sim, V)$ as follows:

- $S = P_{\psi_i}$;
- the relation $\sim_\sigma$, $\sigma \in N$, is defined by setting $w \sim_\sigma v$ if $v = w\sigma$, for $w, v \in S$;
- $V(p) = \emptyset$ for all $p \in \mathsf{PL}$.

A second model $\mathcal{M}' = (S', \sim', V')$ is defined as the restriction of $\mathcal{M}$ to the set of states

$$S' := S \setminus \{w \in N^* \mid w = \widehat{w}w' \text{ for some } w' \in N^*\}.$$

It is not too hard to show that $\mathcal{M}$ is a model of ψ_i, but $\mathcal{M}'$ is not. Details are left to the reader.

Lemma 3. *$\mathcal{M}, \varepsilon \models \psi_i$ and $\mathcal{M}', \varepsilon \not\models \psi_i$.*

We now show that χ cannot distinguish ε in $\mathcal{M}$ from ε in $\mathcal{M}'$, i.e., $\mathcal{M}, \varepsilon \models \chi$ iff $\mathcal{M}', \varepsilon \models \chi$. This clearly is a contradiction to Lemma 3 and the fact that χ is equivalent to φ_i. Thus, we have established $(*)$. The fact that χ cannot distinguish ε in $\mathcal{M}$ from ε in $\mathcal{M}'$ is an immediate consequence of the following lemma, which is proved in Appendix A.

Lemma 4. *For all $s \in S'$ and $\varphi \in \mathsf{sub}(\chi)$ such that*

$$\{sw \mid w \in P_\varphi\} \subseteq P_\chi,$$

we have $\mathcal{M}, s \models \varphi$ iff $\mathcal{M}', s \models \varphi$.

$$
\begin{array}{rcl}
R(a) & := & \{(\varepsilon, a)\} \\
R(\neg\varphi) & := & R(\varphi) \cup \{(\varepsilon, \neg\varphi)\} \\
R(\varphi \wedge \psi) & := & R(\varphi) \cup R(\psi) \cup \{(\varepsilon, \varphi \wedge \psi)\} \\
R(K_i\varphi) & := & R(\varphi) \cup \{(\varepsilon, K_i\varphi)\} \\
R([\varphi]\psi) & := & R(\varphi) \cup \{(\varphi \cdot \sigma, \vartheta) \mid (\sigma, \vartheta) \in R(\psi)\} \\
 & & \cup\, \{(\varepsilon, [\varphi]\psi)\}
\end{array}
$$

Figure 3: The relevant pairs.

$$
\begin{array}{rcl}
B^\sigma_q & := & p^\sigma_q \leftrightarrow q \\
B^\sigma_{\neg\varphi} & := & p^\sigma_{\neg\varphi} \leftrightarrow \neg p^\sigma_\varphi \\
B^\sigma_{\varphi\wedge\psi} & := & p^\sigma_{\varphi\wedge\psi} \leftrightarrow (p^\sigma_\varphi \wedge p^\sigma_\psi) \\
B^\sigma_{K_i\varphi} & := & p^\sigma_{K_i\varphi} \leftrightarrow K_i\big(\bigwedge_{\nu \in \mathsf{pref}(\sigma)} p^\nu_{\nu/\sigma} \rightarrow p^\sigma_\varphi\big) \\
B^\sigma_{[\varphi]\psi} & := & p^\sigma_{[\varphi]\psi} \leftrightarrow (p^\sigma_\varphi \rightarrow p^{\sigma\cdot\varphi}_\psi)
\end{array}
$$

Figure 4: The biimplications B^σ_φ.

Finally, $|\chi| \geq 2^i$ is a consequence of $(*)$ together with the following lemma and the fact that, as is easily proved by structural induction, we have $|\varphi| \geq |P_\varphi|$ for all formulas φ.

Lemma 5. *For all $i \geq 0$, $|P_{\psi_i}| \geq 2^i$.*

Proof. It is straightforward to prove by induction on i that, for all $i \geq 0$, we have $\{a, b\}^i \subseteq P_{\psi_i}$. □

This finishes the proof of Theorem 2. We believe that PAL is also exponentially more succinct than EL on epistemic structures, but leave the proof as an open problem.

4. UPPER BOUNDS FOR PAL WITHOUT COMMON KNOWLEDGE

Given the succinctness of PAL established in the previous section, the question arises whether a penalty has to be paid for this succinctness in terms of computational complexity: is reasoning in PAL more expensive than reasoning in epistemic logic? Interestingly, this is not the case. We show that satisfiability in PAL is NP-complete in the single-agent case and PSpace-complete in the multi-agent case, just as in epistemic logic [13]. Lower bounds are immediate since PAL contains EL as a fragment. The idea for obtaining the upper bounds is to convert the equivalence-preserving, but exponential translation given in Section 3 into a satisfiability-preserving and polynomial translation. We start with the single-agent case.

4.1 Reducing Single-Agent PAL

We start with introducing some relevant notions. As these notions will also be useful for dealing with multi-agent PAL, we do not restrict ourselves to single-agent formulas here. Let φ be a PAL formula. With $\mathsf{sub}(\varphi)$, we denote the set of all subformulas of φ, including φ. With $\Sigma(\varphi)$, we denote the set of all pairs (σ, ψ), where $\psi \in \mathsf{sub}(\varphi)$ and σ is a finite (and possibly empty) sequence of formulas from $\mathsf{sub}(\varphi)$. The subset $R(\varphi) \subseteq \Sigma(\varphi)$ of *relevant pairs for* φ is defined inductively as in Figure 3.

Intuitively, $R(\varphi)$ gives us a representation of the subordinate translations that occur when inductively translating the PAL formula φ to the EL formula φ^ε. The central observation is that, while there are exponentially many *calls* to subordinate translations ψ^σ while translating φ into φ^ε, there are only polynomially many sub-translations with different arguments ψ, σ. Using structural induction on φ, it is easy to show that the number of relevant pairs is polynomial in $|\varphi|$ and that each pair in $R(\varphi)$ is of size polynomial in $|\varphi|$.

Lemma 6. *For all PAL formulas φ, we have the following:*

1. *$|R(\varphi)| \leq |\varphi|$;*
2. *for all $(\sigma, \psi) \in R(\varphi)$, the length of the sequence σ is bounded by $|\varphi|$.*

We now convert the equivalence-preserving and exponential translation from PAL to EL into a satisfiability-preserving polynomial one. Intuitively, we introduce a propositional letter p^ψ_σ for each subordinate translation ψ^σ and enforce that p^ψ_σ is true precisely where ψ^σ is true. This can be done in an incremental fashion without actually using the (exponentially long) formulas of ψ^σ.

Let φ_0 be a single-agent PAL formula whose satisfiability is to be decided. We introduce a set of propositional letters $L_{\varphi_0} := \{p^\sigma_\varphi \mid (\sigma, \varphi) \in R(\varphi_0)\}$. W.l.o.g., assume that no letter from L_{φ_0} occurs in φ_0. For every $(\sigma, \varphi) \in R(\varphi_0)$, we define a biimplication B^σ_φ as in Figure 4. Note that the right-hand side of the biimplications is derived in a straightforward way from the equivalence-preserving translation given in Figure 2. Now define

$$\varphi_0^* := p^\varepsilon_{\varphi_0} \wedge \bigwedge_{(\sigma,\varphi) \in R(\varphi_0)} KB^\sigma_\varphi.$$

Observe that $|\varphi_0^*|$ is polynomial in φ_0: by Point 2 of Lemma 6, $|B^\sigma_\varphi|$ is linear in $|\varphi_0|$ for each $(\sigma, \varphi) \in R(\varphi_0)$. By Point 1 of Lemma 6, $|\varphi_0^*|$ is thus quadratic in $|\varphi|$. Clearly, φ_0^* is an EL formula. As the next lemma shows, we have obtained a satisfiability-preserving reduction as desired. A proof can be found in Appendix A.

Lemma 7. *φ_0 is satisfiable iff φ_0^* is satisfiable.*

Taking together Lemma 7, the fact that $|\varphi_0^*|$ is quadratic in $|\varphi_0|$, and the known NP upper bound of single-agent EL (i.e., modal $\mathsf{S5}$ [9]), we obtain an NP upper bound for single-agent PAL.

Theorem 8. *Satisfiability in single-agent PAL is* NP*-complete.*

4.2 Reducing Multi-Agent PAL

The general idea for obtaining a PSpace-upper bound for multi-agent PAL is to proceed analogously to the single-agent case. However, there is a complication: in the second conjunct of the formula φ_0^* of Lemma 7, we use K as a *master modality* that allows us to access all states that are (directly or indirectly) reachable from some given state. Alas, a master modality is not available in multi-agent PAL. For this reason, we reduce multi-agent PAL to the extension of EL with the *everybody knows operator*. The addition of this operator provides us with a restricted version of the master modality that is sufficient for our purposes.[2] Since adding

9[2]Adding *everybody knows* does not actually increase the exthe everybody knows operator to EL does not increase the computational complexity, we obtain the desired PSPACE upper bound.

Epistemic logic is extended to *epistemic logic with everybody knows (ELE)* by adding the *everybody knows operator* $E_A\varphi$, where A is a finite set of agents. The semantics of the new operator is as follows:

$$\mathcal{M}, s \models E_A\varphi \text{ iff } s \sim_a t \text{ implies } \mathcal{M}, t \models \varphi, \text{ for all } t \in S \text{ and all } a \in A.$$

PSPACE-completeness of satisfiability in multi-agent ELE is folklore. For the sake of completeness, we prove it explicitly in Appendix B.

We now reduce satisfiability in multi-agent PAL to satisfiability in ELE. Let φ_0 be the PAL formula whose satisfiability is to be decided, and let A be the set of agents used in φ_0. As in the single agent case, we introduce a set of propositional letters $L_{\varphi_0} := \{p^\sigma_\varphi \mid (\sigma, \varphi) \in R(\varphi_0)\}$ that are assumed to be disjoint from the propositional letters used in φ_0. The *modal depth* $\mathsf{md}(\varphi)$ of a PAL formula φ is defined inductively in the usual way:

$$\begin{aligned}
\mathsf{md}(p) &:= 0\\
\mathsf{md}(\neg\varphi) &:= \mathsf{md}(\varphi)\\
\mathsf{md}(\varphi \wedge \psi) &:= \mathsf{md}([\varphi]\psi) := \max(\mathsf{md}(\varphi), \mathsf{md}(\psi))\\
\mathsf{md}(K_a\varphi) &:= \mathsf{md}(\varphi) + 1.
\end{aligned}$$

Define an ELE formula

$$\varphi_0^* := p_{\varphi_0} \wedge \bigwedge_{j \le \mathsf{md}(\varphi_0)} \bigwedge_{(\sigma,\varphi)\in R(\varphi_0)} E^j_A B^\sigma_\varphi,$$

where B^σ_φ is the biimplication as defined in Figure 4, $E^j_A\varphi$ is an abbreviation for the j-fold nesting $E_A \cdots E_A\varphi$ if $j > 0$, and $E^0_A\varphi$ is simply φ. The proof of the following lemma is analogous to the proof of Lemma 7. The only difference concerns the "if" direction, where we cannot assume anymore that the accessibility relations are universal relations. To compensate for this, it is not hard to argue that the second conjunct of φ_0^* ensures that B^σ_φ is satisfied at all relevant states in models of φ_0^*, for all $(\sigma, \varphi) \in R(\varphi_0)$. Details are left to the reader.

Lemma 9. *φ_0 is satisfiable iff φ_0^* is satisfiable.*

By Lemma 6 and the fact that $\mathsf{md}(\varphi_0)$ is bounded by $|\varphi_0|$, $|\varphi_0^*|$ is polynomial in $|\varphi_0|$. From Lemma 9 and PSPACE-completeness of satisfiability in ELE, we thus obtain the following result.

Theorem 10. *Satisfiability in multi-agent PAL is* PSPACE-*complete.*

5. UPPER BOUNDS FOR PAL WITH COMMON KNOWLEDGE

It is known that (multi-agent) PALC is more expressive than the standard epistemic language extended with a common knowledge operator (ELC). For example, it is shown in [2] that there is no formula of ELC that is equivalent to the PALC formula $[p]\neg C^*\neg q$. Thus, to obtain complexity results for PALC we cannot proceed analogous to the PAL case, i.e., first exhibit an equivalence-preserving translation to the logic obtained by dropping public announcements, and then use this translation to devise a decision procedure. The solution is to establish an equivalence-preserving translation from PALC to a more expressive language than ELC: propositional dynamic logic (PDL). Since satisfiability in both ELC (which is a fragment of PALC) and PDL is EXPTIME-complete, we can then continue as in the previous section.

$$\begin{aligned}
p^\sigma &:= p\\
(\neg\varphi)^\sigma &:= \neg\varphi^\sigma\\
(\varphi \wedge \psi)^\sigma &:= \varphi^\sigma \wedge \psi^\sigma\\
(K_i\varphi)^\sigma &:= K_i(\bigwedge_{\nu\in\mathsf{pre}(\sigma)} (\nu/\sigma)^\nu \rightarrow \varphi^\sigma)\\
(C_A\varphi)^\sigma &:= [[(\bigcup_{a\in A} a; (\bigwedge_{\nu\in\mathsf{pre}(\sigma)} (\nu/\sigma)^\nu)?)^*]]\varphi^\sigma\\
([\varphi]\psi)^\sigma &:= \varphi^\sigma \rightarrow \psi^{\sigma\cdot\varphi}
\end{aligned}$$

Figure 5: The formulas φ^σ for PALC.

Recall that PDL formulas and programs are built according to the following syntax rules:[3]

$$\begin{aligned}
\varphi &::= p \mid \varphi \vee \varphi \mid \neg\varphi \mid [[\alpha]]\varphi\\
\alpha &::= a \mid \alpha \cup \alpha \mid \alpha;\alpha \mid \alpha^* \mid \varphi?
\end{aligned}$$

where p ranges over propositional letters and a over agents (usually called *atomic programs* in PDL). We define the semantics of PDL based on epistemic models by simultaneously definining the consequence relation together with accessibility relations $\sim_\alpha$ for complex programs α. Let $\mathcal{M} = (S, \sim, V)$ be an epistemic model. Then, for all $s, t \in S$, we have:

$s \sim_{\alpha\cup\beta} t$	iff	$s \sim_\alpha t$ or $s \sim_\beta t$
$s \sim_{\alpha;\beta} t$	iff	$s \sim_\alpha u$ and $u \sim_\beta t$ for some $u \in S$
$s \sim_{\alpha^*} t$	iff	$\exists u_0, \ldots, u_n \in S$, $n \ge 0$, such that $s = u_0$, $t = u_n$, and $u_i \sim_\alpha u_{i+1}$ for $i < n$
$s \sim_{\varphi?} t$	iff	$s = t$ and $\mathcal{M}, s \models \varphi$
$\mathcal{M}, s \models [[\alpha]]\varphi$	iff	$s \sim_\alpha t$ implies $\mathcal{M}, t \models \varphi$, for all $t \in S$

where the clauses for the Booleans are as in Section 2. To distinguish PDL on unrestricted models from PDL on epistemic models, we will from now on call the latter ePDL.

Figure 5 defines, for each PALC formula φ and finite sequence of PALC formulas σ, an ePDL formula φ^σ. Observe that the only difference to Figure 2 is the additional line dealing with the common knowledge operator. As in the case of PAL, we use φ^ε as the ePDL-translation of the PALC formula φ. The following lemma shows that φ^ε is indeed equivalent to φ. The proof is a straightforward extension of the proof of Lemma 1. Details are given in Appendix A.

Lemma 11. *For all models $\mathcal{M} = (S, \sim, V)$, PALC formulas φ, finite sequences of PALC formulas σ, and states $s \in S(\mathcal{M}|\sigma)$, we have $\mathcal{M}, s \models \varphi^\sigma$ iff $\mathcal{M}|\sigma, s \models \varphi$.*

pressive power. However, it allows us to formulate the restricted master modality exponentially more succinct than in standard EL, which is crucial for obtaining a polynomial translation.

9[3]We use the notation $[[\alpha]]\varphi$ instead of the more familiar $[\alpha]\varphi$ to distinguish this operator from the public announcement operator.

Our aim is to show ExpTime-completeness of PALC. As already the fragment ELC of PALC is ExpTime-hard [10], it remains to establish an upper bound. In the following, we prove this bound by a satisfiability-preserving and polynomial reduction to ePDL. Let us first fix the complexity of this logic.

Lemma 12. *Satisfiability in ePDL is* ExpTime*-complete.*

Proof. The lower bound stems from ELC [10]. The upper bound is easily obtained by a reduction to converse-PDL, which is ExpTime-complete [10, 17, 12]: to decide whether a PDL formula φ is satisfiable in an epistemic model, simply replace all atomic programs a in φ with $(a \cup a^-)^*$, and check whether the resulting IPDL formula is satisfiable in an unrestricted model. $\square$

We may now reduce satisfiability in PALC to satisfiability in ePDL using the same approach as in the previous section. Let φ_0 be the PALC formula whose satisfiability is to be decided. The definition of the set of relevant pairs $R(\varphi_0)$ can be extended to PALC formulas by adding the clause

$$R(C_A\varphi) \quad := \quad R(\varphi) \cup \{(\varepsilon, C_A\varphi)\}.$$

As usual, we then introduce a set of propositional letters $L_{\varphi_0} := \{p^\sigma_\varphi \mid (\sigma, \varphi) \in R(\varphi_0)\}$ that are disjoint from the propositional letters used in φ_0. Let $a_1, \ldots, a_k$ be the agents referred to in φ_0. We define an ePDL-formula

$$\varphi_0^* := p_{\varphi_0} \wedge \bigwedge_{(\sigma,\varphi)\in R(\varphi_0)} [[(a_1 \cup \cdots \cup a_k)^*]]B^\sigma_\varphi,$$

where the biimplications B^σ_φ are defined as in Figure 4, with the following additional clause for common knowledge:

$$B^\sigma_{C_A\varphi} \quad := \quad p^\sigma_{C_A\varphi} \leftrightarrow [[(\bigcup_{a\in A} a; (\bigwedge_{\nu\in\mathsf{pre}(\sigma)} p^\nu_{\nu/\sigma})?)^*]]p^\sigma_\psi.$$

The proof of the following lemma is analogous to the proof of Lemmas 7 and 9.

Lemma 13. φ_0 *is satisfiable iff* φ_0^* *is satisfiable.*

Since Lemma 6 can easily be extended to the PALC case, $|\varphi_0^*|$ is polynomial in $|\varphi_0|$. From Lemmas 12 and 13, we thus obtain the following result.

Theorem 14. *Satisfiability in PALC is* ExpTime*-complete.*

In [4], van Benthem et al. introduce a generalization of the common knowledge operator $C_A(\varphi, \psi)$ that is called the *relativized common knowledge operator* and has the following semantics:

$$\mathcal{M}, w \models C_A(\varphi, \psi) \quad \text{iff} \quad \mathcal{M}, w \models \varphi \to [[(\bigcup_{a\in B} a; \varphi?)^*]]\psi,$$

where the formula on the right-hand side is to be read as a PDL formula. Clearly, $C_A\varphi$ is equivalent to $C_A(\top, \varphi)$. Intuitively, relativized common knowledge resembles the until operator from temporal logic.

Let PAL-RC denote the variant of PALC in which common knowledge is replaced with relativized common knowledge, and let EL-RC be the extension of EL with the relativized common knowledge operator. The introduction of the new operator is a reaction to the fact that there exists no equivalence-preserving translation from PALC to ELC, i.e, to the logic obtained by dropping the common knowledge operator. By moving from common knowledge to the stronger relativized common knowledge, such a translation is recovered: as shown in [4], there exists an equivalence-preserving translation from PAL-RC to EL-RC.

We can easily modify the translation given in Figure 5 such that it maps formulas of PAL-RC to formulas of ePDL:

$$C_A(\varphi, \psi)^\sigma \;:=\; \varphi^\sigma \to [[(\bigcup_{a\in A} a; (\bigwedge_{\nu\in\mathsf{pre}(\sigma)} (\nu/\sigma)^\nu)?; \varphi^\sigma?)^*]]\psi^\sigma$$

Then, we can continue along the lines of the reduction from PALC to ePDL to obtain a reduction from PAL-RC to ePDL. As in the PALC case, we obtain the following theorem.

Theorem 15. *Satisfiability in PAL-RC is* ExpTime*-complete.*

6. CONCLUSION

We have analyzed the succinctness and computational complexity of (several variations of) epistemic logic extended with a public announcement operator. The main results are that, first, there are certain properties that can be expressed exponentially more succinct in PAL than in EL, and, second, despite this succinctness the computational complexity of PAL and EL coincides. As future work, it would be nice to prove the exponential succinctness of PAL as compared with EL also on epistemic structures. It is not clear whether a relatively simple argument such as the one given in Section 3 can be used in this case. Moreover, it would be interesting to analyze the computational complexity of other announcement operators, in particular of private announcements as considered, e.g., in [2, 11].

Acknowledgements

I'd like to thank Johan van Benthem, whose talk in Dresden started this work. I am also grateful to the following people for explanations and clarifying discussions: Johan van Benthem, Jelle Gerbrandy, Mitch Harris, Wiebe van der Hoek, and Mehrnoosh Sadrzadeh.

7. REFERENCES

[1] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change. *Journal of Symbolic Logic*, 50:510–530, 1985.

[2] A. Baltag, L. S. Moss, and S. Solecki. The logic of public announcements and common knowledge for distributed applications (extended abstract). In I. Gilboa, editor, *Theoretical Aspects of Reasoning about Knowledge: Proceedings of the Seventh Conference (TARK 1998)*, pages 43–56. Morgan Kaufmann, 1998.

[3] J. van Benthem. Logics for information update. In J. van Benthem, editor, *Proceedings of TARK-VIII*, pages 51–88. Morgan Kaufmann, 2001.

[4] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. Unpublished Manuscript.

[5] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001.

[6] H. van Ditmarsch, W. van der Hoek, and B. Kooi. Dynamic epistemic logic with assignment. In F. Dignum, V. Dignum, S. Koenig, S. Kraus,

M. Singh, and M. Wooldridge, editors, *Proceedings of AAMAS 2005 (Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems)*, pages 141–148. ACM Inc., 2005.

[7] H. van Ditmarsch, W. van der Hoek, and B. Kooi. Public announcements and belief expansion. In R. Schmidt, I. Pratt-Hartmann, M. Reynolds, and H. Wansing, editors, *Advances in Modal Logic, Volume 5*, pages 335–346. King's College Publications, 2005.

[8] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Kluwer academic publishers. To appear.

[9] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.

[10] M. J. Fischer and R. E. Ladner. Propositional modal logic of programs. In *Conference record of the ninth annual ACM Symposium on Theory of Computing*, pages 286–294. ACM Press, 1977.

[11] J. Gerbrandy. *Bisimulations on Planet Kripke*. ILLC Dissertation Series, Amsterdam, 1999.

[12] D. Giacomo and Massacci. Combining deduction and model checking into tableaux and algorithms for converse-PDL. *Information and Computation*, 162, 2000.

[13] J. Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54(3):319–380, 1992.

[14] J. Hintikka. *Knowledge and Belief*. Cornell University Press, Ithaca, New York, 1962.

[15] R. E. Ladner. The computational complexity of provability in systems of modal propositional logic. *SIAM Journal on Computing*, 6(3):467–480, 1977.

[16] J. A. Plaza. Logics of public communications. In M. L. Emrich, M. Pfeifer, M. Hadzikadic, and Z. M. Ras, editors, *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, pages 201–216, 1989.

[17] V. R. Pratt. A near optimal method for reasoning about action. *Journal of Computer and System Sciences*, 20:231–254, 1980.

APPENDIX

A. PROOF DETAILS

Lemma 1. For all models $\mathcal{M} = (S, \sim, V)$, PAL formulas φ, finite sequences of PAL formulas σ, and states $s \in S(\mathcal{M}|\sigma)$, we have $\mathcal{M}, s \models \varphi^\sigma$ iff $\mathcal{M}|\sigma, s \models \varphi$.

Proof. The proof is by induction on $|\varphi| + |\sigma|$. The base case is $|\varphi| + |\sigma| = 1$. Then $\varphi = p \in \mathsf{PL}$ and $\sigma = \varepsilon$. We have $\mathcal{M}|\sigma = \mathcal{M}$ and $\varphi^\sigma = \varphi$ and are done. For the induction step, let $\mathcal{M}|\sigma = (T, \approx, U)$ and make a case distinction on the form of φ:

- $\varphi = p$. Trivial by definition of p^σ.
- $\varphi = \neg\psi$ or $\varphi = (\psi \wedge \vartheta)$. Easy using the definition of $(\neg\psi)^\sigma$, the semantics, and the induction hypothesis.
- $\varphi = K_a\psi$. Let $\mathcal{M}, s \models (K_a\psi)^\sigma$. We have to show that $\mathcal{M}|\sigma, s \models K_a\psi$, i.e. that $\mathcal{M}|\sigma, t \models \psi$ for all $t \in T$ with $s \approx_a t$. Hence, let $t \in T$ with $s \approx_a t$. Then we have $\mathcal{M}, t \models (\nu/\sigma)^\nu$ for all $\nu \in \mathsf{pre}(\sigma)$: assume to the contrary that $\mathcal{M}, t \not\models (\nu/\sigma)^\nu$ for some $\nu \in \mathsf{pre}(\sigma)$. By induction hypothesis, we get $\mathcal{M}|\nu, t \not\models (\nu/\sigma)$. Thus the state t is not present in $\mathcal{M}|\nu\cdot\nu/\sigma$ and consequently $t \notin T$, which is a contradiction. Thus $\mathcal{M}, t \models (\nu/\sigma)^\nu$ for all $\nu \in \mathsf{pre}(\sigma)$, implying $\mathcal{M}, t \models \bigwedge_{\nu\in\mathsf{pre}(\sigma)}(\nu/\sigma)^\nu$. Since $s \approx_a t$, we additionally have $s \sim_a t$. Together with $\mathcal{M}, s \models (K_a\psi)^\sigma$, we get $\mathcal{M}, t \models \psi^\sigma$. By induction hypothesis, we obtain $\mathcal{M}|\sigma, t \models \psi$ as required.

  Now let $\mathcal{M}|\sigma, s \models K_a\psi$. We have to show that $\mathcal{M}, s \models (K_a\psi)^\sigma$, i.e. that $\mathcal{M}, t \models \psi^\sigma$ for all $t \in S$ with (i) $s \sim_a t$ and (ii) $\mathcal{M}, t \models \bigwedge_{\nu\in\mathsf{pre}(\sigma)}(\nu/\sigma)^\nu$. Hence, let $t \in S$ such that (i) and (ii) are satisfied. We show by induction on the length of ν that t is a state in $\mathcal{M}|\nu$ for all $\nu \in \mathsf{pre}(\sigma) \cup \{\sigma\}$.

  - $\nu = \varepsilon$. Trivial since $\mathcal{M}|\varepsilon = \mathcal{M}$.
  - $\nu = \nu' \cdot \vartheta$. By induction hypothesis, t is a state in $\mathcal{M}|\nu'$. By (ii), we have $\mathcal{M}, t \models (\nu'/\sigma)^{\nu'}$. By (outer) induction hypothesis, this yields $\mathcal{M}|\nu', t \models (\nu'/\sigma)$. Thus, t is a state in $(\mathcal{M}|\nu')|(\nu'/\sigma) = \mathcal{M}|\nu$.

  Thus, $s, t \in T$. Hence, (i) yields $s \approx_a t$. Together with $\mathcal{M}|\sigma, s \models K_a\psi$, we get $\mathcal{M}|\sigma, t \models \psi$. By induction hypothesis, we get $\mathcal{M}, t \models \psi^\sigma$ as required.
- $\varphi = [\psi]\vartheta$. Then $\mathcal{M}, s \models ([\psi]\vartheta)^\sigma$ iff $\mathcal{M}, s \not\models \psi^\sigma$ or $\mathcal{M}, s \models \vartheta^{\sigma\cdot\psi}$ iff $\mathcal{M}|\sigma, s \not\models \psi$ or $\mathcal{M}|\sigma\cdot\psi, s \models \vartheta$ iff $\mathcal{M}|\sigma, s \not\models \psi$ or $(\mathcal{M}|\sigma)|\psi, s \models \vartheta$ iff $\mathcal{M}|\sigma, s \models [\psi]\vartheta$.

  The first "iff" holds by definition of $([\psi]\vartheta)^\sigma$, the second by induction hypothesis, the third since $\mathcal{M}|\sigma\cdot\psi = (\mathcal{M}|\sigma)|\psi$, and the fourth by the semantics.

$\square$

Lemma 3. For all $s \in S'$ and $\varphi \in \mathsf{sub}(\chi)$ such that

$$\{sw \mid w \in P_\varphi\} \subseteq P_\chi,$$

we have $\mathcal{M}, s \models \varphi$ iff $\mathcal{M}', s \models \varphi$.

Proof. The proof is by induction on the structure of φ. As the induction start and the Boolean cases are trivial, we only treat the case $\varphi = K_\sigma\psi$.

"$\Leftarrow$". Let $\mathcal{M}', s \models \varphi$. We have to show that, for all $t \in S$, $s \sim_\sigma t$ implies $\mathcal{M}, t \models \varphi$. Hence, let $t \in S$ with $s \sim_\sigma t$. Then $t = s\sigma$. We first show that t does no have prefix $\widehat{w}$. Assume that the contrary holds. Since $s \in S'$ and thus s does not have prefix $\widehat{w}$, we have $t = \widehat{w}$. Since $\sigma \in P_\varphi$ and $\{sw \mid w \in P_\varphi\} \subseteq P_\chi$, the latter yields $\widehat{w} \in P_\chi$, in contradiction to $\widehat{w} \notin P_\chi$. Thus, t does not have prefix $\widehat{w}$, i.e., $t \in S'$. We have $P_\psi = \{w \mid \sigma w \in P_\varphi\}$. Thus, $t = s\sigma$ and $\{sw \mid w \in P_\varphi\} \subseteq P_\chi$ yield $\{tw \mid w \in P_\psi\} \subseteq P_\chi$. Finally, IH, $\mathcal{M}', s \models \varphi$, and $s \sim'_\sigma t$ yield $\mathcal{M}, t \models \psi$ as required.

"$\Rightarrow$". Let $\mathcal{M}, s \models \varphi$. We have to show that, for all $t \in S'$, $s \sim'_\sigma t$ implies $\mathcal{M}', t \models \varphi$. Hence, let $t \in S'$ with $s \sim'_\sigma t$. Then $t = s\sigma$. We have $P_\psi = \{w \mid \sigma w \in P_\varphi\}$. Thus, $t = s\sigma$ and $\{sw \mid w \in P_\varphi\} \subseteq P_\chi$ yield $\{tw \mid w \in P_\psi\} \subseteq P_\chi$. Finally, IH, $\mathcal{M}, s \models \varphi$, and $s \sim_\sigma t$ yield $\mathcal{M}', t \models \psi$ as required. $\square$

Lemma 7. The single-agent PAL formula φ_0 is satisfiable iff the single-agent EL-formula φ_0^* is satisfiable.

Proof. "if". Let $\mathcal{M} = (S, \sim, V)$ be a model of φ_0^*, and let $s_0 \in S$ with $\mathcal{M}, s_0 \models \varphi_0^*$. By standard results on the unimodal logic $\mathsf{S5}$, we may w.l.o.g. assume that $\sim = S \times S$ [5]. Thus, the second conjunct of φ_0^* implies that $K, s \models B_\varphi^\sigma$ for all $s \in S$ and all $(\varphi, \sigma) \in R(\varphi_0)$.

We show that, for all $s \in S$ and all $(\sigma, \varphi) \in R(\varphi_0)$, we have

$$\mathcal{M}, s \models \varphi^\sigma \text{ iff } \mathcal{M}, s \models p^\sigma_\varphi.$$

By Lemma 1, from this we get $\mathcal{M}|\sigma, s \models \varphi$ iff $\mathcal{M}, s \models p^\sigma_\varphi$. Since $\mathcal{M}, s_0 \models p^\varepsilon_{\varphi_0}$, this yields $\mathcal{M}, s_0 \models \varphi_0$ as required. The proof is by induction on $|\varphi| + |\sigma|$. For the induction start, we have $\varphi = q$ and $\sigma = \varepsilon$. Then $\varphi^\sigma = q$. Since $\mathcal{M}, s \models B^\sigma_\varphi = p^\sigma_q \leftrightarrow q$, we are done. For the induction step, we make a case distinction according to the structure of φ:

- $\varphi = q$. Identical to the induction start.
- $\varphi = \neg\psi$. Then $\mathcal{M}, s \models (\neg\psi)^\sigma$ iff $\mathcal{M}, s \models \neg\psi^\sigma$ iff $\mathcal{M}, s \not\models \psi^\sigma$ iff $\mathcal{M}, s \not\models p^\sigma_\psi$ iff $\mathcal{M}, s \models \neg p^\sigma_\psi$ iff $\mathcal{M}, s \models p^\sigma_{\neg\psi}$.
  The first "iff" holds by definition of $(\neg\psi)^\sigma$, the second by the semantics, the third by induction hypothesis, the fourth by the semantics, and the last since $\mathcal{M}, s \models B^\sigma_{\neg\psi} = p^\sigma_{\neg\psi} \leftrightarrow \neg p^\sigma_\psi$.
- $\varphi = \psi \wedge \vartheta$. Similar to the previous case.
- $\varphi = K\psi$. We have $\mathcal{M}, s \models (K\psi)^\sigma$ iff

  $$\mathcal{M}, s \models K(\bigwedge_{\nu \in \mathsf{pre}(\sigma)} (\nu/\sigma)^\nu \to \varphi^\sigma),$$

  which is the case iff, for all $t \in S$,

  $$s \sim t \text{ implies } \mathcal{M}, t \models (\bigwedge_{\nu \in \mathsf{pre}(\sigma)} (\nu/\sigma)^\nu \to \varphi^\sigma). \qquad (\dagger)$$

  By induction hypothesis, we get that, for all $t \in S$, (i) $\mathcal{M}, t \models \nu/\sigma^\nu$ iff $\mathcal{M}, t \models p^\nu_{\nu/\sigma}$ for all $\nu \in \mathsf{pre}(\sigma)$ and (ii) $\mathcal{M}, t \models \varphi^\sigma$ iff $\mathcal{M}, t \models p^\sigma_\varphi$. Thus, $(\dagger)$ holds iff, for all $t \in S$,

  $$s \sim t \text{ implies } \mathcal{M}, t \models (\bigwedge_{\nu \in \mathsf{pre}(\sigma)} p^\nu_{\nu/\sigma} \to p^\sigma_\varphi). \qquad (\ddagger)$$

  Since $\mathcal{M}, s \models p^\sigma_{K\varphi} \leftrightarrow K(\bigwedge_{\nu \in \mathsf{pref}(\sigma)} p^\nu_{\nu/\sigma} \to p^\sigma_\varphi)$, $(\ddagger)$ holds iff $\mathcal{M}, s \models p^\sigma_{K\varphi}$ and we are done.
- $\varphi = [\psi]\vartheta$. We have $\mathcal{M}, s \models ([\varphi]\psi)^\sigma$ iff $\mathcal{M}, s \models \varphi^\sigma \to \psi^{\sigma\cdot\varphi}$ iff $\mathcal{M}, s \models p^\sigma_\varphi \to p^{\sigma\cdot\varphi}_\psi$ iff $\mathcal{M}, s \models p^\sigma_{[\varphi]\psi}$.
  The first "iff" holds by definition of $([\varphi]\psi)^\sigma$, the second by induction hypothesis, and the third since $\mathcal{M}, s \models p^\sigma_{[\varphi]\psi} \leftrightarrow (p^\sigma_\varphi \to p^{\sigma\cdot\varphi}_\psi)$.

"only if". Let $\mathcal{M} = (S, \sim, V)$ be a model of φ_0, and let $s_0 \in S$ with $\mathcal{M}, s_0 \models \varphi_0$. Define a model $\mathcal{M}'$ as $\mathcal{M}$ by additionally setting, for $(\sigma, \varphi) \in R(\varphi_0)$,

$$V(p^\sigma_\varphi) := \{s \in S \mid \mathcal{M}, s \models \varphi^\sigma\}.$$

Using the definition of the translation φ^σ and the implications B^σ_φ, it is straightforward to show by induction on $|\varphi| + |\sigma|$ that

$$\mathcal{M}', s \models B^\sigma_\varphi \text{ for all } (\sigma, \varphi) \in R(\varphi_0) \text{ and } s \in S.$$

Since $\mathcal{M}, s_0 \models \varphi_0$, we have $\mathcal{M}', s_0 \models p^\varepsilon_{\varphi_0}$. Thus, $\mathcal{M}', s_0 \models \varphi_0^*$. □

Lemma 11. For all models $\mathcal{M} = (S, \sim, V)$, PALC formulas φ, finite sequences of PALC formulas σ, and states $s \in S(\mathcal{M}|\sigma)$, we have $\mathcal{M}, s \models \varphi^\sigma$ iff $\mathcal{M}|\sigma, s \models \varphi$.

Proof. The proof is analogous to that of Lemma 1. We only treat the additional case for common knowledge:

- $\varphi = C_A\psi$. Let $\mathcal{M}, s \models (C_A\psi)^\sigma$. We have to show that $\mathcal{M}|\sigma, s \models C_A\psi$, i.e. that $\mathcal{M}|\sigma, t \models \psi$ for all $t \in T$ with $s\,(\bigcup_{a \in A} \approx_a)^*\, t$. Hence, let $t \in T$ with $s\,(\bigcup_{a \in A} \approx_a)^*\, t$. Then there are $s_1, \ldots, s_k \in T$ with $s_1 = s$, $s_k = t$, and $(s_i, s_{i+1}) \in \bigcup_{a \in A} \approx_a$ for $1 \le i < k$. We have $\mathcal{M}, s_i \models (\nu/\sigma)^\nu$ for all $\nu \in \mathsf{pre}(\sigma)$ and $1 < i \le k$: assume to the contrary of what is to be shown that $\mathcal{M}, s_i \not\models (\nu/\sigma)^\nu$ for some $\nu \in \mathsf{pre}(\sigma)$ and some i with $1 < i \le k$. By induction hypothesis, we get $\mathcal{M}|\nu, s_i \not\models (\nu/\sigma)$. Thus the state s_i is not present in $\mathcal{M}|\nu{\cdot}\nu/\sigma$ and consequently $s_i \notin T$, which is a contradiction. Thus $\mathcal{M}, s_i \models (\nu/\sigma)^\nu$ for all $\nu \in \mathsf{pre}(\sigma)$ and $1 < i \le k$. Moreover, we clearly have $(s_i, s_{i+1}) \in \bigcup_{a \in A} \sim_a$ for $1 \le i < k$. It follows that $(s, t) \in \sim_{\alpha^*}$, where

  $$\alpha := \bigcup_{a \in A} a; (\bigwedge_{\nu \in \mathsf{pre}(\sigma)} (\nu/\sigma)^\nu)?.$$

  Together with $\mathcal{M}, s \models (C_A\psi)^\sigma$, we get $\mathcal{M}, t \models \psi^\sigma$. By induction hypothesis, we obtain $\mathcal{M}|\sigma, t \models \psi$ as required.

  Now let $\mathcal{M}|\sigma, s \models C_A\psi$. We have to show that $\mathcal{M}, s \models (C_A\psi)^\sigma$, i.e. that $\mathcal{M}, t \models \psi^\sigma$ for all $t \in S$ with $(s, t) \in \sim_{\alpha^*}$. Hence, let $t \in S$ with $(s, t) \in \sim_{\alpha^*}$. Then there are $s_1, \ldots, s_k \in T$ with $s_1 = s$, $s_k = t$, and $(s_i, s_{i+1}) \in \sim_\alpha$ for $1 \le i < k$. We show by induction on the length of ν that s_i is a state in $\mathcal{M}|\nu$ for all $\nu \in \mathsf{pre}(\sigma) \cup \{\sigma\}$ and all i with $1 < i \le k$.

  - $\nu = \varepsilon$. Trivial since $\mathcal{M}|\varepsilon = \mathcal{M}$.
  - $\nu = \nu' \cdot \vartheta$. By induction hypothesis, s_i is a state in $\mathcal{M}|\nu'$. Since $(s_{i-1}, s_i) \in \sim_\alpha$, we have $\mathcal{M}, s_i \models (\nu'/\sigma)^{\nu'}$. By (outer) induction hypothesis, this yields $\mathcal{M}|\nu', s_i \models (\nu'/\sigma)$. Thus, s_i is a state in $(\mathcal{M}|\nu')|(\nu'/\sigma) = \mathcal{M}|\nu$.

  Thus, we have $s_1, \ldots, s_k \in T$. As $(s_i, s_{i+1}) \in \sim_\alpha$ implies $(s_0, s_{i+1}) \in \bigcup_{a \in A} \sim_a$, this yields

  $$(s, t) \in (\bigcup_{a \in A} \approx_a)^*.$$

  Together with $\mathcal{M}|\sigma, s \models C_A\psi$, we get $\mathcal{M}|\sigma, t \models \psi$. By induction hypothesis, we get $\mathcal{M}, t \models \psi^\sigma$ as required.

□

B. EPISTEMIC LOGIC WITH "EVERYBODY KNOWS" IS IN PSPACE

To obtain a PSPACE algorithm for satisfiability in ELE, which is EL extended with the "everybody knows" operator, we devise a variation of the K-worlds style algorithm as first described in [15].

Definition 16 (Type). Let Γ be a set of ELE-formulas. We use $\mathsf{cl}(\Gamma)$ to denote the smallest set of ELE-formulas that satisfies the following properties:

- $\Gamma \subseteq \mathsf{cl}(\Gamma)$;
- $\mathsf{cl}(\Gamma)$ is closed under taking subformulas and single negations;
- if $E_{a_1,\ldots,a_k}\psi \in \mathsf{cl}(\Gamma)$, then $K_{a_1}\psi, \ldots, K_{a_k}\psi \in \mathsf{cl}(\Gamma)$.

A *type for* Γ is a subset $t \subseteq \mathsf{cl}(\Gamma)$ satisfying the following properties:

1. $\neg\psi \in t$ iff $\psi \notin t$, for all $\neg\psi \in \mathsf{cl}(\Gamma)$;

```
define procedure ELE-World(Δ, Γ, â)
  if Δ is not a type for Γ then
    return false
  for all ¬K_a φ ∈ Δ with a ≠ â do
    set Ψ := {¬φ} ∪ {ψ, K_a ψ | K_a ψ ∈ Δ}
               ∪ {¬K_a ψ | ¬K_a ψ ∈ Δ}
    non-deterministically choose a subset Δ' ⊆ cl(Ψ)
    if Ψ ⊈ Δ' or ELE-World(Δ', Ψ, a) = false then
      return false
  return true
```

Figure 6: The Procedure ELE-World.

2. $\psi \wedge \vartheta \in t$ iff $\psi, \vartheta \in t$, for all $\psi \wedge \vartheta \in \mathsf{cl}(\Gamma)$;
3. $E_{a_1,\ldots,a_k}\psi \in t$ iff $K_{a_1}\psi, \ldots, K_{a_k}\psi \in t$, for all $E_{a_1,\ldots,a_k}\psi \in \mathsf{cl}(\Gamma)$;
4. $K_a\psi \in t$ implies $\psi \in t$, for all $K_a\psi \in \mathsf{cl}(\Gamma)$.

A type t is called *realizable* if there exists a model $\mathcal{M}$ and a state s of $\mathcal{M}$ such that $\mathcal{M}, s \models \varphi$ for all $\varphi \in t$. ◇

The algorithm for deciding satisfiability in ELE is based in the procedure ELE-World given in Figure 6. This procedure gets as arguments two sets of formulas Δ and Γ, and an agent $\widehat{a}$. It checks whether Δ is a realizable type for Γ by trying to construct a tree-shaped model. Intuitively, every recursive call of the algorithm corresponds to one state of this model, and the relational structure is the reflexive-transitive closure of the recursion tree. The agent $\widehat{a}$ is passed as an argument to ensure termination: if Δ contains a formula $\neg K_a\varphi$, then we will generate a direct a-successor y of the current state x such that y satisfies φ. Since a is an equivalence relation, the type of y will contain exactly the same formulas of the form $\neg K_a\psi$ that we find in the type of x. However, there is no need to introduce successors of y as witnesses for these formulas since relations are equivalence relations and we have already generated witnesses for x. Such situations are checked by explicitly passing the agent a as an argument if the current state is an a-successor of its predecessor.

To decide the satisfiability of the input formula φ_0, we guess a subset $\Psi \subseteq \mathsf{cl}(\varphi_0)$ such that $\varphi_0 \in \Psi$ and call ELE-World$(\Psi, \{\varphi_0\}, \bot)$, where "$\bot$" is simply a dummy value. We claim that ELE-World always terminates, and that its recursion depth is bounded linearly in the length of the input formula. To prove this, we need a few notions. The *modal depth* $\mathsf{md}(\varphi)$ of an ELE-formula φ is defined inductively in the usual way:

$$\begin{aligned} \mathsf{md}(p) &:= 0 \\ \mathsf{md}(\varphi \wedge \psi) &:= \max(\mathsf{md}(\varphi), \mathsf{md}(\psi)) \\ \mathsf{md}(K_a\varphi) &:= \mathsf{md}(E_{a_1,\ldots,a_k}\varphi) := \mathsf{md}(\varphi) + 1 \end{aligned}$$

Let Δ be a set of formulas. Then we use $\mathsf{md}(\Delta)$ to denote $\max\{\mathsf{md}(\varphi) \mid \varphi \in \Delta\}$ if Δ is non-empty and 0 otherwise. For $b \in N$, we use Δ^b to denote the set $\{\varphi \in \Delta \mid \varphi$ of the form $K_b\psi$ or $\neg K_b\psi\}$.

Lemma 17. *The recursion depth of ELE-*World$(\Psi, \{\varphi_0\}, \bot)$ *is bounded by* $\mathsf{md}(\varphi_0)$.

Proof. Consider a path of length k in the recursion tree generated by the call ELE-World$(\Psi, \{\varphi_0\}, \bot)$, and let

$$(\Delta_1, \Gamma_1, a_1), \ldots, (\Delta_k, \Gamma_k, a_k)$$

be the arguments to ELE-World on this path, with

$$(\Delta_1, \Gamma_1, a_1) = (\Psi, \{\varphi_0\}, \bot).$$

It is easily seen that $\mathsf{md}(\Delta_1) \leq \mathsf{md}(\Delta_2)$. We additionally show that $\mathsf{md}(\Delta_{i+1}) < \mathsf{md}(\Delta_i)$ for $2 \leq i < k$. Together with $\mathsf{md}(\Delta_1) \leq \mathsf{md}(\varphi_0)$, it follows that $k \leq \mathsf{md}(\varphi_0) + 1$. Thus, the length of the path is bounded by $\mathsf{md}(\varphi_0)$.

We first establish the following property: by construction of the set Δ' that is used as an argument in recursive calls, it is readily checked that we have, for $2 \leq i \leq k$,

$$\mathsf{md}(\Delta_i^b) < \mathsf{md}(\Delta_i^{a_i}) \text{ for all } b \neq a_i. \qquad (*)$$

We can now show that, for $2 < i \leq k$,

$$\mathsf{md}(\Delta_i) = \mathsf{md}(\Delta_i^{a_i}) > \mathsf{md}(\Delta_i^{a_{i+1}}) = \mathsf{md}(\Delta_{i+1}^{a_{i+1}}) = \mathsf{md}(\Delta_{i+1}).$$

The first equality is implied by $(*)$. The inequality follows from $(*)$ and the fact that, by definition of ELE-World, we have $a_i \neq a_{i+1}$ for $1 \leq i < k$. The last but one equality holds since, also by definition of ELE-World, we have $\Delta_i^{a_{i+1}} = \Delta_{i+1}^{a_{i+1}}$ for $1 \leq i < k$. Finally, the last equality is again due to $(*)$. □

Concerning correctness of the algorithm, it is a matter of routine to establish the following lemma.

Lemma 18. *An ELE-formula φ_0 is satisfiable in an epistemic model iff there exists a set $\Psi \subseteq \mathsf{cl}(\varphi_0)$ such that $\varphi_0 \in \Psi$ and ELE-*World$(\Psi, \{\varphi_0\}, \bot)$ *returns* true.

Since $\mathsf{md}(\varphi_0)$ is linearly bounded by the length of φ_0 and the space consumption of ELE-World is bounded linearly by its recursion depth, Lemmas 17 and 18 together with Savitch's Theorem yield the following result.

Theorem 19. *Satisfiability of (multi-agent) ELE-formulas is in* PSpace.