The Complexity of Model Checking for Knowledge Update

Chitta Baral*
Dept. of Computer Science & Engineering
Arizona State University
Tempe, AZ 85287, USA

Yan Zhang†
School of Computing & Information Technology
University of Western Sydney
Penrith South DC, NSW 1797, Australia

Abstract

The authors have recently proposed a formal theory of knowledge update based on the semantics of modal logic S5 (Baral & Zhang 2001). In that system, an agent's knowledge set is represented as a S5 formula and update on agent's knowledge is implemented by updating the corresponding Kripke models of the agent's knowledge set. In this paper, we further investigate the computational complexity of model checking for knowledge update. We first show that in general the model checking for knowledge update is Σ_3^P-complete, which places the problem at one layer higher in the polynomial hierarchy than the traditional model based belief update (e.g. PMA). We then identify a subclass of knowledge update problems that has polynomial time complexity for model checking. We point out that some important knowledge update problems belong to this subclass. We also address other subclasses of knowledge update problems for which the complexity of model checking ranges from NP-complete to Σ_2^P-complete.

1 Introduction

Different from traditional belief revision and update, knowledge update is concerned with the issue of changing an agent's knowledge where a modal logic is usually used to specify the agent's knowledge. The authors have recently proposed a formal theory of knowledge update based on the semantics of the modal logic S5 (Baral & Zhang 2001). In that formulation, an agent's knowledge is represented as a S5 formula and updating the agent's knowledge is implemented by updating the corresponding Kripke models of the agent's knowledge (Baral & Zhang 2001). In this paper, we investigate the computational complexity of model checking for knowledge update. We first show that in general the model checking problem for knowledge update is Σ_3^P-complete, which places the problem at one layer higher in the polynomial hierarchy than the traditional model based belief update (e.g. PMA) (Liberatore & Schaerf 2001). We then identify a subclass of knowledge update problems for which model checking can be achieved in polynomial time. We show that some important knowledge update problems belong to this subclass. We also identify other subclasses of knowledge updates for which the complexity for model checking ranges from NP-complete to Σ_2^P-complete.

The rest of this paper is organized as follows. In Section 2 we introduce the necessary preliminaries that are needed in the paper. We first review the semantics of knowledge update, and then summarize the basic concepts and notations of complexity theory. In Section 3 we investigate the complexity of model checking for general knowledge updates, and in Section 4 we identify an important subclass of knowledge update problems for which model checking can be done in polynomial time. In Section 5 we examine other subclasses of knowledge update problems for which the complexity of model checking is lower than the general case, but still intractable (i.e. NP-complete or beyond). We show that several interesting knowledge update cases belong to these subclasses. Finally, in Section 6 we conclude the paper with some remarks.

2 Preliminaries

In this section, we briefly present the formal semantics of knowledge update that was developed by the authors in (Baral & Zhang 2001) and the background

*E-mail: chitta@asu.edu
†Corresponding author. E-mail: yan@cit.uws.edu.au

and terminology needed to understand our complexity results described in this paper. Throughout the paper we restrict our attention to a finite propositional language.

2.1 Knowledge Update: Semantics

Our formalization of knowledge update will be based on the semantics of the propositional modal logic S5 with a single agent. In general, under Kripke semantics, a *Kripke structure* is a triple (W, R, π), where W is a set of possible worlds, R is an equivalence relation on W, and π is a truth assignment function that assigns a propositional valuation to each world in W. Given a Kripke structure $S = (W, R, \pi)$, a *Kripke interpretation* is a pair $M = (S, w)$, where $w \in W$ is referred to as the *actual world* of M. The entailment relation $\models$ between Kripke interpretations and formulas is defined to provide semantics for S5 formulas (Fagin *et al.* 1995).

In the case of single agent, however, we may restrict ourselves to those S5 structures in which the relation R is universal, i.e. each world is accessible from every world, and worlds are identified with the set of atoms true at the worlds (Meyer & van der Hoek 1995). To simplify a comparison between two worlds (e.g. Definition 2), we may view an atom $p \in w$ iff $w \models p$. Therefore, in our context a Kripke structure (W, R, π) is uniquely characterized by W and we simplify a Kripke interpretation as a pair (W, w) which we call a *k-model*, where w indicates the actual world of the agent and W presents all possible worlds that the agent may access. Note that w is in W for any k-model (W, w).

In our following description, we use $a, b, c, p, \cdots$, $x, y, z, \cdots$ to denote primitive propositional atoms; $P, Q, R, X, X, Z, \cdots$ to denote sets of primitive propositional atoms; $\phi, \psi, \upsilon, \cdots$ to denote propositional formulas without including modalities (we also call them *objective* formulas); and $\alpha, \beta, \gamma, \mu, \cdots$ to denote formulas that may contain the modal operator K. For convenience, we also use $T = \alpha_1 \wedge \cdots \wedge \alpha_k$ to represent a finite set of formulas $\{\alpha_1, \cdots, \alpha_k\}$ and call T a *(knowledge) set*.

Definition 1 *(***Simplified S5 Semantics***) Let P be the set of all primitive propositions in the language. The entailment relation $\models$ under simplified S5 semantics is defined as follows:*

1. *$(W, w) \models p$ iff p is primitive (i.e. $p \in P$) and $w \models p$;*
2. *$(W, w) \models \alpha \wedge \beta$ iff $(W, w) \models \alpha$ and $(W, w) \models \beta$;*
3. *$(W, w) \models \neg\alpha$ iff it is not the case that $(W, w) \models \alpha$;*
4. *$(W, w) \models K\alpha$ iff $(W, w') \models \alpha$ for all $w' \in W$.*

Given a formula α, $M = (W, w)$ is called a k-model of α if $M \models \alpha$. We use $Mod(\alpha)$ to denote the set of all k-models of α. For an objective formula ϕ, $Mod(\phi)$ simply denotes the set of models of ϕ, and $w \models \phi$ denotes that w is a *model* of ϕ, i.e. w is a world in which ϕ is true.

Now the basic problem of knowledge update that we would like to investigate is described as follows: given a *k-model* $M = (W, w)$, that is viewed as a *knowledge state* of an agent (or simply called a state of the agent), and a formula μ - the agent's new knowledge that may contain modal operator K, how do we update M to another k-model $M' = (W', w')$ such that $M' \models \mu$ and M' is *minimally different* from M with respect to some criterion. To approach this problem, we first need to provide a definition of *closeness* between two k-models with respect to a given k-model.

Definition 2 *(Baral & Zhang 2001) Let $M = (W, w)$, $M_1 = (W_1, w_1)$ and $M_2 = (W_2, w_2)$ be three k-models. We say M_1 is* as close to M as *M_2, denoted as $M_1 \leq_M M_2$, if:*

1. *$(w_1 \setminus w \cup w \setminus w_1) \subseteq (w_2 \setminus w \cup w \setminus w_2)$; or*
2. *$w_1 = w_2$ and one of the following conditions holds:*
   - *(i) if $W \subseteq W_1$, then (a) there exist some ϕ and ψ such that $M \models K\phi$ and $M_2 \not\models K\phi$ and $M \not\models K\psi$ and $M_2 \models K\psi$, or (b) for any ϕ if $M \models K\phi$ and $M_1 \not\models K\phi$, then $M_2 \not\models K\phi$;*
   - *(ii) if $W_1 \subset W$, then condition (a) above is satisfied, or (c) for any ϕ if $M \not\models K\phi$ and $M_1 \models K\phi$, then $M_2 \models K\phi$;*
   - *(iii) if $W \not\subset W_1$ and $W_1 \not\subset W$, then both conditions (b) and (c) above are satisfied.*

We denote $M_1 <_M M_2$ if $M_1 \leq_M M_2$ and $M_2 \not\leq_M M_1$.

In the above definition, condition 1 simply says that the symmetric differences between w and w_1 is not bigger than that between w and w_2, while in condition 2, (i), (ii) and (iii) express that different preferences are applied to compare knowledge between M_1 and M_2 with respect to M. For convenience, given a k-model M, if we denote $KM = \{\phi \mid \text{for all } w \in W, w \models \phi\}$, then (a) is equivalent to $KM \setminus KM_2 \neq \emptyset$ and $KM_2 \setminus KM \neq \emptyset$; (b) is equivalent to $KM \setminus KM_1 \subseteq KM \setminus$

KM_2; and (c) is equivalent to $KM_1 \setminus KM \subseteq KM_2 \setminus KM$. Also (b) and (c) together present a difference on both knowledge decrease and increase between M_1 and M_2 in terms of M. It is also easy to see that $\leq_M$ is a partial ordering.

Note that during the comparison between two k-models, we give preference to the change of the actual world over the change of the knowledge about the world. (The closeness criterion between actual worlds that we use is the commonly used criterion (Winslett 1988) based on symmetric difference). For instance, if the actual world of a k-model M_1 is closer to the actual world of M than the actual world of another k-model M_2, we will think M_1 is closer to M and the comparison of knowledge between M_1 and M_2 is ignored. Only when both M_1 and M_2 have the same actual world, we will compare the knowledge of M_1 and M_2 in terms of M. This seems to be intuitive to us. In fact, the comparison between actual worlds determines the *actual distance* between two k-models to the given k-model M. If M_1 and M_2 have the same actual distance to M, the *knowledge distance* is then taken into account.

Basically, condition (i) (or (ii) resp.) in Definition 2 defines a knowledge preference based on knowledge decrease (or increase resp). That is, if M_1 *only* looses knowledge from M (or *only* gains some knowledge to M, resp.), then M_1 is preferred over those k-models that have both knowledge decrease and increase from M, i.e. (a), and also preferred over those k-models that only loose more knowledge from M (or add more knowledge to M, resp.) than M_1 does, i.e. (b) or (c) respectively. Condition (iii), on the other hand, deals with the mixed situation that M_1 has both knowledge decrease and increase from M. In this case, a combined difference on knowledge decrease and increase is applied to determine the knowledge distance, i.e. (b) *and* (c).

Definition 3 *(Baral & Zhang 2001) Let $M = (W, w)$ be a k-model and μ a formula. A k-model $M' = (W', w')$ is called a* possible resulting k-model *after updating M with μ if and only if the following conditions hold:*

1. *$M' \models \mu$;*
2. *there does not exist another k-model $M'' = (W'', w'')$ such that $M'' \models \mu$ and $M'' <_M M'$.*

We denote the set of all possible resulting k-models after updating M with μ as $Res(M, \mu)$.

Example 1 Let $T = Kc \wedge \neg Ka \wedge \neg Kb \wedge K(a \vee b)$ and $\mu = K\neg c$. We denote

$$w_0 = \{a, b, c\},\ w_1 = \{a, c\},\ w_2 = \{b, c\},$$
$$w_3 = \{c\},\ w_4 = \{a, b\},\ w_5 = \{a\},$$
$$w_6 = \{b\},\ w_7 = \emptyset.$$

Clearly, $M_0 = (\{w_0, w_1, w_2\}, w_0)$ is a k-model of T. Consider the update of M_0 with μ. Let $M_1 = (\{w_4, w_5, w_6\}, w_4)$. Now we show that M_1 is a possible resulting k-model after updating M_0 with μ.

Since $(w_0 \setminus w_4 \cup w_4 \setminus w_0) = \{c\}$, we first consider any possible k-model $M' = (W', w')$ such that $(w_0 \setminus w' \cup w' \setminus w_0) \subset \{c\}$. Clearly, the only possible w' would be w_0 itself. Let $M' = (W', w_0)$, where W' is a subset of $\{w_0, \cdots, w_7\}$. However, since $c \in w_0$, there does not exist any W' such that $M' \models K\neg c$. Therefore, from Definition 2, only condition 2 can be used to find a possible M' such that $M' <_M M_1$. So we assume $M' = (W', w_4)$. On the other hand, from M_0 and M_1, it is easy to see that $KM_0 = \{c, a \vee b\}$ and $KM_1 = \{\neg c, a \vee b\}$[1]. Then we have $KM_0 \setminus KM_1 = \{c\}$ and $KM_1 \setminus KM_0 = \{\neg c\}$. Ignoring the detailed verifications, we can show that there does not exist such $M' = (W', w_4)$ satisfying $KM_0 \setminus KM' = KM' \setminus KM_0 = \emptyset$. ∎

Based on the k-model update, updating a formula (knowledge set) T in terms of another formula μ is then achieved by updating every k-model of T with μ.

Definition 4 *(Baral & Zhang 2001) Let T and μ be two S5 formulas. The* update *of T with μ, denoted as $T \diamond \mu$, is defined by $Mod(T \diamond \mu) = \bigcup_{M \in Mod(T)} Res(M, \mu)$.*

Given a set of k-models $\mathcal{S}$ and a k-model M, let $\leq_M$ be an ordering on $\mathcal{S}$ as we defined in Definition 2. By $Min(\mathcal{S}, \leq_M)$ we mean the set of all elements in $\mathcal{S}$ that are minimal with respect to the ordering $\leq_M$. The following theorem shows that our knowledge update can be characterized by a minimal change criterion based on an ordering on k-models.

Theorem 1 *(Baral & Zhang 2001) Let T and μ be two formulas. Then*

$$Mod(T \diamond \mu) = \bigcup_{M \in Mod(T)} Min(Mod(\mu), \leq_M).$$

2.2 Computational Complexity

We now introduce some basic notions from complexity theory and refer to (Garey & Johnson 1979) for

[1] For simplicity, here we only consider the *prime* formulas ϕ in KM in the sense that if $\phi \in KM$, then there is not another ψ such that $\models \psi \supset \phi$ and $\psi \in KM$.

further details. Two important complexity classes are P and NP. The class of P includes those decision problems solvable by a polynomial-time deterministic Turing machine. The class of NP, on the other hand, consists of those decision problems solvable by a polynomial-time nondeterministic Turing machine.

Let $\mathcal{C}$ be a class of decision problems. The class $P^{\mathcal{C}}$ consists of the problems solvable by a polynomial-time deterministic Truing machine with an oracle for a problem from $\mathcal{C}$, while the class $NP^{\mathcal{C}}$ includes the problems solvable by a nondeterministic Turing machine with an oracle for a problem in $\mathcal{C}$. By co-$\mathcal{C}$ we mean the class consisting of the complements of the problems in $\mathcal{C}$.

The classes Σ_k^P and Π_k^P of the *polynomial hierarchy* are defined as follows:

$\Sigma_0^P = \Pi_0^P = P$, and
$\Sigma_k^P = NP^{\Sigma_{k-1}^P}$, Π_k^P =co-Σ_k^P for all $k > 1$.

It is easy to see that $NP = \Sigma_1^P$ and co-$NP = \Pi_1^P$. A problem A is *complete* for a class $\mathcal{C}$ if $A \in \mathcal{C}$ and for every problem B in $\mathcal{C}$ there is a polynomial transformation of B to A.

The prototypical Σ_k^P-complete and Π_k^P-complete problems are deciding the validity of quantified Boolean formulas (QBFs) of the form:

$$Q_1X_1Q_2X_2\cdots Q_kX_kE, k \geq 1, \tag{1}$$

where E is a Boolean expression using propositional atoms over alphabets $X_1, X_2, \cdots$, and X_k, and the Q_i's are alternating qualifiers from $\{\forall, \exists\}$ $(1 \leq i \leq k)$. If $Q_1 = \exists$, then deciding the validity of (1) is Σ_k^P-complete, while deciding the validity of (1) is Π_k^P-complete if $Q_1 = \forall$.

Let X and Y be two finite set of propositional atoms where X and Y have the same cardinality, i.e. $|X| = |Y|$. For convenience, we use notion $X \equiv Y$ to stand for formula $(x_1 \equiv y_1) \wedge (x_2 \equiv y_2) \wedge \cdots \wedge (x_m \equiv y_m)$. Consequently, $X \equiv \neg Y$ stands for formula $(x_1 \equiv \neg y_1) \wedge$ $(x_2 \equiv \neg y_2) \wedge \cdots \wedge (x_m \equiv \neg y_m)$. We also use $\neg X$ to denote the set $\{\neg x_i \mid x_i \in X\}$ (or formula $\bigwedge_{x_i \in X} \neg x_i$), and use notion $\bigvee \neg X$ to stand for formula $\bigvee_{x_i \in X} \neg x_i$. For a given formula α, we use $|\alpha|$ to denote the length of α.

The problem of *model checking* for knowledge update is described as follows: Given a knowledge set T, a formula μ, and a k-model M, deciding whether $M \in Mod(T \diamond \mu)$. It is well known that the model checking problem for traditional belief revision and update is located at the lower end of the polynomial hierarchy from P to Σ_2^P depending on specific revision/update operators and additional restrictions (if any) (Liberatore & Schaerf 2001).

3 The General Case

In this section, we investigate the complexity of model checking for the general case of knowledge update.

Lemma 1 *Let $M = (w, W), M_1 = (w_1, W_1)$ and $M_2 = (w_2, W_2)$ be three k-models.*

1. *Deciding whether $KM \setminus KM_2 \neq \emptyset$ and $KM_2 \setminus KM \neq \emptyset$ has time complexity $\mathcal{O}(|W| \times |W_2|)$.*
2. *Deciding whether $KM \setminus KM_1 \subseteq KM \setminus KM_2$ has time complexity $\mathcal{O}(|W_1| \times (|W| + |W_2|))$.*
3. *Deciding whether $KM_1 \setminus KM \subseteq KM_2 \setminus KM$ is in co-NP.*

Proof: Result 1 is equivalently to deciding whether $W \not\subset W_2$ and $W_2 \not\subset W$ (proper set inclusion). Obviously, this can be verified in $\mathcal{O}(|W| \times |W_2|)$ time.

Now we prove Result 2. From set inclusion and intersection properties, it is easy to see that $KM \setminus KM_1 \subseteq KM \setminus KM_2$ iff $KM_2 \cap KM \subseteq KM_1 \cap KM$. Then from Proposition 1 (Results 2 and 4) of (Baral & Zhang 2001), it follows that $KM_2 \cap KM \subseteq KM_1 \cap KM$ iff $W_1 \cup W \subseteq W_2 \cup W$. Obviously, checking whether $W_1 \cup W \subseteq W_2 \cup W$ can be done in $\mathcal{O}(|W_1| \times (|W| + |W_2|))$ time.

Finally we prove Result 3. Similarly, from set inclusion and union properties, we have that $KM_1 \setminus KM \subseteq KM_2 \setminus KM$ iff $KM_1 \cup KM \subseteq KM_2 \cup KM$ iff $KM_1 \subseteq KM_2 \cup KM$. Now we show that checking whether $KM_1 \subseteq KM_2 \cup KM$ is in co-NP. To show $KM_1 \not\subseteq KM_2 \cup KM$ we need to find some ϕ such that $\phi \in KM_1$ but $\phi \notin KM_2 \cup KM$. A guess of such a formula ϕ can be done in polynomial time. From the definition of KM_1, we know that $\phi \in KM_1$ iff for each world $w' \in W_1$, $w' \models \phi$. Clearly, this can be verified in polynomial time. Similarly, to check whether $\phi \notin KM_2 \cup KM$, we only need to find a world w'' from W_2 and a world w from W such that $w'' \not\models \phi$ and $w \not\models \phi$. Also, this can be done in polynomial time. ■

Lemma 2 *Let M, M_1 and M_2 be three k-models. Deciding whether $M_1 \leq_M M_2$ is in co-NP.*

Proof: According to Definition 2, if $w_1 \neq w_2$, then $M_1 \leq_M M_2$ iff $(w_1 \setminus w \cup w \setminus w_1) \subseteq (w_2 \setminus w \cup w \setminus w_2)$. Clearly, this can be verified in polynomial time. If $w_1 = w_2$, then we need to check the following conditions: (i) If $W \subseteq W_1$, then $M_1 \leq_M M_2$ iff condition (a) or (b) in Definition 2 is satisfied. From Lemma 1, we know that deciding whether (a) and (b) to be true can be done in polynomial time. (ii) If $W_1 \subseteq W$, then $M_1 \leq_M M_2$ iff condition (a) or (c) in Definition 2 is satisfied. From Lemma 1, deciding whether (c) to be true is in co-NP. (iii) If $W \not\subset W_1$ and $W_1 \not\subset W$, then $M_1 \leq_M M_2$ iff conditions (b) and (c) should be satisfied. Again, deciding whether condition (c) to be true is in co-NP. So, the problem is in co-NP. ■

Lemma 3 *Let M, M' be two k-models and μ a S5 formula. Deciding whether $M' \in Res(M, \mu)$ is in Π_2^P.*

Proof: According to Theorem 1, if $M' \notin Res(M, \mu)$, there must exist another k-model M'' such that $M'' <_M M'$. A guess of a k-model M'' can be done in polynomial time. From Lemma 2, deciding whether $M'' \leq_M M'$ is in co-NP. Since $M'' <_M M'$ iff $M'' \leq_M M'$ and $M' \not\leq_M M''$, this follows that checking whether $M'' <_M M'$ can be polynomially decided with two queries to a co-NP oracle. So the problem is in Π_2^P (co-Σ_2^P). ■

Theorem 2 *Model checking for knowledge update is Σ_3^P-complete.*

Proof: Membership proof. From Definition 4, $M \in Mod(T \diamond \mu)$ iff for some $M' \in Mod(T)$, $M \in Res(M', \mu)$. A guess of M' and check whether $M' \in Mod(T)$, i.e. $M' \models T$, can be achieved in polynomial time. According to Lemma 3, deciding whether $M \in Res(M', \mu)$ can be solved with one call to a co-Σ_2^P oracle. So the problem is in Σ_3^P.

Hardness proof. We prove the hardness by giving a polynomial transformation from deciding the validity of $\exists X \forall Y \exists Z E$, where E is a Boolean expression using propositional atoms over $X \cup Y \cup Z$. We construct formulas T, μ and a k-model M over propositional atoms $X \cup Y \cup Z \cup \hat{X} \cup \hat{Y} \cup \hat{Z} \cup \{a, b\}$, where $|\hat{X}| = |X|$, $|\hat{Y}| = |Y|$ and $|\hat{Z}| = |Z|$, and any two sets among X, Y, Z, $\hat{X}$, $\hat{Y}$, $\hat{Z}$, and $\{a, b\}$ are disjoint.

$$\begin{aligned}
T &= K(X \equiv \hat{X}) \wedge Y \wedge \neg\hat{Y} \wedge K(Z \equiv \hat{Z}) \wedge a \wedge b, \\
\mu &= [K\neg X \wedge K\hat{X} \wedge K\neg Y \wedge K\neg Z \wedge K\hat{Z} \wedge \\
&\quad K(a \wedge \neg b)] \vee \\
&\quad [K(X \equiv \hat{X}) \wedge K\neg\hat{Y} \wedge K(Z \equiv \hat{Z}) \wedge \\
&\quad K(a \equiv b) \wedge K(b \equiv \neg E)], \\
M &= (\{w\}, w), \text{ where } w = \hat{X} \cup \hat{Z} \cup \{a\}.
\end{aligned}$$

For simplicity, we denote $\mu = \gamma_1 \vee \gamma_2$, where

$$\begin{aligned}
\gamma_1 &= K\neg X \wedge K\hat{X} \wedge K\neg Y \wedge K\neg Z \wedge K\hat{Z} \wedge \\
&\quad K(a \wedge \neg b), \text{ and} \\
\gamma_2 &= K(X \equiv \hat{X}) \wedge K\neg\hat{Y} \wedge K(Z \equiv \hat{Z}) \wedge \\
&\quad K(a \equiv b) \wedge K(b \equiv \neg E).
\end{aligned}$$

Clearly, M is a k-model of μ since $M \models \gamma_1$. Now we show that $\exists X \forall Y \exists Z E$ is valid iff M is a k-model of $T \diamond \mu$.

($\Rightarrow$) Suppose $\exists X \forall Y \exists Z E$ is valid. Let $X_1 \subseteq X$ be one of the valuations of X that makes $\forall Y \exists Z E$ valid. That is, for any $Y_1 \subseteq Y$, there exists some $Z_1 \subseteq Z$ such that $X_1 \cup Y_1 \cup Z_1 \models E$. Let $X_2 \subseteq X$ be one of the valuations of X that makes $\forall Y \exists Z E$ not valid[2]. That is, for some $Y_2 \subseteq Y$, no matter what $Z_2 \subseteq Z$ we choose E evaluates to false. This implies that γ_2 has two kinds of k-models as showed by M_1 and M_2 below:

$M_1 = (W_1, w_1)$,
$w_1 = X_1 \cup Y_1 \cup Z_1 \cup \hat{X_1} \cup \hat{Z_1}$[3],
where each world $w_1' \in W_1$ has the same form of w_1:
$w_1' = X_1' \cup Y_1' \cup Z_1' \cup \hat{X_1'} \cup \hat{Z_1'}$, where $X_1' \cup Y_1' \cup Z_1' \models E$,

$M_2 = (W_2, w_2)$,
$w_2 = X_2 \cup Y_2 \cup Z_2 \cup \hat{X_2} \cup \hat{Z_2} \cup \{a, b\}$,
where each world $w_2' \in W_2$ has the same form of w_2:
$w_2' = X_2' \cup Y_2' \cup Z_2' \cup \hat{X_2'} \cup \hat{Z_2'} \cup \{a, b\}$, where $X_2' \cup Y_2' \cup Z_2' \not\models E$.

Note that $M_1 \models KE$, and $M_2 \models K\neg E$. We consider a k-model as follows:

$M^* = (\{w^*\}, w^*)$, where
$w^* = X_1 \cup \hat{X_1} \cup Y \cup Z_1 \cup \hat{Z_1} \cup \{a, b\}$.

It is easy to verify that M^* is a k-model of T. Now we prove $M \in Res(M^*, \mu)$. According to Theorem 1, we only need to show $M \in Min(Mod(\mu), \leq_{M^*})$. We use notion $diff(w, w')$ to denote the symmetric difference between sets w and w', i.e. $w \setminus w' \cup w' \setminus w$.

[2] Note that if there does not exist such valuation of X, then γ_2 will only has one type of k-model M_1.

[3] Here we let $\hat{X_1} = \{\hat{x_i} \mid x_i \in X_1\}$. This kind of denotation is adopted through out this paper whenever there is no explicit declaration.

$$diff(w^*, w) = X_1 \cup (\hat{X} \setminus \hat{X_1}) \cup Y \cup Z_1 \cup (\hat{Z} \setminus \hat{Z_1}) \cup \{b\},$$
$$diff(w^*, w_1) = (Y \setminus Y_1) \cup \{a, b\},$$
$$diff(w^*, w_2) = diff(X_1, X_2) \cup (Y \setminus Y_2) \cup diff(Z_1, Z_2) \cup diff(\hat{X_1}, \hat{X_2}) \cup diff(\hat{Z_1}, \hat{Z_2}).$$

Since $a \in diff(w^*, w_1)$ but $a \notin diff(w^*, w)$, $diff(w^*, w_1)$ is not a proper subset of $diff(w^*, w)$. On the other hand, as $\hat{X} \setminus \hat{X_1} \subseteq diff(w^*, w)$ and $\hat{X_1} \setminus \hat{X_2} \subseteq diff(w^*, w_2)$[4], there exists some $\hat{x_i} \in \hat{X_1}$ such that $\hat{x_i} \in diff(w^*, w_2)$ but $\hat{x_i} \notin diff(w^*, w)$. So $diff(w^*, w_2)$ is not a proper subset of $diff(w^*, w)$ either. This follows that $M = (\{w\}, w) \in Min(Mod(\mu), \leq_{M^*})$, i.e. M is a k-model of $T \diamond \mu$.

($\Leftarrow$) Now suppose $\exists X \forall Y \exists Z E$ is not valid. That is, $\forall X \exists Y \forall Z \neg E$ is valid. Consider any k-model of T:

$$M^* = (W^*, w^*), \text{ where}$$
$$w^* = X' \cup \hat{X}' \cup Y \cup Z' \cup \hat{Z}' \cup \{a, b\},$$

there always exists a k-model of μ given by

$$M' = (\{w'\}, w'), \text{ where}$$
$$w' = X' \cup \hat{X}' \cup Y'' \cup Z' \cup \hat{Z}' \cup \{a, b\}.$$

We can assume that $X' \subseteq X$ and $Y'' \subseteq Y$ is a valuation of Y that makes $\exists Z \neg E$ valid on X', i.e. $X' \cup Y'' \cup Z' \models \neg E$. It is easy to see that $M' \models \gamma_2$ and $M' \models K \neg E$. Now we show $diff(w^*, w') \subset diff(w^*, w)$. Indeed, we have

$$diff(w^*, w) = X' \cup (\hat{X} \setminus \hat{X}') \cup Y \cup Z' \cup (\hat{Z} \setminus \hat{Z}') \cup \{b\}, \text{ and}$$
$$diff(w^*, w') = (Y \setminus Y'') \subseteq Y \subset diff(w^*, w).$$

This concludes that M is not a k-model of $T \diamond \mu$. ■

4 A Tractable Subclass - Knowledge Decreased Update

In this section, we identify a subclass of knowledge update problems for which model checking can be achieved in polynomial time. We first introduce a useful notation. Let α be a S5 formula and ϕ^α be an objective formula (i.e. no K occurs in it) occurring in α. We then say ϕ^α is an *objective sub-formula* of α. We denote the set of all objective sub-formulas of α as $Sub^o(\alpha)$. For instance, given $\alpha = Ka \vee K\neg b$, $Sub^o(\alpha) = \{a, b, \neg b\}$.

[4] Without loss of generality, we assume $\hat{X_1} \setminus \hat{X_2} \neq \emptyset$.

Definition 5 *Given S5 formulas T and μ, updating T with μ is called* knowledge decreased *if for any k-model $M' = (W', w')$ of $T \diamond \mu$, there exists a k-model $M = (W, w)$ of T such that (i) $W \subseteq W'$, and $w = w'$; and (ii) for any $w^* \in W'$, $w^* \in W$ iff $w^* \models \phi^\mu$ or $w^* \models \neg\phi^\mu$ for some ϕ^μ in $Sub^o(\mu)$.*

From the above definition, it is easy to see that if an update is knowledge decreased, then the actual world of the agent's state will not change, and the agent's knowledge can only be decreased. Furthermore, the set of posible worlds in the agent's resulting state can be specifically computed from her previous state. We have the following important result on the model checking for knowledge decreased update.

Theorem 3 *Model checking for knowledge decreased update can be achieved in polynomial time.*

Proof: Given T, μ and a k-mode $M' = (W', w')$. Suppose $T \diamond \mu$ be knowledge decreased. To check whether $M' \in Mod(T \diamond \mu)$, we need to do the following things:

1. Check whether $M' \models \mu$,
2. Compute a subset W of W' such that for any $w^* \in W'$, $w^* \in W$ iff $w^* \models \phi^\mu$ or $w^* \models \neg\phi^\mu$ for some $\phi^\mu \in Sub^o(\mu)$,
3. Check whether $(W, w) \models T$.

Clearly, Steps 1 and 3 can be done in polynomial time. As $|Sub^o(\mu)| \leq |\mu|$, it follows that Step 2 can be also done in polynomial time. ■

Now it is worthwhile to examine some concrete forms of knowledge decreased update which have important applications in practical domains.

Ignorance update

We consider how an agent ignores a fact from her knowledge set. This procedure is so-called *ignorance update*, i.e. updating T with $\neg K\phi$. From Definition 1, it is easy to see that $T \diamond \neg\phi \models \neg K\phi$. However, it should be noted that updating $T \diamond \neg\phi$ can *not* be used to achieve $T \diamond \neg K\phi$. Consider a k-model $M = (\{\{a, b\}, \{a\}\}, \{a, b\})$. Updating M with $\neg Ka$ we have a possible resulting k-model $M' = (\{\{a, b\}, \{a\}, \{b\}\}, \{a, b\})$, while updating M with $\neg a$ will lead to a possible result $M'' = (\{\{a, b\}, \{a\}, \{b\}\}, \{b\})$. Note that both M' and M'' entail $\neg Ka$, but $M' <_M M''$ according to Definition 2. In general, we have the following characteristic result for ignorance update.

Lemma 4 *(Baral & Zhang 2001) Given T and $\mu = \neg K\phi$ where ϕ is objective. $M' = (W', w')$ is a k-model of $T \diamond \neg K\phi$ if and only if there exists a k-model $M = (W, w)$ of T such that*

(i) if $M \models K\phi$, then $w = w'$ and $W = W' \setminus \{w^*\}$, where $w^* \models \neg\phi$;

(ii) otherwise $w = w'$ and $W = W'$.

Theorem 4 *Ignorance update is knowledge decreased.*

Forgetting update

As another important type of knowledge update, we consider the update of T with $\mu = \neg K\phi \wedge \neg K\neg\phi$. This update can be thought of as the result of an agent *forgetting* her knowledge about the fact ϕ. We will refer to such an update as a *forgetting update*. Forgetting knowledge is an important issue in reasoning about knowledge dynamics, and forgetting update has a potential application to overcome the unrealistic assumption unlimited memory that has been adopted in many knowledge based system formalizations (Fagin *et al.* 1995). The following lemma shows that in order to forget ϕ from T, for each k-model of the current knowledge set, the agent only needs to expand the set of possible worlds of this model with exactly *one specific* world.

Lemma 5 *(Baral & Zhang 2001) Given T and $\mu = \neg K\phi \wedge \neg K\neg\phi$ where ϕ is objective. $M' = (W', w')$ is a k-model of $T \diamond \mu$ if and only if there exists a k-model $M = (W, w)$ of T such that*

(i) if $M \models K\phi$, then $w = w'$ and $W = W' \setminus \{w^*\}$, where $w^* \models \neg\phi$;

(ii) if $M \models K\neg\phi$, then $w = w'$ and $W = W' \setminus \{w^*\}$, where $w^* \models \phi$;

(iii) otherwise $w = w'$ and $W = W'$.

Example 2 Suppose $T = Kb \wedge (Ka \vee K\neg a)$ represents the current knowledge of an agent. After executing a forgetting action the agent now would like to update her knowledge with $\mu = \neg Ka \wedge \neg K\neg a$. Let $w_0 = \{a, b\}, w_1 = \{b\}, w_2 = \{a\}, w_3 = \emptyset$. It is easy to see that $M_0 = (\{w_0\}, w_0)$ and $M_1 = (\{w_1\}, w_1)$ are the two k-models of T. Then using Proposition 5, we conclude that $M'_0 = (\{w_0, w_1\}, w_0)$, $M'_1 = (\{w_0, w_3\}, w_0)$, $M'_2 = (\{w_1, w_0\}, w_1)$, and $M'_3 = (\{w_1, w_2\}, w_1)$ are the four k-models of $T \diamond \mu$. Note that $(\{w_0, w_2\}, w_0)$ cannot be a k-model of $T \diamond \mu$ according to Lemma 5. ■

Theorem 5 *Forgetting update is knowledge decreased.*

Corollary 1 *Model checking for ignorance and forgetting updates can be achieved in polynomial time.*

5 Other Intractable Subclasses - Knowledge Increased and Gradual Updates

In this section, we address some other subclasses of knowledge update problems whose model checking complexity are intractable but still lower than the general case. Such investigation will be useful for us to design more optimal model checking algorithms for these subclasses of update problems.

5.1 Knowledge Increased Update

First, as a contrary case to the knowledge decreased update, the knowledge increased update is defined as follows.

Definition 6 *Given T and μ, updating T with μ is called* knowledge increased *if for any k-model $M' = (W', w')$ of $T \diamond \mu$, there exists a k-model $M = (W, w)$ of T such that (i) $W' \subset W$, and $w = w'$; and (ii) for any $w^* \in W$, $w^* \in W'$ iff $w^* \models \phi^\mu$ or $w^* \models \neg\phi^\mu$ for some ϕ^μ in $Sub^o(\mu)$.*

It is clear that if a knowledge increased update is performed to an agent's knowledge set, it only increases the agent's knowledge and does not change the agent's actual world. Unfortunately, different from the knowledge decreased update, the model checking problem for knowledge increased update is not tractable. Before we prove this result, we first consider two specific forms of knowledge update that are knowledge increased.

Gaining knowledge and sensing updates

Given T and $K\phi$, where ϕ is objective and $T \models \phi$. We call $T \diamond K\phi$ is *gaining knowledge update*, because it is easy to observe that the purpose of updating T with $K\phi$ is only to let the agent *know* ϕ.

Another interesting form of knowledge update is so-called *sensing update*, that is, updating T with $K\phi \vee K\neg\phi$. It has been proved that sensing update is particularly useful in reasoning about sensing action and planning (Scherl & Levesque 1993; Son & Baral 2001).

Now we introduce a notion that will be useful for us to provide alternative characterizations on these two kinds of knowledge updates. Let W be a set of worlds and $w \in W$ and ϕ an objective formula. By $W^{(w,\phi)}$, we denote the subset of W: $\{w' \mid w' \in W \text{ and } w' \models \phi \text{ iff } w \models \phi\}$.

Lemma 6 *(Baral & Zhang 2001) Given T and $K\phi$ where ϕ is objective and $T \models \phi$. Then $M' = (W', w')$ is a k-model of $T \diamond K\phi$ if and only if there exists a k-model $M = (W, w)$ of T such that $w = w'$ and $W' = W^{(w,\phi)}$.*

Lemma 7 *(Baral & Zhang 2001) Given T and $K\phi \vee K\neg\phi$ where ϕ is objective and $T \models \phi$. Then $M' = (W', w')$ is a k-model of $T \diamond K\phi$ if and only if there exists a k-model $M = (W, w)$ of T such that $w = w'$ and $W' = W^{(w,\phi)}$, or $w = w'$ and $W' = W^{(w,\neg\phi)}$*

Theorem 6 *Gaining knowledge and sensing updates are knowledge increased.*

Theorem 7 *Model checking for knowledge increased update is NP-complete.*

Proof: Membership proof. Given T, μ and $M' = (W', w')$. To deciding whether $M' \in Mod(T \diamond \mu)$, we only need to show that for some $M \in Mod(T)$, $M' \in Res(M, \mu)$. A guess of $M = (W, w)$ and verifying $M \models T$ can be done in polynomial time. Since $T \diamond \mu$ is knowledge increased, to decide $M' \in Res(M, \mu)$, we only need to check: (1) $w = w'$, and (2) for any $w^* \in W$, $w^* \in W'$ iff $w^* \models \phi^\mu$ or $w^* \models \neg\phi^\mu$ for some $\phi^\mu \in Sub^o(\mu)$. Obviously, both (1) and (2) can be checked in polynomial time. So the problem is in NP.

Hardness proof. The hardness is proved by transforming the NP-complete SAT problem to a gaining knowledge update that has been showed to be knowledge increased. Let E be a CNF on the set of propositional atoms X. We construct formulas T, μ and a k-model M' over two disjoint sets X and $\hat{X}$ where $|X| = |\hat{X}|$.

$T = (X \equiv \hat{X}) \wedge \neg K(X \equiv \hat{X})$,
$\mu \equiv K(X \equiv \hat{X} \vee \neg E)$, and
$M' = (W', w')$, where
$W' = \{w'\}$, $w' = X \cup \hat{X}$.

Clearly, $M \models \mu$. We will show that E is satisfiable iff $M' \in Mod(T \diamond \mu)$. Note that since $T \models X \equiv \hat{X} \vee \neg E$ and $\mu = K(X \equiv \hat{X} \vee \neg E)$, $T \diamond \mu$ is a gaining knowledge update that is knowledge increased according to Theorem 6.

($\Rightarrow$) Suppose E is satisfiable. Let $X_1 \subseteq X$ such that $X_1 \models E$. We specify a k-model as follows:

$M^* = (W^*, w^*)$,
$W^* = \{w^*, w''\}$,
$w^* = w' = X \cup \hat{X}$, and
$w'' = X_1 \cup \hat{X}_1$, where $\hat{X}_1 = \{\hat{x}_i \mid \hat{x}_i \in \hat{X}$ and $x_i \notin X_1\}$.

Since $w'' \not\models X \equiv \hat{X}$, it is easy to see that $M^* \models \neg K(X \equiv \hat{X})$. Therefore, M^* is a k-model of T. On the other hand, since $w'' \models E$ and $w'' \not\models X \equiv \hat{X}$, it follows that $W' = \{w'\} = \{w^*\} = W^{*(w^*,\phi)}$, where $\phi = (X \equiv \hat{X}) \vee \neg E$. From Lemma 6, $M' \in Res(M^*, \mu)$, so $M' \in Mod(T \diamond \mu)$.

($\Leftarrow$) Now suppose E is not satisfiable. That is, for any $X_1 \subseteq X$, $X_1 \models \neg E$. Then from Lemma 6, for any k-model of T of the form $M = (W, w)$, where $w \neq w'$, $M' \notin Res(M, \mu)$. We consider k-models of T of the form $M = (W, w)$ where $w = w'$ (note $w' \in W$). Without loss of generality, we assume that there is one world $w^* \in W$ such that $w^* \not\models X \equiv \hat{X}$, otherwise $M \models K(X \equiv \hat{X})$ and M cannot be a k-model of T. On the other hand, since E is not satisfiable, $\neg E$ must be true in each world in W. So $M \models K\neg E$ and hence $M \models K(X \equiv \hat{X} \vee \neg E)$. This implies that $W' \neq W^{(w,\phi)}$, where $\phi = (X \equiv \hat{X}) \vee \neg E$. So M' is not a k-model of $T \diamond \mu$. ■

5.2 Knowledge Gradual Update

Now we consider a more general subclass of knowledge update problems that subsume both knowledge decreased and increased updates. Given T and μ, we say the update of T with μ is *knowledge gradual* if for any k-model $M' = (W', w')$ of $T \diamond \mu$, there exists a k-model $M = (W, w)$ of T such that either $W \subseteq W'$ or $W' \subset W$. Note that, after performing a knowledge gradual update, the agent's knowledge may be decreased or increased (or without change), and the agent's actual world may be changed as well.

Example 3 Let $T = a \wedge \neg Ka$ and $\mu = K\neg a$. Obviously, T has a unique k-model $M = (\{\{a\}, \emptyset\}, \{a\})$. Then updating M with μ generates a unique k-model of $T \diamond \mu$: $M' = (\{\emptyset\}, \emptyset)$. Obviously, M' has increased knowledge from M and the actual world of M' is also different from M's. ■

Lemma 8 *Let $M = (W, w)$, $M_1 = (W_1, w_1)$ and $M_2 = (W_2, w_2)$ be three k-models where $W \subseteq W_2$ or $W_2 \subseteq W$. Deciding whether $M_1 \leq_M M_2$ can be done in polynomial time.*

Proof: Obviously, if $w_1 \neq w_2$, then deciding whether $M_1 \leq_M M_2$ is in polynomial time according to condition 1 in Definition 2. Now suppose $w_1 = w_2$. If $W \subseteq W_2$, then from Definition 2, it is easy to see that

$M_1 \leq_M M_2$ iff $W \subseteq W_1 \subseteq W_2$. To see why this is the case, first consider $M_1 \leq_M M_2$. Since $W \subseteq W_2$, it is obviously that $W \not\subset W1$ and $W_1 \not\subset W$, otherwise $M_2 <_M M_1$. It cannot be $W_1 \subset W$ either since according to condition (ii) in Definition, condition (c) is not satisfied. So it must be the case that $W \subseteq W_1$ and condition (b) is satisfied. This implies that $W \subseteq W_1 \subseteq W_2$. On the ther hand, suppose $W \subseteq W_1 \subseteq W_2$. From condition (i) in Definition, this directly follows that $M_1 \leq_M M_2$.

Clearly, under the condition that $W \subseteq W_2$, checking whether $W \subseteq W_1 \subseteq W_2$ can be done in polynomial time.

Now consider the case that $W_2 \subseteq W$. We prove that under the condition $W_2 \subseteq W$, $M_1 \leq_M M_2$ iff $W_2 \subseteq W_1 \subseteq W$, from which it is clear that checking whether $M_1 \leq_M M_2$ can be verified in polynomial time. Obviously, if $W_2 \subseteq W_1 \subseteq W$, it directly follows that $M_1 \leq_M M_2$ from condition (ii) in Definition 2. Consider the other direction. If $M_1 \leq_M M_2$, then we have $W_1 \not\subset W$ and $W \not\subset W_1$ due to the fact $W_2 \subseteq W$, otherwise we will have $M_2 <_M M_1$. It is also impossible that $W \subseteq W_1$, since under this condition we can only use condition (b) to verify if $M_1 \leq_M M_2$ and this follows that (b) is not satisfied. So it must be the case $W_1 \subseteq W$. As $W_2 \subseteq W$, condition (c) must be satisfied, which implies $KM_1 \setminus KM \subseteq KM_2 \setminus KM$. From simple set inclusion and union properties, we have $KM_1 \setminus KM \subseteq KM_2 \setminus KM$ iff $KM_1 \cup KM \subseteq KM_2 \cup KM$. Since $W_1 \subseteq W$ and $W_2 \subseteq W$, we have $KM \subseteq KM_1$ and $KM \subseteq KM_2$. So $KM_1 \cup KM \subseteq KM_2 \cup KM$ iff $KM_1 \subseteq KM_2$. From Proposition 1 in (Baral & Zhang 2001), we have $KM_1 \subseteq KM_2$ iff $W_2 \subseteq W_1$. This completes our proof. ■

Lemma 9 *Let μ be a S5 formula and $M = (W, w)$, $M' = (W', w')$ two k-models where $W \subseteq W'$ or $W' \subseteq W$. Deciding whether $M' \in Res(M, \mu)$ is in co-NP.*

Proof: If $M' \notin Res(M, \mu)$, then there must exists a k-model M'' of μ such that $M'' \leq_M M'$ and $M' \not\leq_M M''$. Obviously, a guess of M'' and verifying $M'' \models \mu$ can be done in polynomial time. From Lemma 8, deciding whether $M'' \leq_M M'$ and $M' \not\leq_M M''$ can be done in polynomial time. So the original problem is in co-NP. ■

Now we have the following main result of model checking for gradual knowledge update.

Theorem 8 *Model checking for gradual knowledge update is Σ_2^P-complete.*

To prove Theorem 8, we need first to prove the following lemma.

Lemma 10 *Let $X, Y, \hat{X}, \hat{Y}$ be sets of propositional atoms and a be a propositional atom, where $|X| = |\hat{X}|$, $|Y| = |\hat{Y}|$ and any two sets of $X, Y, \hat{X}, \hat{Y}$ and $\{a\}$ are disjoint. Suppose ϕ is an objective formula formed based on $X \cup Y$. Let T and μ be the following two S5 formulas respectively:*

$$\begin{aligned}
&T = \gamma_1 \vee \gamma_2, \text{ where} \\
&\gamma_1 = (X \equiv \hat{X}) \equiv \phi \equiv a \wedge (Y \wedge \neg\hat{Y}) \wedge \\
&\quad\quad \neg K \neg(a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y}), \\
&\gamma_2 = ((X \wedge \neg\hat{X}) \equiv \neg\phi \equiv \neg a) \wedge \hat{Y}, \text{ and} \\
&\mu = K(a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y}) \vee \\
&\quad\quad K(\neg a \wedge X \wedge \neg\hat{X} \wedge (\bigvee \neg Y) \wedge \hat{Y}) \vee \\
&\quad\quad K(\neg a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y})^5.
\end{aligned}$$

Then $T \diamond \mu$ is knowledge gradual.

Proof: To prove $T \diamond \mu$ to be knowledge gradual, we need to show that for any k-model $M = (W, w) \in Mod(T)$, if $M' = (W', w') \in Res(M, \mu)$, then either $W \subseteq W'$ or $W' \subseteq W$. From the construction of T, it is easy to see that if $M \in Mod(T)$, then either $M \models \gamma_1$ or $M \models \gamma_2$, but $M \not\models \gamma_1 \wedge \gamma_2$. Based on this observation, our proof consists of two cases.

<u>*Case 1*</u>. Let $M = (W, w) \in Mod(\gamma_1)$. Since $M \models \gamma_1$, we have

$w = X_1 \cup \hat{X}_1 \cup Y \cup \{a\}$, where $X_1 \cup Y \models \phi$ for some $X_1 \subseteq X$.

Furthermore, since $M \models \neg K \neg(a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y})$, there exists a world $w^* \in W$ such that

$$w^* = X \cup Y \cup \hat{Y} \cup \{a\}.$$

Now we specify a k-model of μ as follows:

$$M^* = (\{w^*\}, w^*).$$

It is easy to see that $M^* \models K(a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y})$. So $M^* \models \mu$. Furthermore, M^* is the unique k-model of $K(a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y})$. We prove M^* is the unique k-model in $Res(M, \mu)$.

Note $Diff(w, w^*) = (X \setminus X_1) \cup \hat{X}_1 \cup \hat{Y}$. Besides M^*, μ has other two types of k-models:

[5] Recall that $\bigvee \neg Y = \bigvee_{y_i \in Y} \neg y_i$.

$M_1 = (W_1, w_1)$, where $w_1 = X \cup Y_1 \cup \hat{Y}$, where $Y_1 \subset Y$ ($Y_1 \neq Y$), and
$M_2 = (W_2, w_2)$, where $w_2 = X \cup Y \cup \hat{Y}$.

Note that

$M_1 \models K(\neg a \wedge X \wedge \neg\hat{X} \wedge (\bigvee \neg Y) \wedge \hat{Y})$, and
$M_2 \models K(\neg a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y})$.

Consider

$Diff(w, w_1) = (X \setminus X_1) \cup \hat{X}_1 \cup (Y \setminus Y_1) \cup \hat{Y} \cup \{a\}$,
$Diff(w, w_2) = (X \setminus X_1) \cup \hat{X}_1 \cup \hat{Y} \cup \{a\}$.

Clearly, we have

$Diff(w, w^*) \subset Diff(w, w_1)$, and
$Diff(w, w^*) \subset Diff(w, w_1)$.

So $Res(M, \mu) = \{M^*\}$. Also observe that $M^* = (\{w^*\}, w^*)$, $\{w^*\} \subset W$.

<u>*Case 2*</u>. Let $M = (W, w) \in Mod(\gamma_2)$. We have

$w = X \cup Y_1 \cup \hat{Y}$, where $X \cup Y_1 \models \neg\phi$ for some $Y_1 \subseteq Y$.

If $Y_1 \neq Y$, then we have

$$w \models \neg a \wedge X \wedge \neg\hat{X} \wedge (\bigvee \neg Y) \wedge \hat{Y}),$$

this implies that there exists a subset W_1 of W where for each $w_i \in W_1$,

$$w_i \models (\neg a \wedge X \wedge \neg\hat{X} \wedge (\bigvee \neg Y) \wedge \hat{Y}).$$

By specifying W_1 to be the maximal such subset of W, it is easy to note that $M_1 = (W_1, w)$ is a k-model in $Res(M, \mu)$.

On the other hand, if $Y_1 = Y$, then we have

$$w \models (\neg a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y}),$$

this implies that there exists a subset W_2 of W where for each $w_i \in W_2$,

$$w_i \models (\neg a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y}).$$

Similarly, by specifying W_2 to be the maximal such subset of W, it is easy to note that $M_2 = (W_2, w)$ is a k-model in $Res(M, \mu)$.

Since in both cases, we have $W_1 \subseteq W$ and $W_2 \subseteq W$, this follows that for any k-model M of T where $M \models \gamma_2$, every resulting k-model after updating M with μ only increases the knowledge from M. ■

Proof: (**Proof of Theorem 8**) Membership proof. Given T, μ and a k-model M where $T \diamond \mu$ is knowledge gradual. $M \in Mod(T \diamond \mu)$ iff for some $M' \in Mod(T)$ such that $M \in Res(M', \mu)$. Obviously, a guess of M' and verifying $M' \models T$ can be done in polynomial time. From Lemma 9, deciding whether $M \in Res(M', \mu)$ can be solved with one call to a co-NP oracle. So the problem is in Σ_2^P.

Hardness proof. This part is based on a variation of the proof of Lemma 10. We prove the hardness by giving a polynomial transformation from deciding the validity of $\exists X \forall Y E$, where E is a Boolean expression using propositional atoms over $X \cup Y$. We construct T, μ and a k-model M^* over propositional atoms $X \cup Y \cup \hat{X} \cup \hat{Y} \cup \{a\}$, where $|\hat{X}| = |X|$ and $|\hat{Y}| = |Y|$, and any two sets among $X, Y, \hat{X}, \hat{Y}$ and $\{a\}$ are disjoint.

$T = \gamma_1 \vee \gamma_2$, where
$\gamma_1 = (X \equiv \hat{X}) \equiv E \equiv a \wedge (Y \wedge \neg\hat{Y}) \wedge$
$\quad \neg K \neg (a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y})$,
$\gamma_2 = ((X \wedge \neg\hat{X}) \equiv \neg E \equiv \neg a) \wedge \hat{Y}$,
$\mu = K(a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y}) \vee$
$\quad K(\neg a \wedge X \wedge \neg\hat{X} \wedge (\bigvee \neg Y) \wedge \hat{Y}) \vee$
$\quad K(\neg a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y})$,
$M^* = (W^*, w^*)$, where
$W^* = \{w^*\}$, $w^* = X \cup Y \cup \hat{Y} \cup \{a\}$.

Note that

$$M^* \models K(a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y}).$$

So M^* is a k-model of μ. Furthermore, it is the unique k-model of $K(a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y})$. From Lemma 10, we know that $T \diamond \mu$ is knowledge gradual. Now we will show that M^* is a k-model of $T \diamond \mu$ if and only if $\exists X \forall Y E$ is valid.

($\Rightarrow$) Suppose $\exists X \forall Y E$ is valid. Then for some $X_1 \subseteq X$, $X_1 \cup Y \models E$. We specify a k-model of γ_1 as follows:

$$M = (W, w), \text{ where } w = X_1 \cup \hat{X}_1 \cup Y \cup \{a\}.$$

Since $M \models \neg K \neg (a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y})$, it is clear that the world w^* must be in W, i.e. $w^* \in W$. With the same justification as described in the proof of Lemma 10, we conclude that M^* is a k-model of updating M with μ.

($\Leftarrow$) Suppose $\exists X \forall Y E$ is not valid. That is, $\forall X \exists Y \neg E$ is valid. Then $X \cup Y_1 \models \neg E$ for some $Y_1 \subseteq Y$. In this case, T has the following type of k-models:

$$M = (W, w), \text{ where } w = X \cup Y_1 \cup \hat{Y}.$$

Note that $w \models ((X \wedge \neg\hat{X} \equiv \neg E \equiv \neg a) \wedge \hat{Y})$. That is, M is a k-model of γ_2.

If $Y_1 \neq Y$, then we have

$$w \models (\neg a \wedge X \wedge \neg\hat{X} \wedge (\bigvee \neg Y) \wedge \hat{Y}).$$

We now specify a k-model of μ as follows: $M_1 = (W_1, w_1)$, where $w_1 = w$ and W_1 is the maximal subset of W such that for each $w_i \in W_1$,

$$w_i \models (\neg a \wedge X \wedge \neg\hat{X} \wedge (\bigvee \neg Y) \wedge \hat{Y}).$$

Since

$$Diff(w, w^*) = (X \setminus X_1) \cup \hat{X}_1 \cup \hat{Y}, \text{ and}$$
$$Diff(w, w_1) = \emptyset \subset Diff(w, w^*),$$

M^* is not a k-model in $Res(M, \mu)$.

If $Y_1 = Y$, then we have

$$w \models (\neg a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y}).$$

Again, we can specify a k-model of μ as follows: $M_2 = (W_2, w_2)$, where $w_2 = w$ and W_2 is maximal subset of W such that for each $w_i \in W_2$,

$$w_i \models (\neg a \wedge X \wedge \neg\hat{X} \wedge Y \wedge \hat{Y}).$$

Since $Diff(w, w_2) = \emptyset \subset Diff(w, w^*)$, M^* is not a k-model in $Res(M, \mu)$ in this case either.

Finally, suppose for some $X_1 \subseteq X$ and $Y_1 \subseteq Y$, E is evaluated to be true on $X_1 \cup Y_1$, i.e. $X_1 \cup Y_1 \models E$. Without loss of generality, we can assume $Y_1 \neq Y$[6]. This implies that γ_1 does not have a k-model under this situation. Therefore, if $\exists X \forall Y E$ is not valid, all k-models of T must be k-models of γ_2. ■

[6] Note that this assumption is always feasible. For instance, if $X_1 \cup Y \models E$, we can expand Y to be Y' by adding a new atom y' into Y to make $Y \neq Y'$, i.e. $Y' = Y \cup \{y'\}$, and modify E to be $E' = E \wedge \neg y'$ such that $X_1 \cup Y \models E'$ but $X_1 \cup Y' \not\models E'$.

6 Conclusion

The notions of belief revision and update have been well-studied in the literature (Gardenfors 1988; Katsuno & Mendelzon 1991; Winslett 1988; Zhang & Foo 2000) and the computational complexity of model checking and performing such revision/update and its variations were well understood (Eiter & Gottlob 1992; Liberatore 2000; Liberatore & Schaerf 2001; Zhang 2001). On the other hand, although the issue of reasoning about knowledge dynamics has been investigated by some researchers from different perspectives, e.g. (Fagin *et al.* 1995; Groeneveld 1995; Lomuscio & Ryan 1998), the formal account of modeling knowledge change from update viewpoint is far from clear. In (Baral & Zhang 2001) we introduced the notion of knowledge update, defined its formal semantics and related it to the impact of sensing actions on the knowledge of an agent, and to actions corresponding to forgetting and ignorance. We also argued why such a new notion of knowledge update is necessary for the above and why the earlier notions of belief revision/update are not sufficient.

In this paper we undertook a further study about the complexity issue of knowledge update. In particular, we analyzed the complexity of model checking for knowledge update in the general case and in special cases. We identify special subcases where the model checking is either tractable or its complexity is lower than the general case (see the summary in the following table). We expect that these results would be useful for designing more optimal model checking algorithms in the implementation of knowledge update (Baral & Zhang 2002).

Update Type		**Complexity**
General case		Σ_3^P-complete
Knowledge decreased		P
ignorance	forgetting	P
Knowledge increased		NP-complete
gaining knowledge	sensing	NP-complete
Knowledge gradual		Σ_2^P-complete

References

Baral, C., and Zhang, Y. 2001. On the semantics of knowledge update. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, 97–102.

Baral, C., and Zhang, Y. 2002. Knowledge update: Semantics and complexity. *Manuscript.*

Eiter, T., and Gottlob, G. 1992. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence* 57:227–270.

Fagin, R.; Halpern, J.; Moses, Y.; and Vardi, M. 1995. *Reasoning about Knowledge.* MIT Press.

Gardenfors, P. 1988. *Knowledge in Flux.* Cambridge, Massachusetts: MIT Press.

Garey, M., and Johnson, D. 1979. *Computers and Intractability.* Freeman.

Groeneveld, W. 1995. *Logical Investigations into Dynamic Semantics.* PhD Thesis, Universiteit van Amsterdam.

Katsuno, H., and Mendelzon, A. 1991. On the difference between updating a knowledge base and revising it. In *Proceedings of KR-91*, 387–394.

Liberatore, P., and Schaerf, M. 2001. Belief revision and update: complexity of model checking. *Journal of Computer and System Sciences* 62:43–72.

Liberatore, P. 2000. The complexity of belief update. *Artificial Intelligence* 119:141–190.

Lomuscio, A., and Ryan, M. 1998. Ideal agents sharing (some!) knowledge. In *Proceedings of ECAI-98*, 557–561.

Meyer, J.-J., and van der Hoek, W. 1995. Epistemic logics for computer science and artificial intelligence. In *Cambridge Tracts in Theoretical Computer Science, 41.* Cambridge University Press.

Scherl, R., and Levesque, H. 1993. The frame problem and knowledge producing actions. In *Proceedings of AAAI-93*, 689–695.

Son, T., and Baral, C. 2001. Formalizing sensing actions: a transition function based approach. *Artificial Intelligence* 125:19–91.

Winslett, M. 1988. Reasoning about action using a possible models approach. In *Proceedings of the 7th National Conference on AI (AAAI-88)*, 89–93.

Zhang, Y., and Foo, N. 2000. Updates with disjunctive information: From syntactical and semantical perspectives. *Computational Intelligence* 16:29–52.

Zhang, Y. 2001. The complexity of logic program updates. In *Proceedings of 14th Australian Joint Conference on Artificial Intelligence (AI2001)*, 630–643.