Partial Goal Satisfaction and Goal Change

Weak and Strong Partial Implication, Logical Properties, Complexity

Yi Zhou
Intelligent Systems Laboratory
University of Western Sydney
NSW, Australia
yzhou@scm.uws.edu.au

Leendert van der Torre
CSC
University of Luxembourg
Luxembourg
leendert@vandertorre.com

Yan Zhang
Intelligent Systems Laboratory
University of Western Sydney
NSW, Australia
yan@scm.uws.edu.au

ABSTRACT

Partial implication semantics in the context of a background theory has been introduced to formalize partial goal satisfaction in the context of beliefs. In this paper, we introduce strong partial implication prohibiting *redundancies* and weak partial implication allowing *side effects*, we study their semantic as well as complexity properties, and we apply the three notions of partial implication to goal change in the context of beliefs.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Intelligent agents*

General Terms

Theory

Keywords

Goal change, partial implication, non-classical logic

1. INTRODUCTION

Consider the following four scenarios for a shopping agent desiring to buy apples and oranges today, represented by $x \wedge y$, and believing from the only shop on-line today that either there are:

1. apples (x), or
2. apples and bananas ($x \wedge z$), or
3. apples but no oranges ($x \wedge \neg y$), or
4. bananas (z).

In none of the scenarios the agent can achieve its goal, because none of the scenarios classically implies the goal $x \wedge y$. Many agents would therefore drop the goal, but Zhou and Chen [13] argue that it is rational for the agent in scenario 1 and 2 to go to the store. They formalize their intuition by introducing a notion so-called partial implication, where x and $x \wedge z$ partially imply $x \wedge y$ while $x \wedge \neg y$ and z do not.

In this paper we formalize partial goal satisfaction of an agent going to the store only in the first scenario, and we call the corresponding logical notion strong partial implication, which prohibits what we call redundancies, like the bananas in the example. It is an intermediate solution between dropping the goal and using partial implication. Moreover, we formalize partial goal satisfaction of an agent that goes to the shop in the first three scenarios, and we call the corresponding logical notion weak partial implication, which allows what we call side effects, like the absence of oranges. All notions are defined in the context of background knowledge representing the beliefs of the agent.

-	strong partial	partial	weak partial
partial satisfaction	yes	yes	yes
redundancy	no	maybe	maybe
side effect	no	no	maybe

Table 1: The family of partial implication

Another perspective on the example is that the agent is first *changing* its goal using partial implication, from $x \wedge y$ to for example x, $x \wedge z$ or $x \wedge \neg y$, and thereafter using classical implication again. In this paper we therefore address the following questions:

1. How to define strong and weak partial implication in the context of background knowledge?
2. What are their properties and complexity?
3. How to define the problem of changing a goal rather than simply dropping it?
4. How to use partial implication for goal change in the context of beliefs?

The *dynamics of goals* has attracted a lot of research recently [2, 12, 11, 9, 8, 7, 10]. Goal change is inevitable in real domains for several reasons. Firstly, the agent has bounded resources so that it cannot make "perfect" plans. Secondly, the environment is unpredictable and uncertain, which means that it is impossible to predict the future. Finally, the agent's information about the environment maybe ambiguous, even wrong. Therefore, the agent's plans may lead to unexpected results. All of these would possibly result in a situation where the agent's original goal is infeasible or meaningless. Hence, the agent should change its goal.

This paper is organized as follows. In Section 2 we discuss the use of prime implicants for partial goal satisfaction, and in Section 3-5 we formalize partial goal satisfaction semantically as three kinds of partial implication, and investigate the main properties. In Section 6, we propose a framework based on partial implication for the problem of how to change the agent's goal in the context of beliefs. In Section 7, we address the complexity issues for the three kinds of partial implication.

Cite as: Partial Implication Semantics for Goal Change, Yi Zhou, Leendert van der Torre, and Yan Zhang, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. pp. 413-420.

2. PARTIAL GOAL SATISFACTION

Zhou and Chen [13] argue that *partial satisfaction of goals* is an important issue in rational decision making of autonomous agents. In the traditional logical approach, agents always try to find plans to *completely* achieve their goals, which is usually formalized based on classical implication.

> P satisfies a goal Q if P logically implies Q.

However, people often perform plans *partially* achieving their goals. Informally, Zhou and Chen characterize partial goal satisfaction as follows.

> P partially satisfies a goal to achieve Q if for all cases of P, there is a case of Q such that the former is related to a part of the latter.

The reason for this complex definition is that they assume that P and Q are represented by propositional formulas, where inconsistent propositions cannot both be achieved. For example, a goal to achieve p and a goal to achieve $\neg p$ cannot both be satisfied by achieving p and $\neg p$ simultaneously, like switching on and off the light. In such a case, when p stands for having the light on at some points in the future, $\neg p$ would stand for never having the light on. Likewise a goal to achieve p together with a belief in $\neg p$ implies that either the goal cannot be achieved, or the belief is wrong. If P and Q would be represented by sets of literals, for example, then they would say that P partially satisfies goal to achieve Q if P is related to a part of Q. However, in propositional logic they also have to deal with disjunction, for which they introduce the cases in the definition of partial goal satisfaction.

Therefore, they introduce a notion of *partial implication* using propositional language to capture partial satisfaction relationship between two propositional formulas with respect to a background formula set. For instance, $x \wedge z$ does not imply $x \wedge y$ in classical propositional logic. However, it partially implies $x \wedge y$ since x, which can be considered as a *part* of $x \wedge y$, is a logical consequence of $x \wedge z$. Thus, we may treat a proposition as a choice among several cases, where each case consists of a set of literals. We may say, for example, that P partially implies goal Q, where P and Q are propositional formulas, i.e., not necessarily propositional atoms, when every case of P is a part of a case of Q.

Before we make this idea more precise, we introduce some notations. We denote the propositional language by $\mathcal{L}$. Formulas in $\mathcal{L}$ are composed recursively by a finite set $Atom$ of atoms (also called variables) with $\{\top, \bot\}$ and standard connectives $\neg$ and $\rightarrow$. The connectives $\wedge$, $\vee$, $\leftrightarrow$ are defined as usual. Literals are atoms and their negations. We use lower case letters, upper case letters, lower Greek letters and upper Greek letters to denote atoms and literals, formulas, literal sets and formula sets respectively. We write $-l$ to denote the negation of a literal l, $-\pi$ to denote the set of negations of all literals in π. We use Γ itself to denote the conjunction of a set of formulas in Γ if it is clear from the context. We write $Atom(l)$, $Atom(P)$, $Atom(\pi)$, $Atom(\Gamma)$ to denote the sets of atoms occurred in literal l, formula P, literal set π and formula set Γ respectively. We say that a literal set π is an assignment over a set of atoms $A \subseteq Atom$ (assignment for short if $A = Atom$) if for each atom $x \in A$, exactly one of x and $\neg x$ is in π. Notice that both a literal l and an assignment π can also be considered as formulas. Hence, in the rest of this paper, both l and π also denote their corresponding formulas for convenience if it is clear from the context. We write $P|l$ to denote the formula obtained from P by replacing every occurrence of l into $\top$ and every occurrence of $-l$ into $\bot$ simultaneously. Let $\pi = \{l_1, l_2, ..., l_k\}$, we write $P|\pi$ to denote the formula $(((P|l_1)|l_2)|...)|l_k$.

To formalize their notion of case, Zhou and Chen use Quine's notion of a prime implicant of a formula [5] as a minimal set (in the sense of set inclusion) of literals satisfying this formula. Prime implicants are used in many areas of logic, but they also relativize the notion of prime implicant to a belief set Γ.

DEFINITION 1 (PRIME IMPLICANT). *A literal set π is a prime implicant of a formula P with respect to formula set Γ*[1] *if:*

(1) $\Gamma \cup \pi$ *is consistent.*

(2) $\Gamma \cup \pi \models P$.

(3) *There is no literal set π' satisfying the above two conditions and $\pi' \subset \pi$.*

The set of all prime implicants of P with respect to Γ is denoted by $PI(\Gamma, P)$.

The prime implicants of a formula P *w.r.t.* a formula set Γ play two roles.

Cases On the one hand, if there exists an assignment π satisfying both Γ and P, then there exists a subset of π which is a prime implicant of P *w.r.t.* Γ. On the other hand, if π is a prime implicant of P *w.r.t.* Γ, then it can be extended into an assignment satisfying both Γ and P. This means that the prime implicants of P *w.r.t.* Γ are corresponding to the possible worlds satisfying both Γ and P. In other words, intuitively, they represent all the *cases* which make the proposition true *w.r.t.* the background theory.

Part Suppose that π is a prime implicant of P *w.r.t.* Γ and l is a literal such that $l \in \pi$. Then we have that $\Gamma \cup \pi \models P$ and $\Gamma \cup \pi \backslash \{l\} \not\models P$. Intuitively, this means that l plays an essential role for achieving P *w.r.t.* Γ via π. Thus, l can be considered as a *part* of P *w.r.t.* Γ.

Both roles of prime implicants are exploited by the notions of partial implication introduced in the following sections. P strongly partially implies Q means that P implies some parts of Q and nothing else, and this corresponds roughly to a weakening of Q to P. P (weakly) partially implies Q means that P implies some parts of Q, but may have something more (redundancies). In this sense, strong partial implication captures that P is a part of Q, while (weak) partial implication captures that P implies part of Q. Before we are going into the technical details, we mention three general properties of prime implicants which are used in the following sections.

LEMMA 1. *Let P be a formula, Γ a set of formulas and π a literal set such that $\Gamma \cup \pi$ is consistent and $\Gamma \cup \pi \models P$. Then there exists $\pi' \subseteq \pi$ such that π' is a prime implicant with respect to Γ.*

LEMMA 2. *Let P be a formula and Γ a set of formulas. We have that $\Gamma \models \neg P$ if and only if $PI(\Gamma, P) = \emptyset$; $\Gamma \models P$ if and only if $PI(\Gamma, P) = \{\emptyset\}$.*

Lemma 2 also shows the difference between $PI(\Gamma, P) = \emptyset$ and $PI(\Gamma, P) = \{\emptyset\}$. The former means that there is no prime implicant of P *w.r.t.* Γ, while the latter means that there is a unique prime implicant of P *w.r.t.* Γ, which is the empty set. We say that a formula P is *trivial w.r.t.* a formula set Γ if $\Gamma \models P$ or $\Gamma \models \neg P$. Otherwise we say that P is *non-trivial w.r.t.* Γ.

LEMMA 3. *Let P and Q be two formulas and Γ a set of formulas. $\Gamma \models P \leftrightarrow Q$ iff $PI(\Gamma, P) = PI(\Gamma, Q)$.*

[1] In [13], this is called Γ-prime implicant.

3. STRONG PARTIAL IMPLICATION

A formula P strongly partially implies a formula Q when for all the cases in P, there is a case in Q that extends it in the sense that all the literals of the case of P are also appeared in the case of Q. Moreover, this basic definition is extended with respect to a formula set Γ and it eliminates trivial cases by excluding that P is implied by the beliefs, or contradictory with it.

DEFINITION 2 (STRONG PARTIAL IMPLICATION). *We say that a formula P strongly partially implies a formula Q with respect to a formula set Γ, denoted by $P \succ^S_\Gamma Q$, if:*

1. *$PI(\Gamma, P)$ is not empty and $PI(\Gamma, P) \neq \{\emptyset\}$.*
2. *For each $\pi \in PI(\Gamma, P)$, there exists $\pi' \in PI(\Gamma, Q)$, such that $\pi \subseteq \pi'$.*

We write $P \not\succ^S_\Gamma Q$ if it it not the case that $P \succ^S_\Gamma Q$. For convenience, we omit Γ when it is empty.

The following example illustrates the apples and oranges example for strong partial implication.

EXAMPLE 1. *x strongly partially implies $x \wedge y$ while $x \wedge z$, $x \wedge \neg y$ and z do not.*

Strong partial implication should not be considered as a variant of classical implication such as relevant implication, and it should not be applied iteratively. The following example illustrates some of its properties.

EXAMPLE 2. *p partially satisfies both $p \wedge q$ and $p \vee q$, represented by $p \succ^S p \wedge q$ and $p \succ^S p \vee q$ respectively, but for different reasons. In the former, suppose that an agent has a goal to accomplish every element in a set of goals, then accomplishing one of them is helpful to the original goal. In the latter, if it has a goal to accomplish one of the goals in a set of goals, then accomplishing one of them is of course "helpful" in the sense that this behavior achieves the goal.*

$p \wedge q$ partially implies p, but $p \vee q$ does not partially imply p, represented by $p \wedge q \succ^S p$ and $p \vee q \not\succ^S p$ respectively, which illustrates that in some examples conjunction and disjunction are distinct. The latter does not hold, because the case q is in no way helpful for the goal to achieve p.

Theorem 1 shows that a proposition already true or always false *w.r.t.* a background theory does not strongly partially satisfy any other formulas, and it is not strongly partially satisfied by any other formulas.

THEOREM 1 (NON-TRIVIALITY). *Let P and Q be two formulas and Γ a set of formulas. If either P or Q is trivial w.r.t. Γ, then $P \not\succ^S_\Gamma Q$.*

PROOF. This assertion can be easily proved by Lemma 2 since the strong partial implications require the intersections of two prime implicants not to be empty. □

Theorem 2 means that strong partial implications are syntactically independent. Propositions equivalent *w.r.t.* the background theory play the same roles in partial implication semantics. Theorem 2 also shows that we can substitute a proposition with an equivalent one *w.r.t.* the background theory. In classical propositional logic, we can also substitute an atom with a formula. However, this can not be done in partial implications. For instance, observe that $x \succ x \wedge y$ ($x \succ^W x \wedge y$, $x \succ^S x \wedge y$). If we replace x by $\neg y$, then by Theorem 1, this strong partial implication relationship no longer holds.

THEOREM 2 (INDEPENDENCY OF SYNTAX). *Let P, Q and R be three formulas and Γ a set of formulas such that $\Gamma \models P \leftrightarrow Q$. Then, $P \succ^S_\Gamma R$ implies that $Q \succ^S_\Gamma R$; $R \succ^S_\Gamma P$ implies that $R \succ^S_\Gamma Q$.*

PROOF. This assertion follows directly from Lemma 3 since strong partial implication is defined only based on the sets of prime implicants. □

Theorem 3 states that relevancy for partial implication holds.

THEOREM 3 (RELEVANCY). *Let P and Q be two formulas such that $P \succ^S R$. Then, $Atom(P) \cap Atom(Q) \neq \emptyset$.*

PROOF. According to the definitions of strong partial implication, there exists $\pi \in PI(\Gamma, P)$ and $\pi' \in PI(\Gamma, Q)$ such that $\pi \cap \pi' \neq \emptyset$. Thus, $Atom(\pi) \cap Atom(\pi) \neq \emptyset$. Since $Atom(\pi) \subseteq Atom(P)$ and $Atom(\pi') \subseteq Atom(Q)$, $Atom(P) \cap Atom(Q) \neq \emptyset$. □

PROPOSITION 1 (NON MONOTONICITY). *Let P and Q be two propositions, Γ and Γ' two sets of formulas such that $\Gamma \subset \Gamma'$. $P \succ^S_\Gamma Q$ does not imply that $P \succ^S_{\Gamma'} Q$.*

PROOF. For instance, $x \succ^S x \wedge y$. However, $x \not\succ^S_{\{\neg y\}} x \wedge y$. This shows that monotonicity does not hold for strong partial implication in general. □

THEOREM 4 (TRANSITIVITY). *Let P, Q and R be three formulas and Γ a set of formulas. If $P \succ^S_\Gamma Q$ and $Q \succ^S_\Gamma R$, then $P \succ^S_\Gamma R$.*

PROOF. By Theorem 1, P, Q and R are all non-trivial *w.r.t.* Γ. Let π be a prime implicant of P *w.r.t.* Γ. Since $P \succ^S_\Gamma Q$, there exists a literal set π_1 consistent with Γ such that $\pi_1 \in PI(\Gamma, Q)$ and $\pi \subseteq \pi_1$. Moreover, since $Q \succ^S_\Gamma R$, there exists a literal set π_2 consistent with Γ such that $\pi_2 \in PI(\Gamma, R)$ and $\pi_1 \subseteq \pi_2$. Hence, $\pi \subseteq \pi_2$. This shows that $P \succ^S_\Gamma R$. □

4. PARTIAL IMPLICATION

Strong partial implication is an intuitive replacement of material implication, but in some cases it may be too strong, and a weaker notion may be called for. Zhou and Chen [13] say that a formula P partially implies a formula Q when for all the cases in P, there is a case in Q that is not disjoint and does not conflict with it, in the sense that some literals appeared in the case of P are also appeared in the case of Q, and there are no literals in the case of P such that the negated literal is in the case of Q. For example, consider two sets π and π' of consistent literals. On the one hand, $\pi \cap \pi' \neq \emptyset$ means that π achieves some parts of π' since all elements of π' can be considered as parts of π'. On the other hand, $\pi \cap -\pi' = \emptyset$ means that π has no side effect to π' since the side effects of π' can be considered as the negations of all elements in π'. Moreover, this definition is relativized to belief set Γ, and condition 1 ensures that there exists at least one case in P.

DEFINITION 3 (PARTIAL IMPLICATION). *A formula P partially implies a formula Q with respect to a formula set Γ, denoted by $P \succ_\Gamma Q$[2], if:*

1. *$PI(\Gamma, P)$ is not empty.*
2. *For each $\pi \in PI(\Gamma, P)$, there exists $\pi' \in PI(\Gamma, Q)$, such that $\pi \cap \pi' \neq \emptyset$ and $\pi \cap -\pi' = \emptyset$.*

[2] In [13], this is denoted as $\Gamma \models P \succ Q$.

We write $P \not\succ_\Gamma Q$ if it it not the case that $P \succ_\Gamma Q$.

The following example illustrates the apples and oranges example for partial implication.

EXAMPLE 3. *x and $x \wedge z$ partially imply $x \wedge y$ while $x \wedge \neg y$ and z do not.*

Theorem 5 shows that strong partial implication implies partial implication.

THEOREM 5 (STRONG AND PARTIAL IMPLICATION). *Let P, Q be two formulas and Γ a set of formulas. $P \succ^S_\Gamma Q$ implies $P \succ_\Gamma Q$.*

PROOF. Suppose that $P \succ^S_\Gamma Q$. Let $\pi \in PI(\Gamma, P)$. We have that $\pi \neq \emptyset$, otherwise, $\emptyset$ is the only prime implicant of P with respect to Γ, a contradiction. By condition 2 in Definition 2, there exists $\pi' \in PI(\Gamma, Q)$ such that $\pi \subseteq \pi'$. Thus, $\pi \cap \pi' = \pi \neq \emptyset$. Moreover, $\pi \cap -\pi' = \emptyset$, otherwise, there exists $l \in \pi \cap -\pi'$. Thus, $l \notin \pi'$ since π' is a consistent literal set, a contradiction. This shows that $P \succ_\Gamma Q$. □

Non-triviality (Theorem 1, Theorem 4 in [13]), Independency of Syntax (Theorem 2), Relevancy (Theorem 3, Proposition 8 in [13]) hold for partial implication just as for strong partial implication. However, transitivity does not hold for partial implication. For instance, $x \succ x \wedge y$ and $x \wedge y \succ y \wedge z$. However $x \not\succ y \wedge z$. Finally, Zhou and Chen [13] show that partial implication is an extension of classical implication in the nontrivial cases: if P and Q are two formulas non-trivial *w.r.t.* Γ and $\Gamma \models P \rightarrow Q$, then $P \succ_\Gamma Q$. However, this does not hold for strong partial implication in general. For example, $\models x \wedge y \rightarrow x$ but $x \wedge y \not\succ^S x$. For further discussion on the properties of partial implication, see [13].

5. WEAK PARTIAL IMPLICATION

Even partial implication may be too strong for some applications, and we therefore also consider weak notion of partial implication. A formula P weakly partially implies a formula Q when for all the cases in P, there is a case in Q that is not disjoint, in the sense that some literal part of the case of P is also part of the case of Q. Weak partial implication is a weaker notion of partial implication by allowing possible side effects. Thus, it can be captured simply by $\pi \cap \pi' \neq \emptyset$.

DEFINITION 4 (WEAK PARTIAL IMPLICATION). *We say that a formula P weakly partially implies a formula Q with respect to a formula set Γ, denoted by $P \succ^W_\Gamma Q$, if:*

1. *$PI(\Gamma, P)$ is not empty.*
2. *For each $\pi \in PI(\Gamma, P)$, there exists $\pi' \in PI(\Gamma, Q)$, such that $\pi \cap \pi' \neq \emptyset$.*

We write $P \not\succ^W_\Gamma Q$ if it it not the case that $P \succ^W_\Gamma Q$.

The following example illustrates the apples and oranges example for weak partial implication.

EXAMPLE 4. *x, $x \wedge z$ and $x \wedge \neg y$ weakly partially imply $x \wedge y$ while z does not.*

Theorem 6 shows that partial implication implies weak partial implication.

THEOREM 6 (PARTIAL AND WEAK PARTIAL). *Let P, Q be two formulas and Γ a set of formulas. $P \succ_\Gamma Q$ implies $P \succ^W_\Gamma Q$.*

PROOF. That $P \succ_\Gamma Q$ implies $P \succ^W_\Gamma Q$ follows directly from Definition 3 and 4. □

Theorem 5 and Theorem 6 state the basic relationships among the family of partial implications which coincide with our intuitions. However, the converses of Theorem 5 do not hold in general.

Again, non-triviality (Theorem 1), independency of syntax (Theorem 2), and relevancy (Theorem 3) hold for weak partial implication just as for strong and normal partial implication, and transitivity does not hold for weak partial implication like it does not hold for partial implication. For example, $x \succ^W x \wedge y$ and $x \wedge y \succ^W y \wedge z$ but $x \not\succ^W y \wedge z$.

THEOREM 7. *Let P, Q and R be three formulas and Γ a set of formulas. If $P \succ^W_\Gamma Q$ and $Q \succ^S_\Gamma R$, then $P \succ^W_\Gamma R$.*

PROOF. By Theorem 1, P, Q and R are all non-trivial *w.r.t.* Γ. Let π be a prime implicant of P *w.r.t.* Γ. Since $P \succ^W_\Gamma Q$, there exists a literal set π_1 consistent with Γ such that $\pi_1 \in PI(\Gamma, Q)$ and $\pi \cap \pi_1 \neq \emptyset$. Moreover, since $Q \succ^S_\Gamma R$, there exists a literal set π_2 consistent with Γ such that $\pi_2 \in PI(\Gamma, R)$ and $\pi_1 \subseteq \pi_2$. Thus, $\pi \cap \pi_1 \subseteq \pi \cap \pi_2$. Therefore $\pi \cap \pi_2 \neq \emptyset$. This shows that $P \succ^W_\Gamma R$. □

THEOREM 8. *Let P, Q and R be three formulas and Γ a set of formulas. If $P \succ^W_\Gamma Q$ and $P \wedge R$ is non-trivial w.r.t. Γ, then $P \wedge R \succ^W_\Gamma Q$.*

PROOF. Let π be a prime implicant of $P \wedge R$ *w.r.t.* Γ. We have that $\Gamma \cup \pi \models P \wedge R$. Then, $\Gamma \cup \pi \models P$. By Lemma 1, there is a subset π_1 of π, which is a prime implicant of P *w.r.t.* Γ. Moreover, $P \succ^W_\Gamma Q$. Then, there exists a prime implicant π_2 of Q *w.r.t.* Γ such that $\pi_1 \cap \pi_2 \neq \emptyset$. Thus, $\pi \cap \pi_2 \neq \emptyset$. This shows that $P \wedge R \succ^W_\Gamma Q$. □

Theorem 8 does not hold for either partial implication or strong partial implication. For example, we have that $x \succ^S x \wedge y$ and $x \succ x \wedge y$. But $x \wedge \neg y \not\succ^S x \wedge y$ and $x \wedge \neg y \not\succ x \wedge y$.

Weak partial implication is also an extension of classical implication in the nontrivial cases.

COROLLARY 1. *Let Γ be a set of formulas, P and Q two formulas non-trivial w.r.t. Γ. If $\Gamma \models P \rightarrow Q$ then $P \succ^W_\Gamma Q$.*

PROOF. This assertion follows directly from this property for partial implication [13] and Theorem 6. □

The family of partial implications is related to the notion of relevance [3], but also distinct in various ways. There are two major differences between Lakemeyer's relevance and weak partial implication. Firstly, weak partial implication is defined by a $\forall - \exists$ style (See Definition 4), while Lakemeyer's notion of relevance is defined by a $\exists - \exists$ style. Secondly, the background theory is not considered in Lakemeyer's approach. Another related notion is probabilistic positive relevance, introduced in [14]. According to the definition in [14], $x \vee y$ is positive relevant to x, while $x \vee y \not\succ x$ ($x \vee y \not\succ^W x$, $x \vee y \not\succ^S x$).

Finally, not that the following definition does not capture partial satisfaction. A formula P partially implies a formula Q *w.r.t.* a formula set Γ iff there exists a formula R such that $\Gamma \cup \{R\} \not\models Q$ but $\Gamma \cup \{P, R\} \models Q$. Actually, for every formula P such that $\Gamma \not\models P \rightarrow Q$, there always exists such a formula R (Let R be $\neg\Gamma \vee \neg P \vee Q$). Even if we restrict R with set of consistent literals, this definition cannot capture partial satisfaction either. If so, for instance, $x \vee \neg y$ should partially imply $x \wedge y$ according to this definition since $(x \vee \neg y) \wedge y \models x \wedge y$. However, this conclusion is counter-intuitive.

6. GOAL CHANGE

Since goals play a central role in agent theory, goal change is a central problem in, for example, agent theory and programming, planning, learning, communication and coordination. [2, 12, 11, 9, 8]. Several strategies have been proposed *when* the agents' goals ought to be changed. The first one is so called "blind" [11]. That is, an agent will never change its goal unless it believes this goal has been achieved. Another one is that the agent will never change its goal unless it believes this goal has been achieved or is unachievable [1, 6]. The third strategy introduces a trigger condition for each goal [12], if this condition becomes true then the agent should change that goal associated with this condition. In the fourth strategy, the agent may change its goal if it receives a request from another agent to cancel this goal [9].

Few approaches are concerned with how to change the agents' goals. For almost all of them, the answer is quite simple, that is to drop the goal. Huang and Bell [2] propose to represent the agents' goals with preferences, an then to change the goals according to this preference structure. However, a reasonable way is changing the goal instead of simply dropping it, especially in the case where the agent has to drop his goal due to requirements from other agents.

6.1 From partial goal satisfaction to goal change

How to change the goals of agents depends on the context and application, and we therefore propose a logical framework based on partial implication for goal change in the context of beliefs. The inspiration for our approach is that partial goal satisfaction may be seen as combination of goal change and full goal satisfaction.

> P partially satisfies a goal to achieve Q if the goal to achieve Q can be changed into a goal to achieve Q', and P satisfies the goal to achieve Q', i.e., P logically implies Q'.

The relation between partial goal satisfaction and goal change is illustrated by the following example.

EXAMPLE 5. *$p \wedge r$ partially satisfies the goal to achieve $p \wedge q$, because the goal to achieve $p \wedge q$ can be changed to the goal to achieve p, and $p \wedge r$ logically implies p.*

We therefore characterize goal change analogously to partial goal satisfaction.

> A goal to achieve Q can be changed into a goal to achieve Q' if for all cases of Q', there is a case of Q such that a part of the former is related to a part of the latter.

This suggests that the three notions of partial implication can be used to define goal change operators. Before we consider which notion of partial implication is best suited for goal change, we have to discuss the role of beliefs in goal change.

6.2 Introducing beliefs in goal change

Goal change takes place in the context of the agent's beliefs. In particular, a goal may be changed to a logically unrelated proposition, when the agent believes that the propositions are related. The following definition models this role of beliefs in goal change by representing the beliefs as the background knowledge Γ. We only consider strong partial implication, the other definitions are analogous.

DEFINITION 5. *The goal to achieve Q can be changed to the goal to achieve Q' in the context of a set of beliefs Γ, if and only if $Q' \succ^S_\Gamma Q$.*

The following example illustrates how beliefs can be used in goal change using strong partial implication.

EXAMPLE 6. *A goal to achieve $p \wedge q$ may be changed to a goal to achieve p, represented by $p \succ^S_\emptyset p \wedge q$. When the agent believes that s implies p, then the goal to achieve $p \wedge q$ may also be changed to a goal to achieve s, represented by $s \succ^S_{\{s \to q\}} p \wedge q$.*

The following combination of Example 5 and 6 illustrates the use of beliefs in goal change to relate partial goal satisfaction in the context of beliefs to goal change in the context of beliefs.

EXAMPLE 7. *$s \wedge r$ partially satisfies the goal to achieve $p \wedge q$ in the context of $s \to p$, as represented by $s \wedge r \succ_{\{s \to p\}} p \wedge q$. Moreover, the goal to achieve $p \wedge q$ may also be changed to a goal to achieve s, represented by $s \succ^S_{\{s \to p\}} p \wedge q$. Consequently, $s \wedge r$ partially satisfies the goal to achieve $p \wedge q$ in the context of $s \to p$, because the goal to achieve $p \wedge q$ in the context of $s \to p$ can be changed to the goal to achieve s, and $s \wedge r$ logically implies s. However, if the beliefs are not used in the goal change, and the goal to achieve $p \wedge q$ is changed to the goal to achieve p, then we do not have to $s \wedge r$ logically implies p.*

The following example illustrates that strong partial implication cannot be used when a goal conflicts with a belief (the same holds for the other notions of partial implication).

EXAMPLE 8. *A goal for $p \wedge q$ conflicts with the belief for $\neg q$ and, maybe, therefore $p \wedge q$ may be changed to p, but not to q. However, neither $p \succ^S_{\{\neg q\}} p \wedge q$ nor $q \succ^S_{\{\neg q\}} p \wedge q$ holds, because $p \succ^S_\Gamma q$ implies that both $PI(\{\neg q\}, p)$ and $PI(\{\neg q\}, q)$ are nonempty.*

The latter example can be modeled with a contraction operator '$-$' as developed in the context of belief change, and as it has been studied in the AGM framework. For example, the goal to achieve $p \wedge q$ may be changed to the goal to achieve p in the context of belief $\neg q$, since $Cn(p \wedge q) - q = Cn(p)$. However, in most examples of goal contraction it seems dangerous to change a goal because it conflicts with a belief, because the belief may later turn out to be wrong. In case of conflict between a goal to achieve $p \wedge q$ and a belief $\neg q$, it is better to keep the goal $p \wedge q$, and add one of its subgoals like the goal to achieve p.

6.3 Which partial implication?

If an agent has to change its goal to achieve $p \wedge q$, and it can freely choose how to change its goal, it seems most reasonable to change it to either a goal to achieve p or to a goal to achieve q. Deliberately introducing a side effect such as changing its goal to $p \wedge r$, or even side effects such as $p \wedge \neg q$, seems irrational. For example, if the agent changes its goal to $p \wedge r$, it has to make some effort to see to r, which does not seem to serve any purpose. In other words, in an unconstrained situation, it seems best to use strong partial implication.

However, in other circumstances the other notions of partial implication may be used too. For example, consider the fourth strategy to change goals mentioned at the beginning of this section, in which the agent may change its goal if it receives a request from another agent to cancel this goal [9]. If the other agent requesting the cancelation also gives some suggestions how to change the goal, the agent has to decide which alternative to accept. In such constrained cases, in which the agent has to choose from a set of alternatives, and in which none of the alternatives strongly partially implies the goal, weaker notions of partial implication can be adopted too.

7. COMPLEXITY ISSUES

In this section, we analyze the complexity issues related to all three kinds of partial implications. We assume that the readers are familiar with some basic notions of computational theory. More details can be found in [4]. We recall some complexity classes here:

- DP is the complexity class of all languages L such that $L = L_1 \cap L_2$, where L_1 is in NP and L_2 is in coNP. The canonical problem of DP is $SAT - UNSAT$: a pair of propositional formulas $\langle P, Q\rangle$ is in $SAT - UNSAT$ if and only if P is in SAT and Q is in $UNSAT$.

- $\Sigma_2^P = NP^{NP}$ is the complexity class of all languages that are recognizable in polynomial time by a non-deterministic Turing Machine equipped with an NP oracle. The canonical problem of $\Sigma_2^P = NP^{NP}$ is $2 - QBF$: let X and Y be two disjoint sets of atoms and P a formula such that $Atom(P) \subseteq X \cup Y$, a triple $\langle X, Y, P\rangle$ is in $2 - QBF$ if and only if there exists an assignment π over X such that for all assignments π' over Y, $\pi \cup \pi' \models P$.

- $\Pi_2^P = co\Sigma_2^P = coNP^{NP}$ is the complexity class of all languages whose complementary problems are in Σ_2^P. The canonical Π_2^P problems is $2 - \overline{QBF}$: let X and Y be two disjoint sets of atoms and P a formula such that $Atom(P) \subseteq X \cup Y$, a triple $\langle X, Y, P\rangle$ is in $2 - \overline{QBF}$ if and only if for all assignments π over X, there exists an assignment π' over Y such that $\pi \cup \pi' \models P$.

- Π_3^P is the complexity class of all languages whose complementary problems are recognizable in polynomial time by a non-deterministic Turing Machine equipped with a Σ_2^P oracle. The canonical Π_3^P problems is $3 - \overline{QBF}$: let X, Y and Z be three disjoint sets of atoms and P a formula such that $Atom(P) \subseteq X \cup Y \cup Z$, a quadruple $\langle X, Y, Z, P\rangle$ is in $3 - \overline{QBF}$ if and only if for all assignments π over X, there exists an assignment π' over Y such that for all assignments over π'', $\pi \cup \pi' \cup \pi'' \models P$.

We first start our complexity analysis when the background formula set Γ is empty.

LEMMA 4. *A literal l is in one of the prime implicants of a formula P if and only if $\not\models P|l \rightarrow P| - l$.*

PROOF. "$\Rightarrow$:" Assume that $\models P|l \rightarrow P| - l$ and $l \in \pi$, $\pi \in PI(P)$. Let $\pi_1 = \pi \backslash \{l\}$. Then $\pi_1 \cup \{l\} \models (l \wedge P|l) \vee (-l \wedge P|-l)$. It follows that $\pi_1 \cup \{l\} \models l \wedge P|l$. Therefore $\pi_1 \models P| - l$, which means that $\pi_1 \models P$. Hence, π is not a prime implicant of P, a contradiction.

"$\Leftarrow$:" First, $P|l$ can be satisfied. Suppose that π satisfies $P|l$ but not $P| - l$. By Lemma 1, there exists a subset $\pi_1 \subset \pi$ which is a prime implicant of $P|l$. Therefore $\pi_1 \not\models P$ but $\pi_1 \cup \{l\} \models P$. By Lemma 1, there is a subset $\pi_2 \subset \pi_1 \cup \{l\}$ which is a prime implicant of P. Moreover, $l \in \pi_2$ since $\pi_1 \not\models P$. □

THEOREM 9. *To determine whether a literal l is in at least one of the prime implicants of a formula P is NP complete. To determine whether a literal l partially (or weak partially, strong partially) implies a formula P are all NP complete as well.*

PROOF. Membership of point 1 follows directly from Lemma 4. Hardness of point 1 follows from the fact that P is satisfiable iff x is in at leat one of the prime implicants of $x \wedge P$, where x is a new atom not occurred in P. Point 2 is obvious. □

THEOREM 10. *To determine whether a literal l is in all prime implicants of a formula P is DP complete.*

PROOF. It is easy to prove that l occurring in all prime implicants of P iff P can be satisfied and $\models P \rightarrow l$. It immediately follows that the membership of this assertion. Hardness follows from the fact that P is satisfiable and Q is unsatisfiable iff x is in all prime implicants of $(x \wedge P) \vee (\neg x \wedge Q)$, where x is a new atom. □

LEMMA 5. *Let $\pi = \{l_1, ..., l_k\}$ be a consistent literal set and P a formula. π is a prime implicant of P if and only if $\models P|\pi$ and $\not\models P|\pi_i$ $1 \leq i \leq k$, where $\pi_i = \pi \backslash \{l_i\} \cup \{-l_i\}$, $1 \leq i \leq k$.*

PROOF. "$\Rightarrow$:" By the definition, $\pi \models P$. Therefore $\models P|\pi$. Moreover, $\not\models P|\pi_i$. Otherwise, $\pi \backslash \{l_i\} \models P$, which shows that π is not a prime implicant of P, a contradiction.

"$\Leftarrow$:" Firstly, $\pi \models P$ since $\models P|\pi$. Secondly, by Lemma 1, there exists $\pi' \subset \pi$ such that $\pi' \in PI(P)$. Thus, for all i, $l_i \in \pi'$. Otherwise, $\pi \backslash \{l_i\} \models P$, which means that $\pi_i \models P$, a contradiction. Thus, $\pi' = \pi$. □

THEOREM 11. *To determine whether a literal set π is a prime implicant of a formula P is DP complete.*

PROOF. Membership follows from Lemma 5. For hardness, we construct a reduction from $SAT - UNSAT$. $\langle P, Q\rangle$ is in $SAT - UNSAT$ if and only if $\{x\}$ is a prime implicant of $(x \vee \neg P) \wedge (\neg x \vee \neg Q)$, where x is a new atom. □

THEOREM 12. *To determine whether a literal set π weakly partially implies a formula P is NP complete.*

PROOF. Hardness follows directly from Theorem 9. For membership, let $\pi = \{l_1, ..., l_k\}$. Then π weakly partially implies P if and only if π is consistent and there exists $l_i, 1 \leq i \leq k$ such that l_i is in one of the prime implicant of P. By Theorem 9, this problem is in NP. □

LEMMA 6. *A literal set π partially implies a formula P iff there is an assignment π_1 over $Atom \backslash Atom(\pi)$ and an assignment π_2 over $Atom(\pi)$ such that $\pi \cup \pi_1 \models P$ and $\pi_1 \cup \pi_2 \models \neg P$.*

PROOF. "$\Rightarrow$:" By Definition 3, there is a prime implicant π' of P such that $\pi \cap \pi' \neq \emptyset$ and $\pi \cap -\pi' = \emptyset$. Let $l \in \pi' \cap \pi$. Then $\pi' \backslash \{l\} \cup \{-l\} \not\models P$. It can be extended into an assignment π_0 over $Atom$, which satisfies $\neg P$. Let $\pi_1 \subseteq \pi_0$ and $Atom(\pi_1) = Atom \backslash Atom(\pi)$; let $\pi_2 \subseteq \pi_0$ and $Atom(\pi_2) = Atom(\pi)$. Clearly, $\pi \cup \pi_1 \models P$ and $\pi_1 \cup \pi_2 \models \neg P$.

"$\Leftarrow$:" By Lemma 1, there is a prime implicant π' of P such that $\pi' \subseteq \pi \cup \pi_1$. It follows that $\pi' \cap \pi \neq \emptyset$ and $\pi' \cap -\pi = \emptyset$. Hence, π partially implies P. □

THEOREM 13. *To determine whether a literal set π partially implies a formula P is NP complete.*

PROOF. Hardness follows directly from Theorem 9. For membership, the following algorithm determines whether π partially implies P: 1. guess π_1, π_2 in Lemma 6; 2. check the conditions in Lemma 6. Step 2 can be done in polynomial time. □

THEOREM 14. *To determine whether a literal set π strongly partially implies a formula P is Σ_2^P complete.*

PROOF. "Membership:" The following algorithm determines whether π strongly partially implies P: 1. guess a consistent literal set π'; 2. check whether π' is a prime implicant of P; 3. if yes, check that whether π is a subset of π'. By Theorem 11, step 2 requires an NP oracle. Hence, this problem is in Σ_2^P.

"Hardness:" Reduction from $2-QBF$. Let X and Y be two disjoint sets of atoms and P a formula such that $Atom(P) \subseteq X \cup Y$. Let $Y = \{y_1, y_2, ..., y_k\}$; T_1 be $y_1 \vee ... \vee y_k$; T_2 be $\neg y_1 \vee ... \vee \neg y_k$; x and y be two new atoms different with $X \cup Y$; Q be $(x \wedge y \wedge P) \vee (x \wedge \neg y \wedge T_1) \vee (\neg x \wedge y \wedge T_2)$. It is easy to prove that $\exists X \forall Y P$ holds if and only if $x \wedge y$ strongly partially implies Q. □

THEOREM 15. *To determine whether a formula P partially implies another formula Q is Π_2^P complete.*

PROOF. "Membership:" The following algorithm determines whether P does not partially imply Q: 1. guess a consistent literal set π; 2. check whether π is a prime implicant of P; 3. if yes, check that whether π does not partially imply Q. By Theorem 11 and Theorem 13, step 2 and step 3 requires an NP oracle respectively. Hence, this problem is in Π_2^P.

"Hardness:" Reduction from $2-\overline{QBF}$. Let X and Y be two disjoint sets of atoms and P a formula such that $Atom(P) \subseteq X \cup Y$. Let $X = \{x_1, x_2, ..., x_k\}$; $X' = \{x'_1, x'_2, ..., x'_k\}$ be k new atoms different with $Atom$. Formula K is $(x_1 \leftrightarrow x'_1) \wedge (x_2 \leftrightarrow x_2) \wedge ... \wedge (x_k \leftrightarrow x_k)$. It is easy to prove that $\forall X \exists Y P$ holds if and only if $x \wedge K$ partially implies $x \wedge P$, where x is a new atom. □

THEOREM 16. *To determine whether a formula P weakly partially implies another formula Q is Π_2^P complete.*

PROOF. "Membership:" The following algorithm determines whether P does not weakly partially imply Q: 1. guess a consistent literal set π; 2. check whether π is a prime implicant of P; 3. if yes, check that whether π does not weakly partially imply Q. By Theorem 11 and Theorem 12, step 2 and step 3 requires an NP oracle respectively. Hence, this problem is in Π_2^P.

"Hardness:" reduction from $2-\overline{QBF}$. Let X and Y be two disjoint sets of atoms and P a formula such that $Atom(P) \subseteq X \cup Y$. Let $Y = \{y_1, y_2, ..., y_k\}$. Formula T is $(y_1 \vee ... \vee y_k) \wedge (\neg y_1 \vee ... \vee \neg y_k)$. It is easy to prove that $\exists X \forall Y P$ holds if and only if P does not weakly partially imply T. □

THEOREM 17. *To determine whether a formula P strongly partially implies another formula Q is Π_3^P complete.*

PROOF. "Membership:" The following algorithm determines whether P does not strongly partially imply Q: 1. guess a consistent literal set π; 2. check whether π is a prime implicant of P; 3. if yes, check that whether π doesn't strongly partially implies Q. By Theorem 11, step 2 requires an NP oracle; by Theorem 14, step 3 requires an Σ_2^P oracle. Hence, this problem is in Π_3^P.

"Hardness:" Reduction from $3-\overline{QBF}$. Let X, Y and Z be three disjoint sets of atoms and P a formula such that $Atom(P) \subseteq X \cup Y \cup Z$. Let $X = \{x_1, x_2, ..., x_k\}$; $X' = \{x'_1, x'_2, ..., x'_k\}$ be k new atoms; $Z = \{z_1, z_2, ..., z_k\}$; T_1 be $z_1 \vee ... \vee z_k$; T_2 be $\neg z_1 \vee ... \vee \neg z_k$; K be $(x_1 \leftrightarrow x'_1) \wedge (x_2 \leftrightarrow x_2) \wedge ... \wedge (x_k \leftrightarrow x_k)$, where x be y are two new atoms; R be $x \wedge y \wedge K$, Q be $(x \wedge y \wedge P \wedge K) \vee (x \wedge \neg y \wedge T_1 \wedge K) \vee (\neg x \wedge y \wedge T_2 \wedge K)$.

We will prove that $\forall X \exists Y \forall Z P$ holds if and only if R strongly partially implies Q. Suppose that $\forall X \exists Y \forall Z P$ holds. Given a prime implicant of R, which is $\{x, y\} \cup \pi \cup \pi'$. By the assumptions, there exists an assignment π_1 over Y such that $\pi \cup \pi_1 \models P$. Then $\{x, y\} \cup \pi \cup \pi' \cup \pi_1 \models x \wedge y \wedge P \wedge K$. It follows that $\{x, y\} \cup \pi \cup \pi' \cup \pi_1 \models Q$. By Lemma 1, there exists a subset π_2 of $\{x, y\} \cup \pi \cup \pi' \cup \pi_1$, which is a prime implicant of Q. We have that $x \in \pi_2$ (otherwise $\{y\} \cup \pi \cup \pi' \cup \pi_1 \models Q$. Therefore $\{y\} \cup \pi \cup \pi' \cup \pi_1 \models x \vee T_2$, a contradiction). Symmetrically, $y \in \pi_2$. Moreover, for each atom $l \in (\pi \cup \pi')$, we have that $l \in \pi_2$ since $\pi_2 \models K$. Therefore $\{x, y\} \cup \pi \cup \pi' \subseteq \pi_2$. It means that for all prime implicants π_3 of R, there exists a prime implicant π_4 of Q such that $\pi_3 \subseteq \pi_4$. Hence, R strongly partially implies Q.

On the other hand, suppose that R strongly partially implies Q. Then for all assignments π over X, $\{x, y\} \cup \pi \cup \pi'$ is a prime implicant of R. By the assumptions, there is a prime implicant of Q, which contains $\{x, y\} \cup \pi \cup \pi'$. Let it be $\{x, y\} \cup \pi \cup \pi' \cup \pi_1$, where $\pi_1 \cap (X \cup -X) = \emptyset$. Then $\pi_1 \not\models T_1$ (otherwise $\{x\} \cup \pi \cup \pi' \cup \pi_1 \models Q$, a contradiction). Symmetrically, $\pi_1 \not\models T_2$. Hence $\pi_1 \cap (Z \cup -Z) = \emptyset$. It follows that $\pi_1 \subseteq (Y \cup -Y)$. It can be extended to an assignment π_3 over Y, which satisfies $\{x, y\} \cup \pi \cup \pi' \cup \pi_3 \models Q$. Therefore $\{x, y\} \cup \pi \cup \pi' \cup \pi_3 \models x \wedge y \wedge P \wedge K$, $\pi \cup \pi_3 \models P$. It means that for all assignments π over X, there is an assignment π_3 over Y, such that for all assignments π_4 over Z, $\pi \cup \pi_3 \cup \pi_4 \models P$. □

We now face to the cases with background formula sets.

LEMMA 7. *Let Γ be a set of formulas, P a formula and π a set of literals. $\pi \in PI(\Gamma, P)$ if and only if π is consistent with Γ and $\pi \in PI(\Gamma \to P)$.*

PROOF. Suppose that π is consistent with Γ. $\pi \in PI(\Gamma, P)$ iff a) $\Gamma \cup \pi \models P$ and b) there is no subset π' of π such that $\Gamma \cup \pi' \models P$ iff a) π is a model of $\Gamma \to P$ and b) there is no subset π' of π such that $\pi' \models \Gamma \to P$ iff $\pi \in PI(\Gamma \to P)$. □

THEOREM 18. *To determine whether a literal set π is a prime implicant of a formula P w.r.t. a formula set Γ is DP complete.*

PROOF. Membership follows from Lemma 7 and Theorem 11. Hardness follows from Theorem 11. □

THEOREM 19. *To determine whether a literal l is in at least one of the prime implicants of a formula P w.r.t. a formula set Γ is Σ_2^P complete.*

PROOF. Membership is easy by guessing a literal set π and checking if $l \in \pi$ and $\pi \in PI(\Gamma, P)$. For hardness, we will show that $\exists X \forall Y P$ iff x is in one of the elements in $PI(\Gamma, F)$, where $\Gamma = \neg(x \wedge \wedge P \wedge (y_1 \vee \cdots \vee y_k))$, $F = x \wedge P \wedge \neg(y_1 \vee \cdots \vee y_k) \vee \neg x \wedge (\neg y_1 \vee \cdots \vee \neg y_k)$ and x is a new atom.
x is in one of the elements of $PI(\Gamma, F)$
iff
$\exists\, \pi, x \in \pi, \pi \not\models \neg\Gamma, \pi \models \Gamma \to F$ and $\forall \pi' \subset \pi, \pi' \not\models \Gamma \to F$.
iff
$\exists\, \pi_1, \pi_1 \cup \{x\} \not\models \neg\Gamma, \pi_1 \cup \{x\} \models \Gamma \cup \neg\Gamma \vee F$ and $\pi_1 \not\models \neg\Gamma \vee F$.
iff
$\exists\, \pi_1, \pi_1 \not\models (\neg\Gamma)|x, \pi_1 \models (\neg\Gamma)|x \vee F|x$ and $\pi_1 \not\models (\neg\Gamma)|\neg x \vee F|\neg x$.
iff
$\exists\, \pi_1, \pi_1 \not\models F \wedge (y_1 \vee \cdots \vee y_k), \pi_1 \models P$ and $\pi_1 \not\models \neg y_1 \vee \cdots \vee \neg y_k$.
iff
$\exists X \forall Y P$. □

THEOREM 20. *To determine whether a formula P strongly partially implies another formula Q w.r.t. a formula set Γ is Π_3^P complete.*

PROOF. Hardness follows from Theorem 17. Given a literal set π, it is easy to see that checking whether there exists a prime implicant π' of a formula P *w.r.t.* a formula set Γ such that $\pi \subseteq \pi$ is in Σ_2^P. Thus, it follows directly that this problem is in Π_3^P. □

8. CONCLUSION

Partial goal satisfaction has been formalized as three kinds of partial implication with respect to a background theory. Roughly, while ignoring the borderline conditions as well as the role of the background theory, the three kinds of partial implication have been defined as follows.

P **strongly partially implies** Q if for every case of P there is a case of Q such that the case of P is a subset of the case of Q. If there is a part of a case P which does not occur in any case of Q, which we call redundancy, then P does not strongly partially imply Q.

P **partially implies** Q if for every case of P there is a case of Q such that a part of the case of P is a part of the case of Q, and there is no part of the case of P which conflicts with the case of Q. There may be redundancy. If there is a part of a case P which conflicts with any case of Q, which we call a side effect, then P does not partially imply Q.

P **weakly partially implies** Q if for every case of P there is a case of Q such that a part of the case of P is a part of the case of Q. There may be redundancy or side effects in P.

We show that the properties in Table 2 hold for the three notions of partial implication. Moreover, we show relations between the three kinds of partial implication, for example that strong partial implication implies partial implication, and that partial implication implies weak partial implication.

-	strong partial	partial	weak partial
Non triviality	yes	yes	yes
Independence of syntax	yes	yes	yes
Relevancy	yes	yes	yes
Transitivity	yes	no	no
Extension classical	no	yes	yes
Left strengthening	no	no	yes

Table 2: Properties of partial implication

We define the problem of changing a goal by relating the problem to partial goal satisfaction: P partially satisfies a goal to achieve Q if the goal to achieve Q can be changed into a goal to achieve Q', and P satisfies the goal to achieve Q', i.e., P logically implies Q'. We show how to use partial implication for goal change in the context of belief, by defining that the goal to achieve Q can be changed to the goal to achieve Q' in the context of a set of beliefs Γ, if and only if $Q' \succ_{\Gamma}^{S} Q$. We show that beliefs can be used only if they do not conflict with the goals, and we argue that in case of conflict other techniques might be used. We argue also that in unconstrained cases of goal change strong partial implication is the most intuitive alternative, but in more constrained applications the other notions of goal change may be used too.

All the complexity results addressed in this paper are on the first three levels of polynomial heterarchy (i.e. from NP complete to Π_3^P complete). Surprisingly, checking strong equivalence between two formulas is Π_3^P complete, even when the background formula set is empty. Our complexity results shows that checking partial implications with background formula set are often harder than that without backgrounds. One of the unexpected result is Theorem 19, which states that literal checking in general case is even Σ_2^P complete. Our complexity results also shows that checking partial implication and checking weak partial implication have almost the same difficulties, while checking strong partial implication is sometimes more complex. The complexity results of checking weak partial implication and partial implication in general case (i.e. with background formula set) have not been addressed in this paper and are left for further research. We believe both of them are in the third level of polynomial hierarchy. More precisely, we believe that checking weak partial implication in general case is $\Delta_3^P O(log\ n)$ complete while checking partial implication in general case is Π_3^P complete.

For future work, the notion of partial implication can be extended to the case between actions (or action sequences) and propositions. Moreover, the consequences of goal change on plan and intention reconsideration approaches could be considered. At this moment, most of these approaches abort the goals, ignore logical connections between goals, do not consider goals as propositional formulas, and are restricted to an agent language. Since our work can be embedded into those approaches since our work is independent of agent languages (we only use the logic connections). Also the application of the partial implication to other related problems about reasoning about goals should be studied, such as subgoal generation: create subgoals to achieve a main goal, typically using a goal hierarchy. For such other applications of partial implication, the question can be raised about the role of beliefs, how partial implication can be used, which kind of partial implication may be used in which cases, and whether other kind of partial implication can be defined.

9. REFERENCES

[1] P. R. Cohen and H. J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42(2-3):213–261, 1990.
[2] Z. Huang and J. Bell. Dynamic goal hierarchies. In *Working Papers of the AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, pages 9–17, 1997.
[3] G. Lakemeyer. A logical account of relevance. In *Proceedings of IJCAI'95*, pages 853–861, 1995.
[4] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
[5] W. Quine. The problem of simplifying truth functions. *American Mathematical Monthly*, 59(8):521–531, 1952.
[6] A. S. Rao and M. P. Georgeff. An abstract architecture for rational agents. In *Proceedings of KR'92*, pages 439–449, 1992.
[7] S. Sardina and L. Padgham. Goals in the context of bdi plan failure and planning. In *Proceedings of AAMAS'07*, 2007.
[8] S. Shapiro and G. Brewka. Dynamic interactions between goals and beliefs. In *Proceedings of IJCAI'07*, pages 2625–2630, 2007.
[9] S. Shapiro, Y. Lespérance, and H. J. Levesque. Goal change. In *Proceedings of IJCAI'05*, pages 582–588, 2005.
[10] J. Thangarajah, J. Harland, D. Morley, and N. Yorke-Smith. Aborting goals and plans in bdi agents. In *Proceedings of the AAMAS'07*, 2007.
[11] B. van Riemsdijk, M. Dastani, F. Dignum, and J. C. Meyer. Dynamics of declarative goals in agent programming. In *Proceedings of DALT'2004, LNCS 3476*, pages 1–18, 2004.
[12] M. Winikoff, L. Padgham, J. Harland, and J. Thangarajah. Declarative & procedural goals in intelligent agent systems. In *Proceedings of KR'02*, pages 470–481, 2002.
[13] Y. Zhou and X. Chen. Partial implication semantics for desirable propositions. In *Proceedings of KR'04*, pages 606–612, 2004.
[14] Y. Zhou and X. Chen. Toward formalizing usefulness in propositional language. In *Proceedings of KSEM'06, LNAI 4092*, pages 650–661, 2006.