

# On the Semantics of Conditional Commitment

Shakil M. Khan and Yves Lespérance

Department of Computer Science and Engineering  
York University, Toronto, ON, Canada M3J 1P3  
{skhan, lesperan}@cs.yorku.ca

**Abstract.** In this paper, we identify some problems with current formalizations of conditional commitments, i.e. commitments to achieve a goal if some condition becomes true. We present a solution to these problems. We also formalize two types of communicative actions that can be used by an agent to request another agent to achieve a goal or perform an action provided that some condition becomes true. Our account is set within ECASL [7], a framework for modeling communicating agents based on the situation calculus.

## 1 Introduction

In recent years, the importance of agent communication in multiagent systems has been widely recognized. As a result, many researchers have developed communicative multiagent frameworks [3, 5, 6, 14, 17, 22, 23] and attempted to formalize various types of communicative actions in these frameworks. One important concept in these is the notion of conditional commitment. A conditional commitment is a commitment to achieve some goal if some condition becomes true (e.g. a commitment to ship some goods when payment of an agreed to amount arrives). Conditional requests are requests that seek to have the addressee acquire a conditional commitment. Any multiagent framework that deals with negotiation and cooperation ought to handle conditional commitments. Unfortunately, most definitions found in the literature ([5, 26, 22, 1], for example) are inadequate: they either define conditional commitments as disjunctive goals, which makes the agents under-committed to the conditional goal, or define them as conjunctive goals, which renders the agent over-committed.

We will go over some examples to point out the problems associated with the disjunctive and the conjunctive accounts of conditional commitment. In these, we use the following modal operators:  $\diamond\phi$ , i.e.  $\phi$  eventually holds,  $\text{Happens}(\alpha)$ , i.e. the action  $\alpha$  is performed next,  $\phi \text{ Until } \psi$ , i.e. eventually  $\psi$  becomes true, and as long as  $\psi$  is false,  $\phi$  holds, and  $\text{Before}(\psi, \phi)$ , i.e.  $\phi$  becomes true before  $\psi$  does. The formal semantics of these operators are given in Section 2.

In the disjunctive account, a conditional commitment to achieve some goal provided that some condition holds is modeled as a commitment to achieve the goal if the condition holds, i.e. as a simple material implication. For example, consider an online marketplace domain. Suppose that there are two agents, a seller agent *slr*, and a buyer agent *byr*. If we use a disjunctive account, *slr*'s

conditional commitment to ship some goods to *byr* on the condition that *byr* pays can be modeled as follows:

$$\begin{aligned} \text{CondInt}_{\text{dis}}(slr, \text{GetPaid}, \text{Happens}(\text{shipGoods}(slr, byr))) \doteq \\ \text{Int}(slr, \neg \diamond \text{GetPaid} \vee \\ [\neg \text{GetPaid} \text{ Until } (\text{GetPaid} \wedge \text{Happens}(\text{shipGoods}(slr, byr)))])) \end{aligned}$$

This says that *slr*'s conditional commitment to ship the goods when *byr* pays amounts to *slr* having the intention that *byr* eventually pays and after that she ships the goods, if *byr* eventually pays (as mentioned above, we assume that the Until construct above implies that  $\diamond \text{GetPaid}$  and  $\diamond \text{Happens}(\text{shipGoods}(slr, byr))$ ). One problem with this account of conditional intention is that there is a counter-intuitive way to satisfy a conditional intention, namely, the agent may commit to the triggering condition remaining false and deliberately perform some action that makes it remain false. Thus, in the example, to satisfy her conditional intention, *slr* may intentionally perform some action to stop *byr* from paying her, such as blocking debits from *byr*. In other words, there is nothing in this formalization of conditional intention that stops *slr* from intending not to get paid and not to send the goods. However, this is counter-intuitive and a model of conditional commitment should not support this. Thus, with the disjunctive account of conditional commitment, the agent seems under-committed to the goal. Some examples in the literature that formalize conditional commitments as disjunctive goals include [22] and [1].

In the conjunctive account, a conditional commitment to achieve a goal provided that a condition holds is modeled as a temporally ordered conjunctive commitment to the triggering condition and the conditional goal, where the triggering condition is achieved first. Although this model may seem appropriate in many cases, it often leads to problems. For example, suppose that *slr* has the conditional commitment to ship a replacement unit provided that *byr* reports and returns a defective good. If we use a conjunctive account, this can be modeled as follows:

$$\begin{aligned} \text{CondInt}_{\text{con}}(slr, \text{DefGoodRet}, \text{Happens}(\text{shipRepl}(slr, byr))) \doteq \\ \text{Int}(slr, \text{Before}(\text{Happens}(\text{shipRepl}(slr, byr)), \text{DefGoodRet}) \\ \wedge \diamond \text{Happens}(\text{shipRepl}(slr, byr))). \end{aligned}$$

This says that *slr*'s conditional commitment to ship a replacement unit provided that *byr* returns a defective good can be modeled as *slr*'s intention that eventually *byr* returns a defective good, and *slr* ships the replacement unit after that. Note that, according to this definition, since *slr* has the intention that the defective product is returned before she ships the replacement unit, and that she eventually ships the replacement units, it follows that *slr* has the intention that *byr* eventually returns a product, i.e.  $\text{Int}(slr, \diamond \text{DefGoodRet})$ . So *slr* may deliberately perform some action, such as shipping a defective good in the first place, to achieve this intention. Thus, the conjunctive account of conditional commitment results in over-committed agents. Both [5] and [26] formalize conditional commitments as conjunctive goals.

In this paper, we propose a solution to these problems (the *under/over-commitment problems*, henceforth). Our solution involves using an additional constraint with the disjunctive account to eliminate the under-commitment problem. We use the Extended Cognitive Agent Specification Language (ECASL) [7] as our base formal framework for this. Our account is formulated for internal/mental states semantics for communication acts. Nevertheless, similar issues arise for public/social-commitment semantics (as discussed in Section 5). In what follows, we will use the terms ‘intention’ and ‘commitment’ interchangeably. We also discuss some simple properties of conditional commitment, and formalize two communicative action types that can be used by an agent to request another agent to achieve a goal, or perform an action, provided that some condition holds.

The paper is organized as follows: in the next section, we outline the ECASL framework. In Section 3, we present our model of conditional commitment and discuss some of its properties. In Section 4, we present some communicative acts that allow agents to make requests that result in conditional commitments. In Section 5, we compare our approach to previous work on conditional commitments. Finally in Section 6, we summarize our results and discuss possible future work.

## 2 ECASL

The Extended Cognitive Agent Specification Language (ECASL) [7], an extension of CASL [20, 21], is a framework for specifying and verifying complex communicating multiagent systems that incorporates a formal model of means-ends reasoning. In this section, we outline the part of ECASL that is needed for our formalization of conditional commitment.

In ECASL, agents are viewed as entities with mental states, i.e., knowledge and goals, and the specifier can define the behavior of the agents in terms of these mental states. ECASL combines a declarative action theory defined in the situation calculus with a rich programming/process language, ConGolog [4]. Domain dynamics and agents’ mental states are specified declaratively in the theory, while the agents’ behavior is specified procedurally in ConGolog.

In ECASL, a dynamic domain is represented using an action theory [16] formulated in the situation calculus [11], a (mostly) first order language for representing dynamically changing worlds in which all changes are the result of named actions. ECASL uses a theory  $D$  that includes the following set of axioms:

- action precondition axioms, one per action,
- successor state axioms (SSA), one per fluent, that encode both effect and frame axioms and specify exactly when the fluent changes [15],
- initial state axioms describing what is true initially including the mental states of the agents,
- axioms identifying the agent of each action,
- unique name axioms for actions, and
- domain-independent foundational axioms describing the structure of situations [9].

Within ECASL, the behavior of agents is specified using the notation of the logic programming language ConGolog [4], the concurrent version of Golog [10]. A typical ConGolog program is composed of a sequence of procedure declarations, followed by a complex action. Complex actions can be composed using constructs that include primitive actions ( $a$ ), waiting for a condition ( $\phi?$ ), sequence ( $\delta_1; \delta_2$ ), nondeterministic branch ( $\delta_1 \mid \delta_2$ ), nondeterministic choice of arguments ( $\pi x. \delta$ ), conditional branching (If  $\phi$  Then  $\delta_1$  Else  $\delta_2$  EndIf), while loop (While  $\phi$  Do  $\sigma$  EndWhile), and procedure call ( $\beta(\vec{p})$ ). Intuitively,  $\pi x. \delta$  nondeterministically picks a binding for the variable  $x$  and performs the program  $\delta$  for this binding of  $x$ . ConGolog also supports nondeterministic iteration, concurrent execution with and without priorities, and interrupts. To deal with multiagent processes, primitive actions in ECASL take the agent of the action as argument.

The semantics of the ConGolog process description language is defined in terms of *transitions*, in the style of structural operational semantics [13]. Two special predicates *Final* and *Trans* are introduced, and are characterized by defining axioms for each of the above constructs, where  $Final(\delta, s)$  means that program  $\delta$  may legally terminate in situation  $s$ , and where  $Trans(\delta, s, \delta', s')$  means that program  $\delta$  in situation  $s$  may legally execute one step, ending in situation  $s'$  with program  $\delta'$  remaining.<sup>1</sup> The overall semantics of a program is specified by the *Do* relation:

$$Do(\delta, s, s') \doteq \exists \delta' \cdot (Trans^*(\delta, s, \delta', s') \wedge Final(\delta', s')).$$

$Do(\delta, s, s')$  holds if and only if  $s'$  can be reached by performing a sequence of transitions starting with program  $\delta$  in  $s$ , and the remaining program  $\delta'$  may legally terminate in  $s'$ . Here,  $Trans^*$  is the reflexive transitive closure of the transition relation  $Trans$ .

ECASL incorporates a branching time temporal logic, where each situation has a linear past and a branching future. In this framework, one can write both state formulas and path formulas. A state formula  $\phi(s)$  takes a single situation as argument and is evaluated with respect to that situation. On the other hand, a path formula  $\psi(s_1, s_2)$  takes two situations as arguments and is evaluated with respect to the interval (finite path)  $[s_1, s_2]$ . We often use  $\phi$  (and  $\psi$ ) to denote a formula whose fluents may contain a placeholder constant *now* (*now* and *then*, resp.) that stands for the situation in which  $\phi$  ( $\psi$ , resp.) must hold.  $\phi(s)$  (and  $\psi(s_1, s_2)$ ) is the formula that results from replacing *now* with  $s$  (*now* and *then* with  $s_1$  and  $s_2$ , resp.). Where the intended meaning is clear, we sometimes suppress the placeholder(s).

ECASL allows the specifier to model agents in terms of their mental states by including operators to specify agents' information (i.e., their knowledge), and motivation (i.e., their goals or intentions). We usually use state formulas within the scope of knowledge, and path formulas within the scope of intentions.

---

<sup>1</sup> Since we have predicates that take programs as arguments, we need to encode programs and formulae as first-order terms as in [4]. For notational simplicity, we suppress this encoding and use formulae and programs as terms directly.

Following [12, 18], ECASL models knowledge using a possible worlds account adapted to the situation calculus.  $K(agt, s', s)$  is used to denote that in situation  $s$ ,  $agt$  thinks that she could be in situation  $s'$ .  $s'$  is called a  $K$ -*alternative situation* for  $agt$  in  $s$ . Using  $K$ , the knowledge or belief of an agent,  $\text{Know}(agt, \phi, s)$ , is defined as  $\forall s'(K(agt, s', s) \supset \phi(s'))$ , i.e.  $agt$  knows  $\phi$  in  $s$  if  $\phi$  holds in all of  $agt$ 's  $K$ -accessible situations in  $s$ . In ECASL,  $K$  is constrained to be reflexive, transitive, and Euclidean in the initial situation to capture the fact that agents' knowledge is true, and that agents have positive and negative introspection. As shown in [18], these constraints then continue to hold after any sequence of actions since they are preserved by the successor state axiom for  $K$ .

ECASL supports knowledge expansion as a result of sensing actions [18] and some *informing* communicative actions. Here, we restrict our discussion to knowledge expansion as a result of *inform* actions. The preconditions of *inform* are as follows:

$$\begin{aligned} \text{Poss}(\text{inform}(\text{inf}, \text{agt}, \phi), s) &\equiv \text{Know}(\text{inf}, \phi, s) \\ &\wedge \neg \text{Know}(\text{inf}, \text{Know}(\text{agt}, \phi, \text{now}), s). \end{aligned}$$

In other words, the agent *inf* can inform *agt* that  $\phi$ , iff *inf* knows that  $\phi$  currently holds, and does not believe that *agt* currently knows that  $\phi$ . The successor state axiom (SSA) for  $K$  is defined as follows:

$$\begin{aligned} K(agt, s^*, \text{do}(a, s)) &\equiv \exists s'. [K(agt, s', s) \wedge s^* = \text{do}(a, s') \wedge \text{Poss}(a, s') \wedge \\ &\quad \forall \text{inf}, \phi. (a = \text{inform}(\text{inf}, \text{agt}, \phi) \supset \phi(s'))]. \end{aligned}$$

This says that after an action happens, every agent learns that it was possible and has happened. Moreover, if the action involves someone informing *agt* that  $\phi$  holds, then *agt* knows this afterwards. Note that this axiom only handles knowledge expansion, not revision.

ECASL also incorporates goal expansion and a limited form of goal contraction. Goals or intentions are modeled using an accessibility relation  $W$  over possible situations. The  $W$ -accessible situations for an agent are the ones where she thinks that all her goals are satisfied.  $W$ -accessible situations may include situations that the agent thinks are impossible, unlike Cohen and Levesque's [2]  $G$ -accessible worlds. But intentions are defined in terms of the more primitive  $W$  and  $K$  relations so that the intention accessible situations are  $W$ -accessible situations that are also compatible with what the agent knows, in the sense that there is a  $K$ -accessible situation in their history. This guarantees that agents' intentions are realistic, that is, agents can only intend things that they believe are possible. Thus we have:

$$\text{Int}(agt, \psi, s) \doteq \forall s', s^*. [W(agt, s^*, s) \wedge K(agt, s', s) \wedge s' \leq s^*] \supset \psi(s', s^*).$$

This means that the intentions of an agent in  $s$  are those formulas that are true for all intervals between situations  $s'$  and  $s^*$  where the situations  $s^*$  are  $W$ -accessible from  $s$  and have a  $K$ -accessible situation  $s'$  in their history. Intentions are future oriented, and any goal formula will be evaluated with respect to a

finite path defined by a pair of situations, a current situation *now* and an ending situation *then*. This formalization of goals can deal with both achievement goals and maintenance goals. An achievement goal  $\phi$  is said to be eventually satisfied if  $\phi$  holds in some situation between *now* and *then*, i.e., if  $\diamond(\phi, now, then)$ , which is defined as  $\exists s'. (now \leq s' \leq then \wedge \phi(s'))$ .<sup>2</sup> In [19], Shapiro showed how positive and negative introspection of intentions can be modeled by placing some constraints on  $K$  and  $W$ . To make sure that agents' wishes and intentions are consistent,  $W$  is also constrained to be serial.

ECASL provides an intention transfer communication action, *request*, which is defined in terms of *inform*. This is similar to Herzig and Longin's account [6], where a request is defined as informing about one's intentions, and the requested goals are adopted via cooperation principles. The *request* action can be used by an agent to request another agent to achieve some state of affairs. Formally, we have:

$$request(req, agt, \phi) \doteq inform(req, agt, Int(req, \phi, now)).$$

The SSA for  $W$  which handles intention change in ECASL, has the same structure as a SSA for a domain dependent fluent. In the following,  $W^+(agt, a, s^*, s)$  ( $W^-(agt, a, s^*, s)$ , respectively) denotes the conditions under which  $s^*$  is added to (dropped from, respectively)  $W$  as a result of the action  $a$ :

$$W(agt, s^*, do(a, s)) \equiv W^+(agt, a, s^*, s) \vee (W(agt, s^*, s) \wedge \neg W^-(agt, a, s^*, s)).$$

An agent's intentions are expanded when it is requested something by another agent. After the  $request(req, agt, \psi)$  action, *agt* adopts the goal that  $\psi$ , unless she has a conflicting goal or is not willing to serve *req* for  $\psi$ . Therefore, this action should cause *agt* to drop any paths in  $W$  where  $\psi$  does not hold. This is handled in  $W^-$ :

$$\begin{aligned} W^-(agt, a, s^*, s) \doteq & [\exists req, \psi. a = request(req, agt, \psi) \\ & \wedge Serves(agt, req, \psi, s) \wedge \neg Int(agt, \neg\psi, s) \\ & \wedge \exists s'. K(agt, s', s) \wedge s' \leq s^* \wedge \neg\psi(do(a, s'), s^*)]. \end{aligned}$$

A limited form of intention contraction is also handled in ECASL. Agents intentions are contracted as a result of a *cancelRequest* action. ECASL also incorporates a formal model of means-ends reasoning and commitment to rational plans to achieve intentions. See [7] for details.

Table 1 shows some abbreviations that will be used throughout the paper.

---

<sup>2</sup> We sometimes use  $\diamond$  with a path formula  $\psi$  argument, in which case, we mean that  $\psi$  holds over some interval  $[s, then]$  that starts at some situation  $s$  between *now* and *then*; see Table 1 for the formal definition.

**Table 1.** Some Definitions of Temporal Operators

1.  $\diamond(\psi, now, then) \doteq \exists s'. now \leq s' \leq then \wedge \psi(s', then)$ ,
2.  $\square(\psi, now, then) \doteq \neg \diamond(\neg \psi, now, then)$ ,
3.  $[\phi \text{ Until } \psi](now, then) \doteq \exists s'. now \leq s' \leq then \wedge \psi(s', then)$   
 $\quad \wedge \forall s''. now \leq s'' < s' \supset \phi(s'')$ ,
4.  $\text{Before}(\psi, \phi, now, then) \doteq \exists s'. now \leq s' \leq then \wedge \psi(s', then) \supset$   
 $\quad \exists s''. now \leq s'' < s' \wedge \phi(s'')$ ,
5.  $E_{\diamond}(\phi, now) \doteq \exists s. now \leq s \wedge \phi(s)$ ,
6.  $A_{\square}(\phi, now) \doteq \neg E_{\diamond}(\neg \phi, now)$ ,
7.  $\text{Happens}(a, now, then) \doteq do(a, now) \leq then$ ,
8.  $\text{Happens}_C(\delta, now, then) \doteq \exists s'. s' \leq then \wedge Do(\delta, s', then)$ .

### 3 Conditional Commitments

Having presented our framework, we now return to our discussion about conditional commitments. Informally, an agent *agt* has a conditional commitment or intention that  $\psi$  on the condition that  $\phi$  if *agt* intends to achieve  $\psi$  as soon as the condition  $\phi$  holds. In our specification, we assume that  $\phi$  is an achievement goal, whereas  $\psi$  can be any kind of goal (achievement, maintenance, etc.). In other words, the trigger condition  $\phi$  of a conditional intention is a state formula that takes a single situation *now* as argument, unlike the goal formula  $\psi$ , which is a path formula and takes two situations *now* and *then* as arguments.<sup>3</sup> If one wishes to use an achievement goal  $\phi'$  for  $\psi$ , one can use  $\diamond(\phi', now, then)$ . For simplicity, we also assume that the trigger condition  $\phi$  is a one-time goal, i.e. once  $\phi$  becomes true, it remains true forever.

So we now propose a formalization of conditional intentions that avoids the under/over-commitment problem:

$$\begin{aligned}
 \text{CondInt}(agt, \phi, \psi, s) &\doteq \\
 &\text{Int}(agt, \text{DisjGoal}(\phi, \psi, now, then) \\
 &\quad \wedge \text{NoUnderComm}(agt, \phi, now, then), s), \\
 \text{DisjGoal}(\phi, \psi, now, then) &\doteq \\
 &[\neg \phi \text{ Until } (\phi \wedge \psi)](now, then) \vee \neg \diamond(\phi, now, then), \\
 \text{NoUnderComm}(agt, \phi, now, then) &\doteq \\
 &\square([\text{Int}(agt, \square(\neg \phi, now, then), now) \supset \\
 &\quad \text{Know}(agt, A_{\square}(\neg \phi, now), now)], now, then).
 \end{aligned}$$

That is, *agt* conditionally intends that  $\psi$  provided that  $\phi$ , iff *agt* intends that the following conditions hold:

<sup>3</sup> We could also handle trigger conditions that are not achievement goals. However, in these cases, since the trigger condition holds over a time interval, it is not always clear when exactly the triggering of the commitment to the conditional goal should occur. To avoid these complications, we stick to state formulas as triggers.

1. either (a)  $\phi$  eventually holds, and  $\psi$  holds immediately from the time  $\phi$  comes to hold, or (b)  $\phi$  never holds, and
2. if in any situation  $agt$  intends that  $\phi$  never comes to hold, she must also know in that situation that it can never become true.

Intuitively, this says that one way to fulfill an agent's conditional intention is to (1a) satisfy  $\psi$  after  $\phi$  comes to hold, and a second way is that (1b)  $\phi$  never comes to hold in the future. This part of our account is as in the disjunctive approach. However, we add to this that (2) the agent does not intend that  $\phi$  never comes to hold unless she knows that it can never hold. Thus we require that if at some situation,  $agt$  intends that  $\phi$  never comes true, it must be the case that she knows in that situation that  $\phi$  can never become true, and she only intends this because it has become inevitable. So the additional constraint that  $NoUnderComm(agt, \phi, now, then)$  ensures that  $agt$  will not do anything intentionally to make the triggering condition  $\phi$  remain false. One might be tempted to define  $NoUnderComm(agt, \phi, now, then)$  as  $\Box(\neg Int(agt, \Box(\neg\phi, now, then), now), now, then)$ , i.e.  $agt$  never intends that  $\phi$  never holds. However, since some event may make  $\phi$  impossible to achieve, there is a possibility that  $agt$  may come to intend that  $\phi$  always be false, if this becomes inevitable. The only case in which  $agt$  intends that  $\phi$  always be false is when she knows that it can never become true.

Consider once again our online marketplace example given in Section 1 for the disjunctive account. Using this definition of conditional commitment, a seller  $slr$ 's intention to send the goods when a buyer  $byr$  pays,  $CondInt(slr, -GetPaid(byr, slr), Happens(shipGoods(slr, byr), now, then), s)$  can be formalized as follows:

$$\begin{aligned} &Int(slr, DisjGoal(GetPaid(byr, slr), \\ &\quad Happens(shipGoods(slr, byr), now, then), now, then) \\ &\quad \wedge NoUnderComm(slr, GetPaid(byr, slr), now, then), s). \end{aligned}$$

$slr$ 's intention can be further expanded to:

$$\begin{aligned} &Int(slr, [GetPaidAndThenSendGoods(byr, slr, now, then) \\ &\quad \vee \neg\Diamond(GetPaid(byr, slr), now, then)] \wedge \\ &\quad [\Box((Int(slr, \Box\neg GetPaid(byr, slr), now) \supset \\ &\quad Know(slr, A_{\Box}(\neg GetPaid(byr, slr), now), now)), now, then)], s), \end{aligned}$$

where,

$$\begin{aligned} &GetPaidAndThenSendGoods(byr, slr, now, then) \doteq \\ &[\neg GetPaid(byr, slr) \text{ Until} \\ &\quad (GetPaid(byr, slr) \wedge Happens(shipGoods(slr, byr), now, then))](now, then). \end{aligned}$$

From this, we can see that there are only two ways by which  $slr$  can satisfy this conditional intention: either at some future or current situation  $byr$  pays  $slr$  and

then *slr* sends the goods to *byr*, or, *byr* never pays *slr*, and as long as *slr* does not know that *byr* will never pay her, she does not intend not to get paid. Since *slr* cannot intend not to get paid, she cannot deliberately perform anything (for example block the payment from *byr*) to make the triggering condition remain false. If at a later situation, *slr* learns that it has become impossible for *byr* to ever pay her, *slr* will inevitably intend that *byr* never pays her, but otherwise she cannot acquire this intention. Thus, our formalization of conditional commitment does not suffer from the under-commitment problem.

Moreover, since we use the disjunctive approach, our account does not suffer from the over-commitment problem associated with the conjunctive approach. Consider the second example given in Section 1, where *slr* has the intention to ship a replacement unit when *byr* returns a defective good. Using our definition, this can be expanded to *slr*'s intention that either *byr* never returns a defective product, or *byr* returns a defective product and *slr* ships the replacement unit after that. Thus *slr* is not over-committed and will not perform something deliberately so that *byr* returns a product. The additional constraint that *slr* never intends that *byr* never return a product unless she knows that *byr* will never return a product does not seem to lead to any over-commitment. Thus our formalization of conditional intention is also free from the over-commitment problem.

Note that, our account allows the agent who has a conditional intention to intend not to know whether the condition holds. We could easily strengthen the definition to rule this out, but it is not clear that this is always appropriate.

Next, we show a simple property of conditional intention. Assume that the domain theory  $D$  (as discussed in Section 2) includes our definition of conditional commitment given above. Then we have the following theorem that says that if an agent *agt* conditionally intends that  $\psi$  provided that  $\phi$  in situation  $s$ , and if she knows that  $\phi$  holds in  $s$ , then *agt* intends that  $\psi$  in  $s$ .

**Theorem 1.**

$$D \models \text{CondInt}(agt, \phi, \psi, s) \wedge \text{Know}(agt, \phi, s) \supset \text{Int}(agt, \psi, s).$$

So when the agent knows that the condition has become true, a conditional intention reduces to an ordinary intention.

It would be interesting to prove additional results about conditional intentions, for instance, that a conditional intention persists as long as its condition is known to remain false and not known to have become impossible. We leave this for future work.

## 4 Conditional Requests

We now discuss two communicative acts, *requestWhen* and *reqActWhen*, that can be used by an agent to request someone to achieve  $\psi$  or to execute a program  $\delta$  respectively, on the condition that  $\phi$  becomes true. Recall that, in ECASL the SSA for  $W$  determines whether an agent adopts a goal when requested; the

requested goal is adopted by the requestee via cooperation principles encoded in the SSA for  $W$ . Thus, we model requests as informing about intentions, rather than as primitives. In the following, we use  $CondIntCont(agt, \phi, \psi)$  as an abbreviation for the content of a conditional intention  $DisjGoal(\phi, \psi, now, then) \wedge NoUnderComm(agt, \phi, now, then)$ . Now, one simple way to model a requester  $req$ 's request to requestee  $agt$  to achieve  $\psi$  on the condition that  $\phi$  is as follows:

$$\begin{aligned} requestWhen_{sim}(req, agt, \phi, \psi) \doteq \\ request(req, agt, CondIntCont(req, \phi, \psi)). \end{aligned}$$

This says that,  $req$ 's conditional request to  $agt$  to achieve  $\psi$  provided that  $\phi$  amounts to  $req$ 's request to  $agt$  to fulfill the content  $CondIntCont(req, \phi, \psi)$  of her own conditional intention. Using the definition of  $request$ , this conditional request amounts to  $req$  informing  $agt$  that she currently intends to achieve  $\psi$  provided that  $\phi$ . However, note that the content  $CondIntCont(req, \phi, \psi)$  of this conditional intention includes mental attitudes that refer to  $req$ , rather than  $agt$ . Since the SSA for  $W$  does not automatically replace the agent parameters of mental state operators used in a goal formula, if we model conditional request as above, given appropriate conditions (i.e., when  $agt$  agrees to serve  $req$  on  $CondIntCont(req, \phi, \psi)$  and does not currently have the intention that  $\neg CondIntCont(req, \phi, \psi)$ ),  $agt$  will adopt the intention that  $CondIntCont(req, \phi, \psi)$ , but not that  $CondIntCont(agt, \phi, \psi)$ . Thus she will not have the conditional intention to achieve  $\psi$  provided that  $\phi$  after the conditional request is performed, and this simple definition is not quite correct.

For example, suppose that the manager agent  $mgr$  wants to conditionally request the seller  $slr$  in situation  $s$  to ship the goods when the buyer  $byr$  pays her. So  $mgr$  can do this by performing the following action in  $s$ :

$$\begin{aligned} requestWhen_{sim}(mgr, slr, GetPaid(byr, slr), \\ Happens(shipGoods(slr, byr), now, then)), \end{aligned}$$

which can be expanded to:

$$\begin{aligned} request(mgr, slr, CondIntCont(mgr, GetPaid(byr, slr), \\ Happens(shipGoods(slr, byr), now, then))). \end{aligned}$$

After the request is performed, if  $slr$  agrees to serve  $mgr$  on  $CondIntCont(mgr, GetPaid(byr, slr), Happens(shipGoods(slr, byr), now, then))$ , and does not intend that  $\neg CondIntCont(mgr, GetPaid(byr, slr), Happens(shipGoods(slr, byr), now, then))$ , the SSA for  $W$  will ensure that:

$$\begin{aligned} Int(slr, CondIntCont(mgr, GetPaid(byr, slr), \\ Happens(shipGoods(slr, byr), now, then)), s_r), \end{aligned}$$

which can be expanded to:

$$\begin{aligned} Int(slr, DisjGoal(GetPaid(byr, slr), \\ Happens(shipGoods(slr, byr), now, then), now, then) \\ \wedge NoUnderComm(mgr, GetPaid(byr, slr), now, then), s_r), \end{aligned}$$

where  $s_r$  is the situation that results from performing the *requestWhen* action in  $s$ . Now, using the definition of conditional intention, we can see that in  $s_r$ ,  $slr$  does not have the conditional intention of sending the goods provided that  $byr$  pays her. The problem is with the mental state operators in the *NoUnderComm*( $mgr, \dots$ ) part of  $slr$ 's intention: they say that  $mgr$  will not intend that the payment not occur unless she knows it can never occur. What we need is for this constraint to hold for  $slr$ .

To deal with this problem, we propose the following model of conditional requests:

$$\begin{aligned} requestWhen(req, agt, \phi, \psi) &\doteq \\ request(req, agt, CondIntCont(agt, \phi, \psi)). \end{aligned}$$

This says that  $req$ 's request to  $agt$  to conditionally achieve  $\psi$  provided that  $\phi$  amounts to  $req$ 's request to  $agt$  to fulfill the content of  $agt$ 's conditional intention to achieve  $\psi$  provided that  $\phi$ , i.e.,  $CondIntCont(agt, \phi, \psi)$ . Using the definition of *request*, this can be further expanded to:

$$\begin{aligned} requestWhen(req, agt, \phi, \psi) &\doteq \\ inform(req, agt, Int(req, CondIntCont(agt, \phi, \psi), now)). \end{aligned}$$

That is,  $req$  can request  $agt$  to achieve  $\psi$  on the condition that  $\phi$  by informing  $agt$  that she intends that  $CondIntCont(agt, \phi, \psi)$ . Note that, the agent parameter of  $CondIntCont(agt, \phi, \psi)$  is now the requestee  $agt$ , rather than the requester  $req$ . This guarantees that given that  $agt$  serves  $req$  and does not has the opposite intention, she will conditionally intend to achieve  $\psi$  provided that  $\phi$  after  $req$  conditionally requests her this. Thus this formalization of conditional request does not suffer from the above mentioned problem.

We also define a special type of conditional request, namely, the request to perform an action when some condition holds:

$$\begin{aligned} reqActWhen(req, agt, \phi, \delta) &\doteq \\ requestWhen(req, agt, \phi, Happen_C(\delta, now, then)). \end{aligned}$$

This states that  $req$ 's conditional request to  $agt$  to execute the program  $\delta$  provided that  $\phi$  amounts to  $req$ 's conditional request to  $agt$  to execute  $\delta$  starting in the situation where  $\phi$  holds.

Now consider what happens when  $mgr$  conditionally requests  $slr$  to ship the goods when  $byr$  pays her, that is, when  $mgr$  performs the *reqActWhen*( $mgr, slr, GetPaid$ ( $byr, slr$ ), *shipGoods*( $slr, byr$ )) action. Given that  $slr$  agrees to serve  $mgr$  and does not have the opposite intention, the SSA for  $W$  will make  $slr$  adopt the following intention:

$$\begin{aligned} Int(slr, CondIntCont(slr, GetPaid(byr, slr), \\ Happen_C(sendGoods(slr, byr), when, then)), \end{aligned}$$

and thus, by the definition of conditional intention, she will conditionally intend to send the goods when  $byr$  pays her. Thus, our formalization of conditional

requests allows the proper transfer of conditional intention from the requester to the requestee.

We next present a theorem that shows how agents' intentions are affected by the *requestWhen* action. Assume that the domain theory  $D$  includes our definition of these new communicative actions. We can show that:

**Theorem 2.**

$$\begin{aligned}
 D \models & [\neg \text{Int}(agt, \neg \text{CondIntCont}(agt, \phi, \psi), s) \\
 & \wedge \text{Serves}(agt, req, \text{CondIntCont}(agt, \phi, \psi), s) \\
 & \wedge \text{Poss}(\text{requestWhen}(req, agt, \phi, \psi), s)] \supset \\
 & \text{CondInt}(agt, \phi, \psi, \text{do}(\text{requestWhen}(req, agt, \phi, \psi), s)).
 \end{aligned}$$

This says that if in some situation  $s$ , an agent  $agt$  does not intend not to fulfill the content of her conditional intention to achieve  $\psi$  provided that  $\phi$ , and if she serves another agent  $req$  on the content of this conditional commitment in  $s$ , then she will have the conditional intention to achieve  $\psi$  given that  $\phi$  after  $req$  conditionally request her this in  $s$ , provided that the request is possible in  $s$ .

It would be useful to extend our framework with a communication act that allows a conditional commitment created as a result of a *requestWhen* to be cancelled. We believe that the existing ECASL *cancelRequest* action can be used to define such a conditional commitment cancelling act. We leave this for future work.

## 5 Related Work

The under-commitment problem that we pointed out in Section 1 is related to another problem involving intentions discussed by Cohen and Levesque [2]. In that paper, they consider a robot who drops the intention of bringing a bottle of beer by breaking the last available bottle and thus making the intention impossible to achieve. Their solution was twofold: (1) they formalize intentions as persistent goals and (2) they assume that existing intentions act as a screen of admissibility over new intentions. In their framework, an agent's intentions persist until she knows that they have been achieved, or knows that it has become impossible to achieve them. Since the robot intends to bring a bottle of beer, she will not drop this goal until she achieves it or gets to know that it has become impossible to achieve. However, the robot can break the last available bottle to make her goal unachievable. But since an agent's current intentions provide a screen of admissibility for adopting new intentions, she cannot have these two conflicting intentions at the same time. Thus since she intends to bring a bottle, she cannot adopt the intention to break the only available bottle. Note that while this problem has similarities with the one addressed here, it does not involve conditional intentions.

In the literature, there has been some work on conditional intention. However, as mentioned earlier, all of the proposed treatments that we are aware

of seem to suffer from the under- or over-commitment problems. Although it does not explicitly address conditional intentions, the FIPA agent communication language specification [5] defines a type of communication act that leads to conditional intentions. In that framework, an agent can conditionally request another agent to execute an action when some condition holds. This is modeled as follows: *req*'s conditional request to *agt* to perform *act* when  $\phi$  holds amounts to *req* informing *agt* that she has the intention that *agt* execute *act* and that  $\phi$  be true just before that. Note that *req*'s intention amounts to the conjunction that  $\phi$  be true at some point and *agt* executes *act* right after that. Thus, this treatment of conditional intention can be viewed as a conjunctive account where the intention is to first achieve the triggering condition  $\phi$ , and then to achieve the conditional goal. As discussed in Section 1, this leads to the over-commitment problem.

Yolum and Singh [26] present a different model of conditional commitment that relies on a social obligation-based semantics rather than a traditional one based on mental states. Their main concern was the study of communication protocols that accommodate exceptions and take advantage of opportunities. They model interaction protocols using *commitment machines* that supply a content to protocol states and actions in terms of the social commitments of the participating agents. In their formal semantics, which is only briefly described, they adopt a branching time temporal model. The semantics for commitments involves a modal accessibility relation for commitments  $C$  that relates a state of the protocol (i.e. a time-point)  $s$ , a debtor agent  $x$ , and a creditor agent  $y$  to a set of paths  $P$ . Intuitively,  $x$  is responsible to  $y$  for satisfying  $\phi$  in state  $s$  iff  $\phi$  holds at time-point  $s$  along all paths  $p$  that are  $C$ -accessible from  $(x, y, s)$ . To model conditional commitment, they introduce a strict implication operator (denoted by  $\rightsquigarrow$ ) that requires the consequent to hold when the antecedent holds. The strict implication is false when the antecedent is false. Their semantics says that  $\phi \rightsquigarrow \phi'$  holds in a state  $s$  iff  $\phi$  holds in  $s$  and for all  $s'$  that satisfy  $\phi$ , every  $s''$  that is similar to  $s'$  (i.e.  $s \approx s'$ ) also satisfies  $\phi'$ . What they mean by the similarity relation  $\approx$  is not explained. Thus for them, a conditional commitment  $C(x, y, \phi \rightsquigarrow \phi')$  holds in state  $s$  iff on all  $C$ -accessible paths  $p$ ,  $\phi$  holds at  $s$ , and whenever some  $s'$  satisfies  $\phi$ , every  $s''$  that is similar to  $s'$  satisfies  $\phi'$ . Since they model conditional commitments using the  $\rightsquigarrow$  operator, which behaves like a conjunction with some additional constraints, it appears that their formalization suffers from the over-commitment problem. It is also not clear how their formalization ensures that the goal is achieved *after* the condition along the paths.

Both [1] and [22] model conditional commitment as a disjunctive goal. In their social commitment and argument network based framework, Bentahar et al. [1] define conditional commitments as a simple implication. Their semantics of conditional commitment goes as follows:  $M, s \models \text{CondInt}_{Ben}(agt_1, agt_2, \phi, \psi)$  iff  $M, s \models \text{EF}^+\phi \Rightarrow M, s \models \text{ABC}(agt_1, agt_2, \psi)$ , where  $s, E, F^+$ , and ABC denotes a timepoint, there exists a path, sometime in the future, and absolute commitment, respectively. This says that  $agt_1$  is committed to  $agt_2$  to achieve

$\psi$  on the condition that  $\phi$  means that  $agt_1$  is unconditionally committed to  $agt_2$  to achieve  $\psi$  if  $\phi$  holds at some timepoint over some path in the future. Besides suffering from the under-commitment problem associated with disjunctive accounts, this seems to require commitment to the goal too early, before the condition becomes true.

In [22], Shapiro et al. describe a framework for specifying communicative multiagent systems using ConGolog [4] within the situation calculus, an early version of CASL. Since they were lacking a goal-revision mechanism at that point, they introduced a type of conditional request, the *requestUnless* action, in an attempt to avoid the need for goal-revision. *requestUnless*(*req*, *agt*,  $\phi$ ,  $\psi$ ) means that *req* is requesting *agt* to adopt the goal that  $\psi$  unless  $\phi$  is obtained. The execution of *requestUnless*(*req*, *agt*,  $\phi$ ,  $\psi$ ) makes *agt* adopt the goal that  $\phi \vee \psi$ . This amounts to modeling conditional intentions as disjunctive goals, and hence the account suffers from the under-commitment problem.

## 6 Conclusion

In this paper, we identified some problems with many existing formalizations of conditional commitments. These seem to either have the agents over-committed, intending to achieve the condition under which the goal would have to be achieved, or under-committed, possibly intending that this condition remain false forever. We could not find any problem-free account in the literature. We presented a definition of conditional intentions that does not suffer from these problems. We then formalized two types of communicative actions that allow agents to make requests that lead to conditional commitments. We also proved some properties of conditional commitments and conditional requests. Finally, we discussed previous work on conditional commitments.

The theory presented here is a part of our ongoing research on the semantics of speech acts and agent communication in the situation calculus. In [8], we present an extended version of this work where we model some simple communication protocols that deal with conditional requests. Much work remains. In the future, we would like to prove other properties of conditional commitments, for example, about the persistence and revision of such commitments. We also plan to formalize complex interaction protocols, such as the Contract Net protocol [25] and the Net Bill protocol [24], using our formalization of conditional intention.

## Acknowledgements

We thank Hector Levesque and the reviewers for useful comments on this work.

## References

1. Bentahar, J., Moulin, B., Meyer, J.-J. Ch., Chaib-draa, B.: A Logical Model of Commitment and Argument Network for Agent Communication. In: Proceedings of

- the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-04), New York, USA (2004) 792–799
2. Cohen, P., Levesque, H.: Intention is Choice with Commitment. *Artificial Intelligence*, 42:(2-3), (1990) 213–361
  3. Cohen, P., Levesque, H.: Rational Interaction as the Basis for Communication. In: Cohen P., Morgan, J., Pollack, M. (eds.): *Intentions in Communication*. Cambridge, MA, MIT Press (1990) 221–255
  4. De Giacomo, G., Lespérance, Y., Levesque, H.: ConGolog, a Concurrent Programming Language Based on the Situation Calculus. *Artificial Intelligence*, 121 (2000) 109–169
  5. Foundations for Intelligent Physical Agents: FIPA Communicative Act Library Specification. Document 37 (1997-2002)
  6. Herzig, A., Longin, D.: A Logic of Intention with Cooperation Principles and with Assertive Speech Acts as Communication Primitives. In: *Proc. of AAMAS 02* (2002)
  7. Khan, S., Lespérance, Y.: ECASL: A Model of Rational Agency for Communicating Agents. In: *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-05)*, Utrecht University, The Netherlands (2005)
  8. Khan, S.: A Situation Calculus Account of Multiagent Planning, Speech Acts, and Communication. M.Sc Thesis. Dept. of Computer Science, York University. In preparation (2005)
  9. Lakemeyer, G., and Levesque, H.: AOL: A Logic of Acting, Sensing, Knowing, and Only-Knowing. In: *Proc. of KR 98* (1998) 316–327
  10. Levesque, H., Reiter, R., Lespérance, Y., Lin, F., Scherl, R.: GOLOG: A Logic Programming Language for Dynamic Domains. *J. of Logic Programming*, 31 (1997)
  11. McCarthy, J., Hayes, P.: Some Philosophical Problems from the Standpoint of Artificial Intelligence. *Machine Intelligence*, 4 (1969) 463–502
  12. Moore, R.: A Formal Theory of Knowledge and Action. In: Hobbs J., Moore, R. (eds.): *Formal Theories of the Commonsense World*. Ablex (1985) 319–358
  13. Plotkin, G.: A Structural Approach to Operational Semantics. Technical Report DAIMI-FN-19, Computer Science Dept., Aarhus University, Denmark (1981)
  14. Rao, A., Georgeff, M.: Modeling Rational Agents within a BDI-architecture. In: Fikes, R., Sandewall, E. (eds.): *Proc. of KR&R 91* (1991) 473–484
  15. Reiter, R.: The Frame Problem in the Situation Calculus: A Simple Solution (Sometimes) and a Completeness Result for Goal Regression. In: Lifschitz, V. (ed.): *Artificial Intelligence and Mathematical Theory of Computation: Papers in the Honor of John McCarthy*. San Diego, CA, Academic Press (1991)
  16. Reiter, R.: *Knowledge in Action. Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press (2001)
  17. Sadek, D.: Communication Theory = Rationality Principles + Communicative Act Models. In: *Proc. of AAAI 94 Workshop on Planning for Interagent Comm.* (1994)
  18. Scherl, R., and Levesque, H.: Knowledge, Action, and the Frame Problem. *Artificial Intelligence*, vol. 144(1-2) (2003)
  19. Shapiro, S.: Specifying and Verifying Multiagent Systems Using CASL. PhD Thesis. Dept. of Computer Science, University of Toronto. (2004)
  20. Shapiro, S., Lespérance, Y.: Modeling Multiagent Systems with the Cognitive Agents Specification Language - A Feature Interaction Resolution Application. In: Castelfranchi, C., Lespérance, Y. (eds.): *Intelligent Agents Vol. VII - Proc. of ATAL 00, LNAI 1986* (2001) 244–259

21. Shapiro, S., Lespérance, Y., and Levesque, H.: The Cognitive Agents Specification Language and Verification Environment for Multiagent Systems. In: Castelfranchi, C., and Johnson, W. (eds.): Proc. of AAMAS 02 (2002) 19–26
22. Shapiro, S., Lespérance, Y., Levesque, H.: Specifying Communicative Multi-Agent Systems. In: W. Wobcke, M. Pagnucco, C. Zhang, (eds.): Agents and Multi-Agent Systems – Formalisms, Methodologies, and Applications, LNAI 1441, Springer-Verlag (1998) 1–14
23. Singh, M.: Multiagent Systems: A Theoretical Framework for Intentions, Know-How, and Communications. LNAI 799 (1994)
24. Sirbu, M. A.: Credits and debits on the Internet. In: M. N. Huhns, M. P. Singh (eds.): Readings in Agents, Morgan Kaufmann (1998) 299–305
25. Smith, R. G.: The Contract Net Protocol: High Level Communication and Control in a Distributed Problem Solver. In: IEEE Transactions on Computers, C-29(12). December (1980) 1104–1113
26. Yolum, P., Singh, M. P.: Commitment Machines. In: J.-J. Ch. Meyer and M. Tambe (eds.): Intelligent Agents VIII : 8th International Workshop, ATAL '01, LNAI 2333 (2002) 235–247