

# Abstract

CHOPRA, AMIT KHUSHWANT. Nonmonotonic Commitment Machines. (Under the direction of Dr. Munindar P. Singh).

Protocols in open and dynamic multiagent systems cannot be modeled via formalisms such as finite state machines (FSMs). FSMs though formal, lack content, that is, a protocol-level semantics, which leads to rigidity in protocols. We propose formalism called Nonmonotonic Commitment Machines (NCMs) for representing protocols. An NCM gives a commitment-based semantics to actions and states in a protocol. Protocols represented as NCMs afford the agent flexibility in interactions with other agents. In particular, situations in protocols when nonmonotonic reasoning is required can be efficiently represented in NCMs. For applications that need efficient execution, we show how to compile an NCM into an FSM and prove the compilation to be sound and complete.

# NONMONOTONIC COMMITMENT MACHINES

BY

AMIT KHUSHWANT CHOPRA

A THESIS SUBMITTED TO THE GRADUATE FACULTY OF  
NORTH CAROLINA STATE UNIVERSITY  
IN PARTIAL FULFILLMENT OF THE  
REQUIREMENTS FOR THE DEGREE OF  
MASTER OF SCIENCE

DEPARTMENT OF COMPUTER SCIENCE

RALEIGH

NOVEMBER 2003

APPROVED BY:

---

R. MICHAEL YOUNG

RADA CHIRKOVA

---

MUNINDAR P. SINGH

CHAIR OF ADVISORY COMMITTEE

*To my mother.*

# **Biography**

Amit Chopra was born on June 13, 1977 in Mumbai, India. He received his Bachelor's Degree in Computer Engineering in 1999 from the University of Pune, India. He has been with the Department of Computer Science at North Carolina State University since Fall 2001.

# Acknowledgement

I am indebted to my advisor, Dr. Munindar Singh for his constant guidance, attention and patience. I thank Dr. Jon Doyle and Dr. Rada Chirkova for their numerous helpful suggestions that helped improve this document. I am grateful to Pinar, who has since graduated, for mentoring me. I thank my roommates, Gopi and JD, for bearing with my eccentricities and not minding my minding theirs. Ashok, Ruben, Vishwas, Rakesh and Gotu have been such good company throughout that I seldom felt homesick. I thank Hemant, Veena, Sonu and Manjusha who are all in India, for many a delightful long distance conversations.

I thank my family for supporting me in this endeavor. And the person who selflessly let me go.

This research was supported by the National Science Foundation under grant DST-0139037.

# Table of Contents

<b>List of Figures</b>	<b>vii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Motivation . . . . .	2
1.2 Approach . . . . .	4
1.3 Outline . . . . .	5
<b>2 Nonmonotonic Causal Logic</b>	<b>6</b>
2.1 Introduction . . . . .	7
2.2 Semantics . . . . .	8
2.3 Action Descriptions in C+ . . . . .	10
2.4 Translating C+ to NCL . . . . .	11
<b>3 Commitments</b>	<b>14</b>
3.1 Definition . . . . .	14
3.2 Operations . . . . .	15
<b>4 Nonmonotonic Commitment Machine Representation of Protocols</b>	<b>16</b>
4.1 Commitments in Causal Logic . . . . .	16
4.2 NetBill Specification in Causal Logic . . . . .	18
4.3 Executing NetBill in CCalc . . . . .	20
<b>5 Formal Definition of Nonmonotonic Commitment Machines</b>	<b>23</b>
<b>6 Compiling Nonmonotonic Commitment Machines</b>	<b>26</b>
6.1 Compilation Procedure . . . . .	27
6.2 Correctness of Compilation . . . . .	28

<b>7</b>	<b>Discussion</b>	<b>30</b>
7.1	Conventional Protocols and Protocol Modeling . . . . .	31
7.2	Commitments . . . . .	32
7.3	Protocols . . . . .	33
7.4	Agent Communication Languages . . . . .	34
<b>8</b>	<b>Future Directions</b>	<b>36</b>
8.1	Protocol Specification . . . . .	36
8.2	Protocol Distribution . . . . .	37
8.3	Creating Role Skeletons from BPEL Flows . . . . .	37
8.4	Verifying Strategies . . . . .	38
	<b>Bibliography</b>	<b>39</b>

# List of Figures

1.1	FSM representation of the simplified NetBill protocol . . . . .	3
4.1	Example Query . . . . .	21
4.2	Answer to Example Query . . . . .	22
6.1	FSM representation of the extended NetBill protocol . . . . .	27



# Chapter 1

## Introduction

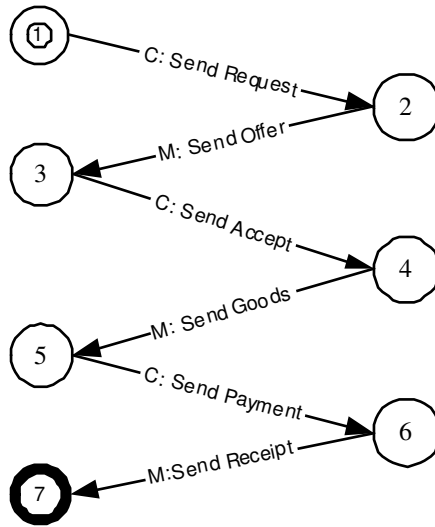
A protocol is a means of achieving meaningful interaction. Agents that constitute a multiagent system use protocols to guide their interactions with each other. Protocols have traditionally been specified as FSMs that specify sequences of states. The protocol designers have certain scenarios in mind that they directly incorporate in an FSM. As a result, agents using a protocol specified as FSMs are limited to behaving in a rigid manner. Such agents cannot handle exceptions or take advantage of opportunities that might arise during interactions with other agents. In this thesis, we present an alternative way of specifying protocols that is based on commitments which we formalize below. Our approach is based on the general notion that an agent does not violate a given protocol as long the agent does not violate the commitments prescribed by the protocol. Using commitments makes the protocol flexible and enables the agent to handle exceptions and opportunities without violating the given protocol.

## 1.1 Motivation

Protocols for interaction in multiagent systems often resemble protocols routinely used by humans in their social interactions. The Contract Net [Smith, 1980] and NetBill [Cox et al., 1995] are examples of such protocols. Such protocols have traditionally been represented by FSMs that represent sequences of states and transitions. Since FSMs are a low level representation, it becomes cumbersome to capture multiple scenarios in an FSM. Thus FSMs designed by hand tend to be rigid and do not allow scenarios other than the specified “normal” ones. A protocol transitions from state to state as a result of the actions of the interacting agents. A transition is usually labeled with the actions that cause it. In an FSM, the states and the actions in the protocol are meaningless tokens. The agent is limited to executing one of the sequences of actions hard-coded by its designer. These sequences represent the only legal behaviors of the agent. Anything the agent does outside of this protocol is considered a violation. This makes the protocol inflexible and, therefore, undesirable in open multiagent systems where agents are autonomous and heterogeneous, and opportunities and exceptions need to be handled appropriately.

Acting flexibly presupposes reasoning about the protocols. Reasoning formally presupposes that the protocols have a formal semantics. We base our semantics on the notion of commitments. Protocols can naturally be seen as an exchange and manipulation of commitments. A commitment is a directed obligation from one agent to another for achieving or maintaining a state of affairs. A commitment is social because it involves two parties and is publicly observable by all the agents in the agent society. Since a commitment is public, it is also possible to verify whether an agent has fulfilled its commitment, thereby making it possible to check an agent’s compliance with a protocol.

We use the NetBill e-commerce protocol [Cox et al., 1995] as a running example



**Figure 1.1:** FSM representation of the simplified NetBill protocol

throughout the thesis. Figure 1.1 shows an FSM representation of a NetBill simplified to focus on the core part. The customer, represented by  $c$ , sends a request for offers to the merchant, represented by  $m$ . The merchant sends an offer in response. If the customer accepts the offer, the merchant sends the goods. The customer then sends the payment for the goods in return for which the merchant sends a receipt. The only execution scenario possible in the protocol starts with the customer sending a request and ends with the merchant sending a receipt. The FSM in Figure 1.1 does not accommodate scenarios that would arise naturally in open and dynamic multiagent systems and is, therefore, unnecessarily rigid. Protocols for multiagent interaction should be flexible in the following ways:

- *Autonomy:* A protocol specification should not impinge on the autonomy of an agent beyond the essential nature of the interaction it describes. Consider a scenario where a customer wants to buy goods from a merchant. A desirable specification should not limit the autonomy of an merchant by preventing him from advertising his wares by sending an *offer* message prior to receiving a

request for offers.

- *Opportunities*: A protocol should enable an agent to take advantage of opportunities that may arise. For example, if a merchant advertises an attractive deal to a customer, the customer should be able to entertain this offer.
- *Exceptions*: A protocol should enable an agent to deal with exceptions instead of aborting the interaction altogether. For example, a customer who doesn't have enough money might delegate a commitment to pay the merchant to some other agent. This is not allowed in the NetBill protocol as specified above.

## 1.2 Approach

The above forms of flexibility can be achieved only if we are able to reason about the content of the states and actions in a protocol. Often, the essential element of content in many protocols is the commitments of the different parties in a protocol. Specifically, we claim that if the protocol representation uses commitments and the criterion for protocol compliance is the satisfaction of commitments, then the above scenarios would be valid behaviors in the protocol.

We propose a formalism for specifying protocols called Nonmonotonic Commitment Machines (NCMs) that uses commitments for representing states and actions. The meaning of a state is given by the commitments that hold in that state; a state is a description of the world. The meaning of an action is given by how it manipulates commitments. An NCM does not directly specify sequences of states and transitions. Instead, it specifies rules in Nonmonotonic Causal Logic (NCL) [Giunchiglia et al., 2003]. These rules model the changes in the state of a protocol as a result of execution of actions. The inference mechanism of NCL computes new states at runtime. Yolum and Singh [2002a] first studied

commitment machines. They did not consider situations during the execution of a protocol when agents must act with incomplete information. In such situations, the agent would need nonmonotonic or defeasible reasoning. Since NCL supports nonmonotonic reasoning, NCMs can express defaults in a protocol in a natural manner. Protocols represented as NCMs are more elaboration tolerant McCarthy [1999] than those represented using classical logic or FSMs. We develop a causal theory of commitments and represent the NetBill protocol as an NCM using that theory. We also formally define NCMs in terms of NCL and show how an NCM can be compiled into an FSM. This is helpful for applications where efficient execution is important.

### **1.3 Outline**

The rest of the thesis is organized as follows. Chapter 2 motivates the need for nonmonotonic logic for protocol representation and describes the NCL that we employ for this purpose. Chapter 3 formalizes commitments. Chapter 4 formalizes commitments in this logic and presents the NetBill NCM. Chapter 5 presents a formal definition of NCMs. Chapter 6 shows how an NCM can be compiled into an FSM. Chapter 7 discusses the relevant literature and chapter 8 discusses future directions.

## Chapter 2

# Nonmonotonic Causal Logic

In logic, the consequence relation  $\vdash$  is a relation between sets of propositions and individual propositions.  $A \vdash x$ , where  $A$  is a set of propositions and  $x$  is a proposition, means that  $x$  is a logical consequence of  $A$ . Classical logic is monotonic meaning that if  $A \subseteq B$ , where  $A \vdash x$  and  $B$  is a set of propositions, then  $B \vdash x$ . Informally, monotonicity means that the addition of new information does not invalidate old information. Therefore, making rules defeasible in the face of change poses difficulties. Consider Example 1.

**Example 1** A customer may not return goods received and should pay for them. However, if the goods are damaged, then the customer may return the goods and then cancel his commitment to pay. ■

Example 1 involves defeasible reasoning. The *default* rule is that the commitment to pay cannot be canceled. (A more accurate rule is that the debtor of a commitment cannot cancel his commitment. However, the above simplified version is adequate for this example.) The general rule is defeasible, i.e., if a special condition applies, like when the goods are damaged and they are returned, the commitment can be canceled. The default rule should be applicable when no information about the condition is available. This is

the essence of nonmonotonic reasoning. Such defeasible reasoning is beyond the realm of classical logic.

## 2.1 Introduction

To overcome the aforementioned difficulties in specifying protocols, we use NCL. We choose NCL because it has an intuitive syntax and semantics. The language  $C+$  [Giunchiglia et al., 2003], which is based on NCL, has a semantics based on state transition systems which agrees with our intuition about protocols. Further, it has been shown to be elaboration tolerant [Lifschitz, 2000]. NCL is a logic of *universal causation* meaning that every fact that is caused holds and every fact that holds is caused. Universal causation is not so much a philosophical stance as a practical one, as universal causation yields a uniform semantics for causal theories. Taking this stance makes NCL suitable for simulation and planning since everything can be explained. Also, in our domain, a commitment holds or does not hold only because there is a reason for it to hold or not hold. Moreover, as we show below, universal causation can be disabled for selected formulas.

The *signature* of a causal theory is the set  $\alpha$  of symbols called *constants*. Each constant  $c$  is assigned a nonempty finite domain  $Dom(c)$  of symbols. An *atom* is of the form  $c = v$  where  $v \in Dom(c)$ . An *interpretation* of  $\alpha$  is an assignment  $c = v$  for each  $c \in \alpha$  where  $v \in Dom(c)$ . Since we consider only boolean atoms, either  $c = true$  or  $c = false$ . A *formula* in NCL is a combination of atoms using the connectives of classical logic. A causal rule is of the form  $F \Leftarrow G$ , where  $F$  and  $G$  are formulas of classical logic and are called the head and the body of the rule, respectively. This means that there is cause for  $F$  to be true if  $G$  is true. It does not say that  $G$  is the cause for  $F$ . This reflects the intuition that it is sufficient to know the conditions under which a fact is caused. As an example, consider a switch  $S$

that, when closed, lights two bulbs  $A$  and  $B$ . Even though  $A$  being lit is not the cause for  $B$  being lit, it is correct to say that there is a cause for  $B$  to be lit when  $A$  is lit. A theory in NCL consists of a set of causal rules.

Constants in causal theories are either *fluents* or *actions*. A causal theory describes histories of length  $m+1$ , ( $m \geq 0$ ), by creating for all  $i$ , ( $i \in \{0, \dots, m\}$ ), a copy of every fluent and, for all  $i$ , ( $i \in \{0, \dots, m-1\}$ ), a copy of every action. The interpretation of fluents for a particular  $i$  represents state  $s_i$  and the interpretation of actions in state  $s_i$  represent the transition to state  $s_{i+1}$ .

## 2.2 Semantics

Models for formulas in NCL are defined in the same way as classical logic. An interpretation is a model of a set  $X$  of formulas iff it satisfies all the formulas in  $X$ . If every model of  $X$  satisfies a formula  $F$ , then  $X$  entails  $F$ , or symbolically  $X \models F$ . We can now define models of a causal theory. Let  $I$  be an interpretation of the signature  $\alpha$  of a theory  $T$ . The reduct  $T^I$  is the set of all heads whose bodies are satisfied by the interpretation  $I$ . If  $I$  is also a unique model of  $T^I$ , then  $I$  is a model of  $T$ . If  $I$  is not a unique model of the reduct, then some constant is missing from the reduct, and therefore there is no explanation for that constant. But NCL is a logic of universal causation, therefore, no constant should be unexplained.

As an example, consider the following theory  $T_1$  consisting of rules  $R1$  and  $R2$ .

R1.  $p \Leftarrow q$

R2.  $q \Leftarrow q$

R1 says that there is a cause for  $p$  if  $q$  is true. R2 says that there is a cause for  $q$ . Reasoning informally using the principle of universal causation, we see that there is no



cause for  $\neg p$  to be true. Therefore,  $p$  has to be true. Therefore,  $p$  must be caused. The only way  $p$  can be caused is if  $q$  is true. And,  $q$  is caused by R2. Therefore, the only possible interpretation for this theory is  $I = \{p = \text{true}, q = \text{true}\}$ .

Intuitively,  $T^I$  represents facts that are caused, according to theory  $T$  under interpretation  $I$ . If a causal theory  $T$  has a model  $I$ , we say that it is consistent or satisfiable. If all models of  $T$  satisfy a formula  $F$ , that means  $T$  entails  $F$  or  $T \models F$ .

Coming back to our example theory  $T_1$  with rules R1 and R2, we consider all the possible interpretations to see which, if any, is a model of  $T_1$ :

1.  $I^1 = \{p = \text{true}, q = \text{true}\}$ :  $T_1^{I^1} = \{p, q\}$ .  $I_1$  is a unique model of  $T_1^{I^1}$ . Therefore,  $I_1$  is a model of  $T_1$ .
2.  $I^2 = \{p = \text{false}, q = \text{true}\}$ :  $T_1^{I^2} = \{p, q\}$ .  $I_2$  is not a model of  $T_1^{I^2}$ . Therefore,  $I_2$  is not a model of  $T_1$ .
3.  $I^3 = \{p = \text{true}, q = \text{false}\}$ :  $T_1^{I^3} = \{\}$ .  $T_1^{I^3}$  has no models. Therefore,  $I_3$  is not a model of  $T_1$ .
4.  $I^4 = \{p = \text{false}, q = \text{false}\}$ :  $T_1^{I^4} = \{\}$ .  $T_1^{I^4}$  has no models. Therefore,  $I_4$  is not a model of  $T_1$ .

Note that  $I^1 = \{p = \text{true}, q = \text{true}\}$  is the only model of this theory that matches the result of our informal reasoning.

To see how NCL is nonmonotonic, consider a theory  $T_2$  consisting of the rule  $c = 1 \Leftarrow c = 1$ . This rule is like a default rule. The only model of  $T_2$  is  $I(c) = 1$ . Now consider a theory  $T_3$  such that it had two rules,  $c = 1 \Leftarrow c = 1$  and  $c = 2 \Leftarrow \text{true}$ . Note that  $T_2 \subseteq T_3$ . However, the only model of  $T_3$  is  $I(c) = 2$ .

## 2.3 Action Descriptions in C+

Recall that  $C+$  is a high level action description language based on NCL. It is easier to specify theories in  $C+$  than directly in NCL because of its concise notation. Before we describe the syntax and semantics of NCL and  $C+$  formally, we describe informally the meanings of some of  $C+$  rules that we use later. A formula in  $C+$  is a propositional combination of constants which could either be *action* constants or *fluent* constants. Actions in NCL are interpreted to be unit-length. This thesis is restricted to boolean constants. An action constant being true represents the execution of the action. The meanings of the rules we use follow.

- $a$  causes  $b$ , where  $a$  and  $b$  are actions.

This means that action  $a$  causes action  $b$  and both happen concurrently.

- $a$  causes  $f$ , where  $a$  is an action and  $f$  is a fluent.

This means that action  $a$  causes  $f$  to hold in the next state.

- $A \wedge F$  causes  $b$ , where  $A$  is a conjunction of actions,  $b$  is an action and  $F$  is a conjunction of fluents.

This means that in a state where  $F$  is true and actions in  $A$  happen, then action  $b$  happens concurrently.

- $a$  causes  $a$ , together with the rule  $\neg a$  causes  $\neg a$ , where  $a$  is an action.

These two rules mean that there is a cause for  $a$  and there is a cause for  $\neg a$  respectively. In other words the  $a$  is exogenous, it simply happens or does not happen. Without these rules, the action is not exogenous. Universal causation is disabled for these rules.

- $a$  may cause  $f$ , where  $a$  is an action and  $f$  is a fluent.

This means that  $f$  may hold after  $a$ 's execution if it does not already hold. Thus,

this rule expresses nondeterminism.

- *caused a after f*, where *a* is an action and *f* is a fluent.

This means that *f* causes *a* in the same state. Notice that this rule uses the *caused* form and not the *causes* form because no suitable formulation in terms of *causes* exists. This rule expresses the causation of an action and differs from the rule *a causes f* above which expresses the causation of a fluent.

- Fluents are declared as *inertialFluents* meaning that their assignment persists from one state to the next unless changed by some other rule.

## 2.4 Translating C+ to NCL

Let's describe *C+* formally and show the translation from rules in *C+* to rules in NCL. *C+* includes three kinds of rules, namely, static rules, fluent dynamic rules and action dynamic rules. A fluent can be either a *statically determined fluent* or a *simple (dynamic) fluent*. Static fluents can appear in the heads of only static rules. A *fluent formula* consists only of fluents. An *action formula* consists only of actions. A *static rule* is an expression of the form

R3. *caused F if G*

where *F* and *G* are fluent formulas. Static rules express indirect effects of actions that are instantaneous with respect to the causal fluent formula. A dynamic rule is an expression of the form

R4. *caused F if G after H*

It is called an *action dynamic rule* if *F* and *G* are both action formulas. It is called a *fluent dynamic rule* if *F* and *G* are both fluent formulas. Action dynamic rules express the

causation of an action and fluent dynamic rules the causation of a fluent.

An action description  $D$  consisting of such rules is turned into a causal theory  $D_m$  where  $m$  is the length of the history. The signature  $\alpha$  of  $D_m$  then consists of

1.  $i:c$  for every fluent  $c \in \alpha$  for every  $i \in \{0, \dots, m\}$
2.  $i:c$  for every action  $c \in \alpha$  for every  $i \in \{0, \dots, m-1\}$

The domain of  $i:c$  is the same as  $Dom(c)$  and  $i:F$  means  $i$  is inserted in front of every occurrence of every constant in  $F$ . The rules of  $D_m$  are then:

- R5.  $i:F \Leftarrow i:G$ , for every static rule R3 in  $D$  and every  $i \in \{0, \dots, m\}$ ;
- R6.  $i:F \Leftarrow i:G \wedge i:H$ , for every action dynamic rule R4 in  $D$  and every  $i \in \{0, \dots, m-1\}$ ;
- R7.  $i+1:F \Leftarrow i+1:G \wedge i:H$ , for every fluent dynamic rule R4 in  $D$  and every  $i \in \{0, \dots, m-1\}$ ;
- R8.  $0:c = v \Leftarrow 0:c = v$ , for every simple fluent constant  $c \in \alpha$  and every  $v \in Dom(c)$ .  
(Notice that every simple fluent has all possible values in the initial state and therefore, they are exogenous in the initial state. Thus, universal causation is disabled for them.)

All examples that follow are in the abbreviated C+ notation. We list the relevant C+ abbreviations below, extracted from [Giunchiglia et al., 2003].

- A1. A dynamic rule of the form *caused F if true after H* abbreviates to *H causes F*.
- A2. An action dynamic rules of the form *caused F if G after H* abbreviates to *G  $\wedge$  H causes F*

- A3. The action dynamic rules *caused a if a* and *caused  $\neg a$  if  $\neg a$*  where *a* is an action, together abbreviate to *exogenous a*. In *C+* such an action is called an *exogenousAction*.
- A4. The fluent dynamic rules *caused p if p after p* and *caused  $\neg p$  if  $\neg p$  after  $\neg p$*  where *p* is a fluent together abbreviate to *inertial p*
- A5. The fluent dynamic rule *caused F if F after H* abbreviates to *H may cause F*.

## Chapter 3

# Commitments

Commitments among agents have been recognized as a fundamental notion in cooperative problem solving [Bratman, 1992; Shoham, 1993; Grosz and Kraus, 1996]. Castelfranchi [1993] and Krogh [1995] present, respectively, a social and logical perspective on commitments. In our work, we do not reason about the commitments from the point of view of cooperation among agents. We use commitments to specify protocols. As agents interact with each other using some protocol, they create and manipulate commitments. The breach of a commitment represents a violation of a protocol.

### 3.1 Definition

The agent that is bound to fulfill the commitment is called the debtor of the commitment.

The agent that is the beneficiary of the commitment is called the creditor.

**Definition 1** A base-level commitment  $C(x,y,G,p)$  binds a debtor  $x$  to a creditor  $y$  for fulfilling the condition  $p$  in context  $G$ .

**Definition 2** A conditional commitment  $CC(x,y,G,p,q)$  denotes that if a condition  $p$  is brought about, then the commitment  $C(x,y,G,q)$  will hold.

Both commitments and conditional commitments are created in a context  $G$ , which can be thought of as an institution or society whose rules are binding on the agents that join it. The context also defines the meanings of the terms used in the context. Henceforth, we omit  $G$  to reduce clutter.

## 3.2 Operations

Singh[1999] lists operations for the creation and manipulation of commitments. These operations cannot be arbitrarily carried out. They are subject to metacommitments that are rules that govern the commitment operations and are part of the context  $G$ . The operations are listed below.

- $Create(x,y,p)$  creates a new commitment  $C(x,y,p)$ .
- $Discharge(x,y,p)$  discharges the existing commitment  $C(x,y,p)$  so that it no longer holds.
- $Cancel(x,y,p)$  cancels the existing commitment  $C(x,y,p)$  so that it no longer holds.
- $Delegate(x,y,p,z)$  delegates the commitment  $C(x,y,p)$  to a new debtor  $z$ . More specifically, the original commitment  $C(x,y,p)$  no longer holds and a new commitment  $C(z,y,p)$  is created in its place.
- $Assign(x,y,p,z)$  assigns the commitment  $C(x,y,p)$  to a new creditor  $z$ . More specifically, the original commitment  $C(x,y,p)$  no longer holds and a new commitment  $C(x,z,p)$  is created in its place.
- $Release(x,y,p)$  releases the debtor  $x$  from the commitment  $C(x,y,p)$  so that the commitment no longer holds.

## Chapter 4

# Nonmonotonic Commitment Machine

## Representation of Protocols

In our approach, we represent protocols as NCMs. An NCM is a causal theory in  $C+$  that consists of two parts. The first part is a protocol-independent causal theory of commitments in which we capture the representation of commitments and the operations on them. The second part is protocol specific and includes constants and rules describing the given protocol's domain. The distinction between the two parts is only to separate out the domain independent part, logically they form a complete causal theory as we shall see later. We first present the theory of commitments and then model the NetBill protocol in causal logic. Together they represent the NetBill NCM.

### 4.1 Commitments in Causal Logic

We represent commitments and operations on them in the causal logic. Commitments in causal logic are declared to be constants of the type *inertialFluents*.

- $C(x,y,p), CC(x,y,p,q) :: inertialFluents$



where  $x$  and  $y$  are variables of the sort *agent* and  $p$  and  $q$  are variables of the sort *condition*. By declaring commitments to be *inertialFluents*, we include rules of the kind A4 for each commitment. Conditional commitments are declared as  $CC(x,y,p,q)$  where  $q$  is also a variable of sort *condition*. Conditional commitments are also declared as *inertialFluents*. For each of the operations on commitments listed in Chapter 3, there is a declaration of the form

- $\langle \textit{Operation} \rangle :: \textit{action}$

Constants of type *action* are not exogenous, that is, rules of the form A3 are not included in the theory. Their execution, therefore, has to be *caused* by other actions or fluents. We add two more operations for handling conditional commitments.

- $CDischarge(x,y,p,q), CCreate(x,y,p,q) :: \textit{action}$

The following rules capture the meaning of the operations:

R9.  $Create(x,y,p)$  causes  $C(x,y,p)$

R10.  $Discharge(x,y,p)$  causes  $\neg C(x,y,p)$

R11.  $Cancel(x,y,p)$  causes  $\neg C(x,y,p)$

R12.  $Delegate(x,y,p,z)$  causes  $\neg C(x,y,p) \ \& \ C(z,y,p)$

R13.  $Release(x,y,p)$  causes  $\neg C(x,y,p)$

R14.  $Assign(x,y,p,z)$  causes  $\neg C(x,y,p) \ \& \ C(x,z,p)$

R15.  $CCreate(x,y,p,q)$  causes  $CC(x,y,p,q)$

R16.  $CDischarge(x,y,p,q)$  causes  $\neg CC(x,y,p,q) \ \& \ C(x,y,q)$

All the variables are grounded such that  $x \neq y$  and  $p \neq q$ . We omit the rules specifying the grounding of the variables. Since we want the operations to be *caused* by other things, then in those states where an operation is not caused, there must be a reason for it to be not caused. In other words, we want the operations to be partially exogenous. So for each of the commitment operations, we include rules of the form

R17.  $\neg\langle \text{Operation} \rangle \text{ causes } \neg\langle \text{Operation} \rangle$

An example is the rule  $\neg\text{Create}(x,y,p) \text{ causes } \neg\text{Create}(x,y,p)$ . The specification also includes rules to capture the restriction that no two commitment operations are concurrent.

## 4.2 NetBill Specification in Causal Logic

We now represent the NetBill protocol in causal logic. The following rules together with the theory of commitments given above represent the specification of NetBill NCM. We declare

- $m, c$  to be of the sort *agent*
- $\text{goodsc}, \text{payc}, \text{acceptc}, \text{receiptc}$  to be of the sort *condition*
- $\text{request}, \text{offer}, \text{accept}, \text{goods}, \text{pay}, \text{receipt}$  to be *inertialFluents*
- $\text{SendRequest}, \text{SendOffer}, \text{SendAccept}, \text{SendGoods}, \text{SendPayment}, \text{SendReceipt}$  to be *exogenousActions*.

The meanings of the above constants are as their name indicates. Conditions are an artifact of conditional commitments. We assume that all fluents and actions have unique identifiers. We have the following rules.

R18.  $\text{SendRequest} \text{ causes } \text{request}$

R19.  $\text{SendOffer} \text{ causes } \text{offer}$

- R20. *SendOffer causes CCreate(m, c, acceptc, goodsc)*
- R21. *SendAccept causes accept*
- R22. *SendAccept & CC(m,c,acceptc,goodsc)*  
*causes CDischarge(m,c,acceptc,goodsc)*
- R23. *SendAccept causes CCreate(c,m,goodsc,payc)*
- R24. *SendGoods causes goods*
- R25. *SendGoods causes CCreate(m, c, payc, receiptc)*
- R26. *SendGoods & CC(c,m,goodsc,payc) causes CDischarge(c,m,goodsc,payc)*
- R27. *SendGoods & C(m,c,goodsc) causes Discharge(m,c,goodsc)*
- R28. *SendPayment causes pay*
- R29. *SendPayment & CC(m,c,payc,receiptc) causes CDischarge(m,c,payc,receiptc)*
- R30. *SendPayment & C(c,m,payc) causes CDischarge(c,m,payc)*
- R31. *SendReceipt causes receipt*
- R32. *SendReceipt & C(m,c,receiptc) causes Discharge(m,c,receiptc)*

In our representation no two commitment operations are concurrent. Also, no two protocol actions are concurrent. However, when a protocol action causes a commitment operation, they are concurrent. By rule R6, the interpretation of *ActionA causes ActionB* is such that *ActionA* and *ActionB* are concurrent. This ensures that the protocol action is concurrent with the commitment operation it causes is satisfied. There could be other concurrency models possible for NetBill. We choose this one because of its simplicity.

We now add rules for our motivating example, Example 1, to this protocol specification. We introduce *SendReturn* and *SendGoods* as exogenous actions. We introduce a new action *Ab* and a new fluent *damagedGoods* to indicate that the goods are damaged. We also include the nondeterministic rule R34 to say that as a result of the *SendGoods* action, the goods may be damaged. Rule R33 captures the condition that the cancel operation is not allowed for any commitment. Rule R35 however says that a commitment can be canceled under abnormal conditions. Rules R36 and R37 ensure that *Ab* is false, except when *damagedGoods* is true. We add the Rules R35 – R37 to accommodate Example 1. Rules R33 and R34 are already in the theory. *Ab* is an action because it has no meaning in the states. It acts as a qualifier for the exogenous action *SendReturn*.

R33.  $\neg \text{Cancel}(x,y,p)$  causes  $\neg \text{Cancel}(x,y,p)$

R34. *SendGoods* may cause *damagedGoods*

R35.  $\text{SendReturn} \wedge C(c,m,\text{pay}) \wedge \text{Ab}$  causes  $\text{Cancel}(c,m,\text{pay})$

R36. *caused Ab if damagedGoods*

R37.  $\neg \text{Ab}$  causes  $\neg \text{Ab}$

The theory also contains rules that place constraints on the actions so that the execution of actions makes sense. For example, we specify that the *SendRequest* action cannot happen after the payment has been made.

### 4.3 Executing NetBill in CCalc

*CCalc* (Causal Calculator) is a reasoning tool that implements causal logic. Given a causal theory and a goal in the form of a query, *CCalc* finds paths to the goal. We load the NetBill

causal theory into *CCalc* and pose queries one by one. *CCalc* then finds paths to satisfy each query.

```
:- query
label::0;
maxstep:: 3;
0: -offer,
   -accept,
   -returned,
   -goods,
   -C(x,y,p),
   -CC(x,y,p,q),
   -pay,
   -request,
   -receipt,
   -damagedGoods;
maxstep: returned.
```

**Figure 4.1:** Example Query

Figure 4.1 shows an example query. This query asks for an execution sequence of three or fewer steps, beginning from a state in which all fluents are false, in which the goods have been returned. Running this query in *CCalc* produces the output as shown in Figure 4.2 (formatted for readability).

The action *SendAccept* is caused which in turn causes the *CCreate* which creates the conditional commitment that if the goods are sent then the customer will pay. The result of these actions is reflected in state 1. *SendGoods* is then caused which reflects the fact that the merchant has sent the goods. *SendGoods* causes the discharge of the conditional commitment created by the customer resulting in the customer's commitment to pay. *SendGoods* also creates a new conditional commitment that if the customer pays then the

Solution:

State 0:

ACTIONS: CCreate (c, m, goodsc, payc)  
SendAccept

State 1: CC (c, m, goodsc, payc) accept

ACTIONS: CCreate (m, c, payc, receiptc)  
CDischarge (c, m, goodsc, payc)  
SendGoods

State 2: C (c, m, payc) CC (m, c, payc, receiptc)  
accept goods damagedGoods

ACTIONS: Cancel (c, m, payc) Ab Return

State 3: CC (m, c, payc, receiptc) accept  
goods returned damagedGoods

**Figure 4.2:** Answer to Example Query

merchant will send the receipt. This example is interesting as *SendGoods* also causes *damagedGoods*, resulting in state 2. *damagedGoods* causes *Ab* in state 2. So the *SendReturn* action is successful, which in turn causes the cancellation of the commitment to pay. State 3 is the resulting state which also satisfies the goal state of our query.

## Chapter 5

# Formal Definition of Nonmonotonic Commitment Machines

We formalize NCMs in terms of the state machine semantics of  $C+$  [Giunchiglia et al., 2003]. Consider an action description  $D$  in  $C+$ . Let  $\sigma^{fl}$  and  $\sigma^{act}$  represent the set of fluents and actions in  $D$  respectively. An action description  $D$  represents a sequence of causal theories  $D_0, D_1, \dots$ . A causal theory with a history of length  $m$  is denoted by  $D_0, D_1, \dots, D_{m-1}$ , where  $m \geq 0$ . In Chapter 2.3, we saw how causal rules in  $C+$  are translated into causal logic to represent histories.  $D_m$  contains all those causal rules that have state indices (represented by  $i$ ) less than or equal to  $m$ . An interpretation of the signature of  $D_m$  is represented by

$$(0: s_0) \cup (0: e_0) \cup (1: s_1) \cup (1: e_1) \cup \dots \cup (m: s_m)$$

where  $s_0, \dots, s_m$  are interpretations of  $\sigma^{fl}$  and  $e_0, \dots, e_{m-1}$  are interpretations of  $\sigma^{act}$ .

**Definition 3** A *state* is an interpretation  $s$  of  $\sigma^{fl}$  such that  $0:s$  is a model of  $D_0$ .

**Definition 4** A *transition* is a triple  $\langle s, e, s' \rangle$ , where  $s$  and  $s'$  are interpretations of  $\sigma^{fl}$  and  $e$  is an interpretation of  $\sigma^{act}$ , such that  $(0: s) \cup (0: e) \cup (1: s')$  is a model of  $D_1$ .

States are the vertices and transitions correspond to the edges of the transition system. Giunchiglia *et al.* [2003] show that for any transition  $\langle s, e, s' \rangle$ ,  $s$  and  $s'$  are states. We denote the set of states in  $D$  by  $\mathbf{M}_D$ . With this background, we can now define an NCM in terms of the transition semantics of action descriptions.

**Definition 5** An NCM based on an action description  $D$  is a 4-tuple  $\langle \mathbf{M}, \mathbf{A}, \mathbf{F}, \Delta \rangle$  where

- $\mathbf{M}$  is the set of states and is defined as  $\mathbf{M} = \{m \mid m \text{ is an interpretation of } \sigma^{fl}\}$
- $\mathbf{A}$  is the set of actions and is defined as  $\mathbf{A} = \sigma^{act}$
- $\mathbf{F}$  is the set of final states ( $\mathbf{F} \subseteq \mathbf{M}$ )
- $\Delta = \{\langle m_i, a, m_j \rangle \mid \langle m_i, e, m_j \rangle \text{ is a transition in } D \text{ where } m_i, m_j \in \mathbf{M}, m_i \neq m_j \text{ and } e \text{ is an interpretation of } \sigma^{act} \text{ such that it assigns true only to the actions } a\}$

We have  $m_i \neq m_j$  to eliminate redundant transitions.

**Lemma 1**  $\mathbf{M} \subseteq \mathbf{M}_D$ .

*Proof.* If there exists some  $m$  such that  $m \in \mathbf{M}$ , but  $m \notin \mathbf{M}_D$ , then  $D$  is not a correct action description of the protocol. It needs to be corrected. ■

To show that  $\mathbf{M}_D$  can have more states than  $\mathbf{M}$ , let  $\sigma^{simfl}$  denote the set of *simple fluents*. By rule R8, we observe that simple fluents are exogenous in the initial state. Therefore, all possible interpretations of  $\sigma^{simfl}$  are all models of  $D_0$ . By Definition 3, they are also states of  $D$ . However, not all interpretations of  $\sigma^{simfl}$  may be meanings of an NCM. Let  $\mathbf{M}_U$  denote  $\mathbf{M}_D - \mathbf{M}$ .

**Lemma 2** There exists no state  $m \in \mathbf{M}_U$  such that  $\langle s, a, m \rangle \in \Delta$  for any  $s$  and any  $a$ .

*Proof.* Consider any  $m \in \mathbf{M}_U$ . Assume that  $\langle s, a, m \rangle \in \Delta$ . Then by Definition 5,  $m \in \mathbf{M}$  which is a contradiction. ■

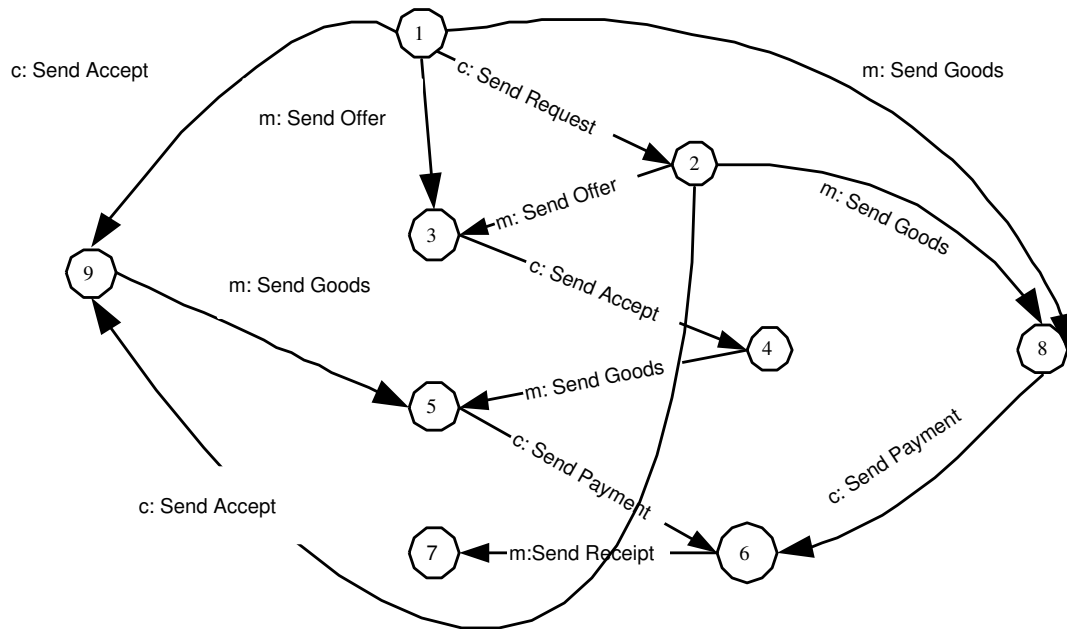


The states in  $\mathbf{M}_U$  are not in  $\mathbf{M}$  and they cannot be reached by any transition. Intuitively, they are there only because of the models of  $D_0$  allowed by simple fluents. Recall that all possible interpretations of  $\sigma^{simfl}$  are all models of  $D_0$  and are therefore states. They can safely be dropped from all consideration without curtailing the functionality of the NCM in any way.

## Chapter 6

# Compiling Nonmonotonic Commitment Machines

The preceding chapters present commitment-based semantics for interaction protocols and show how an agent can reason with commitments. Such a semantics is necessary for verifying compliance. However, it is not necessary that agents be able to reason about commitments to execute a protocol. For some applications efficient execution might be important. It may also be the case that an agent designer wants to exclude certain risky or lengthy paths (behaviors) in the protocol. For such agents, it would be useful if we could compile an NCM into an FSM. The complete NCM need not be compiled into an FSM. A designer could selectively compile behaviors (sequences of transitions) from an NCM into an FSM. The FSM can then be executed without an inference engine. Another advantage of compilation instead of directly designing or extending an FSM is that it is not always clear how to add states and transitions to an FSM. Compilation from an NCM makes this process automatic. Figure 6.1 shows an example FSM for NetBill, that can be compiled from the action description presented in Chapter 4.2. Selective compilation is a subject of future research.



**Figure 6.1:** FSM representation of the extended NetBill protocol

## 6.1 Compilation Procedure

We now present the procedure for compiling an NCM into its corresponding FSM that has no representation of commitments.

**Definition 6** An FSM is defined as a 5-tuple  $\langle \mathbf{S}, \Sigma, s_0, \mathbf{Q}, \delta \rangle$  where  $\mathbf{S}$  is a finite set of states,  $\Sigma$  is the input alphabet,  $s_0 \in \mathbf{S}$  is the start state,  $\mathbf{Q} \subseteq \mathbf{S}$  is the set of final states, and  $\delta \subseteq \mathbf{S} \times \Sigma \times \mathbf{S}$  is the transition relation.

The procedure for compiling an NCM into an FSM is described by Definition 7.

**Definition 7** Let  $X = \langle \mathbf{M}, \mathbf{A}, \mathbf{F}, \Delta \rangle$  be an NCM. Construct an FSM  $Y = \langle \mathbf{S}, \Sigma, s_0, \mathbf{Q}, \delta \rangle$  where  $\mathbf{S} = \mathbf{M}$ ,  $\Sigma = \mathbf{A}$ ,  $s_0 = \text{nothing}$ ,  $\mathbf{Q} = \mathbf{F}$ ,  $\delta = \langle m_i, a, m_j \rangle \mid \langle m_i, a, m_j \rangle \in \Delta$

We defined *nothing* as the initial state  $s_0$  of an FSM. This is simply the state where no actions have happened. In the NetBill NCM, this state would be the one specified as the initial state in the query in Figure 4.1.

## 6.2 Correctness of Compilation

Correctness of a compilation procedure depends on the computations that can be realized as a result of it. Two important aspects of the correctness of a procedure are *soundness* and *completeness*. Soundness means that only originally allowed computations are realized. Completeness means that all allowed computations are realized. Lemma 3 below shows that the FSM resulting from Definition 7 is sound. Lemma 4 shows that the FSM is complete under some restrictions. We start by defining the notion of the computation formally. We introduce sets of indices  $\mathbf{I}, \mathbf{J}$ . Let  $f \in \mathbf{N}$  be an ordinal.  $\mathbf{I}$  is the set  $(0, 1, \dots, f - 1)$ .  $\mathbf{J}$  is the set  $(0, 1, \dots, f)$ .

**Definition 8**  $\tau = \langle m_0, a_0, m_1, a_1, \dots, a_{f-1}, m_f \rangle$  is a computation if  $(\forall i \in \mathbf{I}, j \in \mathbf{J} : a_i \in \Sigma$  and  $m_j \in \mathbf{M})$ .

**Definition 9**  $\tau = \langle m_0, a_0, m_1, a_1, \dots, m_f \rangle$  is a computation *generated* by an NCM  $X = \langle \mathbf{M}, \mathbf{A}, \mathbf{F}, \Delta \rangle$  iff  $m_f \in \mathbf{F}$ , and  $(\forall i \in \mathbf{I} : m_i, m_{i+1} \in \mathbf{M})$  and  $(\exists a_i : \langle m_i, a_i, m_{i+1} \rangle \in \Delta)$ .

**Definition 10**  $\tau = \langle m_0, a_0, m_1, a_1, \dots, m_f \rangle$  is a computation *realized* by an FSM  $Y = \langle \mathbf{S}, \Sigma, s_0, \mathbf{Q}, \delta \rangle$  iff  $m_0 = s_0$ ,  $m_f \in \mathbf{Q}$ , and  $(\forall i \in \mathbf{I}, \langle m_i, a_i, m_{i+1} \rangle \in \delta)$ .

**Lemma 3** Let  $X = \langle \mathbf{M}, \mathbf{A}, \mathbf{F}, \Delta \rangle$  be compiled into  $Y = \langle \mathbf{S}, \Sigma, s_0, \mathbf{Q}, \delta \rangle$  according to Procedure 7. Then any computation realized by  $Y$  can be generated by  $X$ .

*Proof.* Let  $\tau = \langle m_0, a_0, m_1, a_1, \dots, m_f \rangle$  be a computation realized by  $Y$ . By Procedure 7,  $\mathbf{S} = \mathbf{M}$ ,  $(\forall i \in I, m_i \in \mathbf{M})$  and  $m_f \in \mathbf{F}$ . Consider the  $i^{th}$  transition in  $\tau$ ,  $\langle m_i, a_i, m_{i+1} \rangle \in \delta$ . From the Definition of  $\delta$  in Procedure 7,  $\langle m_i, a, m_j \rangle \in \Delta$ . By Definition 9,  $X$  generates  $\tau$ . ■

Since an NCM can start from any state and an FSM has a single start state, it is not possible for a compiled FSM to realize all the allowed computations. We instead define a weaker notion of completeness.

**Lemma 4** Let  $X = \langle \mathbf{M}, \mathbf{A}, \mathbf{F}, \Delta \rangle$  be compiled into  $Y = \langle \mathbf{S}, \Sigma, s_0, \mathbf{Q}, \delta \rangle$  according to Definition 7. Then every computation generated by  $X$  that begins with *nothing* can be realized by  $Y$ .

*Proof.* Let  $m_{st}, m_{st} \in \mathbf{M}$  represent the start state *nothing*. Let  $\tau = \langle m_0, a_0, m_1, a_1, \dots, m_f \rangle$  be any computation realized by  $X$  such that  $m_{st} = m_0$ . Let  $\langle m_i, a_i, a_{i+1} \rangle \in \Delta$  be the  $i^{th}$  transition in  $\tau$ . By Definition 7,  $\langle m_i, a, m_j \rangle \in \delta$ . By Definition 10,  $Y$  realizes  $\tau$ . ■

## Chapter 7

### Discussion

Our focus in this work is to develop meaningful representations of agent communication protocols. We do so by using commitments to declaratively represent states and actions. This gives our representation a verifiable semantics [Singh, 2000]. By using commitments to model protocols, we constrain protocols no more than is necessary. We have highlighted the need for a nonmonotonic logic for commonsense reasoning in protocols. We used NCL towards this end and showed how a protocol can be represented in NCL. Like the NCL, there are a few other noteworthy formalisms for reasoning about action and change. Dynamic Logic [Harel et al., 2000] is a modal logic augmented with an algebra of regular events. However, it is monotonic and therefore not suitable for our purposes. Event Calculus [Kowalski and Sergot, 1986] and situation calculus [McCarthy, 1963] have been extended with circumscription to enable nonmonotonic reasoning.

Both FSMs and NCMs are formal representations of protocols. Both approaches are verifiable and in the case of an FSM, trivially so. An NCM, though, represents meaning. Agents that can reason about commitments take actions accordingly. For example, if, as part of a particular protocol, an agent enters into a commitment, then the agent can plan its actions, even those not directly related to the protocol, so that the commitment is never

violated. Alternate paths through the protocol may be selected based on criteria like safety or number of messages exchanged. For example, an agent can adopt an approach where it does not commit unless another agent also commits for some desirable condition. Also it is not convenient to express defaults in FSMs. In fact, the defaults wouldn't be obvious at all in an FSM. FSMs are also not as elaboration tolerant as NCMs.

## **7.1 Conventional Protocols and Protocol Modeling**

Conventional protocols like TCP/IP, RPC, HTTP, and so on have a well-defined environment and scope. Their focus is on the correct delivery of data and they are therefore strict in the sense that they prescribe all paths for correct execution as well as for error recovery. Modeling such protocols as FSMs is usually sufficient. Notable among other formalisms for modeling protocols are Petri Nets [Girault and Valk, 2003] and statecharts [Harel, 1987]. Petri Nets have proven to be especially useful in modeling concurrency. Petri nets specify transitions (T-elements) between places (P-elements) which are sets of conditions. Statecharts is a visual formalism that extends state machines by adding support for hierarchy, concurrency and communication. Statecharts provide the designer the power to cluster states into super-states as well as refine states, thereby leading to compact representations for complex behavior. Statecharts can also represent defaults. However, statecharts also specify the transitions between states. As such, neither petri nets nor statecharts afford much flexibility. Statecharts though, are easier to comprehend because of the ability to cluster and refine states unlike NCMs where the specification may become difficult to manage as the number of rules increase.

## 7.2 Commitments

Commitments have been studied in the context of distributed problem solving and coordination. Bratman [1992] argues that for shared cooperative activity, among other things, commitment to joint activity and commitment to mutual support are required. Grosz and Kraus [1996] investigate the formulation of shared plans for coordinating group action. In their framework, an agent can adopt two types of intentions, intend-to and intend-that, that commit the agent to an action and state of affairs respectively. Jennings [1993] presents commitments as a fundamental notion for efficient coordination in distributed systems. Jennings also mentions conventions which monitor the commitments and state when a commitment may be reassessed. Jennings further reformulates different models of coordination in terms of commitments. A distinguishing feature of all of the above work is that they present commitments as a mentalistic notion, assuming a system of cooperative agents. Shoham's agent oriented computing paradigm [1993] introduces obligation as a modality required to describe the mental state of an agent. Sandholm and Lesser [1998] study automated negotiation among self-interested agents whose computations are resource bounded. They argue that protocols that have leveled commitments, that is, when commitments vary from breakable to unbreakable in a continuum by assigning a function to evaluate the cost of breach of each commitment, are more suitable for contracts than full commitment protocols. Krogh [1995] examines the possibility of using deontic logic for analyzing multiagent systems. Castelfranchi [1993] presents an ontology of commitments with the aim of understanding organizational activity. He defines social commitment as a relation between two or more agents and discusses its various aspects. Singh [2000] defends a commitment-based social semantics for agent communication.

We do not study commitments from the point of view of coordination. However as we



pointed out earlier in this chapter, an agent can reason about its future actions depending upon the commitments it already has or ones that the agent might have to make in the future. Also, we do not specify that commitments necessarily have to be represented in an agent's state, since a commitment based representation can be compiled into an FSM. We are also not concerned directly with the economic impact of the breach of a commitment. Though we specify protocols in a logical language using commitments, we do not present a deontic logic. Social commitments, that is, directed obligation from one agent to another, represent the cornerstone of our research. Our scope is limited to the flexible specification of protocols and their verification.

### **7.3 Protocols**

Yolum and Singh[2002a] proposed commitment machines and also showed how event calculus can be used to represent protocols and generate new paths in the protocols [2002b]. They do not consider commonsense reasoning situations. Koning and Huget [2001] describe a methodology for designing interaction protocols for multiagent systems. In their work, the focus is on reusability and modularity of protocols. They achieve this by composing protocols out of microprotocols using a formal language called *Communication Protocol Description Language* (CPDL). A formula in CPDL corresponds to an edge going from an initial state to a final state. The edge is labeled with a sequence of microprotocols which makes the protocol quite rigid. Koning and Huget also do not consider which microprotocols can be composed. This job is presumably left to the designer. With our commitment-based approach, it is possible for the agent to determine which protocols can be composed by looking at the legal states in the protocol.

## 7.4 Agent Communication Languages

An agent communication language (ACL) allows agents of heterogeneous designs to interact. Developing a semantics for agent communication languages (ACLs) that is expressive and verifiable has been a long standing goal of the agent community. To be verifiable, an ACL should have social semantics [Singh, 1998]. Earlier efforts at standardization of ACLs like Arcol [Sadek, 1997] and KQML [Finin et al., 1994] promoted mental agency, that is, they were based on mental concepts like beliefs and intentions, and were therefore, not verifiable. The ongoing efforts at standardizing ACLs are based on communicative acts. The problem with giving a communicative acts based semantics is that it is not clear what the meanings of the communicative acts should be. Also it is not clear how many communicative acts are needed. Another challenge is relating the communicative acts to the conversation in which they occur. By giving semantics to protocols directly we are able to give a simple, operational characterization of protocols without getting bogged down with the above issues.

Dignum and van Linder [2002], and Guerin and Pitt [2001] define ACLs in term of communicative acts. Dignum and van Linder's framework considers four components, the *information* component, the *action* component, the *motivational* component and the *social* component as constituting an agent framework and formally describes and relates them. The framework is developed in dynamic logic which is monotonic. They postulate a COMMIT communicative act, but other communicative acts have mentalistic preconditions which makes the framework suitable only for a system of cooperative agents.

Guerin and Pitt propose an ACL specification in which declarative ACL specifications are given procedural interpretations. An ACL specification in their approach, consists of three parts: a *Converse Function* that specifies permissions and obligations for subsequent

speech acts based on the conversation state, a *Protocol Semantics* that captures protocol dependent meanings of speech acts, and a *Speech-Act Semantics* that give the protocol independent part of the meaning. It is not clear how useful it is to model the protocol-independent part of the meaning, since most meaning comes from the protocol.

## **Chapter 8**

### **Future Directions**

Our main aim in this work is to come up with protocols that have verifiable semantics and constrain the agent no more than to the extent necessary to carry out legal interactions. To carry our work further, we have identified the following areas for future work.

#### **8.1 Protocol Specification**

While our formulation of NCM has some desirable properties like declarative rules and elaboration tolerance, it lacks other properties desirable in protocols such as a comprehensible graphical representation, role bindings for agents and temporal model checking [Holzmann, 1991]. For example, we cannot prove satisfactorily if the NetBill NCM is correct and complete. A related piece of work is to compare NCMs with other formalism like statecharts and Petri Nets in more depth. Another interesting avenue to explore is the compilation of NCMs into more expressive graphical formalism like statecharts.

## 8.2 Protocol Distribution

In the NCM representation of NetBill, we did not specify roles in the protocol. However, an agent will be executing the role that it adopts in the protocol. We want to formulate procedures based on symbol manipulation to distribute a centralized protocol among the various roles in the protocol. The result will be roles skeletons that are also NCMs. It is not clear what the specification of a role itself should be. Is a role just a label, or does it specify the required capability of an agent to adopt a role along with normative rules and authorizations that come along with the role? Another technical challenge is proving that a distributed protocol is sound and complete with respect to the centralized protocol.

## 8.3 Creating Role Skeletons from BPEL Flows

BPEL [BPEL], the Business Process Execution Language is a draft standard for specifying and coordinating business processes. Intuitively, it is a process flow graph with nodes as tasks and edges as messages. Business flows, as they are currently specified, are like FSMs in the sense that they have been designed with certain scenarios in mind. As such, they are quite rigid. To make business processes more flexible, we envisage the following design-time methodology.

1. The designers start with an interaction diagram or state machine or some other graphical representation for the protocol that is to be modeled as a BPEL flow.
2. The designers build NCM representation of the protocol with enhancements to make it more flexible and then partition the NCM into the role skeletons.
3. The role skeletons are compiled into FSMs

4. Compile FSMs into a BPEL flow. This should be easier than compiling an NCM into a BPEL flow.

We plan to build a tool which incorporates this methodology. The tool would suggest enhancements to the designer for a given protocol, build an NCM based on the choices of the designer and compile it into an FSM and then, perhaps with annotations from the user, compile it into a BPEL flow.

## **8.4 Verifying Strategies**

It will often be the case that an agent is confronted with the problem of selecting between multiple paths that it can take to reach a goal. The agent selects a path based on some strategy. For example, if the customer does not trust a merchant, the customer might adopt a strategy where it never pays before receiving the goods. On the other hand, if it does, it could adopt a strategy where it accepts to buy goods for a certain price without even asking the merchant for offers or pays before getting the goods. It is possible to imagine more complex strategies. An interesting problem is the specification of strategies with respect to commitments and determining which protocols (more specifically, paths in the protocol) satisfy a given strategy. An agent can then inspect a role to check whether it satisfies its strategy.

We conclude by saying that this work represents the first step in the development of a comprehensive methodology for designing flexible multiagent interaction protocols. Development of the protocol design tool that incorporates this methodology is the primary objective of this research.

# Bibliography

Michael E. Bratman. Shared cooperative activity. *The Philosophical Review*, 101:327–341, 1992.

Cristiano Castelfranchi. Commitments: From individual intentions to groups and organizations. In *Proceedings of the AAI-93 Workshop on AI and Theories of Groups and Organizations: Conceptual and Empirical Research*, 1993.

Benjamin Cox, J.D. Tygar, and Marvin Sirbu. Netbill security and transaction protocol. In *Proceedings of the First USENIX Workshop on Electronic Commerce*, pages 77–88, 1995.

Frank Dignum and Bernd van Linder. Modelling social agents: Towards deliberate communication. In *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, pages 357–380. Kluwer, 2002.

Tim Finin, Richard Fritzson, Don McKay, and Robin McEntire. KQML as an agent communication language. In *Proceedings of the International Conference on Information and Knowledge Management*, pages 456–463. ACM Press, 1994.

Claude Girault and Rudiger Valk. *Petri Nets for System Engineering*. Springer, 2003.

- Enrico Giunchiglia, Joohyung Lee, Vladimir Lifschitz, Norman McCain, and Hudson Turner. Nonmonotonic causal theories. *Artificial Intelligence*, 2003. To appear.
- Barbara J. Grosz and Sarit Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, October 1996.
- Frank Guerin and Jeremy Pitt. Denotational semantics for agent communication languages. In *Proceedings of the Fifth International Conference on Autonomous Agents*, pages 497–504. ACM Press, 2001.
- David Harel. Statecharts: A visual formalism for complex systems. *Science of Computer Programming*, 8:231–274, 1987.
- David Harel, Dexter Kozen, and Jerzy Tiuryn. *Dynamic Logic*. MIT Press, Cambridge, MA, 2000.
- Gerard J. Holzmann. *Design and Validation of Computer Protocols*. Prentice-Hall, London, 1991.
- Michael N. Huhns and Munindar P. Singh, editors. *Readings in Agents*. Morgan Kaufmann, San Francisco, 1998.
- N. R. Jennings. Commitments and conventions: The foundation of coordination in multi-agent systems. *Knowledge Engineering Review*, 2(3):223–250, 1993.
- Jean-Luc Koning and Marc-Philippe Huget. A semi-formal specification language dedicated to interaction protocols. In H. Kangassalo, H. Jaakkola, and E. Kawaguchi, editors, *Information Modeling and Knowledge Bases XII, Frontiers in Artificial Intelligence and Applications*. IOS Press, 2001.



- Robert Kowalski and Marik Sergot. Logic-based calculus of events. *New Generation Computing*, 4(1):67–95, 1986.
- Christen Krogh. Obligations in multiagent systems. In Agnar Åmodt and Jan Komorowski, editors, *Scandinavian Artificial Intelligence Conference 1995 (SCAI'95)*, pages 19–30, Trondheim, May 1995.
- Vladimir Lifschitz. Missionaries and cannibals in the causal calculator. In *Proceedings of the 7th International Conference on Principles of Knowledge Representation and Reasoning*, pages 85–96, 2000.
- John McCarthy. Situations, actions and causal laws. TR, Stanford University, 1963.
- John McCarthy. Elaboration tolerance. *In progress*, 1999. <http://www-formal.stanford.edu/jmc/elaboration.html>.
- David Sadek. Compliance in Arcol, April 1997. Personal communication.
- Tuomas Sandholm and Victor Lesser. Issues in automated negotiation and electronic commerce: Extending the contract net framework. In *Huhns and Singh [1998]*, pages 66–73. 1998. (Reprinted from *Proceedings of the International Conference on Multiagent Systems, 1995*).
- Yoav Shoham. Agent-oriented programming. *Artificial Intelligence*, 60(1):51–92, 1993.
- Munindar P. Singh. Agent communication languages: Rethinking the principles. *IEEE Computer*, 31(12):40–47, December 1998.
- Munindar P. Singh. An ontology for commitments in multiagent systems: Toward a unification of normative concepts. *Artificial Intelligence and Law*, 7:97–113, 1999.

- Munindar P. Singh. A social semantics for agent communication languages. In *Proceedings of the 1999 IJCAI Workshop on Agent Communication Languages*. Springer-Verlag, 2000.
- R. G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, 29(12):1104–1113, 1980.
- BPEL. *Business process execution language for web services, version 1.1*. W3C, May 2003. URL: <http://www-106.ibm.com/developerworks/webservices/library/ws-bpel/>.
- Pınar Yolum and Munindar P. Singh. Commitment machines. In *Proceedings of the 8th International Workshop on Agent Theories, Architectures, and Languages (ATAL-01)*, pages 235–247. Springer-Verlag, 2002a.
- Pınar Yolum and Munindar P. Singh. Flexible protocol specification and execution: Applying event calculus planning using commitments. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, pages 527–534. ACM Press, July 2002b.