

An Algebra for Commitment Protocols *

Ashok U. Mallya and Munindar P. Singh

Department of Computer Science

North Carolina State University

Raleigh, NC 27695-7535 USA

{aumallya, singh}@ncsu.edu

Abstract. Protocols enable unambiguous, smooth interactions among agents. Commitments among agents are a powerful means of developing protocols. Commitments enable flexible execution of protocols and help agents reason about protocols and plan their actions accordingly, while at the same time providing a basis for compliance checking. Multiagent systems based on commitments can conveniently and effectively model business interactions because the autonomy and heterogeneity of agents mirrors real-world businesses. Such modeling, however, requires multiagent systems to host a rich variety of protocols that can capture the needs of different applications. We show how a commitment-based semantics provides a basis for refining and aggregating protocols. We propose an approach for designing commitment protocols wherein traditional software engineering notions such as refinement and aggregation are extended to apply to protocols. We present an algebra of protocols that can be used to compose protocols by refining and merging existing ones, and does this at a level of abstraction high enough to be useful for real-world applications.

Keywords: Commitments, Interaction Protocols

1. Introduction

Multiagent systems composed of autonomous and heterogeneous agents provide a convenient and accurate model for describing and enacting many real-life processes and interactions. While autonomy and heterogeneity are what make the multiagent paradigm attractive, heterogeneity gives rise to incompatibility and autonomy to unpredictability. Agents need to understand each other and behave in predictable ways for their interactions to be fruitful. To achieve consensus and facilitate interaction between agents, standards are required, as in most distributed systems. Web Services are an example of how standards enable heterogeneous systems to interact with each other. Recent efforts for Web Service *choreography*—which deals with the way services interact—and *orchestration*—which deals with the way services are composed using other services—address service interactions (Peltz, 2003) similar in spirit to agent interaction protocols. Agent interaction, however, requires higher-level abstractions to deal with the rich variety of interactions found in multiagent systems.

* We thank Amit Chopra and Nimit Desai for valuable comments. This research was supported partly by the NSF under grant DST-0139037 and partly by a DARPA contract.

INTERACTION PROTOCOLS

A protocol is a description of the steps involved in an interaction. Protocols make interactions coherent and easy to implement. The use of protocols has successfully solved the problem of standardization in areas such as computer networks. Likewise, the heterogeneous and distributed structure of multiagent systems necessitates clear protocols to govern any interaction. Network protocols explain the steps to be taken in great detail, sometimes even enumerating all possible events that can occur. For example, the Session Initiation Protocol (SIP), which is used to set up phone calls over the Internet, describes every message that needs to be sent for setting up and tearing down calls and also every possible resultant reply for the message (Rosenberg et al., 2002). By contrast, multiagent systems require protocols to be specified at a high level of abstraction, to accommodate the complexity of agent systems, and to not overwhelm protocol designers with unnecessary details.

While protocols are needed to force an agent to behave in a predictable manner, they should also allow flexibility of execution. A protocol that allows only one sequence of steps does not let its participants leverage their autonomy. A restrictive protocol, however, is not always bad. If a protocol allows only a single computation, checking whether the participants are compliant with the protocol is trivial. Any step that does not agree with the protocol signals a violation. As protocols become more flexible, however, compliance verification becomes harder, since many choices are offered to the participants at any step of the protocol. Consequently, protocol design is an exercise in finding the right balance between flexibility of execution and ease of compliance checking.

MOTIVATION

The tradeoffs between execution and verification to be borne in mind make protocol design a nontrivial undertaking. It requires human expertise and knowledge of the application domain. To reduce unnecessary effort and to prevent reinventing the wheel, designers should be able to create new protocols by refining or combining existing protocols whose properties are well understood. A sound theory of refining and composing protocols would assist designers in ascertaining the properties of protocols. An algebra of protocols that includes operators for merging and refinement is needed as the basis for protocol composition.

Our central claim is that protocols can be characterized in terms of their content, not just their sequence of steps. We develop a protocol algebra which is at once a high-level abstraction of protocols and a useful tool for composing protocols and reasoning about them, as we demonstrate with an example.

CONTRIBUTION

Our main contribution is in developing an algebra for composing protocols. Just as conceptual modeling in general involves abstractions such as refinement and aggregation, so must the conceptual modeling of protocols. We develop an algebra that provides the underpinnings of such abstractions for protocols. The algebra is a high-level abstraction that relates to real-world protocols, and hence is easy for protocol designers to understand. We also demonstrate how the use of commitments allows reasoning about protocols that leads to richer interaction patterns from existing ones. Further, we outline how a hierarchy of protocols can be generated based on commitments. This hierarchy aids reasoning about which protocol is the most general for a given business process.

ORGANIZATION

The rest of this paper is organized as follows. Section 2 introduces the technical background, and some illustrative examples that are used throughout the paper. Section 3 develops our theory of semantics of protocol subsumptions, introduces the protocol algebra and demonstrates its utility in composing protocols. Section 4 summarizes the paper, identifies related work in the field, and charts out future directions.

2. Technical Framework

We represent protocols as transition systems similar in spirit to finite state machines. These protocols generate computations or *runs*, which are sequences of states that a valid protocol execution can go through. We devise a hierarchical classification based on the runs generated by protocols. Runs are composed of *states* that the protocol computation (execution) goes through based on the *actions* that the participants in the given protocol perform. This classification forms the basis of our work. Next, we introduce commitments, discuss some scenarios from our running example, and then define the basic technical concepts needed for our semantics.

2.1. COMMITMENTS IN PROTOCOLS

Commitments among agents are an abstraction of contracts that exist in the real world (Castelfranchi, 1993; Singh, 1991). Commitments lend coherence to interactions because they help agents plan based on the actions of others, and they are, in principle, enforceable. Commitment-based protocols are more flexible than traditional formalisms like finite state machines and Petri

nets (Verdicchio and Colombetti, 2002; Yolum and Singh, 2002). By specifying the states that need to be reached in terms of commitments, they can allow multiple paths to achieve a state, and consequently create a flexible protocol specification.

A commitment $C(x, y, p)$ denotes that the agent x is responsible to the agent y for bringing about the condition p . Here x is the *debtor*, y the *creditor*, and p the *condition* of the commitment, expressed in a suitable formal language. Commitments can also be *conditional*, denoted by $CC(x, y, p, q)$, meaning that x is committed to y to bring about p if q holds.

Commitment Operations

Commitments are created, satisfied, and transformed in certain ways. Conventionally, six operations are defined on commitments. These are the $CREATE(x, C)$, the $CANCEL(x, C)$, the $RELEASE(y, C)$, the $ASSIGN(y, z, C)$, the $DELEGATE(x, z, C)$, and the $DISCHARGE(x, C)$. The $ASSIGN(y, z, C)$ operation replaces y with z as C 's creditor and the $DELEGATE(x, z, C)$ operation makes z the new debtor of the commitment C . A detailed exposition of these operations is given in (Singh, 1999) and is omitted here, for brevity.

The $DISCHARGE(\cdot, \cdot, \cdot)$ operation satisfies a commitment. A commitment is said to be *active* if it has been created, but not yet discharged.

2.2. RUNNING EXAMPLE

As a real-world example, we consider a variant of the NetBill protocol (Sirbu, 1997) used by a customer's agent to purchase a book from an online bookstore's agent. We identify four distinct, but related, scenarios that can arise during this purchase interaction. Each of these scenarios requires a different amount of effort from the participants in terms of protocol execution, planning, and coordination. Both agents would benefit from being able to compare scenarios to choose the one that best serves their interests.

1. The customer asks the bookstore for a price quote on a book, and upon receiving a quote from the bookstore, accepts the bookstore's offer. The bookstore sends the book, and the customer pays for it. Figure 1.a shows this interaction. This interaction sequence belongs to the *purchase* protocol.
2. The bookstore is willing to refund the price of returned books. This scenario is similar to the previous scenario till the book is delivered to the customer, but is longer, since the customer then returns the book for a refund. Figure 1.b. shows this interaction.
3. The customer delegates the payment to a third party, e.g., a bank. Such a situation is not very different from using a credit card to pay for goods, and is shown in Figure 1.c.

4. The customer wants insured shipping, and the bookstore's existing shipper does not insure goods. The bookstore negotiates with and contracts out the shipping to a shipper. Here, the shipper delivers the books to the customer, after which the shipper is paid by the bookstore. To complicate matters, the customer pays the bookstore via its bank like in the previous scenario. This scenario is shown in Figure 2.

In Figures 1.a, 1.b, 1.c, 2, 4.a, and 4.b, ellipses represent states, named s_i . Solid arrows are labeled by the messages that are passed between the participating agents. These messages correspond to actions that the agents take. Note that each of these figures represent a possible scenario, i.e., a *run* of the protocol. Also, states of the runs are drawn in different columns (also called swimlanes in UML parlance) to show the interacting agents clearly even though states are maintained by all interacting agents.

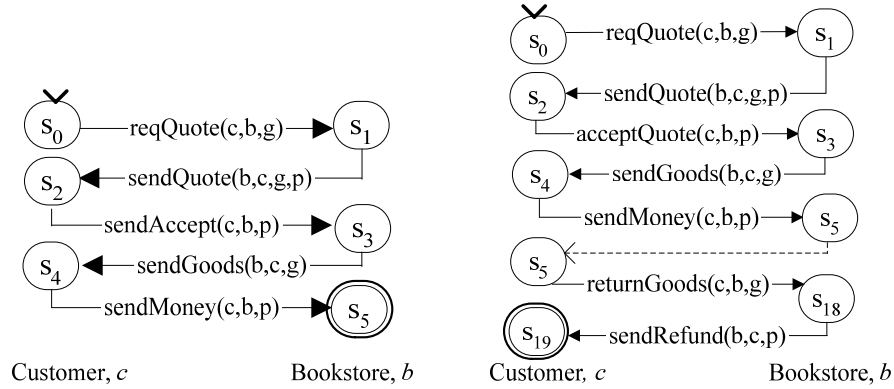
Table I explains the meanings of the states that the first scenario runs through. Table II shows the meanings of the messages passed, where c represents the customer, b , the bookstore, g , the book that the customer is interested in buying, and k , the customer's bank. The *delegate* message relates to corresponding commitment operation.

Table I. Meaning of states in the purchase protocol

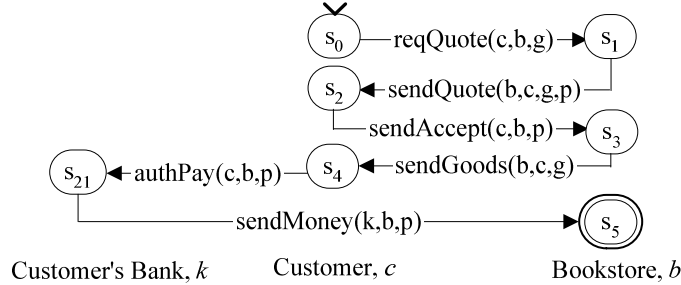
State	Meaning
s_1	Customer has asked the bookstore the price of the goods. No commitments made.
s_2	Bookstore has quoted a price for the said goods. The bookstore is now willing to send the goods if the customer promises to pay for them
s_3	Customer has agreed to the bookstores price. The customer is willing to pay the price if the books are delivered.
s_4	Bookstore has delivered the book.
s_5	Customer has paid for the book.

2.3. PROPOSITIONS

Propositions capture facts about what conditions hold, what commitments have been made, and whether these commitments have been fulfilled. The propositions used in a protocol are assumed to be understood by agents involved in the protocol. In the *purchase* example, we use the propositions given in Table III. In addition to these, active commitments are also represented as propositions, as we shall explain when discussing states.



a. Purchase protocol scenario 1: normal purchase b. Purchase protocol scenario 2: goods returned for a refund



c. Purchase protocol scenario 3: customer pays via bank

Figure 1. Three scenarios of the purchase example

2.4. ACTIONS

Agents perform actions to bring about changes in the world. In our framework, actions are modeled as messages sent by an agent to other agents. Just like an action, a message sent by an agent can affect the state of a protocol in which the agent participates. Messages may be implemented in different ways. For example, filling a form with credit card information and submitting it over the web is a message that represents a transfer of funds. The set of actions is denoted by \mathbb{A} . The meanings of the actions used in our *purchase* example are given in Table II. In addition to the actions shown, \mathbb{A} also contains actions corresponding to the commitment operations applied to each commitment. For example, \mathbb{A} contains an action $delegate(c, k, C)$ corresponding to the operation $delegate(c, k, C)$, where C is a commitment made by c .

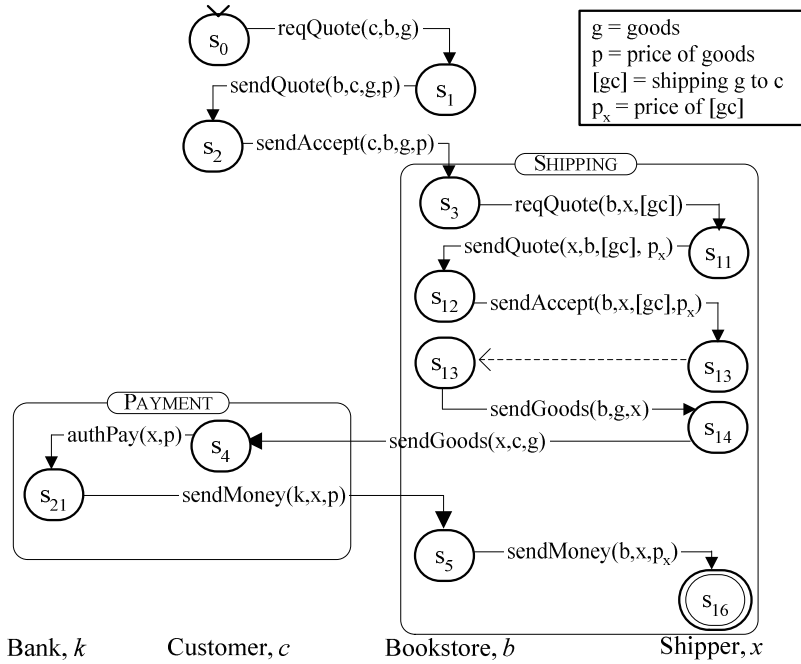


Figure 2. Purchase protocol scenario 4: shipping via a separate shipper and payment via bank

2.5. STATES

A protocol has many *states* that it goes through, during the course of its execution. A state is a snapshot of the world and is labeled by the set of propositions that are true in it. The propositions in the universe of discourse are termed the *frame*. A frame serves as a common ontology for the propositions used by a protocol. Frames provide the universe of discourse of a protocol. A state is an assignment of truth values to propositions. For example, state s_1 of the *purchase* example is labeled by the set $\{reqQuote(c, b, g)\}$ and state s_0 by $\{true\}$. We denote the label of a state s by $[s]$. Table IV shows the labels that are assigned to states in the *purchase* protocol. The set of states is denoted by \mathbb{S} . We include in this set a unique start state s_ϕ , which is labeled by the set $\{true\}$. In the *purchase* example, $s_0 = s_\phi$.

2.6. RUNS

A run is one possible execution sequence of a protocol. A protocol can allow many computations, or *runs*. A run is a sequence of states $\langle s_0 \dots s_i \dots \rangle$. We use \in to indicate the occurrence of a state on in a run. For example, $s_i \in \langle s_0 \dots s_i \dots \rangle$. In this paper, we consider only finite runs. The empty run is allowed.

Table II. Meanings of actions (modeled as messages) in the purchase protocol

Message	Meaning
$reqQuote(c, b, g)$	c asks b what the price of g is
$sendQuote(b, c, g, p)$	b quotes price p to the c , for g
$sendAccept(c, b, g, p)$	c accepts the price p quoted by b for g . c is now committed to pay if the book is sent to it
$sendGoods(b, c, g)$	b sends g to c
$sendMoney(c, b, p)$	c sends the money p to b
$delegate(c, k, C)$	c delegates the commitment C to k
$returnGoods(c, b, g)$	c returns g to b
$sendRefund(b, c, p)$	b refunds the money p to c
$authPay(c, b, p)$	c authorizes its bank to pay the amount p to b ; essentially c delegates $C(c, b, p)$ to k

Table III. Meanings of propositions used in the purchase protocol

Proposition	Meaning
$reqQuote(c, b, g)$	c has requested a quote for g from b
$quote(b, c, g, p)$	b quotes to c price p for g , i.e., b will deliver if c commits to pay upon delivery. This is represented by $CC(b, c, goods(b, c, g), acceptQuote(c, b, g, p))$
$acceptQuote(c, b, g, p)$	c has accepted the price p that b quoted for g , i.e. c commits to pay if the goods are delivered. This is represented by $CC(c, b, pay(c, b, p), goods(b, c, g))$
$goods(b, c, g)$	g has been delivered to c by b
$pay(c, b, p)$	The amount p has been paid to b by c
$return(c, b, g)$	g has been returned to b by c
$refund(b, c, p)$	The amount p has been refunded to c by b

The operator \prec_τ orders states temporally with respect to a run τ , so that $s_i \prec_\tau s_j$ implies that s_i occurs before s_j in the run τ . The concatenation of a state s_n to a run $\tau = \langle s_0 s_1 \dots s_k \rangle$, written as $\tau \circ s_n$, is given by appending the state to the run, i.e., by $\langle s_0 s_1 \dots s_k s_n \rangle$. The sequence of states in a *subrun* of a run $\tau = \langle s_0 \dots s_n \rangle$ is a possibly noncontiguous subsequence of $s_0 \dots s_n$. We denote the first state of a run τ by $[\tau]^0$, the n th state by $[\tau]^n$, and the last state by $[\tau]^T$.

Table IV. State labels in the purchase protocol

State	Associated Label
s_0	$\{\text{true}\}$
s_1	$\{\text{reqQuote}(c, b, g)\}$
s_2	$\{\text{quote}(b, c, g, p)\}$
s_3	$\{\text{C}(b, c, \text{goods}(b, c, g)), \text{CC}(c, b, \text{pay}(c, b, p), \text{goods}(b, c, g))\}$
s_4	$\{\text{goods}(b, c, g), \text{C}(c, b, \text{pay}(c, b, p))\}$
s_5	$\{\text{goods}(b, c, g), \text{pay}(c, b, p)\}$
s_{21}	$\{\text{goods}(b, c, g), \text{C}(k, b, \text{pay}(k, b, p))\}$
s_{17}	$\{\text{goods}(b, c, g), \text{pay}(c, b, p)\}$
s_{18}	$\{\text{goods}(b, c, g), \text{return}(c, b, g), \text{C}(b, c, \text{refund}(b, c, p))\}$
s_{19}	$\{\text{goods}(b, c, g), \text{return}(c, b, g), \text{refund}(b, c, p)\}$

2.7. PROTOCOLS

Computationally, a protocol corresponds to a set of computations that it allows. These can be captured as a set of runs where any of the runs that *subsume* the given runs may be realized. That is, each run in a protocol defines a sequence of steps that must be performed in the same order relative to each other. The subsumption of runs is defined below. A protocol is represented as a transition system as defined by a tuple $\langle \mathbb{A}, \mathbb{S}, S_0, \Delta, \mathbb{F}, \mathbb{R} \rangle$ where \mathbb{A} is a set of actions, \mathbb{S} is a set of states, S_0 is the set of initial states ($S_0 \subseteq \mathbb{S}$), Δ is a set of transitions ($\Delta \subseteq \mathbb{S} \times \mathbb{A} \times \mathbb{S}$), \mathbb{F} is a set of final states, ($\mathbb{F} \subseteq \mathbb{S}$), and \mathbb{R} is a set of roles (or participants).

Δ contains transitions of the form $\langle s_i, a, s_j \rangle$, where $s_i, s_j \in \mathbb{S}$ and $a \in \mathbb{A}$. Here s_i is the source of the transition and s_j its destination. Such a transition advances a computation from state s_i to state s_j when an action a is performed, i.e., when the message corresponding to a is sent (and received, assuming synchronous message passing, for convenience). In other words, a run can be generated from a protocol by the successive concatenation of transitions beginning from the initial state of the transition system. The concatenation of a transition to a run appends the destination of transition to the run if the source of the transition matches the last state of the run. Consequently, a run $\langle s_0 s_1 s_2 \dots s_n \rangle$ can be generated by a protocol whose initial state is s_0 , and whose transition set contains the elements $\langle s_0, -, s_1 \rangle, \langle s_1, -, s_2 \rangle$ and so on till $\langle s_{n-1}, -, s_n \rangle$, where $s_0 \in S_0$ and $s_n \in \mathbb{F}$. The set of all such runs is denoted by $[P]$.

Protocols are specified by propositions and actions that cause states to change. The semantics of actions are given in terms of commitments such as those shown in Table II. Given the actions and their semantics, the formal-

ization of a protocol is straightforward. The transition function of a protocol can be specified explicitly as state-action-state triples or as a set of rules that are compiled into such triples for runtime efficiency. Two example transition mechanisms for commitment-based protocols are commitment machines (Yolum and Singh, 2002) and nonmonotonic commitment machines (Chopra and Singh, 2003). For example, Tables III, II, and IV, along with a set of rules for determining the new state given the old state and the action taken would define the purchase protocol.

3. Reasoning about Protocols

This section describes our theory of comparing and refining protocols. Section 3.1 defines how states are deemed similar to one another, Section 3.2 defines what it means for a run to subsume another or be similar to another, Section 3.3 defines subsumption and similarity of protocols, and Section 3.4 uses comparisons of commitment-operation based propositions to relate different protocols.

3.1. SIMILARITY OF STATES

States form the fundamental components of runs, and are identified (and labeled) by sets of propositions. Any comparison of states, therefore, must be based on comparing propositions. This section introduces three state-similarity functions ι , σ , and $\alpha_{A,P}$, all based on commitment propositions, and shows how these help relate different runs.

A *state-similarity function* f is a mapping from a state to a set of states, i.e., $f : \mathbb{S} \mapsto 2^{\mathbb{S}}$. From such a function, we can induce a binary relation $\approx_f \subseteq \mathbb{S} \times \mathbb{S}$, which is defined as $\approx_f = \{(s, f(s)) : s \in \mathbb{S}\}$. That is,

$$s_i \approx_f s_j \iff s_j \in f(s_i) \quad (1)$$

For a well formed state-similarity function f , \approx_f must be reflexive, symmetric, and transitive, i.e., an equivalence relation.

We constrain all runs to prevent *stuttering*, i.e., to not have consecutive states that are similar, i.e., for every run $\tau = \langle s_0 \dots s_n \rangle$, s_i is not similar under f to s_{i+1} , for all $0 \leq i < n$.

Identity State-Similarity. ι is the identity state-similarity function. That is, $s_i \approx_\iota s_j$ if and only if s_i and s_j are labeled by same set of propositions. $\iota(s_i) = \{s_j \mid [s_i] = [s_j]\}$. \approx_ι is an equivalence relation.

Creditor State-Similarity. As another state-similarity function, consider σ . Under σ , a state s_i is similar to a state s_j if in the two states all the participants of the protocol have the same commitments being made towards

them, regardless of which agent makes it. Since the creditor of a commitment is immaterial under σ and a $delegate(\cdot, \cdot, \cdot)$ action changes the creditor of a commitment, σ can be defined as $\sigma(s_i) = \{s_j | s_j \text{ can be reached by finite number of } delegate(\cdot, \cdot, \cdot) \text{ actions from } s_i\}$

As an example, consider states s_4 and s_{21} from the of the example scenarios. These states are similar under σ because, as described in Table IV, these states have propositions representing commitments that differ only in their creditors. \approx_σ is an equivalence relation.

Role-and-Commitment State-Similarity. A state s_i is similar to a state s_j under $\alpha_{A,P}$, where A is a set of roles and P is a set of propositions, if the commitments made by any role in A to any other role in A , and the propositions in P that hold at s_i , also hold at s_j . If, for example, A represents all the roles in a protocol and P represents all the proposition used by that protocol, then $\alpha_{A,P}$ can be used as a similarity function to detect cases where the protocol has been merged with other protocols. This is explained in greater detail in Section 3.4.2.

3.2. SUBSUMPTION AND SIMILARITY OF RUNS

Comparisons among protocols are based on a notion of *subsumption* of runs. $\llbracket f \rrbracket$ denotes subsumption operator over runs. The operator $\llbracket f \rrbracket$ is an order-preserving mapping from one run to another, and depends on the function f .

Definition 1. A run τ_j *subsumes* a run τ_i under function f if and only if, for every state s_i that occurs in τ_i , there occurs a state s_j in τ_j that is similar under f , and s_j has the same temporal order relative to other states in τ_j as s_i does with states in τ_i .

$$\begin{aligned} \tau_j \llbracket f \rrbracket \tau_i \iff & (\forall s_i, s'_i \in \tau_i : (\exists s_j, s'_j \in \tau_j \\ & \text{and } s_i \approx_f s_j \text{ and } s'_i \approx_f s'_j \\ & \text{and } (s_i \prec_{\tau_i} s'_i \Rightarrow s_j \prec_{\tau_j} s'_j))) \end{aligned} \quad (2)$$

That is, longer runs subsume shorter ones, provided they have similar states occurring in the same order. Before we describe properties of run subsumption, we define *run-similarity*.

Definition 2. A run τ_i is similar to a run τ_j under a well-formed state-similarity function f if and only if the two runs are of equal length and every k th state of τ_i is similar under f to the k th state in τ_j . In notation, $[\tau_i]^k \approx_f [\tau_j]^k, 0 \leq k < |\tau_i|$ and $|\tau_i| = |\tau_j|$.

Lemma 1. If a run τ_j subsumes a run τ_i under f , then, for all subruns β_i of τ_i , there exists a subrun β_j of τ_j such that $\beta_i \approx_f \beta_j$. That is, $\tau_j \llbracket f \rrbracket \tau_i \Rightarrow (\forall \beta_i \text{ subrun of } \tau_i, \exists \beta_j \text{ subrun of } \tau_j : \beta_i \approx_f \beta_j)$.

Proof. This lemma holds trivially for subruns of lengths 0 or 1. We prove this lemma by induction on the length of subruns. As the base case, we note that by the definition of run subsumption, there must be a pair of states in both runs for which the above property holds, i.e., $(\forall \beta_i \text{ subrun of } \tau_i : |\beta_i| = 2 \Rightarrow (\exists \beta_j \text{ subrun of } \tau_j \text{ and } \beta_i \approx_f \beta_j))$. Assume this property holds for subruns of a particular length $k > 2$. Next, consider $|\beta'_i| = k + 1$, where $[\beta'_i]^k = s'_i$. The first k states of β'_i form a subrun of length k . Therefore, there exists a subrun β_j of τ_j such that $\beta_i \approx_f \beta_j$ and β_j is the earliest such subrun in τ_j . Consider the k pairwise temporal precedence relationships (in τ_i) between the first k states of β'_i and s'_i . By the definition of run subsumption, we know that since $\tau_j \llbracket f \rrbracket \tau_i$, $\exists s'_j \in \tau_j$ such that $[\beta_j]^\top \prec_{\tau_j} s'_j$ and $s'_j \approx_f s'_i$. Let s'_j be the earliest such state in τ_j , so that $\beta'_j = \beta_j \circ s'_j$. Thus, the inductive hypothesis holds for subruns of length $k + 1$. By induction, the lemma holds. ■

Theorem 1. Under any well-formed state similarity function, run subsumption is reflexive, transitive, and antisymmetric up to state similarity.

Proof. The reflexivity of run subsumption follows from the reflexivity of state-similarity functions. The transitivity of run subsumption follow from its definition. The antisymmetry of run subsumption means that if two runs subsume each other under some well-formed state similarity function f , then those runs are similar under f . For proof, consider Lemma 1. By Lemma 1 we know that if $\tau_j \llbracket f \rrbracket \tau_i$, then for each subrun of τ_i , there is a similar subrun of τ_j . Let the entire run τ_i be treated as a subrun of itself. Then, there is a similar subrun in τ_j , i.e., $|\tau_j| \geq |\tau_i|$. In a similar manner, from $\tau_i \llbracket f \rrbracket \tau_j$, we infer that $|\tau_i| \geq |\tau_j|$. Hence, $|\tau_i| = |\tau_j|$. Since the runs are of equal length and all the subruns of one have similar subruns in the other, $\tau_i \approx_f \tau_j$ if τ_i and τ_j subsume each other under f . ■

Let us consider an example to better explain the above concepts. Let τ_1 , τ_2 , and τ_3 be the runs shown in Figures 1.a, 1.b, and 1.c respectively. We then have $\tau_2 \llbracket \iota \rrbracket \tau_1$. Also, $\tau_3 \llbracket \iota \rrbracket \tau_1$. However, τ_1 subsumes neither τ_2 nor τ_3 , because τ_3 has a state s_{21} , whose label does not match any state label in τ_1 , and τ_2 has states s_{18} and s_{19} , whose labels do not match any state label in τ_1 . Run τ_3 does not subsume τ_2 because of s_{21} , and τ_2 does not subsume τ_3 because of s_{18} and s_{19} .

Next, we consider run subsumption under the creditor state-similarity function σ . Here, $\tau_2 \llbracket \sigma \rrbracket \tau_1$. Also, since $s_{21} \in \sigma(s_4)$, $\tau_1 \llbracket \sigma \rrbracket \tau_2$. Run τ_3 subsumes

both τ_1 and τ_2 for the same reason, but neither τ_1 nor τ_2 subsumes τ_3 , since neither of them have states that are σ -similar to s_{18} and s_{19} .

Figure 3.a shows the subsumption relation between the runs τ_1 , τ_2 , and τ_3 under the identity function ι and Figure 3.b, under σ .



a. ι subsumption

b. σ subsumption

Figure 3. Subsumption of purchase protocol runs under ι and σ

3.3. SUBSUMPTION AND SIMILARITY OF PROTOCOLS

Closure of Protocol Runs In line with our intuitions that generic protocols allow many variations, we constrain all protocols to have a set of runs that is closed under run subsumption. That is, if any run occurs in the set of runs of a protocol, all the runs that subsume it (under a well-formed state-similarity function) are also found in that set. Given a well formed state-similarity function f , a protocol P contains all runs that subsume any run in $[P]$. We define this new set of runs as the set of runs allowed by the protocol, and denote this set by $\llbracket P \rrbracket = \{\tau \mid \forall \tau' \in [P] : \tau \llbracket f \rrbracket \tau'\}$.

Operationally, the runs allowed by a protocol completely characterize that protocol by naturally illustrating the key tradeoff in protocol design, that of *flexibility versus compliance*: A protocol that allows many runs is better than one that allows a few runs, since the many-run protocol affords more choice and flexibility in protocol execution to the participants. However, checking for compliance is computationally more demanding when protocols have numerous runs. The definition of the subsumption of protocols reflects these intuitions.

Definition 3. A protocol P_j subsumes a protocol P_i under a state-similarity function f if and only if, every run in $\llbracket P_i \rrbracket$ subsumes, under f , a run in $\llbracket P_j \rrbracket$.

$$P_j \llbracket f \rrbracket P_i \iff \forall \tau_i \in \llbracket P_i \rrbracket \exists \tau_j \in \llbracket P_j \rrbracket : \tau_i \llbracket f \rrbracket \tau_j \quad (3)$$

If P_j is a protocol that has short runs (and consequently all runs that subsume them) and P_i is a protocol that has long runs only, then P_j subsumes P_i . Since long runs subsume shorter ones, protocols with long runs only are subsumed by protocols with short runs as well.

Before we describe the properties of protocol subsumption, we define protocol similarity as follows:

Definition 4. Two protocols are similar under a state similarity function f if and only if for every run in one protocol, there exists at least one run in the other protocol that is similar under f and vice versa.

$$P_j \approx_f P_i \Leftrightarrow \forall \tau_i \in \llbracket P_i \rrbracket, \exists \tau_j \in \llbracket P_j \rrbracket : \tau_i \approx_f \tau_j \quad (4)$$

and $\forall \tau_j \in \llbracket P_j \rrbracket, \exists \tau_i \in \llbracket P_i \rrbracket : \tau_j \approx_f \tau_i$

Theorem 2. For any state-similarity function f , the protocol subsumption relation $\llbracket f \rrbracket$ is a partial order, i.e., reflexive, transitive, and antisymmetric (up to state similarity).

Proof. Reflexivity and transitivity of protocol subsumption follow from its definition. Because of the property of closure of the set of runs of a protocol under run subsumption and the definition of protocol subsumption, we know that if $P_i \llbracket f \rrbracket P_j$, then $\forall \tau_j \in \llbracket P_j \rrbracket : \exists \tau_i \in \llbracket P_i \rrbracket : \tau_j \approx_f \tau_i$. Conversely, if $P_j \llbracket f \rrbracket P_i$, then $\forall \tau_i \in \llbracket P_i \rrbracket : \exists \tau_j \in \llbracket P_j \rrbracket : \tau_i \approx_f \tau_j$. Together, these are equivalent to $P_j \approx_f P_i$. ■

3.4. THE PROTOCOL ALGEBRA

We now introduce our protocol algebra as consisting of two operators (*merge* and *choice*), their identity elements ($\mathbb{1}$ and $\mathbb{0}$, respectively), and an ordering relationship (protocol subsumption, as defined above).

Merge. The merge operator, denoted by \otimes_f , splices two protocols (under a state-similarity function f) to produce a new protocol. A merge of two protocols involves interleaving their runs.

Definition 5. The merge of two protocols P and Q under a well-formed state-similarity function f creates a protocol R , each of whose runs subsume some run from $\llbracket P \rrbracket$ and some run from $\llbracket Q \rrbracket$.

$$\llbracket P \otimes_f Q \rrbracket = \{r \mid \exists r_p \in \llbracket P \rrbracket, \exists r_q \in \llbracket Q \rrbracket : r \llbracket f \rrbracket r_p \text{ and } r \llbracket f \rrbracket r_q\} \quad (5)$$

Choice. The choice operator, denoted by \oplus_f , combines two protocols P and Q (under a state similarity function f) to create a protocol R , whose set of runs $\llbracket R \rrbracket$ contains exactly those runs that exist in the sets $\llbracket P \rrbracket$ and $\llbracket Q \rrbracket$.

Definition 6. The choice of two protocols P and Q under a well-formed state-similarity function f creates a protocol R , each run of which subsumes either a run from $\llbracket P \rrbracket$ or a run from $\llbracket Q \rrbracket$

$$\llbracket P \oplus_f Q \rrbracket = \{r \mid \exists r_p \in \llbracket P \rrbracket, \exists r_q \in \llbracket Q \rrbracket : r \llbracket f \rrbracket r_p \text{ or } r \llbracket f \rrbracket r_q\} \quad (6)$$

Choosing a run from $\llbracket P \oplus_f Q \rrbracket$ is equivalent to choosing a run from either $\llbracket P \rrbracket$ or a run from $\llbracket Q \rrbracket$.

Constants. The properties of the merge operator lead us to define two protocols, $\mathbb{0}$ and $\mathbb{1}$. The $\mathbb{0}$ protocol is an “impossible” protocol, which does not have any runs. $\llbracket \mathbb{0} \rrbracket = \{\}$. The $\mathbb{1}$ protocol is a “trivial” protocol which contains the zero-length run in its run-set, and consequently contains all possible runs. $\llbracket \mathbb{1} \rrbracket = \{\tau : \tau \llbracket f \rrbracket \tau_\phi\}$, where $\tau_\phi = \langle \rangle$, and f is any well-formed state-similarity function. The $\mathbb{0}$ and the $\mathbb{1}$ protocols form the bottom and the top element, respectively, of a protocol hierarchy based on the merge function. All protocols are subsumed by the $\mathbb{1}$ protocol and all protocols subsume the $\mathbb{0}$ protocol.

3.4.1. Formal Results

We briefly present some formal results, which simplify reasoning about protocols using our algebra. First, we present a formulation for the operators in terms of protocol subsumption. Since definitions 5 and 6 state the properties of the operators only in terms of the runs of the protocols, this new formulation simplifies the understanding of the behavior of the merge and choice operators, and likens them to set intersection and union, respectively.

Theorem 3. Given the run-subsumption relation $\llbracket f \rrbracket$ under a well-formed state-similarity function f , the merge \otimes_f of two protocols P and Q under f is the greatest lower bound of P and Q under the protocol subsumption relation $\llbracket f \rrbracket$.

Proof. Let $R = P \otimes_f Q$. This implies that

$$\llbracket R \rrbracket = \{r \mid \exists p \in \llbracket P \rrbracket, \exists q \in \llbracket Q \rrbracket : r \llbracket f \rrbracket p \text{ and } r \llbracket f \rrbracket q\} \quad (7)$$

Consider some protocol R' such that $P \llbracket f \rrbracket R'$ and $Q \llbracket f \rrbracket R'$. By definition of protocol subsumption,

$$\forall r'_p \in \llbracket R' \rrbracket, \exists p \in \llbracket P \rrbracket : r'_p \llbracket f \rrbracket p \quad (8)$$

$$\forall r'_q \in \llbracket R' \rrbracket, \exists q \in \llbracket Q \rrbracket : r'_q \llbracket f \rrbracket q \quad (9)$$

From equation 7 we know that all runs that subsume a run each from $\llbracket P \rrbracket$ and $\llbracket Q \rrbracket$ are in $\llbracket R \rrbracket$. From equations 8 and 9, we know that every run in $\llbracket R' \rrbracket$ is just such a run. Hence, $\forall r' \in \llbracket R' \rrbracket, r' \in \llbracket R \rrbracket$. From the closure property of protocol run-sets, we further infer that $\forall r' \in \llbracket R' \rrbracket, \exists r \in \llbracket R \rrbracket : r \llbracket f \rrbracket r'$, which means that $R \llbracket f \rrbracket R'$. Since any R' is subsumed by $P \otimes_f Q$ under f , \otimes_f gives the greatest lower bound of P and Q under protocol subsumption under f . ■

Essentially, the merge is analogous to the intersection of the sets of runs of the protocols being merged, where run similarity is used to determine equality of elements of the sets.

Theorem 4. Given the run-subsumption relation $\llbracket f \rrbracket$ under a well-formed state-similarity function f , the choice \oplus_f of two protocols P and Q under f is the least upper bound of P and Q under the protocol subsumption relation $\llbracket f \rrbracket$.

Proof. Let $R = P \oplus_f Q$. This implies that

$$\llbracket R \rrbracket = \{r \mid \exists p \in \llbracket P \rrbracket : r \llbracket f \rrbracket p \text{ or } \exists q \in \llbracket Q \rrbracket : r \llbracket f \rrbracket q\} \quad (10)$$

Consider some protocol R' such that $R' \llbracket f \rrbracket P$ and $R' \llbracket f \rrbracket Q$. By definition of protocol subsumption,

$$\forall p \in \llbracket P \rrbracket, \exists r'_p \in \llbracket R' \rrbracket : p \llbracket f \rrbracket r'_p \quad (11)$$

$$\forall q \in \llbracket Q \rrbracket, \exists r'_q \in \llbracket R' \rrbracket : q \llbracket f \rrbracket r'_q \quad (12)$$

Let $r \in \llbracket R \rrbracket$. From equation 10, without loss of generality, we can assume that $\exists p \in \llbracket P \rrbracket : r \llbracket f \rrbracket p$. From this, equation 11, and the transitivity of run subsumption, we infer that $\forall r \in \llbracket R \rrbracket, \exists r'_p \in \llbracket R' \rrbracket : r \llbracket f \rrbracket r'_p$, which is the definition of protocol subsumption under f . Hence $R' \llbracket f \rrbracket R$. Since any R' subsumes $P \oplus_f Q$ under f , \oplus_f gives the least upper bound of P and Q under protocol subsumption under f . ■

Choice is analogous to the union of the sets of the runs, using run similarity to compare elements of the sets.

Next, we state some simple properties of the protocol-algebraic operators that are immediate consequences of the properties of the operators described in the theorems above.

Since each run in $\llbracket P \otimes_f Q \rrbracket$ subsumes a run each from $\llbracket P \rrbracket$ and from $\llbracket Q \rrbracket$, the following properties apply to the merge operator

1. The merge of a protocol with another refines the given protocol.

$$\begin{aligned} P \llbracket f \rrbracket (P \otimes_f Q) \\ Q \llbracket f \rrbracket (P \otimes_f Q) \end{aligned}$$
2. The merge of a protocol with itself yields the same protocol (idempotence).

$$\begin{aligned} P \llbracket f \rrbracket (P \otimes_f P) \\ (P \otimes_f P) \llbracket f \rrbracket P \\ (P \otimes_f P) = P \end{aligned}$$

3. The merge operator is commutative and associative.

$$P \otimes_f Q = Q \otimes_f P$$

$$P \otimes_f (Q \otimes_f R) = (P \otimes_f Q) \otimes_f R$$

4. Merge distributes over choice.

$$P \otimes_f (Q \oplus_f M) = (P \otimes_f Q) \oplus_f (P \otimes_f M)$$

5. The merge of any protocol with $\mathbb{1}$ gives that protocol and the merge of any protocol with $\mathbb{0}$ gives $\mathbb{0}$. In this way, the $\mathbb{1}$ protocol is the identity element and the $\mathbb{0}$ is the nil element for merge.

$$P \otimes_f \mathbb{1} = P$$

$$P \otimes_f \mathbb{0} = \mathbb{0}$$

Choice also supports idempotence, commutativity, and associativity. The choice of a protocol with $\mathbb{1}$ yields $\mathbb{1}$ and the choice of a protocol with $\mathbb{0}$ yields that protocol itself.

1. Idempotence.

$$(P \oplus_f P) = P$$

2. Commutativity.

$$(P \oplus_f Q) = (Q \oplus_f P)$$

3. Associativity.

$$P \oplus_f (Q \oplus_f R) = (P \oplus_f Q) \oplus_f R$$

4. Choice distributes over merge.

$$P \oplus_f (Q \otimes_f R) = (P \oplus_f Q) \otimes_f (P \oplus_f R)$$

5. Choice with $\mathbb{1}$ and $\mathbb{0}$

$$P \oplus_f \mathbb{1} = \mathbb{1}$$

$$P \oplus_f \mathbb{0} = P$$

3.4.2. Applying the Algebra

Now we discuss how the algebra can be applied to create new protocols. The choice operator \oplus_f allows us to choose between runs belonging to different protocols. This operator can be used, for example, when multiple ways of payment exist, such as payment by credit card, or payment by personal check. The result of the choice operator is a protocol whose set of runs is larger and thus offers more choices than the individual protocols to which the choice was applied.

The merge operator is more interesting. As an example of its application, consider the run shown in Figure 1.c. This run belongs to the merge of the simple purchase shown in Figure 1.a. and payment, shown in Figure 4.a. The

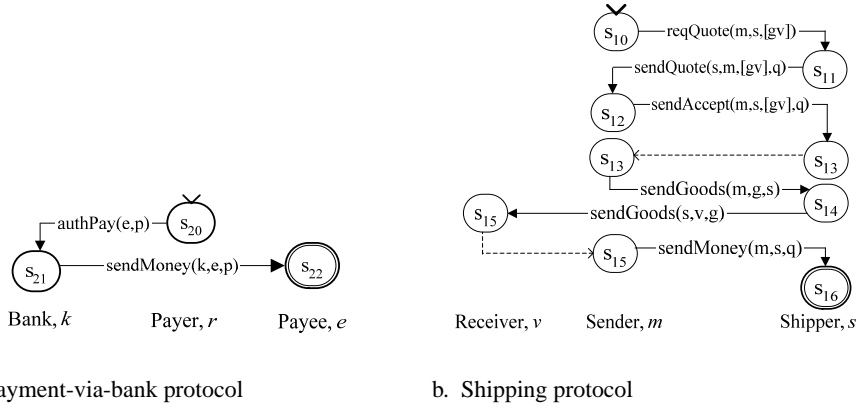


Figure 4. A payment and a shipping protocol

merge is performed under the creditor state-similarity function σ . As a more complicated, consider a run of the refined purchase example as shown in Figure 2. This run belongs to the refined purchase protocol, which is the result of a merge of the simple purchase, the shipping, and the payment protocols. The state-similarity function used here is $\alpha_{A,P}$. Under $\alpha_{A,P}$, A is a set of agents, P , a set of propositions, and two states are similar if all commitments between agents in A and all propositions in P that exist in one state also exist in the other. Under $\alpha_{A,P}$, where A denotes the set containing the participants of *Shipping*, i.e., $\{b, x, c\}$, and P denotes the set of all propositions that are used in *Shipping*, we see that the *Shipping* run shown in Figure 4.b. is subsumed by the refined *Purchase* run shown in Figure 2. Specifically, the states $s_3, s_{11}, s_{12}, s_{13}, s_{14}, s_5$, and s_{16} of refined *Purchase* are $\alpha_{A,P}$ -similar to the states $s_{10}, s_{11}, s_{12}, s_{13}, s_{14}, s_{15}$, and s_{16} of *Shipping* respectively. Similarly, the states s_4 and s_{21} of the refined *Purchase* are similar to states s_{20} and s_{21} of *Payment* under $\alpha_{A,P}$, where A denotes $\{c, k\}$ and P denotes the set of all propositions used in *Payment*. Consequently, the refined *Purchase* run subsumes *Payment*. Note that Figure 2 shows only one run of the refined purchase protocol. Given the semantics of the merge operator, the refined purchase protocol allows more runs, since all valid interleavings of runs of the merged protocols are allowed. One such run could be where the shipping protocol is started before the first step of the purchase protocol. In practice, data dependencies and temporal state ordering are specified to filter the set of runs generated by interleaving.

4. Discussion

Our research program seeks to develop rich abstractions methodologies that will ease the development of large-scale open systems. This paper is part of

our ongoing research in that direction. Singh *et al.* present an overview of the motivations of this program and of the applications envisioned for it (Singh *et al.*, 2004).

The framework presented in this paper can serve as the foundation for developing design-time tools, and possibly for automated, runtime composition. Complete automation of protocol composition requires a complete specification of the behavior of a protocol. This is rarely the case when dealing with complex agent interaction protocols that can find applications in business processes modeling or Web Service composition. Most realistic settings require considerable context-sensitive information, which may be encoded as policies local to the agent. Such contexts may be based on motivations such as trust and economics, which can change unpredictably. The autonomy of agents in a multi agent system allows agents to behave differently under different contexts. It is this dynamic behavior that makes the agent paradigm attractive for application to open systems, and cannot be statically specified for all but the simplest of agents. Our framework helps develop tools that aid protocol designers in tailoring an existing protocol to meet their requirements by automatically verifying properties of the designed protocol.

Our understanding of protocols as specifications of the minimum states that a computation should contain is analogous to the minimal process execution semantics as defined in the MIT Process Handbook (Malone *et al.*, 2003). There is an underlying assumption that concise specifications are better than elaborate ones, since flexibility of a protocol is desirable in business applications where opportunities can be profitably exploited. In some cases, however, a maximal execution semantics might be applicable, e.g., in a protocol for which compliance checking is costly or difficult or where unexpected actions are undesirable.

4.1. LITERATURE

Our work relates to and draws both from well established and emerging fields. Business processes have received much attention lately because of the economic benefits of cross-enterprise business. Coordination and interaction protocols have been studied by the agent research community so that agent conversations and interaction can be computationally realized. We list here selected literature from both areas.

4.1.1. *Workflows and Processes.*

Business processes have been traditionally automated as workflows. Recently, the web service model has been applied to process automation.

Workflows have been studied extensively as Petri-net based models of business processes (van der Aalst *et al.*, 2000; van der Aalst and van Hee, 2002). The Workflow Management Coalition (WfMC) is a standards body

that has created a reference model for workflows (Fischer, 2004). This model has two basic parts, a modeling and an enactment part. The model prescribes a workflow engine as the system that executes the workflow. These models specify a rigid sequence of steps. Workflows require human intervention to handle most exceptions. Because of their inflexibility, workflows have had only limited success.

The MIT Process Handbook (Malone et al., 2003) is a project that aims to create a hierarchy of commonly used business processes. Based on this hierarchy, Grosf and Poon (2003) develop a system to represent and execute business rules.

Of late, web services have been touted as the solution to the business interoperability problem. The need for process composition and interoperability has led to the development of standards for orchestration and choreography of web services (Peltz, 2003). Orchestration refers to intra-service planning and choreography to an overall view of inter-service coordination. Here, we shall mention only two important standards, WSCI and BPEL4WS. The Web Services Choreography Interface (WSCI) is an XML-based language that describes a service interface by the flow of messages sent and received by the service. The standard, however, looks at protocols one level lower than our view, since each WSCI specification corresponds to a role in our scheme. The Business Process Execution Language for Web Services (BPEL4WS) is currently the most widely used web services standard for describing business processes (Andrews et al., 2003). However, BPEL4WS is no more than a procedural script encoded in XML.

Fu et al. (2004) develop methods to verify if a given web service will adhere to a given conversation protocol. Their work develops formal results about verification of protocol compliance for protocols based on finite state machines. Hamadi and Benatallah (2003) develop a protocol algebra for petri nets and show its applicability to workflows and web services. However, this approach suffers from the same pitfalls as workflows modeled using Petri nets.

4.1.2. *Interaction Protocols.*

Yolum and Singh (2002) give one of the first accounts of the use of commitments in modeling agent interaction protocols and the flexibility that it affords the participating agents. Fornara and Colombetti (2003) describe how commitments relate to FIPA-ACL messages and demonstrate with an example. Both approaches highlight the benefits of a commitment-based approach to interaction protocol design.

Johnson et al. (2003) develop a scheme for identifying when two commitment-based protocols are equivalent. Their scheme, however, is simplistic, classifying protocols based solely on their syntactic structure. Our work provides

stronger results about the relationships between protocols from an application point of view and relates better to the Web Services approach.

Bussmann et al. (2003) present a design methodology to aid in the selection of a protocol from a library of existing protocols to apply to agent-based control applications. They identify criteria like the number of agents, the number of roles, and the number and kind of commitments and use these to select a protocol from an existing pool of interaction protocols. This approach is quantitative, and lacks a formal semantics to base the methodology on.

Pitt and Mamdani (2000) describe a semantics for agent interaction protocols using the Belief-Desire-Intention (BDI) theory. Using this semantics, they outline the design of a system of agent plans that are instantiated by agents to carry on conversations with other BDI agents. In our work, an agent's beliefs, desires, and intentions are private to that agent. We work with social commitments which are observable by all agents and whose breach is easier to verify.

In more recent work, Vitteau and Huget (2004) describe an approach for designing agent interaction protocols using modular *micro-protocols*. This scheme is similar to our protocol design proposal in spirit. However, Vitteau and Huget do not provide a formal basis for putting protocols together.

4.2. CONCLUSIONS AND DIRECTIONS

The above is a semantic approach to commitment protocols that yields a simple algebra for protocols. This algebra provides a basis for conceptual reasoning about protocols in terms of refinement and aggregation, which is essential if we are to engineer protocols that way other software systems are engineered. To our knowledge, this work is unique in formulating the problem of problem design at a conceptual level. Partly, it derives its uniqueness from a careful consideration of the commitments that underlie protocols in multiagent settings.

This work opens up some interesting challenges. One, it would help consider how the algebra will work with more subtle kinds of state similarity functions. Two, the abstractions supported by our algebra must be woven into a methodology for designing protocols. Three, such methodologies should be supported by tools that give appropriate reasoning assistance to designers.

References

- Andrews, T., F. Curbera, H. Dholakia, Y. Golland, J. Klein, F. Leymann, K. Liu, D. Roller, D. Smith, S. Thatte, I. Trickovic, and S. Weerawarana: 2003, 'Business Process Execution Language for Web Services, (BPEL 1.1)'. www-106.ibm.com/developerworks/webservices/library/ws-bpel.

- Bussmann, S., N. R. Jennings, and M. Wooldridge: 2003, 'Re-use of Interaction Protocols for Agent-Based Applications'. In: F. Giunchiglia, J. Odell, and G. Weiß (eds.): *Third International Workshop on Agent-Oriented Software Engineering (AOSE 2002)*, Vol. 2585 of *Lecture Notes in Computer Science*. pp. 73–87, Springer.
- Castelfranchi, C.: 1993, '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*.
- Chopra, A. and M. P. Singh: 2003, 'Nonmonotonic Commitment Machines'. In: M.-P. Huget and F. Dignum (eds.): *Proceedings of the AAMAS-03 Workshop on Agent Communication Languages and Conversation Policies*.
- Fischer, L. (ed.): 2004, *The Workflow Handbook 2004*. Lighthouse Point, FL: Future Strategies Inc. Compendium of standards and articles by the Workflow Management Coalition (WfMC).
- Fornara, N. and M. Colombetti: 2003, 'Defining Interaction Protocols using a Commitment-based Agent Communication Language'. In: *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. pp. 520–527, ACM Press.
- Fu, X., T. Bultan, and J. Su: 2004, 'Realizability of Conversation Protocols with Message Contents'. In: *Proceedings of the 2004 International Conference on Web Services*. pp. 96–105, IEEE Computer Press.
- Grosz, B. N. and T. C. Poon: 2003, 'SweetDeal: Representing Agent Contracts with Exceptions using XML Rules, Ontologies, and Process Descriptions'. In: *Proceedings 12th International Conference on the World Wide Web*. pp. 340–349.
- Hamadi, R. and B. Benatallah: 2003, 'Proceedings of the 14th Australasian Database Conference'. In: X. Zhou and K.-D. Schewe (eds.): *ADC*, Vol. 17 of *Conferences in Research and Practice in Information Technology*. pp. 191–200, Australian Computer Society.
- Johnson, M. W., P. McBurney, and S. Parsons: 2003, 'When are Two Protocols the Same?'. In: M.-P. Huget (ed.): *Communication in Multiagent Systems: Agent Communication Languages and Conversation Policies*, Vol. 2650 of *Lecture Notes in Artificial Intelligence*. Berlin: Springer-Verlag, pp. 253–268.
- Malone, T. W., K. Crowston, and G. A. Herman (eds.): 2003, *Organizing Business Knowledge: The MIT Process Handbook*. Cambridge, MA: MIT Press.
- Peltz, C.: 2003, 'Web Service Orchestration and Choreography'. *IEEE Computer* **36**(10), 46–52.
- Pitt, J. and A. Mamdani: 2000, 'Communication Protocols in Multi-agent Systems'. In: F. Dignum and M. Greaves (eds.): *Issues in Agent Communication*, Vol. 1916 of *Lecture Notes in Artificial Intelligence*. pp. 160–177, Springer.
- Rosenberg, J., H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler: 2002, 'SIP: Session Initiation Protocol'. IETF Proposed Standard. Available as an RFC at <http://www.ietf.org/rfc/rfc3261.txt>.
- Singh, M. P.: 1991, 'Social and Psychological Commitments in Multiagent Systems'. In: *AAAI Fall Symposium on Knowledge and Action at Social and Organizational Levels*. pp. 104–106.
- Singh, M. P.: 1999, 'An Ontology for Commitments in Multiagent Systems: Toward a Unification of Normative Concepts'. *Artificial Intelligence and Law* **7**, 97–113.
- Singh, M. P., A. K. Chopra, N. Desai, and A. U. Mallya: 2004, 'Protocols for Processes: Programming in the Large for Open Systems (extended abstract)'. In: *OOPSLA Companion*. pp. 120–123.
- Sirbu, M. A.: 1997, 'Credits and Debits on the Internet'. *IEEE Spectrum* **34**(2), 23–29.

- van der Aalst, W., J. Desel, and A. Oberwis (eds.): 2000, *Business Process Management: Models, Techniques, and Empirical Studies*, Vol. 1806 of *Lecture Notes in Computer Science*. Berlin: Springer-Verlag.
- van der Aalst, W. and K. van Hee (eds.): 2002, *Workflow Management Models, Methods, and Systems*. Cambridge, MA: MIT Press.
- Verdicchio, M. and M. Colombetti: 2002, 'Commitments for Agent-Based Supply Chain Management'. *ACM SIGecom Exchanges* **3**(1), 13–23.
- Vitteau, B. and M.-P. Huet: 2004, 'Modularity in Interaction Protocols'. In: F. Dignum (ed.): *Advances in Agent Communication*, Vol. 2922 of *Lecture Notes in Artificial Intelligence*. pp. 291–309, Springer.
- Yolum, P. and M. P. Singh: 2002, '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)*. pp. 527–534, ACM Press.