

# Interaction Protocols as Design Abstractions for Business Processes

Nirmit Desai, Ashok U. Mallya, Amit K. Chopra, and Munindar P. Singh

**Abstract**—Business process modeling and enactment are notoriously complex, especially in open settings, where business partners are autonomous, requirements must be continually finessed, and exceptions frequently arise because of real-world or organizational problems. Traditional approaches, which attempt to capture processes as monolithic flows, have proven inadequate in addressing these challenges. We propose (business) protocols as components for developing business processes. A protocol is an abstract, modular, publishable specification of an interaction among different roles to be played by different participants. When instantiated with the participants' internal policies, protocols yield concrete business processes. Protocols are reusable and refinable, thus simplifying business process design. We show how protocols and their composition are theoretically founded in the  $\pi$ -calculus.

**Index Terms**—Multiagent systems, software reuse, interaction-based modeling, software design methodologies, rule-based processing,  $\pi$ -calculus.

## 1 INTRODUCTION

**B**USINESS processes in open settings typically involve complex interactions among autonomous and heterogeneous *business partners*. Conventionally, business processes are modeled as centralized flows, specifying exact steps for each participant. However, because of the exceptions and opportunities that arise in open environments, business relationships cannot be preconfigured to full detail. Flow-based models are inherently ill-suited to development and maintenance in the face of evolving requirements. Also, flows are not amenable to reuse via refinement or aggregation. Further, instead of treating all participants equally, conventional models classify the participants as clients and servers, thereby compromising their autonomy.

We apply multiagent systems to model business interactions. Agents mirror the autonomy and heterogeneity of real-world businesses and support rich interaction models. This paper describes a novel, agent-based approach for business process modeling and enactment. The key idea is to capture meaningful interactions as protocols. Protocols involve two or more roles and address specific purposes such as ordering, payment, shipping, etc. Protocols emphasize the essence of the interactions and omit local details. Such abstract protocol specifications are publishable as components and reusable in different settings.

Crucially, we give protocols a semantics in terms of *commitments*. Commitments among roles provide a basis for modeling the state of the interaction, thus allowing a variety of possible executions depending on the specific circum-

stances of the parties involved. Thus, commitments help capture the essence of the interaction supported by a protocol.

A business process reflects a composition of a set of protocols, to be enacted by agents representing real-world partners. Whereas protocols specify the interaction from a global perspective, *(role) skeletons* specify the interaction from the perspective of a particular role. During the enactment of a process, each agent adopts a role in the corresponding composite protocol, setting its role skeleton accordingly, and integrates the skeleton with its *policies*. A policy is a (typically, private) description of an agent's business logic that controls its participation in a process. For example, a protocol may allow a role to choose among multiple actions; the agent playing that role would select an action based on its policy. Policies help determine the parameters of the messages being sent, typically based on the parameters of the messages previously received.

Our theoretical contributions are developing the concept of commitment protocols and applying the  $\pi$ -calculus to formalize certain aspects of protocols and policies. We show how to formally construct composite protocols and derive processes by integrating protocols and policies. The formalization enables us to reason about properties of protocols such as their incorrectness, compatibility, equivalence, and flexibility.

Our practical contributions include a language and tools for protocols, called OWL-P. OWL-P, captured as an ontology expressed in the Web Ontology Language (OWL) [1], provides the primitives to specify protocols, such as roles, the messages exchanged between them, and declarative rules describing the effects of messages in terms of commitments. Our approach includes:

1. specifying protocols,
2. composing protocols,
3. generating role skeletons from a protocol,
4. compiling protocols into executable rules,
5. instantiating protocols, and
6. enacting processes.

This paper focuses on the first four steps.

- N. Desai, A.K. Chopra, and M.P. Singh are with the Department of Computer Science, North Carolina State University, Suite 110, Box 7535, Raleigh, NC 27695-7535. E-mail: {nvdesai, akchopra, singh}@ncsu.edu.
- A.U. Mallya is with Veraz Networks, 926 Rock Avenue, San Jose, CA 95131. E-mail: aumallya@ncsu.edu.

Manuscript received 6 May 2005; revised 27 Oct. 2005; accepted 7 Nov. 2005; published online 22 Dec. 2005.

Recommended for acceptance by S. Uchitel, M. Broy, I.H. Krüger, and J. Whittle.

For information on obtaining reprints of this article, please send e-mail to: tse@computer.org, and reference IEEECS Log Number TSE-0142-0505.

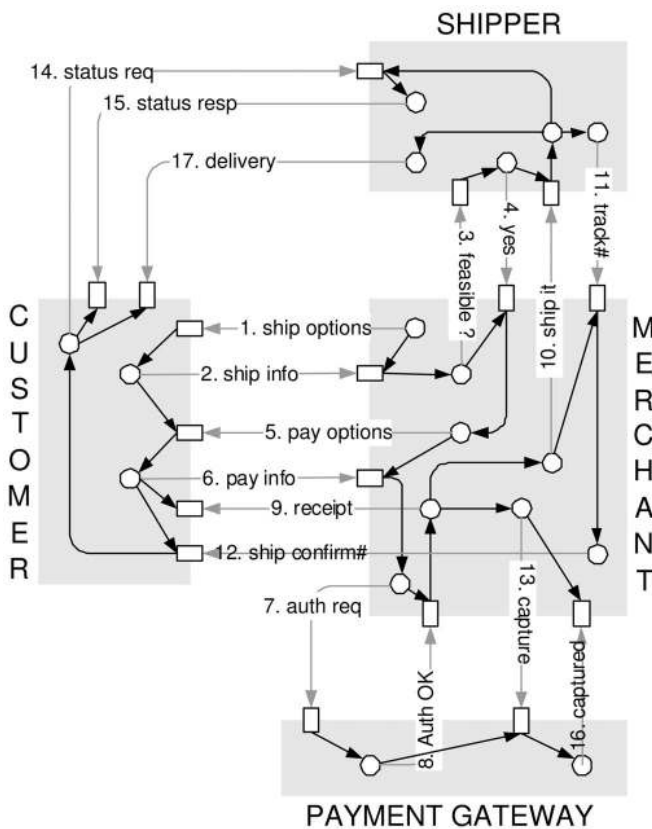


Fig. 1. A purchasing process.

The key benefits of this approach are:

1. a separation of concerns between interaction (protocols) and local decision making (policies),
2. reusability of protocols across processes,
3. evolution and refinement of processes by protocol composition, and
4. flexible process enactment that respects local policies while adapting continually.

### 1.1 Case Study

Let's consider the common business process involving a Customer who wants to procure items, a Merchant who supplies items, a Shipper who is a logistics provider, and a Payment Gateway who authorizes payments. The payment-related interactions are based on the Secure Electronic Transactions (SET) standard [2]. Fig. 1 depicts a variant where the items and price have been agreed upon. Each participant is shown as a separate shaded region, the graph within it denoting its *local process*. Circular nodes represent a participant's internal business logic or policies, e.g., to decide the parameters of an outbound message. Rectangular nodes represent external interfaces through which a participant receives messages. All out-edges from a node are taken concurrently. The messages are labeled with numbers to indicate a possible order in which they might occur. For example, after message 9, the Merchant could send message 10 and message 13 in any order. Just one possible scenario is shown; a realistic process would involve multiple scenarios.

### 1.2 Shortcomings of Traditional Approaches

The above process can be captured via a traditional flow-based approach such as BPEL [3]. Such a representation would be functionally correct, but inadequate from the perspective of software engineering in open environments.

- **Lack of Reusable Components.** Local processes are monolithic in nature, and formed by ad hoc intertwining of internal business logic and external interactions. Since the business logic is proprietary, the local processes of one partner are not usable by another. For instance, if a new customer were to participate in this market, its local process would have to be developed from scratch.
- **Lack of Semantics.** Traditional approaches expose low-level interfaces, but associate no semantics with the participants' actions. This precludes flexible enactment (as needed to handle exceptions) as well as compliance checking. Without semantics, we cannot determine if a deviation from an expected sequence of steps is significant.
- **Inadequate Evolvability.** Suppose the merchant wishes to change the way it interacts with its customers, maybe because of new service features, i.e., say the goods can be shipped before the payment is received. Due to an update in the merchant's process, the customer agent may no longer be able to interact correctly. It is not clear what updates must be made to the process and where.

### 1.3 Organization

Section 2 introduces key concepts and terminology. Section 3 describes our protocol specification language and its semantics. Section 4 discusses protocol composition and demonstrates its applicability in exception handling. Section 5 shows how augmenting policies with protocols can be used to develop processes. Section 6 presents a  $\pi$ -calculus formalization. Section 7 discusses the relevant research. Section 8 outlines directions for future research.

## 2 CONCEPTS AND TERMINOLOGY

Key intuitions underlying protocols and processes and a deeper background were introduced in [4]. A finite automata representation of protocols and their grounding in BPEL [3] were presented in [5]. A state-based model of commitment protocols and a protocol algebra for refinement and aggregation were presented in [6].

Fig. 2 shows our conceptual model for a treatment of business processes based on protocols and policies. Boxed rectangles are abstract entities (interfaces), which must be combined with business policies to yield concrete entities that can be fielded in a running system (rounded rectangles). Abstract entities can be published and reused. They correspond to service specifications in service-oriented computing (SOC). We specify the *protocol logic* via rules that describe interactions among the participating *roles*. Roles are abstract, and are adopted by agents to enable concrete computations.

Whereas a protocol specifies interactions from the global perspective, a skeleton specifies interactions from a participant's perspective. To participate in multiple protocols, an agent would adopt a role in each of them. A *composite*

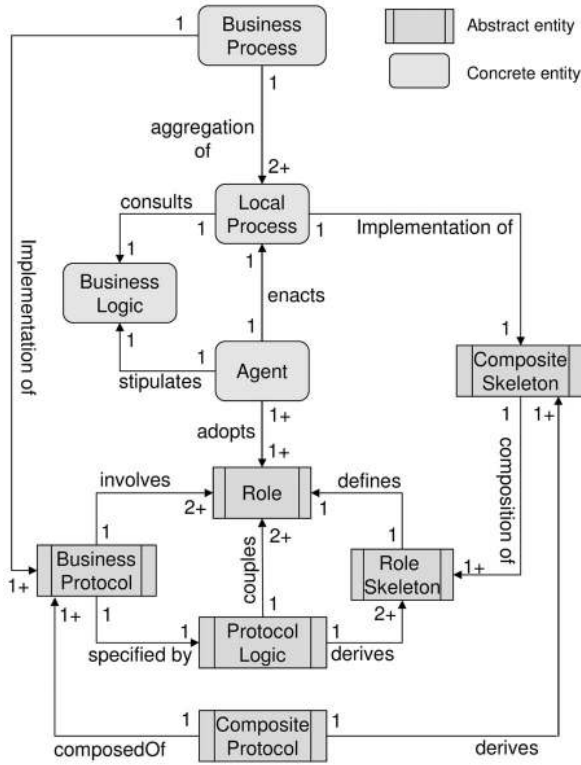


Fig. 2. Conceptual model: business processes based on protocols and policies.

skeleton can be constructed by combining the adopted role skeletons according to some composition constraints. For example, in a supply chain process, a supplier would be a merchant when interacting with a retailer in a trading protocol and an item-sender in a shipping protocol for sending goods to the retailer. A skeleton for such a supplier would combine the skeletons for a trading merchant and a shipping item-sender. Alternatively, as here, a composite protocol for supply chain can be composed from component protocols such as trading and shipping, and the supplier’s skeleton derived from the composite protocol. The resultant composite skeleton could be published and, thus, used by multiple suppliers.

The local process of an agent is an executable realization of its skeleton instantiated with its policies or business logic. A business process aggregates the local processes of the agents participating in it. Equivalently, a business process implements the constituent business protocols.

**2.1 Protocols and Commitments**

Commitments capture a variety of contractual relationships, while enabling manipulations such as delegation and assignment, which are essential for open systems [7]. For example, a customer would commit to the merchant to pay an item’s price upon delivery. Violations of commitments can be detected; in some important settings, violators can be penalized. Verifying compliance is essential in open settings [8].

**Definition 1.** A commitment  $C(x, y, p)$  denotes that agent  $x$  is obliged to agent  $y$  for bringing about condition  $p$ .

Here,  $x$  is called the debtor,  $y$  the creditor, and  $p$  the condition of the commitment. The condition is expressed in a suitable formal language.

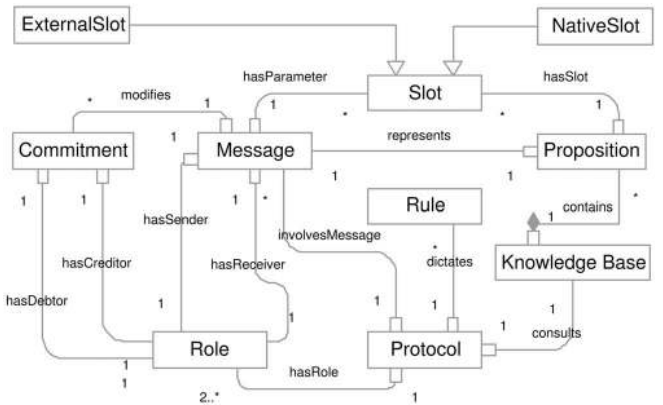


Fig. 3. Basic OWL-P ontology.

Commitments can be conditional, denoted by  $CC(x, y, p, q)$ , meaning that  $x$  is committed to  $y$  to bring about  $q$  if  $p$  holds where  $p$  is called the precondition of the commitment. For example,  $CC(c, b, goods(g), pay(p))$  means that customer  $c$  is committed to pay the bookstore  $b$  an amount  $p$  if  $b$  delivers the book  $g$  to  $c$ . When  $b$  delivers the book, i.e.,  $goods(g)$  holds,  $CC(c, b, goods(g), pay(p))$  yields the base-level commitment  $C(c, b, pay(p))$ .

The following operations describe how a commitment  $c$  is created, satisfied, and manipulated [7]:

1.  $CREATE(x, c)$  establishes  $c$ , if directly performed by  $c$ ’s debtor, or because of a conditional commitment.
2.  $CANCEL(x, c)$  terminates  $c$  as an action of  $c$ ’s debtor  $x$ . Generally, cancellation is constrained and would lead to new commitments.
3.  $RELEASE(y, c)$  terminates  $c$  by releasing  $c$ ’s debtor  $x$ . This only can be performed by the creditor  $y$ .
4.  $ASSIGN(y, z, c)$  replaces  $y$  with  $z$  as  $c$ ’s creditor.
5.  $DELEGATE(x, z, c)$  replaces  $x$  with  $z$  as the  $c$ ’s debtor.
6.  $DISCHARGE(x, c)$   $c$  is terminated by  $x$  fulfilling it.

A commitment is active if it has been created but not yet terminated. Commitments are transformed via the rules below:

1.  $\frac{C(x,y,p)\wedge p}{discharge(x.C(x,y,p))}$ ,
2.  $\frac{CC(x,y,p,q)\wedge p}{create(x.C(x,y,q))\wedge discharge(x.CC(x,y,p,q))}$ , and
3.  $\frac{CC(x,y,p,q)\wedge q}{discharge(x.CC(x,y,p,q))}$ .

**3 PROTOCOL SPECIFICATION**

Let’s now consider how protocols are specified based on the meaning associated with a message.

**3.1 OWL-P**

Fig. 3 shows key OWL-P concepts (nodes) and their properties (edges: the end with a rectangle is the domain), reflecting the conceptual model of Section 2. The core of a protocol is specified using implication rules, as in the Semantic Web Rule Language (SWRL) [9]. A Message has a sender, a receiver, and one or more parameters as Slots. The semantics of a message is given by its effects on commitments, expressed via operations. Slots function as data variables. A slot is defined when it is assigned a value and

used when its value is assigned to another slot. A *native slot* is typed and defined within the protocol. An *external slot* is assigned a value produced by another protocol.

A *protocol* dictates rules that govern the interaction and consults a *knowledge base*, which consists of a set of *propositions*. Propositions are represented as slotted facts in a rule-based system. Propositions express messages sent or received, active commitments, and domain facts.

Fig. 5 shows possible scenarios for protocols. Leading and trailing \*, as in \*amount\*, identify an external slot. Consider the *Order* protocol. The buyer requests a quote for an item; the seller responds with a quote. The semantics of quote is that it creates a commitment from the seller to the buyer guaranteeing delivery if the buyer pays the quoted price. In this scenario, the buyer accepts the quote. The semantics of accept is to create a commitment from the buyer to the seller to pay the quoted price if it receives the requested item. Below are the rules for *Order* in the “antecedent  $\Rightarrow$  consequent” notation:

- 1: startProp  $\Rightarrow$  reqForQuote(?itemID)
- 2: reqForQuoteProp(?itemID)  $\Rightarrow$   
quote(?itemID, ?itemPrice)  $\wedge$   
CC(S, B, pay(?itemPrice), goods(?itemID))
- 3: quoteProp(?itemID, ?itemPrice)  $\Rightarrow$   
acceptQuote(?itemID, ?itemPrice)  $\wedge$   
CC(B, S, goods(?itemID), pay(?itemPrice))
- 4: quoteProp(?itemID, ?itemPrice)  $\Rightarrow$   
rejectQuote(?itemID, ?itemPrice)

A knowledge base contains startProp at the start of a protocol. Here, reqForQuote, quote, and acceptQuote are OWL-P messages, and their corresponding propositions are reqForQuoteProp, quoteProp, and acceptQuoteProp. pay and goods are also propositions. ?itemID and ?itemPrice are native slots. The ?itemID in the first rule is not assigned any value by the antecedent, meaning that the rule is abstract and not executable. As we see in Section 5.2, it can be augmented with business logic that produces such values. The OWL-P ontology and example protocols are available at <http://research.csc.ncsu.edu/mas/OWL-P/>.

### 3.2 Operational Semantics

Protocols are specified from the global perspective with an assumption of an abstract global knowledge base. Later sections show how the abstraction of a global knowledge base maps to the perspectives of the participants’ local knowledge bases. In OWL-P, rules are forward-chained. A message in the consequent of a rule indicates a message exchange. A commitment in the consequent of a rule indicates its creation. By the inference rules in Section 2.1, a commitment is discharged automatically when its condition holds. Other operations on commitments are explicitly specified.

Fig. 4 shows an abstract rule machine that provides operational semantics to the rules. For now, ignore Steps 3, 4, and 5 dealing with policy rules. When a message is received, it is checked against the protocol rules to see if it may be consumed. If so, a corresponding proposition is asserted and any rules thus activated are executed. Doing so may activate other rules, resulting in further propositions being asserted and messages being sent.

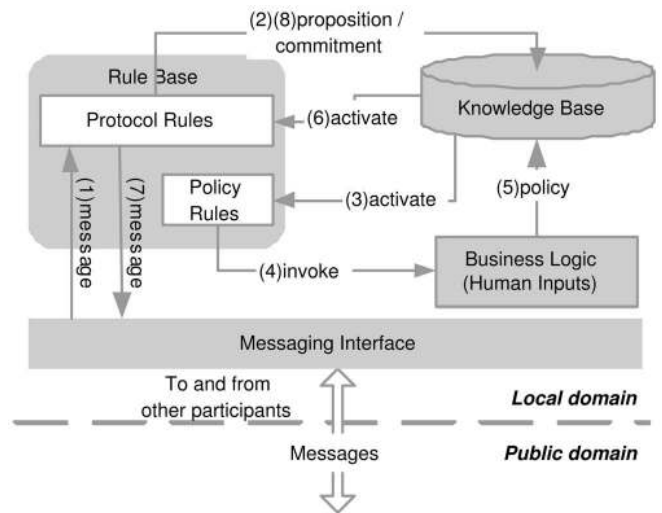


Fig. 4. Agent architecture: protocol and policy interplay.

## 4 COMPOSITE PROTOCOLS

Protocols are rarely used in isolation and must be composed so as to address multiple business goals. Another motivation for composition is that it enables modifying protocols. When a desired modification is captured as a set of rules, it is definitionally identical to a protocol. Thus, a protocol can be modified by composing it with a given modification. Let’s now discuss how we can construct composite protocols and how they facilitate reusability and refinement.

### 4.1 Construction of Composite Protocols

Fig. 6 describes key entities and properties dealing with protocol composition. A *composite protocol* aggregates (component) protocols and is defined by a *composition profile*. A composition profile stipulates several *composition axioms* to constrain the composite protocol. Fig. 5 depicts three protocols, *Order*, *Shipping*, and *Payment*, and a set of composition axioms. We denote the composition axioms as  $axiom_1(property_1 : range_1, property_2 : range_2, \dots)$ , where  $axiom_1$  identifies the axiom class, the  $property_i$  are the properties of  $axiom_1$  and the  $range_i$  are the values of those properties. For brevity, we replace irrelevant slot names by “.”.

A *role definition* axiom defines a (unified) role in the composite protocol in terms of the roles in the component protocols to which it corresponds. In Fig. 5, the first axiom states that customer in *Purchase* unifies buyer in *Order*, payer in *Payment*, and receiver in *Shipping*. Similarly, the second axiom defines a merchant in *Purchase*. Typically, the nonunified roles are played by different agents.

A *data flow* axiom captures a data flow between the protocols. It defines an external slot in one protocol as using an external or a native slot of another. In Fig. 5, the fourth axiom states that slot ?amount in *Payment* gets its value from slot ?itemPrice in *Order*. Such a dependency implicitly ensures that none of the rules using a slot can fire before the slot is assigned a value by the defining rule. Formally:

```
axiom :dataFlow(definition:Order.itemPrice,
               usage:Payment.amount)
def    :reqForQuoteProp(.)  $\Rightarrow$  quote(.,?itemPrice)  $\wedge$  ...
use0  :paymentInfoProp(.,.)  $\Rightarrow$  authReq(.,.,?amount)
```

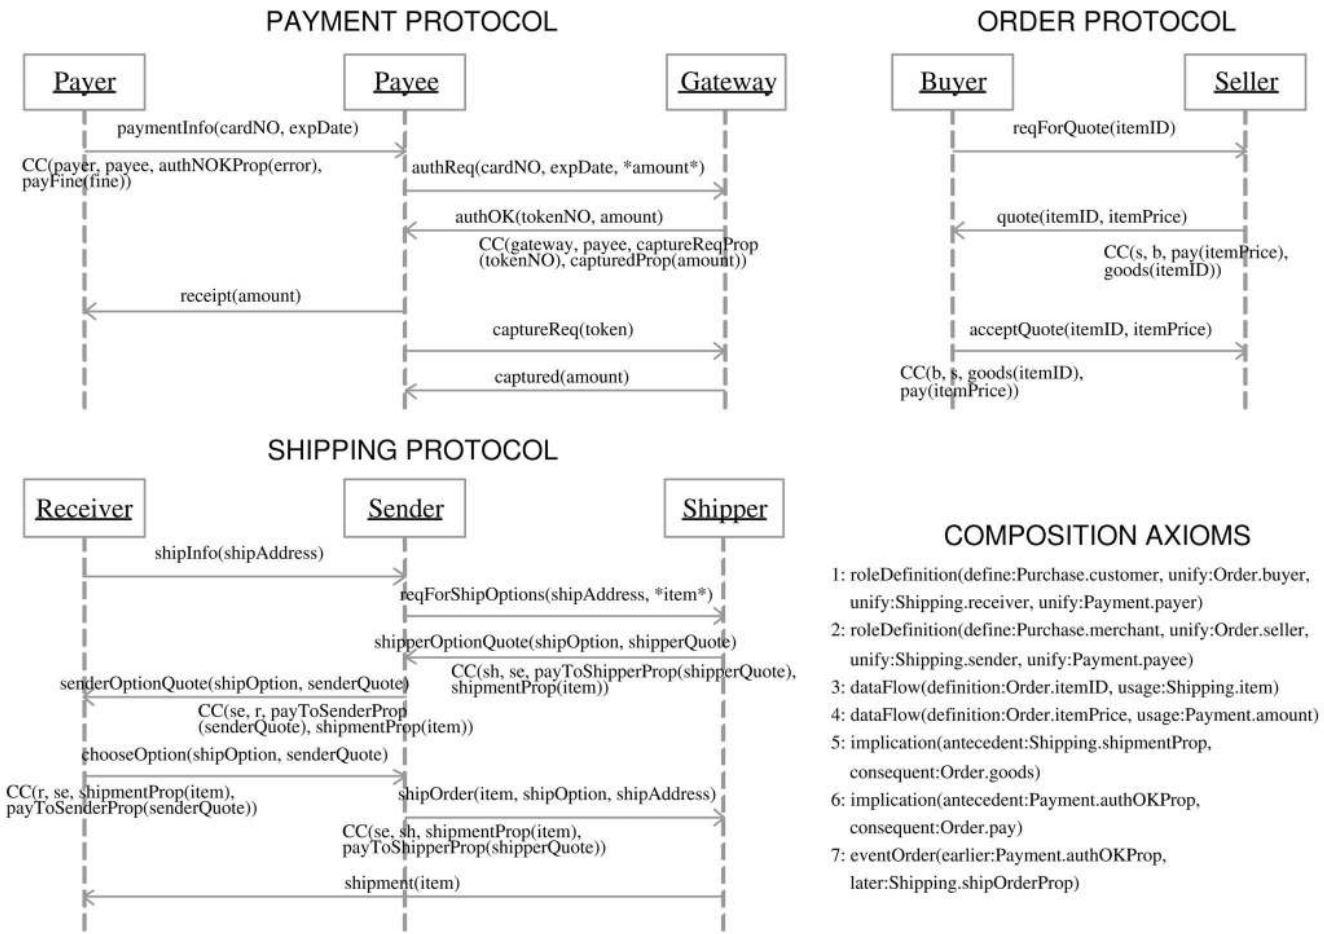


Fig. 5. Example: order, shipping, and payment protocols and their composition.

usei :...  
 used :captureReqProp(.) ⇒ captured(?amount)

**The resulting rules:**

use0 :paymentInfoProp(.,.) ∧ quoteProp(.,?itemPrice) ⇒ authReq(.,.,?amount)

usei :...

used :captureReqProp(.) ∧ quoteProp(.,?itemPrice) ⇒ captured(?itemPrice)

The ordering is achieved by adding the definition of the slot as a premise for the usage of the slot. Since there might be multiple uses of a slot, premises of all of them need to be updated. The usage slot takes the name of the defining slot and the external slot is not needed in the composite protocol.

An *implication* axiom states that proposition *X* in a protocol implies proposition *Y* in another protocol. In Fig. 5, the sixth axiom states that an assertion of authOKProp in *Payment* means an assertion of pay in *Order*. This is achieved by adding an implication rule to the composite rule base. Formally:

axiom:implication(body:Payment.authOKProp, head:Order.pay)

**The resulting rule:**

rule :authOKProp(.,.) ⇒ pay(.)

An *event order* axiom specifies an ordering among selected messages of the component protocols. In Fig. 5, the seventh axiom states that an authOK message from the payment gateway must be received before a shipOrder message is sent to the shipper. This can be ensured by making the rule for the later event depend on the rule for the earlier event, that is:

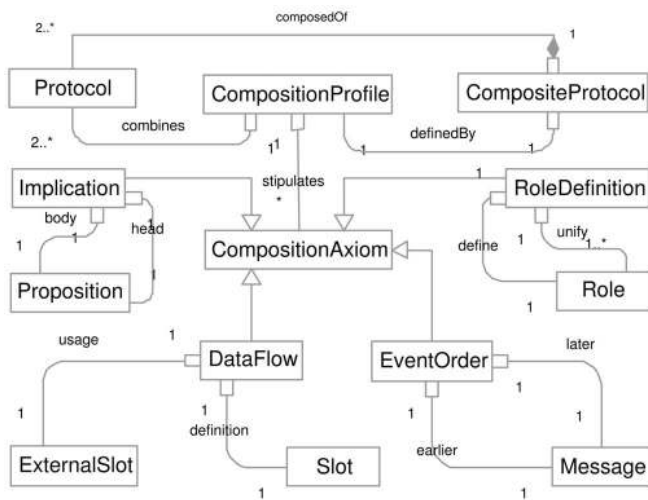


Fig. 6. OWL-P composition classes and properties.

axiom :eventOrder(earlier:Payment.authOK,  
later:Shipping.shipOrder)  
earlier:authReqProp( $\cdot, \cdot, \cdot$ )  $\Rightarrow$  authOK( $\cdot, \cdot$ )  $\wedge$  CC( $\cdot, \cdot, \cdot, \cdot$ )  
later :chooseOptionProp( $\cdot, \cdot$ )  $\Rightarrow$  shipOrder( $\cdot, \cdot, \cdot$ )  $\wedge$   
CC( $\cdot, \cdot, \cdot, \cdot$ )

**The resulting rule:**

later :chooseOptionProp( $\cdot, \cdot$ )  $\wedge$  authOKProp( $\cdot, \cdot$ )  $\Rightarrow$   
shipOrder( $\cdot, \cdot, \cdot$ )  $\wedge$  CC( $\cdot, \cdot, \cdot, \cdot$ )

Composition axioms must be specified by a designer. There may be several ways of composing a protocol. Specifically, if the component protocols are mutually independent, no axioms need be specified; the protocols can simply be unioned, yielding the specification of a composite. If necessary, additional subclasses of *CompositionAxiom* may be introduced by specifying their syntactic properties and operational semantics. A composite protocol is like any other protocol, and can be composed further.

How can we determine whether additional component protocols are needed to obtain an enactable protocol? A protocol may not be enactable due to *nonlocal choice*, which means that an agent's decision depends on information it doesn't have locally [10]. For example, protocols with rules of the form  $a \Rightarrow b$ , where proposition  $a$  is known only to roles  $\rho_1$  and  $\rho_2$ , and message  $b$  is sent by  $\rho_3$ , are not enactable as  $\rho_3$  cannot observe  $a$ . This problem can arise because a protocol is a global specification whereas role skeletons are partial local views. The following properties of a protocol  $P$  avoid this problem and ensure that  $P$  is enactable. Here,  $r$  denotes a rule,  $a$ ,  $b$ ,  $p$ , and  $q$  are propositions or messages as applicable, and  $P$  is interpreted as a set of rules.

**Prop1.**  $\forall r \in P, \forall a \in r.body, a \neq start \rightarrow \exists r' \in Pa \in r'.head.$   
The propositions in the antecedent of a protocol rule must be asserted by some protocol rule.

**Prop2.**  $\forall r \in P, \forall a \in r.head, (start \rightsquigarrow a) \in P.$  All asserted propositions are reachable from *start*. Here,  $p \rightsquigarrow q$  means  $q$  is reachable from  $p$ . Formally,  $p \rightsquigarrow q \stackrel{def}{=} p \Rightarrow q \vee (p \Rightarrow q' \wedge q' \rightsquigarrow q).$

**Prop3.**  $\forall r \in P, \forall a \in r.body, (\exists b \in r.head \wedge \exists \rho = sender(b) \rightarrow \rho = sender(a) \vee \rho = receiver(a)).$  If a message exchange  $b$  has an antecedent  $a$ , then the sender  $\rho$  of  $b$  must be able to verify  $a$ . That is,  $\rho$  must receive a message corresponding to  $a$  or must send a message corresponding to  $a$ , i.e.,  $\rho$  observes  $a$  or causes  $a$ .

A designer's goal is to obtain an enactable protocol by repeated applications of composition. Observe that in Fig. 5, the rules of *Order* do not assert propositions *pay* and *goods* necessary for discharging the commitments created. The rules of *Payment* do not assert the proposition *pay(fineAmount)*. The rules of *Shipping* do not assert *payToSenderProp* and *payToShipperProp*. The rules of the composite *Purchase* do not assert the propositions *payFine(fine)*, *payToSenderProp*, and *payToShipperProp*. A designer would choose protocols that assert the missing propositions and compose them with *Purchase* to obtain an enactable protocol.

## ADJUSTMENT PROTOCOL




Fig. 7. Handling exceptions by composition.

## 4.2 Refinement by Composition

Business protocols evolve continually as new requirements and new features arise. Therefore, the ability to systematically refine protocols is valuable. In the composite *Purchase*, consider a situation in which the customer has already paid the merchant for the goods and, hence, the commitment  $C(m, c, goods(itemID))$  is active. If a fire destroys the merchant's warehouse, the merchant will not be able to honor its commitment to ship the item. How can such exceptions be handled? Of course, we could detect the violation due to an unfulfilled commitment, and the merchant could be held legally responsible. However, a more flexible solution would be to allow the merchant to refund the payment and release the merchant from the commitment, provided the customer agrees to it. We can achieve this flexibility by combining *Purchase* with *Adjustment* as specified below (see Fig. 7):

$$C(d, c, cond) \Rightarrow exception(?reason) \wedge$$

$$cancel(d, C(d, c, cond))$$

$$exceptionProp(?reason) \Rightarrow OK() \wedge C(d, c, newCond)$$

$$exceptionProp(?reason) \Rightarrow NOK() \wedge C(d, c, oldCond)$$

The axioms below yield *New*, a more flexible protocol:

- 1: roleDefinition(define:New.customer,  
unify:Purchase.customer, unify:Adjustment.creditor)
- 2: roleDefinition(define:New.merchant,  
unify:Purchase.merchant, unify:Adjustment.debtor)
- 3: Implication(body:Purchase.C(m,c,goods(itemID)),  
head:Adjustment.C(d, c, cond))
- 4: Implication(body:Adjustment.C(d, c, newCond),  
head:Purchase.C(m, c, refund))
- 5: Implication(body:Adjustment.C(d, c, oldCond),  
head:Purchase.C(m, c, goods(itemID)))

*Adjustment* allows cancellation of the merchant's commitment if the customer deems it reasonable. The semantics of *OK* specifies the creation of a new commitment for compensation. The creditor may not agree, send a *NOK*, and retain the old commitment. Similar protocols for assigning, delegating, and releasing commitments can be defined.

## 5 PROCESSES

To enact its local process, a participant extracts its skeleton from an enactable protocol and instantiates it with its policies.

### 5.1 Role Skeletons

OWL-P describes a protocol assuming a global state. As in all distributed systems, the state of a protocol as seen by a role changes only when a message is sent or received by that role. This observation forms the basis of skeleton-generation as shown in Algorithm 1, where  $P$  is a protocol and  $\rho$  is a role.

Algorithm 1: deriveSkeleton( $P, \rho$ ): Generate the skeleton for role $\rho$
1 $\rho.rules \leftarrow \text{getPertinentRules}(P, \rho);$
2 <b>foreach</b> $r \in \rho.rules \wedge m \in r.head \wedge \rho \neq \text{sender}(m)$ <b>do</b>
3 <b>foreach</b> $a \in r.body$ <b>do</b>
4 $hist \leftarrow \{\};$
5         Replace $a$ with $\text{replace}(P, \rho, a);$
6 <b>Procedure</b> $\text{getPertinentRules}(P, \rho)$ : Get rules for $\rho$
7 $rules \leftarrow \{\};$
8 <b>foreach</b> $r \in P.rules$ <b>do</b>
9 <b>foreach</b> $m \in r.head$ <b>do</b>
10 <b>if</b> $\rho = \text{receiver}(m)$ <b>then</b>
11             Replace $m$ with $\text{RECEIVE}(m);$
12 <b>if</b> $\rho = \text{sender}(m)$ <b>then</b>
13             Replace $m$ with $\text{SEND}(m);$
14 $rules \leftarrow rules \cup \{r\};$
15 <b>return</b> $rules;$
16 <b>Procedure</b> $\text{replace}(P, \rho, a)$ : If $a$ isn't asserted in $\rho.rules$ , replace it by something that is
17 <b>if</b> $a \in hist$ <b>then</b>
18 <b>return</b> "FALSE";
19 <b>if</b> $a = \text{"start"} \vee (\text{rule } r \in \rho.rules \wedge a \in r.head)$ <b>then</b>
20 <b>return</b> $a;$
21 $hist \leftarrow hist \cup a;$
22 $disj \leftarrow \text{"FALSE"};$
23 <b>foreach</b> $r \in P.rules \wedge a \in r.head$ <b>do</b>
24 $conj \leftarrow \text{"TRUE"};$
25 <b>foreach</b> $b \in r.body$ <b>do</b>
26 $conj \leftarrow conj \cdot \text{"\&"} \cdot \text{replace}(P, \rho, b);$
27 $disj \leftarrow disj \cdot \text{"\vee"} \cdot conj;$
28 $hist \leftarrow hist - a;$
29 <b>return</b> $disj;$

The algorithm gathers all the rules from the specification for  $P$  that have  $\rho$  receiving or sending a message. Next, the algorithm invokes  $\text{replace}(P, \rho, a)$ , defined from line 16 to line 29. If a proposition  $a$  enables a rule  $r \in \rho.rules$  but  $a$  is not asserted by any rule in  $\rho.rules$ , it means that  $a$  was not observed by  $\rho$ . This procedure replaces  $a$  with the last proposition that  $\rho$  did observe, i.e., the proposition  $a'$  that was asserted in  $\rho.rules$  and leads to  $a$  being asserted. Such  $a'$  will be found due to the properties of an enactable protocol.

As an example, we show a rule in *Shipping* in Fig. 5, and the same rule in the receiver's skeleton. As the receiver would not be aware of the previous exchanges between the sender and the shipper, the antecedent of the rule for receiving  $\text{senderOptionQuote}$  should be adjusted as shown below:

Protocol rule:

$\text{shipperOptionQuoteProp}(\cdot, \cdot) \Rightarrow \text{senderOptionQuote}(\cdot, \cdot) \wedge \text{CC}(\text{Se, Re, payToSenderProp}(\cdot), \text{shipmentProp}(\cdot))$

Receiver skeleton rule:

$\text{shipInfoProp}(\cdot) \Rightarrow \text{receive}(\text{senderOptionQuote}(\cdot, \cdot)) \wedge \text{CC}(\text{Se, Re, payToSenderProp}(\cdot), \text{shipmentProp}(\cdot))$

Protocol computations are given in terms of the message exchanges  $m$ , whereas skeletons computations are given in terms of  $\text{send}(m)$  and  $\text{receive}(m)$ . We now prove Algorithm 1 correct.

**Theorem 1 (Soundness).** *If condition  $\text{cond}$  enables  $\text{send}(m)$  and  $\text{receive}(m)$  in the skeletons generated from protocol  $P$ , then  $\text{cond}$  enables  $m$  in  $P$ .*

**Proof.** Let  $\text{cond}$  hold in the skeletons, enabling  $\text{send}(m)$  and  $\text{receive}(m)$ . As  $\text{cond}$  enables  $\text{send}(m)$ , there must exist a role  $\rho = \text{sender}(m)$  such that  $(\text{cond} \Rightarrow \text{send}(m)) \in \rho.rules$ . For this rule to be in  $\rho.rules$ , it is necessary that a rule  $(\text{cond} \Rightarrow m) \in P.rules$ , as line 14 in Algorithm 1 is the only way to add such a rule to  $\rho.rules$ . However, if  $(\text{cond} \Rightarrow m) \in P.rules$  then the message exchange  $m$  is enabled in  $P$ .  $\square$

**Theorem 2 (Completeness).** *If condition  $\text{cond}$  enables a message exchange  $m$  in protocol  $P$ , then  $\text{cond}$  also enables  $\text{send}(m)$  and  $\text{receive}(m)$  in the skeletons generated from  $P$ .*

**Proof.** Let  $\text{cond}$  hold in  $P$ , enabling  $m$ . This implies a rule  $(\text{cond} \Rightarrow m) \in P.rules$ . By line 14 in Algorithm 1, this means that there exists a role  $\rho_s = \text{sender}(m)$  such that  $(\text{cond} \Rightarrow \text{send}(m)) \in \rho_s.rules$ . If such a rule exists in  $\rho.rules$ , and if  $\text{cond}$  holds, then  $\text{send}(m)$  must be enabled.

Also, by line 14, there exists a role  $\rho_r = \text{receiver}(m)$  such that  $\text{cond}' \Rightarrow \text{receive}(m) \in \rho_r.rules$ , where  $\text{cond}'$  is the result of line 5. However, by Lemma 1 (Appendix D which is available at <http://computer.org/tse/archives.htm>),  $\text{cond} \rightarrow \text{cond}'$ . Hence,  $\text{receive}(m)$  must be enabled.  $\square$

### 5.2 Policies

As mentioned earlier, a rule in a protocol may be abstract, meaning that values of some of its native slots must be produced via the role's business logic. That is, a role skeleton must be augmented with business logic to obtain a local process. To capture this intuition, we define a skeleton as *concrete* if all of its native slots are defined and as *abstract* otherwise.

**Definition 2.** *A local process is a concrete skeleton derived from an enactable protocol.*

An agent's business logic is specified as local policy rules. The skeleton of the seller in *Order* augmented with the policy rules of the seller agent is shown below. The policy rule invokes a business logic operation to decide what price to quote. The operation asserts  $\text{quotePolicy}$ , which activates the second skeleton rule.

Seller skeleton rules:

1:  $\text{startProp} \Rightarrow \text{receive}(\text{reqForQuote}(\text{?itemID}))$

2:  $\text{reqForQuoteProp}(\text{?itemID}) \wedge \text{quotePolicy}(\text{?itemPrice}) \Rightarrow \text{send}(\text{quote}(\text{?itemID}, \text{?itemPrice})) \wedge \text{CC}(\text{S}, \text{B}, \text{pay}(\text{?itemPrice}), \text{goods}(\text{?itemID}))$

3:  $\text{quoteProp}(\text{?itemID}, \text{?itemPrice}) \Rightarrow$

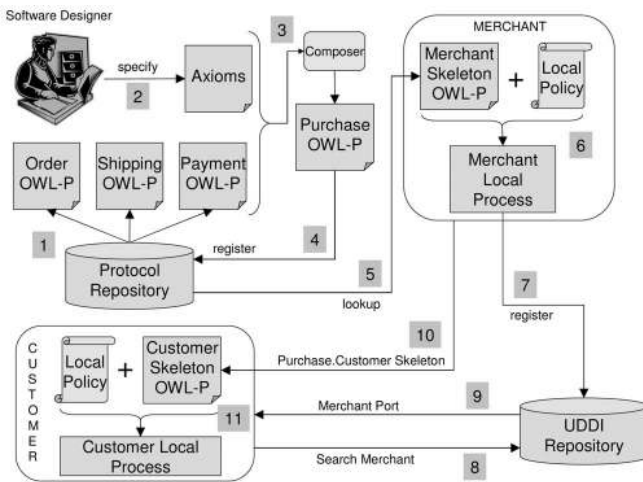


Fig. 8. A usage scenario.

```

receive(acceptQuote(?itemID, ?itemPrice)) ^
4:CC(B, S, goods(?itemID), pay(?itemPrice))
quoteProp(?itemID, ?itemPrice) =>
receive(rejectQuote(?itemID, ?itemPrice))

```

```

Seller's policy for deciding quote:
1:reqForQuoteProp(?itemID) =>
call(policyDecider, quotePolicy(?itemID))

```

This pattern of augmenting policy rules is general and applies to rules where the agent has to make a decision and respond. It also assigns a value to native slots that are not defined. Fig. 4 shows the interplay between the protocol rules and the policy rules of an agent. Steps 3, 4, and 5 show policy rules in action. The business logic could involve looking up a legacy database or waiting for human input.

### 5.3 Usage

Here, we show how our model of business processes can be grounded in Web services. Fig. 8 summarizes our methodology with a scenario involving a customer interested in purchasing goods online. Software designers design protocols and register them with protocol repositories. They may also construct composite protocols by specifying composition axioms and reusing the existing component protocols from the repository. *Composer* is our tool for generating the composite protocol. A merchant wishing to supply goods online looks up *Purchase* in a repository. It generates the skeleton for the merchant role, augments it with its local policies, and deploys the result as a service. The profile for this service contains an OWL-P description of *Purchase*, and may be published to a UDDI registry. If a customer wishes to procure goods online, it searches the UDDI registry, finds the merchant, and acquires the OWL-P skeleton for the customer role from the merchant. The customer enacts its local process by augmenting the skeleton with its local policies and starts interacting with the merchant. Our tools support these development scenarios and a prototype implementation based on the agent architecture of Fig. 4.

## 6 $\pi$ -CALCULUS FORMALIZATION

The foregoing showed how our interaction-based treatment facilitates reusability and evolvability. It would help to reason about the properties of the protocols. To do so, we first introduce the  $\pi$ -calculus, show how protocols and their composition can be specified in it, and then describe some properties of protocols established using the  $\pi$ -calculus formalization.

### 6.1 The $\pi$ -calculus

The  $\pi$ -calculus [11] is a process algebra for modeling concurrent processes whose configurations may change as the processes execute. In the  $\pi$ -calculus, the fundamental unit of computation is the transfer of a communication link between two processes. The simplicity of the  $\pi$ -calculus arises from the fact that it includes only two kinds of entities: names and agents (processes). These are sufficient to rigorously define interactional behavior. This paper uses the synchronous  $\pi$ -calculus in which interaction corresponds to a handshake between two processes and involves the output of a link by one process and simultaneous receipt of the link by another process. The language of the  $\pi$ -calculus consists of prefixes and process expressions, as summarized below.

Prefixes are of the following kinds:

Prefixes  $\alpha ::= \bar{a}(x)$  (Output)  
 $a(x)$  (Input)  
 $\tau$  (Silent).

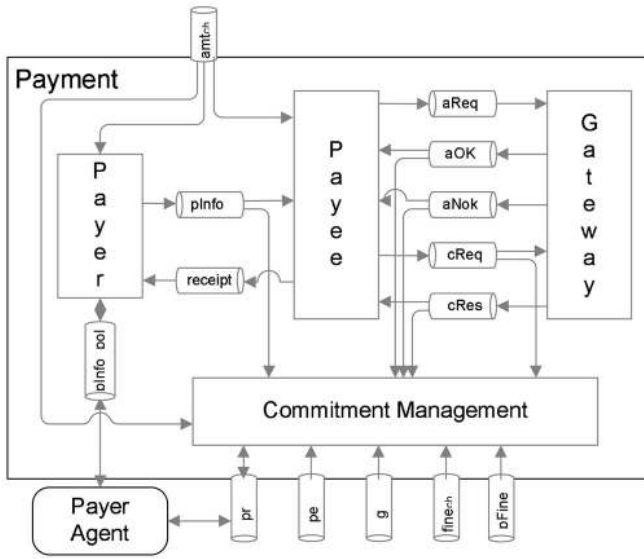
- The output prefix  $\bar{a}(x)$  means that  $x$  is sent along the channel  $a$ .
- The input prefix  $a(x)$  means that the channel  $a$  can be used to receive input and binds this input to  $x$ .
- The silent  $\tau$  means that nothing observable happens.

The process expressions are as follows:

Agents	$P ::=$	$0$	(Nil)
		$\alpha.P$	(Prefix)
		$P + P$	(Sum)
		$P P$	(Parallel)
		$[x = y]P$	(Match)
		$[x \neq y]P$	(Mismatch)
		$(new\ x)P$	(Restriction)
		$!P$	(Replication)
		$A\langle y_1, \dots, y_n \rangle$	(Identifier)
Definitions	$A(x_1, \dots, x_n)$	$\stackrel{def}{=} P,$	(where $i \neq j$ $\Rightarrow x_i \neq x_j$ )

- $0$  represents the nil-process.
- $\alpha.P$  does the action represented by prefix  $\alpha$  and changes to  $P$ .
- $P + Q$  represents the sum-nondeterminism, that is, do either process  $P$  or process  $Q$ .
- $P|Q$  represents that process  $P$  and process  $Q$  execute in parallel.
- $[x = y]P$  represents the process that changes to  $P$  if  $x = y$ . Mismatch is the opposite, i.e., it checks  $x \neq y$ .




 Fig. 9. A  $\pi$ -calculus model of *Payment*.

- $(new\ x)P$  means that the variable  $x$  is declared as a new name local to process  $P$  and bound in  $P$ . It is not visible outside of  $P$ .
- $!P$  represents an unbounded number of copies of the process  $P$ . Formally,  $!P \stackrel{def}{=} P!P$ .
- $A\langle y_1, \dots, y_n \rangle$  represents the instantiation of a defined agent.
- $A(x_1, \dots, x_n) \stackrel{def}{=} P$  represents the declaration of a process  $A$  in terms of process  $P$ . One can think of it as a method declaration in traditional procedural programming.

The input prefix and the *new* operator bind names. For example, in the process  $a(x).P$ , the name  $x$  is bound, but  $a$  is not. It helps to use the *polyadic*  $\pi$ -calculus because it enables sending and receiving tuples of variables [11]. For example,  $\bar{x}(a, b, c)$  outputs three names  $a$ ,  $b$ , and  $c$  over the channel  $x$ . And,  $x(p, q, r)$  receives three names over  $x$  and binds them to  $p$ ,  $q$ , and  $r$  respectively. We introduce a notation for broadcasts. A broadcast on a channel  $x$  of a name  $y$  is indicated by  $\widehat{x}(y)$  and is equivalent to  $!x(y)$ . As an example, consider the following:

$$P \equiv (new\ z)(z(w).\widehat{w}(y)|x(u).\bar{u}(v)|\widehat{x}(z)).$$

The pair  $\widehat{x}(z)$  and  $x(u)$  can interact.  $P_1$  indicates the process left after interaction.

$$P_1 \equiv (new\ z)(z(w).\widehat{w}(y)|\bar{z}(v)|0).$$

In  $P_1$ , another interaction is possible over the channel  $z$  giving process  $P_2$ .

$$P_2 \equiv (new\ z)(\widehat{v}(y)|0|0).$$

Thus, process  $P$  evolves to  $P_2$  over two interactions:  $P \rightarrow P_1$  and  $P_1 \rightarrow P_2$ . We use  $\rightarrow^*$  to abbreviate a series of such interactions, thus writing  $P \rightarrow^* P_2$ .

## 6.2 Protocols

Using *Order* of Fig. 5 as an example, we show how to formalize a protocol in the  $\pi$ -calculus. Fig. 9 shows a model of *Payment* in the  $\pi$ -calculus. Pipes denote the channels,

boxes the  $\pi$ -processes, and rounded rectangles the agents. The Appendix (which is available at <http://computer.org/tse/archives.htm>) presents complete formalizations of *Payment*, *Shipping*, and *Purchase*.

An exchange of messages corresponds to *reactions* in the  $\pi$ -calculus. Thus, a protocol can be specified as a  $\pi$ -process in which each message is represented by a channel. Unique names for channels corresponding to each message type are assumed. Each role skeleton is a  $\pi$ -process having the relevant messages as the  $\pi$ -process parameters. The structure of a protocol  $\pi$ -process would be a parallel composition of the skeletons of each of the roles. Thus, *Order* would be formalized as:

$$\begin{aligned} Order(b, s, r, fq, quote, accept, reject, goods, pay) &\stackrel{def}{=} \\ Buyer(r, fq, quote, accept, reject) &| \\ Seller(r, fq, quote, accept, reject) &| \\ Comm(b, s, r, fq, quote, accept, reject, goods, pay) & \end{aligned}$$

Here, *Comm* is a process that observes the message exchanges and manages commitments according to the message semantics. Each role is also represented as a name, e.g.,  $b$  and  $s$ , and may represent a debtor or a creditor of a commitment. The protocol messages are channels  $r, fq, quote, accept$ , and  $reject$ . The channels  $goods$  and  $pay$  represent external messages, as discussed below. Both skeletons are passed all the message channels as both of the roles are involved in all the interactions. Here is the  $\pi$ -process for the seller skeleton:

$$\begin{aligned} Seller(r, fq, quote, accept, reject) &\stackrel{def}{=} (new\ quote\_pol) \\ r, fq(g).Quote_p.\widehat{quote}(g, p). & \\ (accept(g', p') + reject(g', p')) & \end{aligned}$$

$$Quote_p \equiv \overline{quote\_pol}(g).quote\_pol(g, p)$$

The seller receives  $g$  (*itemID*) on  $r, fq$ , consults its policy for quoting and sends the quote price  $p$  (*itemPrice*) on  $quote$ , and receives either *accept* (*acceptQuote*) or *reject* (*rejectQuote*).  $Quote_p$  is a stub for interacting with the quoting policy of the seller agent, using the channel  $quote\_pol$ . Such policy stubs output necessary names to the corresponding policy channel, and then receive some new names output by the business logic on the same channel. As the policies are private, the channels for interacting with them are restricted to the skeleton  $\pi$ -process, e.g.,  $quote\_pol$ . As  $g$  is bound by  $r, fq$ , a new name  $g'$  should be bound to the input on *accept* or *reject*. Similarly as  $p$  is bound by  $quote\_pol$ ,  $p'$  binds to the input on *accept* or *reject*. Usually,  $g$  and  $g'$  would receive the same input but it cannot be guaranteed. Ideally, the seller should check if they match for the purpose of correlation. The following is the  $\pi$ -process for the buyer skeleton:

$$\begin{aligned} Buyer(r, fq, quote, accept, reject) &\stackrel{def}{=} (new\ r, fq\_pol, accept\_pol) \\ RFQ_p.r, fq(g).quote(g, p'). & \\ (Accept_p.\widehat{accept}(g', p') + Reject_p.\widehat{reject}(g', p')) & \end{aligned}$$

$$RFQ_p \equiv \overline{rfq\_pol}.rfq\_pol(g)$$

$$Accept_p \equiv \overline{accept\_pol}(g', p').accept\_pol$$

$$Reject_p \equiv \overline{reject\_pol}(g', p').reject\_pol$$

It is easy to see that  $(Buyer \mid Seller) \rightarrow^* 0$ , provided the business logic is correct. This means that the skeletons of the buyer and the seller role are compatible for interaction. The commitment semantics of the interaction is given by the observer  $\pi$ -process  $Comm$  as follows:

$$Comm(b, s, rfq, quote, accept, reject, goods, pay) \stackrel{def}{=} \\ (new\ ch_1, ch_2, ch_3, ch_4) \\ quote(g, p).CC\langle s, b, pay, goods, ch_1, ch_2 \rangle.(\overline{ch_1}(p) \mid \overline{ch_2}(g)) \mid \\ accept(g', p').CC\langle b, s, goods, pay, ch_3, ch_4 \rangle.(\overline{ch_3}(g') \mid \overline{ch_4}(p'))$$

On a message exchange involving a commitment operation,  $Comm$  would invoke the corresponding commitment operation  $\pi$ -process. Here, the commitments are created when *quote* or *accept* is observed according to their semantics. After calling the  $\pi$ -process  $CC$  for a conditional commitment, the necessary message parameters are passed to it. The  $\pi$ -process  $CC$  is:

$$CC\langle deb, cred, cond_1, cond_2, ch_1, ch_2 \rangle \stackrel{def}{=} (new\ ch_3, ch_4, ch_5) \\ (ch_1(val_1) \mid ch_2(val_2)).(cond_2(val'_2)).([val_2 \neq val'_2] \\ CC\langle deb, cred, cond_1, cond_2, ch_3, ch_4 \rangle.(\overline{ch_3}(val_1) \mid \overline{ch_4}(val_2)) + \\ [val_2 = val'_2]\tau) + cond_1(val'_1).([val_1 = val'_1] \\ C\langle deb, cred, cond_2, ch_5 \rangle.\overline{ch_5}(val_2) + [val_1 \neq val'_1] \\ CC\langle deb, cred, cond_1, cond_2, ch_3, ch_4 \rangle.(\overline{ch_3}(val_1) \mid \overline{ch_4}(val_2))))$$

First, the parameters passed by  $Comm$  are received, and then the channels related to the commitment conditions are observed. On an exchange, if the commitment condition is satisfied then the process ends. If the precondition is satisfied, the conditional commitment reduces to a commitment, the  $\pi$ -process  $C$  is called, and the necessary parameters passed to it. If neither is satisfied,  $CC$  is called again and the commitment remains active. Private channels  $ch_i$  are used for passing parameters to avoid interference with the protocol channels. The commitment conditions may involve varying number of parameters and  $CC$  needs to be defined accordingly. The definition is for when both the precondition and the condition have one parameter each. Similarly, the  $\pi$ -process  $C$  is:

$$C\langle deb, cred, cond_2, ch_2 \rangle \stackrel{def}{=} (new\ ch_3)ch_2(val_2). \\ (cond_2(val'_2)).([val_2 = val'_2]\tau + [val_2 \neq val'_2] \\ C\langle deb, cred, cond_2, ch_3 \rangle.\overline{ch_3}(val_2)))$$

The separation of the commitment semantics from the interaction specification simplifies the derivation of skeletons, e.g., the buyer skeleton is the parallel composition of the  $\pi$ -processes  $Buyer$ ,  $Comm$ ,  $CC$ , and  $C$ . Also, the loose coupling among the commitment semantics, the interaction specification and the policy definitions (business logic) allows maintainability, reusability, and simplicity of design.

### 6.3 Composite Protocols

A composite protocol is a parallel composition of the component protocols and the composition axioms, e.g, *Purchase* can be composed as follows (details in the Appendix C which is available at <http://computer.org/tse/archives.htm>):

$$Purchase(c, m, g, sh, \\ \dots\ order\ channels\ \dots, \\ \dots\ payment\ channels\ \dots, \\ \dots\ shipping\ channels\ \dots) \stackrel{def}{=} \\ Order\langle c, m, \dots \rangle \mid Payment\langle c, m, g, \dots \rangle \mid \\ Shipping\langle c, m, sh, \dots \rangle \mid Axiom4\langle accept, amt_{ch} \rangle \mid \\ Axiom6\langle aOK, pay \rangle \mid Axiom7\langle aOK, shipment \rangle \mid \\ \dots\ other\ axioms\ \dots$$

$$Axiom4(msg, channel) \stackrel{def}{=} msg(g, p).\widehat{channel}(p) \\ Axiom6(msg_1, msg_2) \stackrel{def}{=} msg_1(tNO, amt).\overline{msg_2}(amt) \\ Axiom7(msg_1, msg_2) \stackrel{def}{=} msg_1(\cdot, \cdot).msg_2(\cdot, \cdot)$$

The channels  $c$ ,  $m$ ,  $g$ , and  $sh$  are the roles of *Purchase*. Notice how passing these channels to *Order*, *Payment*, and *Shipping* accomplishes role unification; role definition axioms are not needed. The axiom numbers relate to the axioms in Fig. 5.

A *Data Flow* axiom is a  $\pi$ -process with two parameters: the channel that defines a name, and the channel that uses the defined name. The axiom receives on the defining channel and forwards it to the using channel. For example, *Axiom4* receives the price on the *accept* channel and outputs the price to the  $amt_{ch}$ , i.e., the channel on which the users of  $amt$  in *Payment* listen for the definition. (Broadcast is needed as there might be multiple listeners.)

An *Implication* axiom is a  $\pi$ -process with two parameters. The first parameter is the antecedent of the implication and the second parameter is the consequent. The axiom listens on the antecedent channel, and on the message exchange, signals by outputting to the consequent channel. For example, *Axiom6* listens on *aOK* to see if the payment is authorized and, signals so to *pay* when it is authorized.

An *Event Order* axiom is a  $\pi$ -process with two parameters: the channels that should exchange messages first and last, respectively. For example, *Axiom7* expresses the condition that *pay* should do its exchange before *shipment*. We take a passive view of event order: monitoring the events to determine if the ordering is satisfied, but not attempting to enforce the condition through planning.

Fig. 10 shows a model of *Purchase*. The dotted pipes are the protocol channels and the axioms listen on them. Agent business logics and the policy channels are not shown. It is encouraging to see how the protocols can be coupled with each other through their external channels via composition axioms. Clearly, *Purchase* itself can be a component protocol in another composition.

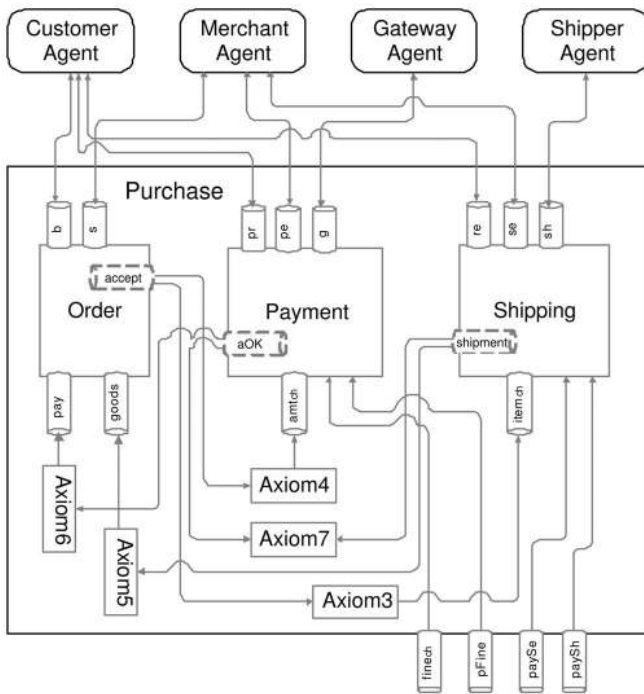


Fig. 10. A  $\pi$ -calculus model of *Purchase*.

## 6.4 Properties of Protocols

We describe some inferences about protocols that are supported by our formalization. Below,  $\mathcal{P}$  is any  $\pi$ -process expression,  $Prot_i$  is a protocol  $\pi$ -process definition,  $Role_{ij}$  is the skeleton definition of role  $j$  in  $Prot_i$ , and  $ch_{ik}$  is an external channel  $k$  for  $Prot_i$ . Subscripts  $i, j, k, m, n$ , etc., range over the cardinality of the respective entity, e.g.,  $Role_{mn}$  could represent a skeleton  $\pi$ -process for role  $n$  in a protocol  $Prot_m$ . Also,  $\mathcal{S}$  and  $\approx$  are the weak simulation and bisimulation relations, respectively. We assume correct functioning of the business logic by replacing each  $Policy_p$  by  $\tau$ .

### 6.4.1 Incorrectness

A protocol specification  $Prot_1$  causes a commitment violation if  $Prot_1 \mid \overline{ch_{11}} \mid \dots \mid \overline{ch_{1k}} \rightarrow^* C \mid \mathcal{P}$ . The debtor of  $C$  is responsible for the violation.

A protocol specification  $Prot_1$  causes an interaction violation if  $Prot_1 \mid \overline{ch_{11}} \mid \dots \mid \overline{ch_{1k}} \rightarrow^* Role_{ij} \mid \mathcal{P}$ .

A protocol specification  $Prot_1$  causes an event ordering violation if  $Prot_1 \mid \overline{ch_{11}} \mid \dots \mid \overline{ch_{1k}} \rightarrow^* Axiom_i \mid \mathcal{P}$ , and  $Axiom_i$  is an event order axiom.

A protocol specification  $Prot_1$  is incorrect if it causes either of the above violations. Note that correctness of a specification can only be verified relative to requirements. The definition of incorrectness given here is general and applies regardless of the requirements.

### 6.4.2 Compatibility

Role skeletons  $Role_{ij}$  and  $Role_{mn}$  are compatible if  $Role_{ij} \mid Role_{mn} \mid \overline{ch_{11}} \mid \dots \mid \overline{ch_{pq}} \rightarrow^* 0$ . Although this definition is intuitive, in case of iterative interactions, the process would never reduce to a  $nil$  process as suggested by Canal et al. [12]. In that case, Skeletons  $Role_{ij}$  and  $Role_{mn}$  are compatible if  $Role_{ij} \mid Role_{mn} \mid \overline{ch_{11}} \mid \dots \mid \overline{ch_{pq}}$  continues making  $\tau$  transitions without interacting with its environment.

### 6.4.3 Equivalence

Protocols  $Prot_1$  and  $Prot_2$  are equivalent if  $Prot_1 \approx Prot_2$ . Due to the presence of policy stubs in the protocol specification, the equivalence is defined by the weak bisimulation with policy interactions modeled as  $\tau$  transitions.

### 6.4.4 Flexibility

A protocol  $Prot_1$  is more flexible than a protocol  $Prot_2$  if  $(Prot_2, Prot_1) \in \mathcal{S}$  and  $(Prot_1, Prot_2) \notin \mathcal{S}$ . Note that this also implies  $Prot_2$  can be substituted by  $Prot_1$ . Canal et al. [12] give a more rigorous treatment of component substitutability.

Other properties of interest might be commitment equivalence of the protocols to see if two protocols are equivalent as far as the commitment semantics of the interaction goes. Also, a relation between a protocol and its refined version can be defined to see if one of them is a specialization of the other. For example, *payment by credit card* is a specialized payment protocol. We believe the formalization is rich enough to allow such inferences.

## 7 RELATED WORK

The uniqueness of our approach arises from 1) our commitment-based semantics for interactions, 2) a focus on a methodology that enables reuse, refinement, and aggregation, and 3) practical grounding of our concepts along with their theoretical foundations. Several relevant research efforts are summarized next.

**Component behavior modeling.** Plasil et al. [13] propose a language based on regular expressions for describing the behavior of software components using protocols. They define software components as agents and the proposed language allows expressions of method invocations over components. This behavior protocols are similar to our protocols but do not support commitments. Similarly, Canal et al. [12] use the  $\pi$ -calculus to describe protocols over the methods provided of CORBA objects. They reason about interaction compatibility and object substitutability, i.e., whether an object can be successfully substituted by another. Studying architectural connectors, Allen et al. [14] express similar intuitions about interactions being a fundamental abstraction for complex systems. Their components correspond to our agents, connectors to protocols, and attachments to role adoption. They formalize components and connectors in CSP [15], and have abstractions of roles and connector composition analogous to ours.

These approaches seek to verify interaction specifications for correctness, compatibility, deadlock freedom, and substitutability for components, whereas our thrust is to exploit protocols to develop processes, thus improving the quality of modeling and reuse of interactions. Previous approaches view interactions as sequences of send and receive actions without considering commitment-based semantics.

**Scenario-based behavior modeling.** This area develops visual formalisms for modeling interaction scenarios of entities (agents, components, or processes). Popular formalisms include AUML [16], Message Sequence Charts [17], and UML sequence diagrams [18].

Whittle et al. [19] present an approach for creating a design model (specified as a statechart) from a collection of interaction scenarios using a state-identification approach to merge scenarios. Uchitel et al. [20] present an MSC language to explicitly specify the assumptions under which scenarios are merged. Whereas the above approaches merge behaviors to produce one protocol, our approach merges several protocols to produce a composite protocol. Further, our methodology incorporates commitments to better capture business requirements.

Uchitel et al. [21] elaborate behavior models by deriving implied scenarios from an existing model, categorizing them as desired or not, and updating the model. Implied scenarios arise as a result of the partial nature of scenario-based specifications. Protocols are naturally extensible, especially nonenactable protocols, which necessarily need to be fleshed out to become enactable.

**Service composition.** BPEL [3] is a language for specifying the static composition of Web services. However, it mixes interaction activities with business logic, thus reducing reusability. WS-CDL [22] specifies the conversational behavior of a service using control flow constructs. However, these specifications lack a flexible semantics, which makes them difficult to compose and reuse. OWL-S [23], which includes a process model for Web services, uses semantic annotations to enable dynamic composition. A composed service is produced at runtime based on stated constraints. Dynamic service composition assumes a perfect markup of the services being composed, and involves ontological matching between inputs and outputs. OWL-S formalizes process models of the ilk of BPEL, but doesn't address protocols as such. An end-to-end service creation approach separates a logical composition (service type) from a physical one (service instances) [24]. Agarwal et al.'s approach combines a service matchmaking similar to OWL-S and a BPEL based service implementation. Due to its close ties with OWL-S and BPEL, it inherits their limitations mentioned above.

Some approaches treat service composition via formal specifications to support verifiability. Solanki et al. [25] employ interval temporal logic to specify and verify ongoing behavior of a composed service. Their use of "assumption" and "commitment" (different meaning than here) assertions yields better compositionality than other approaches. Gerede et al. [26] model services as activity-based finite automata to study the decidability of composition and existence of a lookahead delegator for a set of services. However, these approaches consider neither the autonomy of the partners, nor the flexibility of composition.

**Agent-oriented software engineering.** Agent-oriented software methodologies, e.g., Gaia, KAOS, MaSE, and SADDE [27], apply software engineering principles to agent-based systems. Tropos [28] includes an early requirements stage in the process. Gaia [29] describes roles in the software system being developed, identifies their processes, and the safety and liveness conditions for the processes. It incorporates protocols in its interaction model and can be used with commitments. Baina et al. [30] advocate a model-driven Web service development approach to ensure compliance between a service's implementation and its external protocol specifications. By contrast, our objective is

achieving protocol reusability by separating protocols and policies, and by composing protocols.

**Component catalogs.** Our methodology enables using protocols as building blocks of business processes. Protocols are reused by refining them to yield improved protocols. RosettaNet [31] has put protocols in practice. About 100 published protocols, termed Partner Interface Processes (PIPs), account for billions of dollars of business. However, RosettaNet is limited to two-party request-response PIPs, which lack a formal semantics. The MIT Process Handbook [32] classifies business processes. For example, *sell* is a generic business process, which can be qualified by *sell what*, *sell to whom*, etc.

## 8 CONCLUSIONS

The contributions of this paper include an approach to protocol composition, its grounding in OWL-P, and a formalization of the approach. The charm of this approach is that it recognizes the fundamental interactive nature of open environments where the autonomy of the participants must be preserved. Interactions are treated as first class. Commitments provide a semantics that combines rigor with flexibility. The presented approach will enable the development of an ever-increasing set of protocols for critical business functions.

We demonstrated the practicality of our approach by grounding it in a framework for specifying protocols. Some useful prototype tools supporting this methodology are available. We also demonstrated how our approach is theoretically founded in the  $\pi$ -calculus.

An important direction is to validate this approach in practice on a real-life problem with respect to metrics of reusability, scalability, and cost-effectiveness [33]. Another direction is to adapt protocols or skeletons to the context in which the business partners are interacting.

## ACKNOWLEDGMENTS

This research was partially supported by US National Science Foundation grant DST-0139037 and US Defense Advanced Research Projects Agency contract F3063-00-C-0178.

## REFERENCES

- [1] OWL, Web Ontology Language, Feb. 2004, <http://www.w3.org/TR/2004/REC-owl-features-20040210/>.
- [2] SET, Secure Electronic Transactions (SET) Specifications, 2003, [http://www.setco.org/set\\_specifications.html](http://www.setco.org/set_specifications.html).
- [3] BPEL, Business Process Execution Language for Web Services, Version 1.1, May 2003, [www-106.ibm.com/developerworks/webservices/library/ws-bpel](http://www-106.ibm.com/developerworks/webservices/library/ws-bpel).
- [4] M.P. Singh, A.K. Chopra, N. Desai, and A.U. Mallya, "Protocols for Processes: Programming in the Large for Open Systems," *ACM SIGPLAN Notices*, vol. 39, no. 12, pp. 73-83, 2004.
- [5] N. Desai and M.P. Singh, "Protocol-Based Business Process Modeling and Enactment," *Proc. Int'l Conf. Web Services*, pp. 35-42, 2004.
- [6] A.U. Mallya and M.P. Singh, "A Semantic Approach for Designing Commitment Protocols," *Developments in Agent Comm.*, R.V. Eijk, ed., pp. 37-51, Berlin: Springer, 2005.
- [7] M.P. Singh, "An Ontology for Commitments in Multiagent Systems: toward a Unification of Normative Concepts," *Artificial Intelligence and Law*, vol. 7, pp. 97-113, 1999.
- [8] M. Venkatraman and M.P. Singh, "Verifying Compliance with Commitment Protocols: Enabling Open Web-Based Multiagent

- Systems," *Autonomous Agents and Multi-Agent Systems*, vol. 2, no. 3, pp. 217-236, Sept. 1999.
- [9] I. Horrocks, P.F. Patel-Schneider, H. Boley, S. Tabet, B. Grosz, and M. Dean, "SWRL: A Semantic Web Rule Language Combining OWL and RuleML," May 2004, (W3C Submission), <http://www.w3.org/Submission/2004/SUBM-SWRL-20040521/>.
- [10] P.B. Ladkin and S. Leue, "Interpreting Message Flow Graphs," *Formal Aspects of Computing*, vol. 7, no. 5, pp. 473-509, 1995.
- [11] R. Milner, *Communicating and Mobile Systems: The Pi-Calculus*. Cambridge Univ. Press, 1999.
- [12] C. Canal, L. Fuentes, E. Pimentel, J.M. Troya, and A. Vallecillo, "Adding Roles to Corba Objects," *IEEE Trans. Software Eng.*, vol. 29, no. 3, pp. 242-259, Mar. 2003.
- [13] F. Plasil and S. Visnovsky, "Behavior Protocols for Software Components," *IEEE Trans. Software Eng.*, vol. 28, no. 11, pp. 1056-1076, Nov. 2002.
- [14] R. Allen and D. Garlan, "A Formal Basis for Architectural Connection," *ACM Trans. Software Eng. and Methodology*, July 1997.
- [15] C.A.R. Hoare, *Communicating Sequential Processes*. Prentice Hall, 1985.
- [16] B. Bauer, J. Müller, and J. Odell, *Agent UML: A Formalism for Specifying Multiagent Interaction*, P. Ciancarini and M. Wooldridge, eds. Springer-Verlag, 2001.
- [17] ITU, Message Sequence Charts, Recommendation z.120, Int'l Telecomm. Union, Telecomm. Standardization Sector, 1996.
- [18] I. Jacobson, J. Rumbaugh, and G. Booch, *The Unified Software Development Process*. Addison Wesley, 1999.
- [19] J. Whittle and J. Schumann, "Generating Statechart Designs from Scenarios," *Proc. Int'l Conf. Software Eng.*, pp. 314-323, 2000.
- [20] S. Uchitel, J. Kramer, and J. Magee, "Synthesis of Behavioral Models from Scenarios," *IEEE Trans. Software Eng.*, vol. 29, no. 2, pp. 99-115, Feb. 2003.
- [21] S. Uchitel, J. Kramer, and J. Magee, "Incremental Elaboration of Scenario-Based Specifications and Behavior Models Using Implied Scenarios," *ACM Trans. Software Eng. and Methodology*, vol. 13, no. 1, pp. 37-85, Jan. 2004.
- [22] N. Kavantzaz, D. Burdett, G. Ritzinger, T. Fletcher, and Y. Lafon, "Web Services Choreography Description Language," Dec. 2004, (W3C Working Draft 1. 0), <http://www.w3.org/TR/ws-cdl-10/>.
- [23] "DAML-S: Web Service Description for the Semantic Web," *Proc. First Int'l Semantic Web Conf. (ISWC)*, July 2002.
- [24] V. Agarwal, K. Dasgupta, N. Karnik, A. Kumar, A. Kundu, S. Mittal, and B. Srivastava, "A Service Creation Environment Based on End to End Composition of Web Services," *Proc. Int'l World Wide Web Conf.*, pp. 128-137, 2005.
- [25] M. Solanki, A. Cau, and H. Zedan, "Augmenting Semantic Web Service Descriptions with Compositional Specification," *Proc. Int'l World Wide Web Conf.*, pp. 544-552, 2004.
- [26] Ç.E. Gerede, R. Hull, O. Ibarra, and J. Su, "Automated Composition of e-Services: Lookaheads," *Proc. Int'l Conf. Service Oriented Computing*, 2004.
- [27] *Methodologies and Software Engineering for Agent Systems*, F. Bergenti, M.-P. Gleizes, and F. Zambonelli, eds. Kluwer, 2004.
- [28] P. Bresciani, A. Perini, P. Giorgini, F. Guinchiglia, and J. Mylopoulos, "Tropos: An Agent-Oriented Software Development Methodology," *J. Autonomous Agents and Multi-Agent Systems*, vol. 8, no. 3, pp. 203-236, May 2004.
- [29] F. Zambonelli, N.R. Jennings, and M. Wooldridge, "Developing Multiagent Systems: The Gaia Methodology," *ACM Trans. Software Eng. Methodology*, vol. 12, no. 3, pp. 317-370, 2003.
- [30] K. Baina, B. Benatallah, F. Casati, and F. Toumani, "Model-Driven Web Service Development," *Proc. Advanced Information Systems Eng.: 16th Int'l Conf. (CAiSE)*, June 2004.
- [31] RosettaNet, Home Page, 1998, [www.rosettanet.org](http://www.rosettanet.org).
- [32] *Organizing Business Knowledge: The MIT Process Handbook*, T.W. Malone, K. Crowston, and G.A. Herman, eds. Cambridge, Mass.: MIT Press, 2003.
- [33] P. Vitharana, H. Jain, and F.M. Zahedi, "Strategy-Based Design of Reusable Business Components," *IEEE Trans. Systems, Man, and Cybernetics*, vol. 34, no. 4, pp. 460-474, Nov. 2004.



**Nirmitt Desai** received the BE degree in information technology from Gujarat University, India, in 2001, the MS degree in computer science from North Carolina State University (NCSU) in 2003, and is currently a PhD student in the Department of Computer Science at NCSU. His research interests include software engineering for open environments, business process modeling, and organization modeling.



**Ashok U. Mallya** received the PhD degree in computer science from the North Carolina State University and is a senior engineer at Veraz Networks in San Jose, California. His research interests include agent communication, business process modeling and enactment, and the Semantic Web. His work has sought to achieve a synergistic confluence of these areas to enable dynamic interaction among autonomous enterprises in open environments such as the Internet.



**Amit K. Chopra** is a PhD student in the Department of Computer Science at North Carolina State University (NCSU). His interests span service-oriented computing, business processes, and multiagent systems.



**Munindar P. Singh** is a professor at North Carolina State University. His research interests include multiagent systems and Web services, where he specifically addresses the challenges of trust, service discovery, and business processes and protocols in large-scale open environments. Munindar's recent books include the coauthored *Service-Oriented Computing* and the edited *Practical Handbook of Internet Computing*.

► For more information on this or any other computing topic, please visit our Digital Library at [www.computer.org/publications/dlib](http://www.computer.org/publications/dlib).