

Formalizing and Verifying Protocol Refinements

Scott N. Gerard and Munindar P. Singh
North Carolina State University

A (business) protocol describes, in high-level terms, a pattern of communication between two or more participants, specifically via the creation and manipulation of the commitments between them. In this manner, a protocol offers both flexibility and rigor: a participant may communicate in any way it chooses as long as it discharges all of its activated commitments. Protocols thus promise benefits in engineering cross-organizational business processes. However, software engineering using protocols presupposes a formalization of protocols and a notion of the *refinement* of one protocol by another. Refinement for protocols is both intuitively obvious (e.g., *PayViaCheck* is clearly a kind of *Pay*) and technically nontrivial (e.g., compared to *Pay*, *PayViaCheck* involves different participants exchanging different messages). This paper formalizes protocols and their refinement. It develops Proton, an analysis tool for protocol specifications that overlays a model checker to compute whether one protocol refines another with respect to a stated mapping. Proton and its underlying theory are evaluated by formalizing several protocols from the literature and verifying all and only the expected refinements.

Categories and Subject Descriptors: I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multi-agent Systems*; F.3.1 [Specifying and Verifying and Reasoning about Programs]: Mechanical verification

General Terms: Algorithms, Languages, Theory, Verification

Additional Key Words and Phrases: Commitments, Agent communication, Verification of multiagent systems

1. INTRODUCTION

We focus our attention on business service engagements as realized over the Internet. In current practice, such an engagement is defined rigidly and purely in operational terms. Consequently, the software components of the business partners involved are tightly coupled with each other, and depend closely on the engagement specification. Even small changes in one partner's components must be propagated to others, even when such changes are inconsequential to the business being conducted. Conversely, if the model leaves the engagements unstructured, humans must carry out the necessary interactions manually, with concomitant loss in productivity. We motivate protocols as providing a happy middle ground between rigid automation and flexible manual execution.

Specifically, in contrast with traditional approaches, we model each partner as an autonomous *agent*. The agents participate in a (*business*) *protocol* to realize a service engagement. A protocol describes a pattern of communication between agents. Based on the foregoing, we formulate the following key requirements on a suitable formalization of protocols. First, a protocol is *public*, meaning that it pertains to the messages sent and received by participating agents, not how those agents are implemented. Thus, the semantics of a

Department of Computer Science, North Carolina State University, Raleigh, NC 27695-8206, USA. {sngerard,singh}@ncsu.edu. Gerard is also a Senior Technical Staff Member with IBM.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2011 ACM xxxx-3785/2011/0700-0001 \$5.00

protocol should depend solely on the communications of the agents enacting it, not on their internal policies. Second, the semantics should capture the business meanings of the messages, thereby avoiding operational constraints, and thus enabling the agents to deal better with exceptions and opportunities [Yolum and Singh 2002]. Third, the semantics should be modular: an agent who enacts a protocol correctly may concurrently enact additional protocols. Fourth, designing engagements using protocols presupposes that we support engineering methodologies such as those based on stepwise refinement. We address the above criteria for protocols with an emphasis on their refinement.

We understand a protocol semantically in terms of exactly the set of runs (i.e., computations) that it allows. Following Mallya and Singh [2007], we posit that a putative subprotocol *refines* a putative superprotocol if and only if each run allowed by the subprotocol is also allowed by the superprotocol. In general, a subprotocol would include additional roles and actions: in determining refinement, we disregard those that do not feature in the superprotocol. Doing so facilitates modularity, enhanceability, and reuse of protocols.

Consider a simple protocol *Pay* consisting of two actions where a payer first commits to paying a payee, and next pays. Now consider a protocol *PayViaMM* where the payer first pays a middleman, who in turn pays the payee. Both *Pay* and *PayViaMM* send a payment from the payer to the payee. Even though *PayViaMM* involves an additional role (middleman) and *PayViaMM* uses different messages (two payment messages instead of one), we expect *PayViaMM* refines *Pay*, because *PayViaMM* makes a payment as *Pay* specifies. Similarly, we expect *PayViaCheck* and *PayViaCredit* also refine *Pay*. We imagine a service engagement design exercise where protocol designers begin by identifying the need for payment as *Pay*, then refine it to *PayViaMM*, and then to *PayViaCheck*. The designers may build or find an existing repository of protocols (analogous to taxonomies of business processes [Malone et al. 2003]). The question we address is how can protocols in such a repository be expressed so that their refinements can be rigorously verified.

1.1 Proton: Approach and Contributions

We formulate refinement in technical terms and show how to compute it via a tool called *Proton*. We specify a protocol declaratively in terms of (i) its roles, (ii) the guarded messages the roles exchange, and (iii) the meaning of each message as a set of actions on the public state of the roles, sometimes termed the *social state* [Baldoni et al. 2010]. Commitments between roles are central to our approach [Singh 1999]. Section 2.4 provides additional details. For now, suffice it to say that a state of a protocol is determined by what atomic propositions hold therein (some propositions specify the states of commitments).

We define the semantics of a protocol precisely in terms of the runs (i.e., sequences of actions) it allows. Informally, a *subprotocol* refines a *superprotocol* if and only if the latter allows all the runs the former allows. However, refinement is nontrivial because the protocols may involve different roles and messages, the messages may have different meanings, and the meanings may be at different levels of abstraction. Hence, we define refinement only with respect to a mapping of meanings from the superprotocol to the subprotocol. For example, the payment in *Pay* maps to two payments in *PayViaMM*.

Our approach for verifying refinement takes three inputs: formal descriptions of a putative superprotocol and subprotocol, and a mapping between them. We reduce the protocol descriptions to their canonical forms, taking into account the mapping provided. We generate an input to an existing model checker consisting of (i) a specification of a temporal logic model and (ii) temporal formulae whose truth in the model verifies refinement.

1.2 Contributions

Our main contributions are as follows. One, we offer the *first* approach that computes the refinement for protocols based on static analysis of protocol specifications. Two, we formulate a notion of the serial composition of commitments, which can have broader applications than this paper, e.g., in the treatment of commitments in coalitions.




Fig. 1. Refinements demonstrated by Proton (arrows point from subprotocols to superprotocols).

Further, we have implemented our approach in the Proton tool that overlays the well-known model checker MCMAS (<http://www-lai.doc.ic.ac.uk/mcas/>). Figure 1 summarizes some protocol refinements that Proton verifies (under the obvious mappings) based on the above and other examples known from the literature.

1.3 Organization

Section 2 overviews our syntax and semantics and briefly reviews commitments. Section 3 introduces our running examples for payment and order protocols. Section 4 formalizes our definitions of protocols, mappings between protocols, and protocol refinement. Section 5 describes how Proton generates input for the MCMAS model checker and the CTL formulae that must be satisfied for protocol refinement to hold. Section 6 pulls the previous sections together, illustrating how protocol *PayViaMM* refines, or fails to refine, protocol *Pay* under various mappings. Section 7 shows that the algorithmic implementation in Section 5 is correct with respect to the theoretical framework of Section 4. Section 8 describes the related literature and important future directions.

2. TECHNICAL FRAMEWORK

We adopt the following terminology. A *subprotocol* refines a *superprotocol*. In hyphenated form, *super-x* and *sub-x* refer to element *x* as it occurs in the superprotocol and subprotocol, respectively. For example, a *super-role* is a role defined in the superprotocol and a *sub-commitment* is a commitment defined in the subprotocol.

2.1 Interpreted Systems

We adopt Lomuscio and colleagues’ [Lomuscio and Raimondi 2006; Cohen et al. 2009] formalization of a multiagent system as an *interpreted system*. Importantly, protocols involve roles, not agents. We presume no knowledge of the internals of an agent playing a role and consider all possible strategies that a role may follow in a protocol.

Each role is described by a set of possible local states, a set of local actions, a local strategy listing the legal actions in each local state, and a local progression function defining the progression of the role’s local state based on the actions performed by all the roles. To clarify the terminology, our *role* and *strategy* respectively map to *agent* and *protocol* in work by Lomuscio and colleagues.

Definition 2.1 Interpreted System. An interpreted system \mathcal{I} is

$$\mathcal{I} = \langle \Sigma, P, PV, L^i, Act^i, AP^i, t^i, G, G_0, F \rangle$$

$\Sigma = \{1, \dots, n, e\}$ is a set of three or more roles, including a distinguished role e that stands for the environment. P is a set of atomic propositions. Let $i \in \Sigma$ range over all roles and the environment e . L^i is a nonempty set of possible local states for each i . Act^i is a set of actions for each i . $AP^i : L^i \times L^e \mapsto 2^{Act^i}$ is the local strategy for each i . In local state $l \in L^i, l^e \in L^e$, role i can perform only the actions in $AP^i(l, l^e)$. $G \subseteq L^1 \times \dots \times L^n \times L^e$ is the set of reachable global states. For any global state $g \in G$, we write g^i for the i -th component in g , i.e., the local state of role i in g . $G_0 \subseteq G$ is a nonempty set of initial global states. $PV : P \mapsto 2^G$ is the evaluation function for propositions. The set of joint actions is $Act = Act^1 \times \dots \times Act^n \times Act^e$. $t^i : L^i \times L^e \times Act \mapsto L^i$ is the local progression function for role $i \in \Sigma \setminus e$, and $t^e : L^e \times Act \mapsto L^e$ is the progression function for the environment. All roles progress simultaneously. The global progression function is $T : G \times Act \mapsto G$ and is defined such that $T(g, a) = g'$ iff $\forall i \in \Sigma \setminus e : t^i(g^i, g^e, a^i) = g'^i$ and $t^e(g^e, a^e) = g'^e$. T must be serial ($\forall g \in G, \exists a \in Act, \exists g' \in G : T(g, a) = g'$). F is a set of Boolean fairness conditions, each of which must be true infinitely often on all legal execution paths. A path π in \mathcal{I} is an infinite sequence of global states $\langle g_0, g_1, \dots \rangle$ in G such that every pair of adjacent states is a legal transition, i.e., $\forall i \geq 0 : \exists a \in Act : T(g_i, a) = g_{i+1}$. The i -th state in path π is denoted π_i , and the set of all paths starting at $g \in G$ is denoted $\Pi(g)$.

Given an interpreted system \mathcal{I} , we associate with it a Kripke model $\mathcal{K} = (G, G_0, T, PV)$ where G is the set of possible worlds understood as the reachable states of \mathcal{K} , built from the set of initial states G_0 by iterating the global progression function T ; and, the temporal relation $T \subseteq G \times Act \times G$ connects global states based on the joint actions. The labeling function PV is the propositional labeling function.

The grammar of the temporal language CTL is (*PropName* is an atomic proposition)

$$\phi ::= \text{PropName} \mid \neg\phi \mid \phi \vee \psi \mid \phi \wedge \psi \mid \mathbf{AG}\phi \mid \mathbf{AF}\phi \mid \mathbf{A}[\phi\mathbf{U}\psi] \mid \mathbf{EG}\phi \mid \mathbf{EF}\phi \mid \mathbf{E}[\phi\mathbf{U}\psi]$$

Our temporal formulae use the standard CTL temporal logic operators: **A** (for all paths), **E** (for some path), **G** (on all future states on a path), and **F** (eventually on the path). The following is the semantics for CTL, specified relative to the Kripke structure \mathcal{K} at state g . Given a model \mathcal{K} for an interpreted system, \mathcal{I} satisfies a CTL formula ϕ if and only if $\mathcal{K}, g_0 \models \phi$, where $g_0 \in G_0$ is a starting state.

$$\begin{aligned} \mathcal{K}, g \models p & \quad \text{iff } p \in g \\ \mathcal{K}, g \models \neg\phi & \quad \text{iff it is not the case } \mathcal{K}, g \models \phi \\ \mathcal{K}, g \models \phi \wedge \psi & \quad \text{iff } \mathcal{K}, g \models \phi \text{ and } \mathcal{K}, g \models \psi \\ \mathcal{K}, g \models \phi \vee \psi & \quad \text{iff } \mathcal{K}, g \models \phi \text{ or } \mathcal{K}, g \models \psi \\ \mathcal{K}, g \models \mathbf{AG}\phi & \quad \text{iff } \forall \pi \in \Pi(g), \forall i \geq 0, \mathcal{K}, \pi_i \models \phi \\ \mathcal{K}, g \models \mathbf{AF}\phi & \quad \text{iff } \forall \pi \in \Pi(g), \exists i \geq 0, \mathcal{K}, \pi_i \models \phi \\ \mathcal{K}, g \models \mathbf{A}[\phi\mathbf{U}\psi] & \quad \text{iff } \forall \pi \in \Pi(g), \exists k \geq 0, \mathcal{K}, \pi_k \models \psi \text{ and } \forall 0 \leq j < k, \mathcal{K}, \pi_j \models \phi \\ \mathcal{K}, g \models \mathbf{EG}\phi & \quad \text{iff } \exists \pi \in \Pi(g), \forall i \geq 0, \mathcal{K}, \pi_i \models \phi \\ \mathcal{K}, g \models \mathbf{EF}\phi & \quad \text{iff } \exists \pi \in \Pi(g), \exists i \geq 0, \mathcal{K}, \pi_i \models \phi \\ \mathcal{K}, g \models \mathbf{E}[\phi\mathbf{U}\psi] & \quad \text{iff } \exists \pi \in \Pi(g), \exists k \geq 0, \mathcal{K}, \pi_k \models \psi \text{ and } \forall 0 \leq j < k, \mathcal{K}, \pi_j \models \phi \end{aligned}$$

We write $a \models b$ if and only if for all models \mathcal{K} and states g , $\mathcal{K}, g \models b$ holds whenever $\mathcal{K}, g \models a$ holds.

```

Protocol ::= protocol ProtoName {role ⟨RoleName;⟩* prop ⟨PropName;⟩*
          commitment ⟨Commit;⟩* message ⟨Message;⟩*}
Commit   ::= ComName = C(Debtor, Creditor, Ant, Csq)
Message  ::= Snd → Rcv : ⟨[Guard]⟩?MsgName means {actexp}
Debtor   ::= RoleExp
Creditor ::= RoleExp
Ant      ::= ActExp
Csq     ::= ActExp
Snd     ::= RoleName
Rcv     ::= RoleName
Guard   ::= MsgExp
ActExp  ::= Actions

Map      ::= map MapName : ProtoName ↦ ProtoName{role ⟨RoleMap;⟩*
          prop ⟨PropMap;⟩* commitment ⟨CommitMap;⟩*}
RoleMap  ::= RoleName ↦ RoleExp
PropMap  ::= PropName ↦ ActExp
CommitMap ::= ComName ↦ CommitExp

RoleExp  ::= RoleName | {RoleName ⟨, RoleName⟩*}
MsgExp   ::= MsgName | ¬MsgExp | MsgExp ∨ MsgExp | MsgExp ∧ MsgExp | MsgExp → MsgExp
ActExp   ::= Action | ¬ActExp | ActExp ∨ ActExp | ActExp ∧ ActExp | ActExp → ActExp
CommitExp ::= ComName | CommitExp ⊕ CommitExp
Actions  ::= Action ⟨, Action⟩*
Action   ::= PropName | CREATE(ComName) | TRANSFER(ComName) | RELEASE(ComName)
          | CANCEL(ComName)
    
```

Fig. 2. Proton input syntax in BNF.

2.2 Proton Syntax

Proton verifies whether one protocol refines another under a mapping. Figure 2 shows Proton’s input syntax (here $|$ separates alternatives, $\langle A \rangle^*$ is zero or more repetitions of A , and $\langle A \rangle^?$ is an optional occurrence of A). Notice that $PropName$ is simply an atomic proposition.

The *Protocol* nonterminal describes the syntax for a protocol (as Listing 1 exemplifies). A protocol declares roles, propositions, commitments, and messages. Each message m is sent from a sender ($m.snd$) to a receiver ($m.rcv$), has a guard ($m.guard$) which must be true before the message can be sent, a set of actions ($m.actions$), and means a conjunction of these actions ($m.actexp$). Actions are either propositions (being set true) or a commitment operation (being performed). Boolean negation (\neg) is allowed in antecedent, consequent, and guard expressions to check state, but a message meaning cannot set a proposition to false using a negated action.

The *Map* nonterminal describes the syntax for a mapping between two protocols (as Listing 3 exemplifies). A mapping maps individual roles, propositions, and commitments from the putative superprotocol to expressions in the putative subprotocol. *ProtoName*, *MapName*, *RoleName*, *PropName*, *ComName*, *MsgName*, and *Action* are names.

The serial composition operator, \oplus , chains two commitments together and is described in Section 2.5. We write $\bigoplus_i C_i$ for a left-associated chain $(C_1 \oplus C_2) \oplus \dots \oplus C_n$. In Section 2.6, we compare commitments between superprotocol and subprotocol, under an abstraction mapping M , using commitment covering (\leq_M).

2.3 Proton Semantics

Proton’s semantics is based on interpreted systems: it constructs an interpreted system from an input superprotocol, subprotocol, and mapping.

Each state g is a set of true propositions p_i . All propositions are false in the initial state g_0 . Actions cause state transitions. For this paper, we use a simplified model of actions, assuming actions (i) always succeed, (ii) have definite outcomes (no uncertainty), and (iii) have no side-effects. The actions for role i , Act^i , are the propositional and commitment actions $Act^i = \{p_i\} \cup \{a(C_j) | a \in Act_C\} \cup \{nop\}$ where \mathcal{A} is the set of protocol propositions, \mathcal{C} is the set of protocol commitments, for all $p_i \in \mathcal{A}$, $Act_C = \{\text{CREATE, TRANSFER, RELEASE, CANCEL}\}$, for all $C_j \in \mathcal{C}$, and nop represents no-operation.

Operationally, messaging is point-to-point and synchronous. All protocol state is stored in the “environment” (effectively, a distinguished agent), and is globally accessible by all roles (a current simplification). At each time step, the environment schedules one role to execute next (interleaved execution). When scheduled, the role’s agent (i) determines which of its messages are currently enabled, by accessing the protocol’s global state and evaluating each message’s guard expression, (ii) chooses an enabled message to send or chooses “no-operation”, (iii) performs all actions in message’s meaning, in any order, and (iv) updates the protocol’s state.

In every global state g of the interpreted system, each commitment C_i has a state, $C_i.status \in Status_C$, where $Status_C = \{null, cond, detached, dis, xfer, rel, can\}$, whose value can be any of the states in Figure 3. We define propositions for the expressions $C_i.status = x$ and $C_i.status \neq x$ in each state g . For each commitment C_i , we evaluate the occurrence of the commitment operations using the four propositions:

$$\begin{aligned} \text{CREATE}(C_i) &\triangleq C_i.status \neq null \\ \text{TRANSFER}(C_i) &\triangleq C_i.status = xfer \\ \text{RELEASE}(C_i) &\triangleq C_i.status = rel \\ \text{CANCEL}(C_i) &\triangleq C_i.status = can \end{aligned}$$

As Section 1.1 mentions and Definition 4.6 formalizes, refinement depends upon a mapping between protocols to account for their different levels of abstraction. Specifically, we must map each (i) super-role to a set of sub-roles, (ii) super-proposition to a Boolean expression of sub-propositions, and (iii) super-commitment to an expression of sub-commitments. Proton combines two protocols and a mapping to (i) construct an interpreted system model \mathcal{I} from the subprotocol’s propositions, commitments, and guarded actions, as specified in Definition 4.2 and (ii) generate appropriate CTL formulae as specified in Section 5.5. The refinement in consideration holds if and only if the constructed model satisfies all the CTL formulae that Proton generates.

2.4 Background on Commitments

Commitments are a formal and concise method of describing how agent roles commit to performing future actions [Singh 1999; Yolum and Singh 2002]. We extend previous commitment definitions in two ways. First, we allow both debtors and creditors to be sets of roles, enabling us to handle commitment chains with multiple debtors and multiple creditors. Second, we introduce a new TRANSFER commitment operation to unify and replace prior uses of DELEGATE and ASSIGN.

A commitment $C_{\{\text{debtors}\}, \{\text{creditors}\}}$ (*antecedent, consequent*) means that the debtors com-

mit to the creditors, that if the antecedent holds, they will bring about the consequent. In an active commitment whose antecedent is false, the debtors are conditionally committed to the creditors. When the antecedent becomes true, we say the commitment is *detached*, and the debtors become unconditionally committed to the creditors.

Consider the following two example commitments drawn from the payment scenarios: (i) $C_{\text{Payer, Payee}}(\text{promise}, \text{pay})$: the payer conditionally commits to paying the payee; and (ii) $C_{\{\text{Payer, MM}\}, \text{Payee}}(\text{promise}, \text{pay}_P \wedge \text{pay}_M)$: the payer and middleman commit to paying the payee via pay_P (the payer’s payment) and pay_M (the middleman’s payment).




Fig. 3. State transition diagram for commitments (states in lowercase; transitions in small caps).

Figure 3 shows the state transition diagram for a commitment. The states of a commitment are (i) *null*: where the commitment does not yet exist; (ii) *cond* (active and conditional): after CREATE with antecedent and consequent being false, and with no other operations; (iii) *detached* (active and detached): after CREATE with antecedent true, consequent false, and with no other operations; (iv) *dis* (discharged): after CREATE and consequent being true; (v) *xfer* (transferred): after CREATE and TRANSFER; (vi) *rel* (released): after CREATE and RELEASE; and (vii) *can* (canceled): after CREATE and CANCEL. A commitment in states *cond* or *detached* is said to be *active*. A commitment in states *dis*, *xfer*, *rel*, or *can* is said to be *resolved*. The commitment operations are (i) CREATE, performed only by debtors, creates an active commitment; (ii) DETACH, which occurs implicitly when the antecedent becomes true; (iii) DISCHARGE, which occurs implicitly when the consequent becomes true; (iv) TRANSFER, performed by either debtors or creditors, deactivates the current commitment and replaces it by another commitment; (v) RELEASE, performed only by creditors, deactivates the commitment, thus releasing the debtors; and (vi) CANCEL, performed only by debtors, cancels, deactivates, and “breaks” the debtors’ commitment.

We understand a protocol P as refining a protocol Q if and only if the correct enactments of P are also correct enactments of Q . As in previous work, in a correct enactment of a protocol each detached commitment must eventually resolve. Debtors may act before they are required to do so and the consequent may become true before the antecedent. In general, there is no guarantee that autonomous debtors do not arbitrarily CANCEL. In practice, the creditors would assume the debtors are trustworthy or the setting would include an external mechanism (such as penalties) to ensure the debtors’ compliance.

2.5 Serial Composition of Commitments

In *PayViaMM*, where the payer commits to a middleman who commits to the payee, the two commitments together effectively commit the payer to the payee. We introduce serial composition as a general way to chain commitments over intermediaries, computing a single, resultant commitment. The serial composition of commitments is a static construct,

but the resultant commitment dynamically progresses through the states in Figure 3 as the protocol progresses.

Definition 2.2. Let C_A and C_B be two commitments where $C_A.csq \models C_B.ant$. Then, the serial composition of C_A and C_B is the commitment $C_{\oplus} = C_A \oplus C_B$ whose components are specified precisely as follows:

$$C_{\oplus}.debt := C_A.debt \cup C_B.debt \quad (1)$$

$$C_{\oplus}.cred := C_A.cred \cup C_B.cred \quad (2)$$

$$C_{\oplus}.ant := C_A.ant \quad (3)$$

$$C_{\oplus}.csq := C_A.csq \wedge C_B.ant \wedge C_B.csq \quad (4)$$

The state of C_{\oplus} is defined based on the states of C_A and C_B . C_{\oplus} is created exactly when both C_A and C_B are created. C_{\oplus} is respectively transferred, released, or canceled when at least one of C_A and C_B is transferred, released, or canceled.

Singh's [2008] formalization of commitments includes the similar idea of commitment chaining, but serial composition additionally captures the intuition of a coalition of roles. The above well-definedness condition $C_A.csq \models C_B.ant$ follows Singh's [2008] definition for chaining, although serial composition accommodates different roles. The second commitment becomes active ($C_B.ant$) whenever the first commitment resolves ($C_A.csq$), including the case where the first debtors perform without being required to do so ($C_A.ant$ always false).

Informally, we say debtors are *responsible* for their commitments, and creditors are *beneficiaries* of their commitments. In a detached commitment, the debtors are responsible for eventually making the consequent true. Responsibility can be *several* (each debtor is responsible for just its portion), *joint* (each debtor is individually responsible for the entire commitment), or *joint and several* (the creditors hold one debtor fully responsible, who then pursues other debtors). We use *several* responsibility so that, in serial composition of commitments, a debtor is never compelled to assume additional responsibilities. The result of serial composition is useful for reasoning about multiple commitments, but the original commitment expression, with its individual commitments, must be retained to determine which role(s) failed to perform if the resultant consequent is not produced.

C_{\oplus} states the union of debtors is committed to the union of creditors to bring about the consequent $C_A.csq \wedge C_B.ant \wedge C_B.csq$ when antecedent $C_A.ant$ is true. Debtors are *severally* responsible for C_{\oplus} , so that debtors are never compelled to assume additional responsibilities. Every debtor in $C_A.debt$ is partially responsible for discharging C_A , and thus is partially responsible for discharging C_{\oplus} . Also, every debtor in $C_B.debt$ has some responsibility for C_{\oplus} . Equation 1 captures this intuition for debtors. Equation 2 captures the analogous intuition for creditors.

If we order the states of a commitment as $null < can < rel < xfer < cond < detached < dis$, then the state of a serial composition is the minimum of its constituents' states: $C_{\oplus}.status = \min(C_A.status, C_B.status)$. That is, a serial composition progresses no further than its least constituent.

Because of Equations 3 and 4, serial composition is neither commutative nor associative. However, it creates commitments that are at least as strong as, and typically stronger than, their inputs. $C_A \oplus C_B$ is typically stronger than C_A because, even though both have the same antecedent ($C_A.ant$), in general, $C_A \oplus C_B$ has a stronger consequent ($C_B.csq$ vs.

$C_A.csq \wedge C_B.ant \wedge C_B.csq$). The lemma below shows serial composition is not always stronger, because \oplus is idempotent: a commitment can be usefully added to a commitment chain only once.

LEMMA 2.3. *If C_k is any commitment in a commitment chain $\bigoplus_{1 \leq i \leq n} C_i$, then composing C_k again does not increase the strength.*

$$\left(\bigoplus_{1 \leq i \leq n} C_i \right) \oplus C_k = \bigoplus_{1 \leq i \leq n} C_i$$

PROOF. If $\bigoplus_i C_i$ is well defined, then so is $(\bigoplus_i C_i) \oplus C_k$. By inspection, Equations (1–4) yield the same results for both sides. \square

2.6 Commitment Covering

Because commitments are crucial to our semantics of protocols, commitments are also crucial to refinement. And because we need to compare two protocols, we need a mechanism to compare two commitments. Specifically, each super-commitment must be *covered by*, or make at least the same commitment as, another relevant sub-commitment. The commitment comparison accommodates a mapping to account for the commitments being expressed at different levels of abstraction. Definition 2.4 extends Chopra and Singh’s [2009] notion of *commitment strength*. In addition to the logical relationships between antecedents and consequents, this definition incorporates the mapping of roles and propositions.

Definition 2.4 Commitment Covering. A stronger commitment C_S covers (is stronger than) a weaker commitment C_W with respect to mapping M , written $C_W \leq_M C_S$, if and only if

$$\forall d \in C_W.debt \quad M(d) \cap C_S.debt \neq \emptyset \quad (5)$$

$$\forall c \in C_W.cred \quad M(c) \cap C_S.cred \neq \emptyset \quad (6)$$

$$M(C_W.ant) \models C_S.ant \quad (7)$$

$$C_S.csq \models M(C_W.csq) \quad (8)$$

where $M(x)$ maps (super-) element x in C_W to an expression of (sub-) elements in C_S .

Every super-debtor is partially (severally) responsible for discharging the super-commitment. Each super-role is mapped to (implemented by) a set of sub-roles $M(d)$. We require each super-debtor’s responsibilities be passed to one or more of its sub-debtors. Together these sub-debtors assume the super-debtor’s responsibilities. Equation 5 captures the requirement that every super-debtor’s responsibilities must pass to at least one of its sub-debtors, so that responsibilities are not lost. Similarly, each super-creditor is a partial beneficiary of the super-commitment. Equation 6 captures the requirement that every super-creditor’s benefit pass to at least one of its sub-creditors.


In many situations, multiple sub-commitments must be combined to cover a single super-commitment. In those cases, a super-commitment is covered by the serial composition of multiple sub-commitments.

We visualize our explanations by diagramming a commitment as a labeled arrow.

$$\frac{\text{debt}^{(\cup)} \quad C_{\text{name}}^{(\oplus)} \quad \text{cred}^{(\cup)}}{\text{ant}^{(\wedge)} \quad \text{csq}^{(\wedge)}}$$

The name of the commitment is written in the top center of the arrow. Debtors and creditors are above the arrow and the antecedent and consequent below it. When multiple terms appear in a position, they are implicitly combined using the operator in parentheses.

As a simple example, consider commitments $C_1 = C_{\text{Payer}, \text{Payee}}(\text{promise}, \text{pay})$, $C_2 = C_{\text{Buyer}, \text{Seller}}(\text{order}, \text{pay})$, and an abstraction mapping M that is defined on roles and propositions as follows: (i) $\text{Payer} \mapsto \{\text{Buyer}\}$, (ii) $\text{Payee} \mapsto \{\text{Seller}\}$, (iii) $\text{promise} \mapsto \text{order}$, and (iv) $\text{pay} \mapsto \text{pay}$. The diagram shows C_2 covers C_1 .



As another example, we obtain $C(\text{order}, \text{ship}) \leq C(\text{order} \vee \text{freeCoupon}, \text{ship})$ by Equation 7, since the stronger commitment detaches when order or freeCoupon is true. And, likewise $C(\text{order}, \text{ship}) \leq C(\text{order}, \text{ship} \wedge \text{expressDelivery})$ holds by Equation 8, since to discharge the stronger commitment requires expressDelivery in addition to ship .

3. RUNNING EXAMPLES OF PROTOCOLS

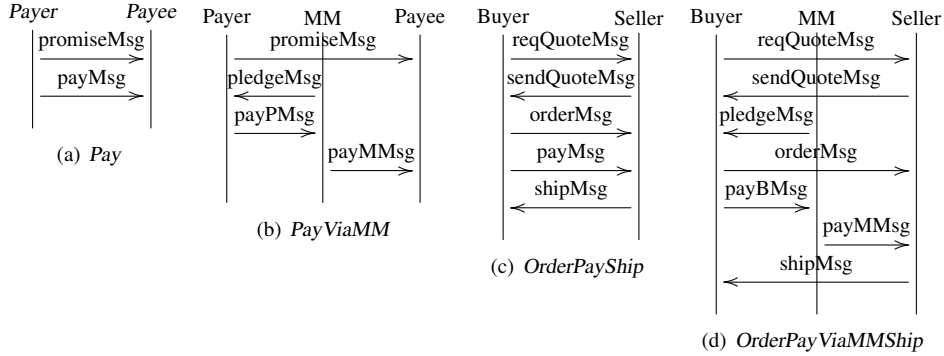


Fig. 4. Suggestive sequence diagrams for selected protocols (in general, alternative sequences may occur).

3.1 Protocol Descriptions

We introduce four running examples. *Pay* and *PayViaMM* are payment protocols whereas *OrderPayShip* and *OrderPayViaMMShip* are order protocols involving payments. Suggestive sequence diagrams are presented in Figure 4. Protocol *Pay* in Figure 4(a) is a basic payment protocol between a payer and a payee. The payer may commit to pay the payee by sending the promise message. Later, it sends a payment message directly to the payee.

Listing 1 shows the Proton specification of protocol *Pay*. Lines 2–5 declare roles Payer and Payee, propositions *promise* and *pay*, and the commitment. Both *promiseMsg* and *payMsg* messages are sent by Payer to Payee. A message may be sent only if its guard (the expression between [and]) is true. The guard for *payMsg* in Line 8 is *promiseMsg*. If no

Listing 1 Pay Protocol

```

1: protocol Pay {
2:   role Payer, Payee;
3:   prop promise; pay;
4:   commitment
5:     Cpay = C(Payer, Payee, promise, pay);
6:   message
7:     Payer → Payee : promiseMsg means {promise, CREATE(Cpay)};
8:     Payer → Payee : [promiseMsg] payMsg means {pay};
9: }

```

guard is explicitly specified, as is the case for *promiseMsg* in Line 7, it is implicitly true. A message's meaning is expressed as a set of actions after *means* and between { and }.

In protocol *PayViaMM* (pay via middleman), if the payer chooses to do so, it commits to paying the payee by sending the promise message. The middleman commits to sending *payM* if the payer sends *payP* by sending *pledgeMsg*. The payer then sends a payment indirectly to the payee, first paying the middleman, who in turn pays the payee. The sequence diagram in Figure 4(b) shows a typical interaction. In this case, other acceptable runs also exist: for example, the middleman may send the *payM* message before the *payP* message. The Proton specification for *PayViaMM* is shown in Listing 2. Middleman commits to Payer to pass along any payment it receives (Line 11). Payer will not pay Middleman without this commitment (Line 12). Since *payMMsg* has an implicit guard of true (Line 14), Middleman is allowed to pay early.

Listing 2 PayViaMM Protocol

```

1: protocol PayViaMM {
2:   role Payer, MM, Payee;
3:   prop promise;
4:     payP; //payment from Payer to MM
5:     payM; //payment from MM to Payee
6:   commitment
7:     CpayP = C(Payer, Payee, promise, payP);
8:     CpayM = C(MM, Payer, payP, payM);
9:   message
10:    Payer → Payee : promiseMsg means {promise, CREATE(CpayP)};
11:    MM → Payer : pledgeMsg means {CREATE(CpayM)};
12:    Payer → MM : [promiseMsg ∧ pledgeMsg]
13:      payPMsg means {payP};
14:    MM → Payee : payMMsg means {payM};
15: }

```

Protocol *OrderPayShip* in Figure 4(c) supports a buyer placing an order with a seller. The buyer requests a price quote for a good from the seller. The seller sends the price quote along with its commitment to ship the good if the buyer orders. The buyer can accept the seller's offer by placing an order, which creates its commitment to pay for the good if it

ships. The seller can ship first, or the buyer can pay first. Protocol *OrderPayViaMMShip* in Figure 4(d) is similar to *OrderPayShip* except that it incorporates *PayViaMM* for payment.

3.2 Mapping Abstractions across Protocols

Since superprotocols represent higher-level abstractions than subprotocols, comparing protocols must address differences in abstraction level. To this end, we map elements (roles, propositions, and commitments) of a putative superprotocol to elements of a putative subprotocol. We map every super-element to an expression of sub-elements, but a subprotocol may contain sub-elements that do not correspond with any super-element.

Listing 3 Mapping M_1 : *Pay* to *PayViaMM*

```

1: map  $M_1$ : Pay  $\mapsto$  PayViaMM {
2:   role
3:     Payer  $\mapsto$  {Payer};
4:     Payee  $\mapsto$  {Payee};
5:   prop
6:     promise  $\mapsto$  promise;
7:     pay  $\mapsto$  payP  $\wedge$  payM;
8:   commitment
9:      $C_{pay} \mapsto C_{payP} \oplus C_{payM};$     //requires  $C_{pay} \leq_{M_1} C_{payP} \oplus C_{payM}$ 
10: }
```

Consider mapping M_1 in Listing 3 from *Pay* to *PayViaMM*. Each super-role is mapped to a set of sub-roles. Line 3 maps the *Payer* super-role in *Pay* to the *Payer* sub-role in *PayViaMM*. Each super-proposition in *Pay* is mapped to a Boolean expression of sub-propositions in *PayViaMM*. Line 7 maps *pay* to the conjunction of *payP* and *payM*. Notice that *payP* and *payM* are messages sent by different roles in *PayViaMM*; thus even the simple Line 7 demonstrates the generality of our mapping approach. Line 9 maps super-commitment C_{pay} to the serial composition of sub-commitments C_{payP} and C_{payM} .

Listing 4 Alternative Mapping M_2 : *Pay* to *PayViaMM*

```

1: map  $M_2$ : Pay  $\mapsto$  PayViaMM {
2:   role
3:     Payer  $\mapsto$  {Payer, MM};
4:     Payee  $\mapsto$  {Payee};
5:   ...
6: }
```

There can be multiple mappings between some protocol pairs. The Middleman role does not appear in mapping M_1 . We can construct alternative mappings that group the Middleman into coalitions with different super-roles. Mapping M_2 in Listing 4 and Mapping M_3 in Listing 5 are each the same as M_1 except for their role mappings: M_2 groups the Middleman into a coalition with *Payer* and M_3 into a coalition with *Payee*. *PayViaMM* refines *Pay* under all three mappings M_1 , M_2 , and M_3 .

Listing 5 Alternative Mapping M_3 : *Pay* to *PayViaMM*

```

1: map  $M_3$ : Pay  $\mapsto$  PayViaMM {
2:   role
3:     Payer  $\mapsto$  {Payer};
4:     Payee  $\mapsto$  {Payee, MM};
5:   ...
6: }
```

We require each commitment to be explicit mapped. A commitment mapping must not violate the role and proposition mappings, but that is not always sufficient to uniquely determine the commitment mapping. It is possible that a super-commitment can be mapped to multiple serial compositions that meet all constraints. In a hypothetical *PayViaTwoMM* protocol, where the payment can be made through either *Middleman*₁ or *Middleman*₂, the super-commitment $C_1 = C_{Payer, Payee}(promise, pay)$ can be mapped to either of two serial compositions (the protocol designer chooses between them based on other factors).

$$C_1 \mapsto C_{Payer, MM_1}(promise, payMM_1) \oplus C_{MM_1, Payee}(payMM_1, pay)$$

$$C_1 \mapsto C_{Payer, MM_2}(promise, payMM_2) \oplus C_{MM_2, Payee}(payMM_2, pay)$$

Mapping B_1 in Listing 6 shows a possible mapping between *Pay* and *PayViaMM*, similar

Listing 6 Nonrefining Mapping B_1 : *Pay* to *PayViaMM*

```

1: map  $B_1$ : Pay  $\mapsto$  PayViaMM {
2:   role
3:     Payer  $\mapsto$  {Payer};
4:     Payee  $\mapsto$  {Payee};
5:   prop
6:     promise  $\mapsto$  promise;
7:     pay  $\mapsto$  payP  $\wedge$  payM;
8:   commitment
9:      $C_{pay} \mapsto C_{payM} \oplus C_{payP};$     //wrong order
10: }
```

to M_1 except the serial composition in Line 9 combines the commitments in the wrong order. In Section 6, we show *PayViaMM* does *not* refine *Pay* under mapping B_1 .

4. FORMALIZING PROTOCOLS AND THEIR REFINEMENT

We assume a set of atomic propositions that describe the state of the world and states of relevant commitments. We define actions as atomic propositions (being made true) and commitment operations (being performed). Messages set propositions true, but not false.

Definition 4.1. A protocol is a sextuple $\langle \mathcal{R}, \mathcal{M}, \mathcal{C}, \mathcal{A}, \mathcal{S}, \mathcal{G} \rangle$ corresponding respectively to (i) a set \mathcal{R} of roles; (ii) a set \mathcal{M} of message names; (iii) a set \mathcal{C} of commitments; (iv) a set \mathcal{A} of Boolean propositions and commitment states; (v) a set \mathcal{S} of states, $\mathcal{S} \subseteq 2^{\mathcal{M}}$ such that $\emptyset \in \mathcal{S}$ and if $s \in \mathcal{S}$, $gs \in \mathcal{G}$, and $s \in gs.guard$ then $s \cup gs.msg \in \mathcal{S}$; and (vi) a set \mathcal{G} of guarded statements of the form $\langle snd, rcv, guard, msg, actions \rangle$ with $snd, rcv \in \mathcal{R}$, $guard \subseteq \mathcal{S}$, $msg \in \mathcal{M}$, and $actions = \{a_i \in \mathcal{A}\} \cup \{Act_{\mathcal{C}}(C_j \in \mathcal{C})\} \cup \{nop\}$. In addition,

we impose the *no overlap* constraint: $\forall gs_1, gs_2 \in \mathcal{G}$, if $gs_1.actions \cap gs_2.actions \neq \emptyset$ then $gs_1.guard \cap gs_2.guard = \emptyset$.

Each message corresponds to an atomic proposition recording whether the message has been sent. Each global state $s \in \mathcal{S}$ is a set of (the atomic propositions corresponding to) the messages that have been sent in that state (Item v). Each guarded statement $gs \in \mathcal{G}$ has a guard $gs.guard$ which is a set of states, and a meaning $gs.actexp$ —a conjunctive expression of actions. A message msg can be sent by the sender ($gs.snd$) to the receiver ($gs.rcv$) in state s only if $s \in gs.guard$. When m is sent, the action expression $gs.actexp$ becomes true in the next state. The actions corresponding to different messages may be interleaved. The *no overlap* constraint ensures that if two or more super-actions contain the same sub-action, and both super-actions are enabled in a state, then the occurrence of the common sub-action in a sub-run is unambiguous as to which super-action it corresponds to, which recall is key to our notion of refinement.

4.1 Protocol Enactment

We introduce a *run*, a possible computation through our model, as a basis for our semantics. A run, notated π , is an alternating sequence of states and actions $\langle s_0, a_1, s_1, a_2, s_2, \dots \rangle$ such that s_{i+1} results from performing a_{i+1} in s_i . The length of π is written $|\pi|$.

We can now express two key intuitions. First, the semantics of a protocol is simply the set of runs it allows. Underlying each run is a coarser *message enactment*: a sequence of states and messages where each message's guard is true in the state where the message occurs. Second, a protocol refines another if and only if the runs of the first are also runs of the second, with the proviso that the putative subprotocol may involve roles and actions that are absent in the putative superprotocol. To capture the above, we need to relate protocols to models. Our approach generates a model from the putative subprotocol and then verifies (using suitable mapping) whether the putative superprotocol relates correctly with the subprotocol in that model, that is, whether the runs of the two protocols relate as explained above. Definition 4.2 specifies such a model.

Definition 4.2 Proton Model. Let $P = \langle \mathcal{R}, \mathcal{M}, \mathcal{C}, \mathcal{A}, \mathcal{S}, \mathcal{G} \rangle$ be a protocol. Then, the *Proton model for P* is $\mathcal{I} = \langle \Sigma, P, PV, L^i, Act^i, AP^i, t^i, G, G_0, F \rangle$ where (i) $\Sigma = \mathcal{R} \cup \{e\}$, with e being the environment, (ii) $P = \mathcal{A}$, (iii) $\forall p \in \mathcal{A} : PV(p) = \{s | p \in s\}$, (iv) $\forall i \in \mathcal{R} : L^i = \{l\}$ and $L^e = \prod_{m_i \in \mathcal{M}} m_i \times \prod_{C_i \in \mathcal{C}} C_i.status$, (v) $\forall i \in \mathcal{R} : Act^i = \{m | m.snd = i\} \cup \{nop\}$, and $Act^e = \{sched = r | r \in \mathcal{R}\}$. (vi) $\forall i \in \mathcal{R}, \forall s \in \mathcal{S} : AP^i(s) = \{m | sched = i \wedge m.snd = i \wedge s \in m.guard\}$, (vii) $\forall i \in \mathcal{R} : t^i(l) = l$, and $t^e = \prod_{m_i \in \mathcal{M}} t^{m_i} \times \prod_{C_i \in \mathcal{C}} t^{C_i}$, (viii) G is the set of all states reachable from G_0 by transition function T in \mathcal{I} , (ix) $G_0 = \emptyset$, and (x) $F = \{C_i.status \neq detached | C_i \in \mathcal{C}\}$.

Where \times is binary cross-product and \prod is set cross-product. The protocol's state is the cross-product of the state of each message and commitment. Since both messages and commitments involve multiple roles, each role has just a single state l and all state is in the environment (iv). Proton supports interleaved rather than concurrent actions with the environment scheduling one role at each step (v). Every role can perform the *nop* action (no-operation) at every step (v). t^{m_i} is the transition function that tracks the past occurrence of message m_i , and t^{C_i} is the transition function that tracks the commitment state of commitment C_i as defined by Figure 3 (vii).

Through a slight abuse of notation, for simplicity, we treat guards and actions as expres-

sions in the following.

Definition 4.3 Enactment. Let $P = \langle \mathcal{R}, \mathcal{M}, \mathcal{C}, \mathcal{A}, \mathcal{S}, \mathcal{G} \rangle$ be a protocol and \mathcal{I} be its Proton model. Then, an alternating sequence of states and messages $\langle h_0, m_1, h_1, m_2, h_2, \dots \rangle$ is a *message enactment of P* if and only if $h_0 = \emptyset$ and $(\forall j \geq 0 : h_j \in \mathcal{S}, m_{j+1} \in \mathcal{M} : \mathcal{I}, h_j \models m_{j+1}.\text{guard}$ and $\mathcal{I}, h_{j+1} \models m_{j+1}.\text{actexp}$).

A message enactment yields one or more runs with different interleavings of each message's actions. We define a function μ that maps each index in the message enactment to the index in the run where the corresponding message expression $m_j.\text{actexp}$ becomes true. Each message expression occurs in the same order in every run, and becomes true precisely at the state where its execution completes.

Definition 4.4 Run. Let $P = \langle \mathcal{R}, \mathcal{M}, \mathcal{C}, \mathcal{A}, \mathcal{S}, \mathcal{G} \rangle$ be a protocol and \mathcal{I} be its Proton model. Then, an alternating sequence of states and actions $\langle s_0, a_1, s_1, a_2, s_2, \dots \rangle$ is a *run of P* if and only if $s_0 = \emptyset$ and $(\forall j \geq 0 : s_j \in \mathcal{S}, a_{j+1} \in \mathcal{A} : \mathcal{I}, s_j \models a_{j+1}.\text{guard}$ and $\mathcal{I}, s_{j+1} \models a_{j+1}.\text{actexp}$).

We say a run is *well defined* to emphasize that it satisfies the guard and action expression conditions above: that it is more than just an alternating sequence of states and actions. The empty run $\langle \emptyset \rangle$ is always well defined, since no agent is required to perform any action.

Definition 4.5 Generated Runs. Let $P = \langle \mathcal{R}, \mathcal{M}, \mathcal{C}, \mathcal{A}, \mathcal{S}, \mathcal{G} \rangle$ be a Proton protocol and \mathcal{I}_P its model. Then, a run $\pi = \langle s_0, a_1, s_1, \dots \rangle$ is *generated by P* if and only if there exists a message enactment $\langle h_0, m_1, h_1, \dots \rangle$, and there exists a strictly increasing function on the natural numbers $\mu : \mathbb{N} \mapsto \mathbb{N}$ such that $(\forall j \geq 0 : \mathcal{I}, s_{\mu(j)} \models m_{j+1}.\text{guard}$ and $\mathcal{I}, s_{\mu(j+1)} \models m_{j+1}.\text{actexp}$).

We write $\mathbf{runs}(P)$ for the set of all runs generated by protocol P in \mathcal{I}_P .

4.2 Protocol Refinement

Definition 4.6 Mapping. M maps protocol $P = \langle \mathcal{R}_P, \mathcal{M}_P, \mathcal{C}_P, \mathcal{A}_P, \mathcal{S}_P, \mathcal{G}_P \rangle$ to protocol $Q = \langle \mathcal{R}_Q, \mathcal{M}_Q, \mathcal{C}_Q, \mathcal{A}_Q, \mathcal{S}_Q, \mathcal{G}_Q \rangle$ if and only if $M = \langle M_R, M_P, M_C \rangle$, where

$$\begin{aligned} M_R &= \{ \langle r, R \rangle \mid r \in \mathcal{R}_P, R \subseteq \mathcal{R}_Q \} \\ M_P &= \{ \langle p, e \rangle \mid p \in \mathcal{A}_P, e \subseteq 2^{\mathcal{A}_Q} \} \\ M_C &= \{ \langle C, \oplus_i C_i \rangle \mid C \in \mathcal{C}_P, C_i \in \mathcal{C}_Q, C \leq_M \oplus_i C_i \} \end{aligned}$$

Informally, a run π_Q *embeds* a run π_P if all of π_P lies within π_Q . In effect, π_Q does everything that π_P does, and possibly more: as Mallya and Singh [2007] propose, a protocol Q refines a protocol P if and only if every run of Q embeds some run of P . This captures the intuition that any computation (run) allowed by Q is allowed by P as well.

Consider the mapping from *Pay* to *OrderPayShip*. In protocol *Pay*, *promiseMsg* means $\{\text{CREATE}(C_{\text{pay}})\}$ and *payMsg* means $\{\text{pay}\}$. In protocol *OrderPayShip*, *orderMsg* means $\{\text{CREATE}(C_{\text{pay}})\}$ and *payMsg* means $\{\text{pay}\}$. Therefore, *promiseMsg* and *payMsg* in *Pay* mean the same, respectively, as *orderMsg* and *payMsg* in *OrderPayShip*.

$$\begin{aligned} \text{Pay} &\mapsto \text{OrderPayShip} \\ \text{promiseMsg} &\mapsto \text{orderMsg} \\ \text{payMsg} &\mapsto \text{payMsg} \end{aligned}$$

Pay has two message enactments: $\langle \rangle$ and $\langle \text{promiseMsg}, \text{payMsg} \rangle$. *OrderPayShip* has five message enactments, which embed *Pay*'s runs as follows.

$$\begin{aligned} \langle \rangle & : \langle \rangle \\ \langle \rangle & : \langle \text{reqQuoteMsg} \rangle \\ \langle \rangle & : \langle \text{reqQuoteMsg}, \text{sendQuoteMsg} \rangle \\ \langle \text{promiseMsg}, \text{payMsg} \rangle & : \langle \text{reqQuoteMsg}, \text{sendQuoteMsg}, \text{orderMsg}, \text{payMsg}, \text{shipMsg} \rangle \\ \langle \text{promiseMsg}, \text{payMsg} \rangle & : \langle \text{reqQuoteMsg}, \text{sendQuoteMsg}, \text{orderMsg}, \text{shipMsg}, \text{payMsg} \rangle \end{aligned}$$

We define a *mapped run* where each sub-state s is enriched to a state $M(s)$ by including values for all super-propositions and super-commitments. We now compare enriched sub-states $M(s)$ in mapped sub-runs with super-states in super-runs. Below, we write $\text{exp}\langle x \mapsto y \rangle$ to mean the expression resulting from the uniform substitution of symbol x by expression y in exp .

Definition 4.7 Mapped Run. Let $\pi = \langle s_0, a_1, s_1, \dots \rangle$ be a run and $M = \langle M_R, M_P, M_C \rangle$ be a protocol mapping. Then the *M-map of π* , $M(\pi) = \langle M(s_0), a_1, M(s_1), \dots \rangle$ is a run where for all s , $M(s) \supseteq s$ and $M(s)$ is the minimal set for which the following conditions hold:

- (Propositions) if $\langle m, E \rangle \in M_P$ and $s \models E$, then $m \in M(s)$.
- (Commitments) if $\langle C, \bigoplus_i C_i \rangle \in M_C$ and $\forall i : s \models C_i.\text{status}$, then $C.\text{status} \in M(s)$ where $C.\text{status} = \min(C_i.\text{status})$.

Continuing with the above discussion, we map each sub-run and verify that it embeds some super-run. The following definition captures the intuition that the embedding sub-run steps through each of the states of the embedded super-run, but may potentially include additional states. We ignore the transitions in each run.

To simplify the notation, we also introduce a projected mapping function $\widehat{M}(q) = M(q) \cap \mathcal{A}_P$ that is the set of just the propositions and states in a (super-)protocol P .

Definition 4.8 Embedding. Let P and Q be two protocols. A run $\pi_Q = \langle q_0, \cdot, q_1, \dots \rangle \in \mathbf{runs}(Q)$ embeds a run $\pi_P = \langle p_0, \cdot, p_1, \dots \rangle \in \mathbf{runs}(P)$, written $\mathbf{emb}(\pi_Q, \pi_P)$, if and only if there exists a strictly increasing function on natural numbers $\tau : \mathbb{N} \mapsto \mathbb{N}$ such that $(\forall i : 0 \leq i \leq |\pi_P| : p_i = \widehat{M}(q_{\tau(i)})$ and $(\forall j : \tau(i) \leq j < \tau(i+1) : \widehat{M}(q_{\tau(i)}) = \widehat{M}(q_j))$, where $\widehat{M}(q) = M(q) \cap \mathcal{A}_P$.

Function τ maps from indices of π_P to indices of π_Q , and the conditions ensure every time $\widehat{M}(q_j)$ changes, the new value matches the next p_i .

Now we can define refinement in purely semantic terms that capture our intuition that each mapped sub-run must embed some super-run. Notice that this definition implicitly uses Proton models \mathcal{I}_P and \mathcal{I}_Q respectively for P and Q .

Definition 4.9 Refinement. Let P and Q be two protocols, and M a mapping from P to Q . Then Q refines P under M if and only if $(\forall \pi_Q \in \mathbf{runs}(Q) : (\exists \pi_P \in \mathbf{runs}(P) : \mathbf{emb}(M(\pi_Q), \pi_P)))$.

5. VERIFYING PROTOCOL REFINEMENT

Figure 5 shows the high-level process flow for verifying protocol refinement. The Proton preprocessor reads the subprotocol, superprotocol, and mapping specifications and con-

structs (as Section 6 details) the input for the MCMAS model checker in the Interpreted Systems Programming Language (ISPL) [Lomuscio et al. 2009].




Fig. 5. Proton process flow. Activities are shown in boxes and specification documents are not. Transformations and comparisons are numbered 1 to 4 in sequence.

The input to MCMAS is a set of guarded statements for each role. MCMAS internally generates a state transition system such as that shown in Figure 6. The system starts in initial state s_0 . Action *requestQuote* transitions to state s_1 , action *sendQuote* transitions to state s_2 , and so on. There is an edge for every action enabled in a state.




Fig. 6. A schematic of some possible enactments of the *OrderPayShip* protocol.

The Proton preprocessor generates an interpreted system model for the subprotocol. There is one ISPL agent definition for each sub-role, and the state of all sub-elements (propositions and commitments) are expressed as model state variables. The model checker simulates the subprotocol’s actions. Because each super-element is mapped to an expression of sub-elements, the state of every super-element can be inferred from the subprotocol’s state. As Section 5.5 shows, protocol refinement conditions are expressed as CTL formulae. If all these CTL formulae are true, the subprotocol refines the superprotocol.

5.1 Intuition: Decomposition

A message can mean multiple things. To better understand and characterize a message, we decompose each *message* into its meaning as a set of primitive, well-defined *actions*. The meaning of a message is then the conjunction of all its constituent actions.

An action is either a Boolean proposition or a commitment operation. A propositional action sets the value of the proposition to true. We do not support setting propositions to false. Commitment actions are the operations CREATE, TRANSFER, RELEASE, and CANCEL that change the state of a commitment.

We replace all message terms with a conjunction of their actions throughout a protocol, *decomposing* protocol messages, converting from a “protocol of guarded messages” to a “protocol of guarded actions.” Each *msg* term is replaced by a conjunction of its actions in

both the guard and the action expression of every guarded statement.

$$\mathbf{means}([guard] msg \mathbf{means}\{act_i\}) \Rightarrow [guard \langle\langle msg \mapsto \bigwedge_i act_i \rangle\rangle] \bigwedge_i act_i$$

5.2 Intuition: Diffusion

The result of decomposition is a set of guarded action expressions, but the model checker executes actions, not action expressions. Therefore, each guarded action expression must be converted to an equivalent set of guarded actions. However, computing the equivalent guarded actions is nontrivial. For example, consider the guarded action expression

$$[guard] payP \wedge payM$$

A naïve approach would be to apply the guard to each action separately: $[guard] payP$ and $[guard] payM$. Doing so would be overly restrictive because neither $payP$ nor $payM$ can occur before the guard becomes true. Greater flexibility is needed.

Given a guarded conjunction of actions $[guard] a_1 \wedge a_2 \wedge \dots \wedge a_n$, the action expression becomes true when the *last* (in time) of the a_i s becomes true. Given a guarded disjunction of actions $[guard] a_1 \vee a_2 \vee \dots \vee a_n$, the action expression becomes true when the *first* (in time) of the a_i s becomes true. We minimally constrain when the a_i become true so the overall expression becomes true at exactly the same point, relative to the other actions, in both the super-runs and the sub-runs. For conjunctions, all the a_i , except the *last*, can move to any *earlier* point in time. For disjunctions, all the a_i , except the *first*, can move to any *later* point in time. The a_i can even move all the way to the run's beginning (conjunction) or end (disjunction). Decomposition (Section 5.1) generates guarded action expressions with conjunctions; abstraction mappings (Section 5.3) can generate guarded action expressions with both conjunctions and disjunctions.

The recursive *diffusion* function **dif** transforms a guarded action expression to a set of guarded actions.

$$\mathbf{dif}([guard] \bigvee_i exp_i) \Rightarrow \{\mathbf{dif}([guard] exp_i)\} \quad (9)$$

$$\mathbf{dif}([guard] \bigwedge_i exp_i) \Rightarrow \{\mathbf{dif}([guard \vee \bigvee_{j \neq i} \neg exp_j] exp_i)\} \quad (10)$$

$$\mathbf{dif}([guard] act) \Rightarrow [guard] act \quad (11)$$

where $guard$ and exp_i are Boolean expressions of actions. Diffusion transforms the guarded expression example above to

$$\begin{aligned} & [guard \vee \neg payM] payP \\ & [guard \vee \neg payP] payM \end{aligned}$$

Both actions can still occur after the guard becomes true, so it allows at least all of the runs allowed in the naïve approach. Additionally, the first action to fire can fire at any time. If $payP$ fires before $payM$, then $\neg payM$ is true when $payP$ fires, so $payP$ can fire at any time. Diffusion thus covers possibilities that the naïve approach omits.

Diffusion can generate multiple guarded statements for the same action, but MCMAS requires a single guarded statement for each action. Therefore, we introduce the *collection* function **col** that converts a set of guarded statements back to canonical form, where

there is a single guarded statement for each action. Consider each individual, input action. Here, **col** collects potentially multiple guarded statements for that action in its input, and generates a single guarded statement for that action in its output. The output guard for the action is the disjunction of all the input guards. If an action appears in only one guarded statement in the input, that guarded statement appears unmodified in the output.

$$\mathbf{col}(\{[guard_i] act_i\}) \Rightarrow \{[\bigvee_j guard_j] act_i \mid act_j = act_i\}$$

Consider a partial protocol containing these two guarded statements. Since the messages overlap on action *ship*, condition *freeCoupon* ensures the no overlap constraint of Definition 4.1.

$$\begin{aligned} & [orderMsg \wedge \neg freeCoupon] paidShipMsg \text{ means } \{bill, ship\}; \\ & [orderMsg \wedge freeCoupon] freeShipMsg \text{ means } \{ship\}; \end{aligned}$$

Diffusion transforms those to the following three guarded statements.

$$\begin{aligned} & [(orderMsg \wedge \neg freeCoupon) \vee \neg ship] bill \\ & [(orderMsg \wedge \neg freeCoupon) \vee \neg bill] ship \\ & [orderMsg \wedge freeCoupon] ship \end{aligned}$$

Collection merges the two statements for *ship*, giving these two, final guarded statements.

$$\begin{aligned} & [(orderMsg \wedge \neg freeCoupon) \vee \neg ship] bill \\ & [(orderMsg \vee \neg bill] ship \end{aligned}$$

Figure 7(a) schematically shows how the foregoing developments of decomposition, diffusion, collection, and embedding combine together to check whether one protocol refines another. However, by themselves, they do not address the fact that different protocols can be written at different layers of abstraction.

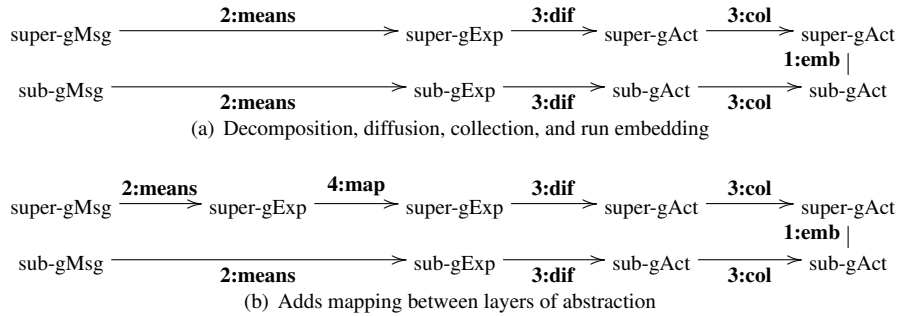


Fig. 7. Protocol refinement defined as transformations (horizontal lines) and comparisons (vertical lines) between protocols of guarded messages (gMsg), guarded action expressions (gExp), and guarded actions (gAct). In each subfigure, the top row refers to a superprotocol and the bottom row to a subprotocol.

5.3 Accommodating Abstraction Mapping

Since superprotocols represent higher-level abstractions than subprotocols, comparing protocols must address differences in levels of abstraction. There is often no one-to-one correspondence between super-elements and sub-elements. Protocol elements (roles, propositions, and commitments) must be mapped between the two protocols to compare them.

An important type of abstraction difference is the introduction of intermediaries or middlemen in lower-level abstractions. Whereas two super-roles may communicate directly with each other using a single message in a high-level protocol, there is a natural tendency for message communication to pass through multiple, intermediary sub-roles as that protocol is refined to lower-level abstractions. Protocol refinement must allow super-elements to span intermediaries. One super-proposition could map to an expression of multiple sub-propositions, each controlled by different sub-roles (intermediaries), and one super-commitment could be fulfilled through multiple sub-commitments and their intermediate sub-debtors.

We say one protocol refines another protocol *under a given mapping*, because mapping functions are an essential element for protocol refinement, and must be an explicit input. A subprotocol might refine a superprotocol under one mapping, but not under a different mapping. Our approach does not determine whether it is impossible for one protocol to refine another protocol under any possible mapping.

A mapping expresses how terms in a putative superprotocol map to expressions in a putative subprotocol. The *mapping* function **map** converts guarded action expressions written with high-level abstractions x_i in a putative superprotocol, to expressions e_i of low-level terms in a putative subprotocol. (Below, $\langle\langle x_i \mapsto e_i \rangle\rangle$ is a mapping assertion.)

$$\mathbf{map}([guard] \exp) \Rightarrow [guard\langle\langle x_i \mapsto e_i \rangle\rangle] \exp\langle\langle x_i \mapsto e_i \rangle\rangle$$

5.4 Verifying Refinement: Summary

Figure 7(b) schematically shows the transformations and comparison required to demonstrate protocol refinement. In both subfigures, horizontal lines show the transformations of a single protocol: decomposition (**means**), mapping (**map**), diffusion (**dif**), and collection (**col**). Vertical lines show the comparison between two protocols: run embedding (**emb**). The nodes in the figure show how guarded messages (gMsg) are transformed first to guarded action expressions (gExp), and then to guarded actions (gAct). In each subfigure, the top row refers to a superprotocol and the bottom row refers to a subprotocol.

5.5 Generating CTL Formulae for Verification

MCMAS checks whether an interpreted system model satisfies specified CTL formulae. In this section, we describe how Proton expresses conditions for commitment resolution, overlapping messages, serial composition, and commitment covering as CTL formulae. All such formulae must be satisfied for protocol refinement to hold.

5.5.1 Verify Run Embedding by Checking Guards. Protocol comparison is fundamentally based on run embedding. Run embedding means, at every state, if the subprotocol can perform an action then the superprotocol must also be able to perform that action. That is, when an action's sub-guard is true, its super-guard must also be true. Since run embedding ignores actions not in the superprotocol, Proton generates CTL formulae for all actions that

result from mapping all super-actions ($\forall a \in M(\mathcal{A}_{super})$):

$$\mathbf{AG}(a.sub\text{-guard} \rightarrow a.super\text{-guard}) \quad (12)$$

5.5.2 *Verify that Messages Do Not Overlap.* So that every action a in a sub-run can be uniquely associated with a message, we verify the no overlap constraint of Definition 4.1. For every pair of guarded statements gs_1 and gs_2 that share a common action meaning a ,

$$\mathbf{AG}(\neg(gs_1.guard \wedge gs_2.guard)) \quad (13)$$

5.5.3 *Verify that Detached Commitments Eventually Resolve.* We require each detached commitment must eventually resolve in every correct protocol enactment. We employ model checker fairness constraints (expressions that must be true infinitely often on any run) to eliminate sub-runs in which the sub-roles fail to act properly and resolve their detached commitments. Doing so restricts our verification to correct enactments of the given protocols, thus avoiding false negatives due to incorrect enactments.

$$\mathbf{Fairness} \ C_{sub.status} \neq \text{detached} \quad (14)$$

The states of super-commitments can be inferred from the states of sub-commitments.

5.5.4 *Verify Commitment Covering.* The truth or falsity of a statement in an unreachable state has no bearing on the enactment of a protocol, so we can replace $a \models b$ statements by the CTL formula $\mathbf{AG}(a \rightarrow b)$. Doing so enables us to use the model checker to verify commitment covering, which would otherwise need to be handled separately, as indeed it was in a previous version of Proton.

Verifying one commitment covers another under map M , $C_W \leq_M C_S$, is done in two parts. First, the preprocessor verifies the debtor and creditor conditions (Equations 5–6). Second, the model checker verifies the antecedent and consequent conditions (Equations 7–8) hold in all (reachable) states with the CTL formulae

$$\mathbf{AG}(M(C_W.ant) \rightarrow C_S.ant) \quad (15)$$

$$\mathbf{AG}(C_S.csq \rightarrow M(C_W.csq)) \quad (16)$$

5.5.5 *Verify Serial Compositions are Well Defined.* For serial compositions $C_{\oplus} = C_A \oplus C_B$, the model checker verifies the well-definedness condition holds in all (reachable) states on all paths with the CTL formula

$$\mathbf{AG}(C_A.csq \rightarrow C_B.ant) \quad (17)$$

6. TOOLING, DETAILED EXAMPLES, AND EXPERIMENTAL RESULTS

In this section, we pull together the many elements: commitments, serial composition of commitments, and commitment covering; the example payment and order protocols, and various mappings between them; and the formal definitions. We concretely demonstrate how *PayViaMM* refines, or fails to refine, *Pay* under various mappings.

Proton verifies protocol refinement using the process flow as shown in Figure 5 and the pseudocode for **refines**(super, map, sub) shown in Listing 7. The inputs P and Q are protocols, which in our syntax are in terms of guarded messages. The first lines of the algorithm transform these into protocols expressed in terms of guarded actions. Proton generates an interpreted system model from the guarded actions of the subprotocol. There is one MC-MAS agent definition for each sub-role, and the state of the sub-elements (propositions and

Listing 7 Calculate $\text{refines}(P, M, Q)$


```

1:  $Q_{gMsg} = Q$                                 ▷ Input Q is a protocol of guarded messages
2:  $P_{gMsg} = P$                                 ▷ Input P is a protocol of guarded messages
3:  $Q_{gAct} = \text{col}(\text{dif}(\text{means}(Q_{gMsg})))$         ▷ protocol of guarded sub-actions
4:  $P_{gAct} = \text{col}(\text{dif}(\text{map}_M(\text{means}(P_{gMsg}))))$     ▷ protocol of guarded super-actions
5:  $\text{model} = \text{genModel}(Q_{gAct})$                 ▷ generate ISPL model
6: for all  $act_P \in P_{gAct}.\text{actions}$  do          ▷ For all super-actions
7:    $\text{genCTL}(\text{AG}(act_P.\text{sub-guard} \rightarrow act_P.\text{super-guard}))$ 
8: end for
9: for all  $C_Q \in Q_{gAct}.\mathcal{C}$  do                ▷ For all sub-commitments
10:   $\text{genFairness}(C_Q.\text{status} \neq \text{detached})$ 
11: end for
12: for all  $C_P \in P_{gAct}.\mathcal{C}$  do                ▷ For all super-commitments
13:   $C_Q = C_P.\text{coveringCommitment}$ 
14:   $\text{genCTL}(\text{AG}(M(C_P.\text{ant}) \rightarrow C_Q.\text{ant}))$ 
15:   $\text{genCTL}(\text{AG}(C_Q.\text{csq} \rightarrow M(C_P.\text{csq})))$ 
16: end for
17: for all  $C_A \oplus C_B$  do                    ▷ For all serial compositions
18:   $\text{genCTL}(\text{AG}(C_A.\text{csq} \rightarrow C_B.\text{ant}))$ 
19: end for
20: for all overlapping guarded statement pairs  $gs_1$  and  $gs_2$  in P and Q do
21:   $\text{genCTL}(\text{AG}(\neg(gs_1.\text{guard} \wedge gs_2.\text{guard})))$ 
22: end for
23:  $\text{ctl} = \text{all generated CTL formulae}$ 
24: return  $\text{MCMAS}(\text{model}, \text{ctl})$             ▷ Are all CTL formulae satisfiable?

```

commitments) are MCMAS state variables or MCMAS evaluation expressions. The MCMAS model checker then simulates the subprotocol's actions. Because each super-element is mapped to an expression of sub-elements, the superprotocol's state can be inferred from the subprotocol's state. Refinement requires the model of the subprotocol to satisfy a set of CTL formulae. If all formulae are true, the subprotocol refines the superprotocol.

$$C_{\text{pay}} \mapsto C_{\text{payP}} \oplus C_{\text{payM}} \quad \text{if } C_{\text{pay}} \leq_{M_1} C_{\text{payP}} \oplus C_{\text{payM}}$$



We now check whether PayViaMM refines Pay under map M_1 . Using the commitment diagrams from Section 2.6, this diagram demonstrates commitment C_{pay} from Pay is covered by the serial composition of C_{payP} and C_{payM} from PayViaMM under mapping M_1 in Listing 3. The bottom-left arrow states that if promise becomes true, Payer commits to making payP true. The bottom-right arrow states that if payP becomes true, MiddleMan

commits to making pay_M true. In the serial composition (middle arrow), if promise becomes true, then Payer and MiddleMan (severally) commit to making both pay_P and pay_M true. The well-definedness condition ensures that the discharge of the first commitment entails the antecedent of the second commitment, thus detaching it. The Payee and Payer are creditors of the input commitments, and are thus creditors of the serial composition. The serial composition covers the commitment in Pay (top arrow).

Proton generates the following six CTL formulae to verify PayViaMM refines Pay under M_1 . Equation 12, which verifies whether sub-guards imply super-guards, for actions promise , pay_P , and pay_M , generates Equations 18–20, respectively. Equation 21 verifies that $C_{\text{pay}_P} \oplus C_{\text{pay}_M}$ is valid. Equations 7–8, which verify the antecedent and consequent conditions of $C_{\text{pay}_P} \oplus C_{\text{pay}_M}$ covers C_{pay} , generates Equation 22–23, respectively (the debtor and creditor conditions in Equations 5–6 are checked directly by the Proton preprocessor, not by MCMAS).

$$\mathbf{AG}(\top \rightarrow \top) \quad (18)$$

$$\mathbf{AG}(\text{promise} \wedge \text{CREATE}(C_{\text{pay}_P}) \wedge \text{CREATE}(C_{\text{pay}_M}) \rightarrow (\text{promise} \wedge \text{CREATE}(C_{\text{pay}_P}) \wedge \text{CREATE}(C_{\text{pay}_M})) \vee \neg \text{pay}_M) \quad (19)$$

$$\mathbf{AG}(\top \rightarrow (\text{promise} \wedge \text{CREATE}(C_{\text{pay}_P}) \wedge \text{CREATE}(C_{\text{pay}_M})) \vee \neg \text{pay}_P) \quad (20)$$

$$\mathbf{AG}(\text{pay}_P \rightarrow \text{pay}_P) \quad (21)$$

$$\mathbf{AG}(\text{promise} \rightarrow \text{promise}) \quad (22)$$

$$\mathbf{AG}(\text{pay}_P \wedge \text{pay}_M \rightarrow \text{pay}_P \wedge \text{pay}_M) \quad (23)$$

All of the above formulae are obviously true, except Equation 20, which can be rewritten $\mathbf{AG}(\text{pay}_P \rightarrow \text{promise} \wedge \text{CREATE}(C_{\text{pay}_P}) \wedge \text{CREATE}(C_{\text{pay}_M}))$. It is true because the progression of the model, controlled by message guards, ensures pay_P becomes true only after promiseMsg . MCMAS verifies each generated CTL formula holds in the model, so PayViaMM refines Pay under map M_1 .

Proton generates exactly the same input to MCMAS when checking whether PayViaMM refines Pay under map M_2 or map M_3 , because the subprotocol models are derived from exactly the same PayViaMM protocol, the superprotocol Pay contains exactly the same propositional and commitment super-elements, and exactly the same CTL conditions must be checked. Therefore, PayViaMM refines Pay under both maps M_2 and M_3 .

$$\mathbf{AG}(\top \rightarrow \top) \quad (24)$$

$$\mathbf{AG}(\text{promise} \wedge \text{CREATE}(C_{\text{pay}_P}) \wedge \text{CREATE}(C_{\text{pay}_M}) \rightarrow (\text{promise} \wedge \text{CREATE}(C_{\text{pay}_P}) \wedge \text{CREATE}(C_{\text{pay}_M})) \vee \neg \text{pay}_M) \quad (25)$$

$$\mathbf{AG}(\top \rightarrow (\text{promise} \wedge \text{CREATE}(C_{\text{pay}_P}) \wedge \text{CREATE}(C_{\text{pay}_M})) \vee \neg \text{pay}_P) \quad (26)$$

$$\mathbf{AG}(\text{pay}_M \rightarrow \text{promise}) \quad (27)$$

$$\mathbf{AG}(\text{promise} \rightarrow \text{pay}_P) \quad (28)$$

$$\mathbf{AG}(\text{promise} \wedge \text{pay}_P \wedge \text{pay}_M \rightarrow \text{pay}_P \wedge \text{pay}_M) \quad (29)$$

Proton correctly reports failures. Proton generates these CTL formulae when checking whether PayViaMM refines Pay under mapping B_1 in Listing 6. Recall that B_1 maps the super-commitment to a serial composition in the wrong order. Equations 24–26 are the same as Equations 18–20, and all hold in the model. Equations 29 obviously holds. But

<i>Superprotocol</i>	<i>Subprotocol</i>	<i>Map</i>	<i>refines</i>	\oplus	<i>cover</i>	<i>formulae</i>	<i>time</i>
<i>Pay</i>	<i>PayViaSpouse</i>	M_1	Yes	0	1	4	0.7
<i>Pay</i>	<i>FullPay</i>	M_1	Yes	0	1	4	0.8
<i>Pay</i>	<i>PayViaMM</i>	M_1	Yes	1	1	6	0.7
<i>Pay</i>	<i>PayViaMM</i>	M_2	Yes	1	1	6	0.7
<i>Pay</i>	<i>PayViaMM</i>	M_3	Yes	1	1	6	0.6
<i>Pay</i>	<i>PayViaMM</i>	B_1	No	1	1	6	0.7
<i>Pay</i>	<i>PayViaCheck</i>	M_1	Yes	2	1	7	1.1
<i>Pay</i>	<i>PayViaCredit</i>	M_1	Yes	2	1	9	3.6
<i>Pay</i>	<i>OrderPayShip</i>	M_1	Yes	1	1	5	1.0
<i>Pay</i>	<i>OrderPayViaMMShip</i>	M_1	Yes	2	1	7	1.2
<i>PayViaMM</i>	<i>PayViaCheck</i>	M_1	Yes	3	2	12	1.7
<i>PayViaMM</i>	<i>PayViaCredit</i>	M_1	Yes	3	2	12	1.7
<i>PayViaMM</i>	<i>OrderPayViaMMShip</i>	M_1	Yes	0	2	7	1.2
<i>OrderPayShip</i>	<i>OrderPayViaMMShip</i>	M_1	Yes	1	2	11	1.4
<i>OrderPayShip</i>	<i>NetBill₂</i>	M_1	Yes	0	2	9	1.2
<i>OrderPayViaMMShip</i>	<i>NetBill₃</i>	M_1	Yes	0	3	12	1.8

Table I. Information about each demonstrated Refinement. The columns are: the name of the superprotocol; the name of the subprotocol; the name of the map; whether subprotocol refines superprotocol under map; \oplus is the number of serial compositions; *covers* is the number of commitment covering checks; *formulae* is the number of CTL formulae verified by the model checker; and *time* is the elapsed time (in seconds) for refinement verification.

Equation 27 does not hold because pay_M has a true guard in Listing 2, so the Middleman can send pay_M at any time, even before *promise*. That leaves Equation 28, which comes from the antecedent of C_{pay} which is *promise* and the antecedent of $C_{payM} \oplus C_{payP}$ which is pay_P . In the states between the Payer promising and the Payer actually paying, the formula does not hold, meaning C_{pay} can become detached without $C_{payM} \oplus C_{payP}$ also becoming detached. The result is C_{pay} is not covered by $C_{payM} \oplus C_{payP}$. MCMAS correctly reports these two formulae are false in the model, and *PayViaMM* does *not* refine *Pay* under mapping B_1 .

As in Mallya and Singh’s [2007] approach, an interesting consequence of our treatment of refinement is that aggregation functions like refinement. For example, consider a protocol *OrderPayShip*. Because all enactments of *OrderPayShip* necessarily include an enactment of *Pay*, *OrderPayShip* is naturally a refinement of *Pay*. The foregoing coheres with the notion of subtype in object-oriented programming.

Proton verified all the refinements shown in Figure 1 and Table I. The first three columns are the superprotocol, subprotocol, and mapping, respectively. *OrderPayShip* is identical to the first NetBill scenario described by Mallya and Singh [2007]. NetBill₂ and NetBill₃ are scenarios 2 and 3 in the same paper. *PayBySpouse* is a new, simple, payment protocol where one person promises and then his or her spouse pays. *FullPay* is similar to *Pay*, but exercises all the commitment operations: CREATE, TRANSFER, RELEASE, and CANCEL.

Each superprotocol-subprotocol pair has a unique set of mapping names, so M_1 for *Pay* and *PayViaMM* is a different mapping than M_1 for *Pay* and *PayViaCheck*. Mappings M_1 , M_2 , M_3 , and B_1 for *Pay* and *PayViaMM* are shown, respectively, in Listings 3, 4, 5, and 6. The *refines* column is whether subprotocol refines superprotocol under map; the \oplus column is the number of serial compositions in the mapping; the *covers* column is the number of commitment covering checks; the *formulae* column is the number of CTL formulae verified by the model checker; and the *time* column is the elapsed time in seconds for

the complete refinement verification, including Proton preprocessing and MCMAS model checking.

7. CORRECTNESS OF OUR REFINEMENT VERIFICATION METHOD

We now establish the correctness of the verification method of Section 5 with respect to the formal definitions of Section 4. Theorem 7.1 is our main correctness result: it ties the algorithm of Listing 7 with Definition 4.9. An interesting subtlety is that whereas Definition 4.9 “maps up,” enriching the states in sub-runs, the algorithm “maps down,” expanding expressions and ignoring elements not in \mathcal{A}_P .

THEOREM 7.1. *Let P and Q be two protocols, and let M be a mapping from P to Q . Then, $\mathbf{refines}(P, M, Q)$ returns true if and only if Q refines P under M .*

PROOF. Let \mathcal{I}_Q be the Proton model generated from Q . Define mappings $M_Q = \mathbf{means}(Q)$ and $M_P = M(\mathbf{means}(P))$.

\Rightarrow Assume $\mathbf{refines}(P, M, Q)$. Then all of the CTL formulae in Listing 7 are true. Because of Line 7, all sub-guards imply their super-guards. Because of the fairness condition at Line 10, all detached commitments eventually resolve. Because of Lines 14–15, all commitment coverings are valid. Because of Line 18, all serial compositions are valid. Because of Line 21, no guarded statement pairs overlap.

Let π_Q be a run in $\mathbf{runs}(Q)$. By Theorem A.7, there is a run $\pi'_Q \in \mathbf{runs}(\mathbf{col}(\mathbf{dif}(M_Q(Q))))$. By Theorem A.2, there exists a run $\pi'_P \in \mathbf{runs}(\mathbf{col}(\mathbf{dif}(M_P(P))))$. And, by Theorem A.7, there exists a run $\pi_P \in \mathbf{runs}(P)$. Therefore, for every π_Q there is a π_P and Q refines P under M by Definition 4.9.

\Leftarrow Assume Q refines P under M . For any $\pi_Q \in \mathbf{runs}(Q)$ there exists a run $\pi_P \in \mathbf{runs}(P)$ such that $\mathbf{emb}(M(\pi_Q), \pi_P)$ by Definition 4.9. By Theorem A.7, there exists a $\pi'_Q \in \mathbf{runs}(\mathbf{col}(\mathbf{dif}(M_Q(Q))))$ and a $\pi'_P \in \mathbf{runs}(\mathbf{col}(\mathbf{dif}(M_P(P))))$. By Theorem A.2, $\mathbf{emb}(M(\pi_Q), \pi_P)$ implies $(\forall a_i \in \mathcal{A}_P : \mathcal{I}, g_0 \models \mathbf{AG}(a_i.\mathit{sub-guard} \rightarrow a_i.\mathit{super-guard}))$, which implies the CTL formulae at Line 7 are true.

Because all detached commitments must eventually resolve, the fairness formulae at Line 10 are true. Because all commitment coverings must be valid, the formulae at Lines 14–15 are true. Because all serial compositions must be valid, the formulae at Line 18 are true. Because protocols must be well defined, the formulae at Line 21 are true. Because all of the CTL formulae evaluate to true, $\mathbf{refines}(P, M, Q)$ returns true. \square

8. DISCUSSION

Commitments support finer guard granularity than propositions can. A proposition divides time into two stages: before and after it holds. A commitment divides time into four stages: null, active and conditional, active and detached, and resolved.

$$\mathit{null} \xrightarrow{\text{CREATE}} \mathit{cond} \xrightarrow{\text{ant}} \mathit{detached} \xrightarrow{\text{csq}} \mathit{resolved}$$

Rather than waiting for final resolution, a protocol can make progress sooner if an action’s guard is enabled after one of the first three stages. Commitments increase protocol flexibility, because guards can specify earlier stages.

For example, suppose the Buyer role in *OrderPayShip* decides whether to pay based on the state of proposition *ship*. Since *ship* has only two stages, the role’s decision can only be “all” (*ship* complete) or “nothing” (*ship* not complete). The “all” choice is represented

by the guarded statement $[ship] \text{ pay}$, and the “nothing” choice is represented by $[true] \text{ pay}$. Using commitments, the protocol can guard pay based on any of the four commitment stages. A guard can enable pay as soon as the debtor commits to make ship true.

$$[\text{CREATE}(\text{C}_{\text{Seller,Buyer}}(\text{pay}, \text{ship}))] \text{ pay}$$

Incorporating commitments can improve flexibility over traditional protocol frameworks.

A necessary prerequisite of employing protocols is that the participants of a service engagement agree on the format and meanings of the messages they exchange. Note that such agreement is unavoidable: it is just that in today’s practice the meanings are not expressed clearly and explicitly and any agreements are hardcoded in implementations.

Our definition of protocol refinement does not mean agents that can participate in a superprotocol can necessarily participate unchanged in a subprotocol. In our model, agents may need to be modified to participate in subprotocols. For example, an agent capable of participating in a basic payment protocol needs to handle the additional messages required in paying via check or credit card.

8.1 Relevant Literature

Proton is the first approach for protocol refinement that incorporates mapping super-elements to expressions of sub-elements. Proton supports mapping super-propositions to Boolean expressions of sub-propositions as well as mapping super-commitments to serial compositions of chains of sub-commitments.

Mallya and Singh [2007] propose a definition of protocol refinement (which they call subsumption) that compares the order of state pairs in state runs. For every pair of states in the superprotocol, there must be *some* matching pair of states in the subprotocol with the same order. However, this definition can create false positives when multiple state pairs in the superprotocol each match the same state pair in the subprotocol or when one super-state matches different sub-states. For example, all state pairs in the super-run $\langle 1, 2, 3 \rangle$ have matching state pairs in the sub-run $\langle 2, 1, 3, 2 \rangle$ even though the two runs are very different. Our definition compares runs step-by-step and thereby so avoids the above problems, even if protocol looping is allowed.

Our definition of commitment covering is an extension of *commitment strength* as defined by Chopra and Singh [2009], who identify the basic requirements in Equations 7 and 8. We extend Chopra and Singh’s definition with the role requirements in Equations 5 and 6, and we allow commitments to be at different levels of abstraction by including an abstraction mapping function.

Singh [2008] states rules for commitment chaining similar to those proposed here, but does not directly state a rule for stronger consequents, and does not directly state a rule similar to serial composition. The concrete commitment created by serial composition provides a midpoint in commitment reasoning, and can potentially make the comparison of commitments across protocols more explicit.

When we say a group of debtors are jointly and severally responsible for eventually making the consequent true, we mean this in the sense of Rescher’s [1998] *legal* responsibility where “individual agents are responsible only for their own individual acts.” We do not mean Rescher’s notion of legal responsibility where the group as a whole becomes a legal person, nor his notion of *moral* responsibility where intentions are crucial (intentions are absent from our formulation). In Norman and Reed [2002], group imperatives can be addressed *distributively* (as a list of individuals) or as a *collective*. In both cases, group

imperatives imply more than just a collection of individual imperatives. While joint and several responsibility is similar to distributive responsibility, because only one member of the group is required to act, it is different, because joint and several responsibility is only a summary of individual responsibilities, and does not impose additional responsibilities on roles that are members of a group.

Our work on protocols builds on the fundamental intuition that protocol states can be effectively characterized in terms of the commitments of the participants, and that such characterization can be used as a basis for correct enactments and for further reasoning. The earliest works that developed the above theme include the commitment machines approach of Yolum and Singh [2002] for business protocols and McBurney and Parsons' [2002] framework for sequencing multiple dialogue games, allowing one dialogue game to be embedded inside another partially completed dialogue game.

We do not propose specific, desirable properties of protocols, but others have. Yolum [2007], Singh and Chopra [2010], and El-Menshawry et al. [2010] describe desirable properties of protocols in general and commitment protocols in particular, including fairness, safety, liveness, operability, and transparency.

We use Boolean guards to constrain actions, but other representations are possible. Baldoni et al. [2010] proposed constraints based on regulative specifications. Regulative specifications constrain the execution flow using special-purpose operators on state values. Gabbay [1987] proposes using past-temporal expressions for controlling when actions can occur and future-temporal expressions for controlling which actions must occur in the future. Past-temporal expressions are more expressive than our guard expressions.

8.2 Directions for Future Research

Constructing a mapping function from a superprotocol to a subprotocol can be a challenging task. Advice to guide protocol designers, in the form of a basic mapping methodology, would be a valuable addition to this work. Winikoff [2006; 2007] proposed a methodology for the related task of designing commitment-based protocols. Some of these ideas could be valuably adapted into a future commitment mapping methodology. The approach begins with an easily understood, but not exhaustive, set of sequence diagrams, and then specifies specific steps to generalize the protocol and expand its set of runs.

Our model and formulation of refinement respects intuitions similar to those of Mallya and Singh [2007]. Mallya and Singh describe protocol composition and establish results relating composition and refinement. We expect that similar results would apply in our framework, and plan to formalize and reason about composition in future work.

Model checkers have been extended to handle epistemic and strategic modal operators [Alur et al. 2002; Fagin et al. 1995]. We have begun investigating the inclusion of such concepts into our definitions. Building on top of a model checker that already handles those concepts, such as MCMAS, will simplify our future extensions. Another important enhancement would be to expand the class of protocols to those that support iteration.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for valuable comments and Amit Chopra for his suggestion of describing complex messages as sets of primitive actions.

REFERENCES

- ALUR, R., HENZINGER, T. A., AND KUPFFERMAN, O. 2002. Alternating-time temporal logic. *Journal of the ACM* 49, 5 (Sept.), 672–713.
- BALDONI, M., BAROGLIO, C., AND MARENGO, E. 2010. Constraints among commitments: Regulative specification of interaction protocols. In *Proceedings of the International Workshop on Agent Communication*. Toronto, 10–29.
- CHOPRA, A. K. AND SINGH, M. P. 2009. Multiagent commitment alignment. In *Proceedings of the 8th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. IFAAMAS, Budapest, 937–944.
- COHEN, M., DAM, M., LOMUSCIO, A., AND RUSSO, F. 2009. Abstraction in model checking multi-agent systems. In *Proceedings of the 8th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. IFAAMAS, Budapest, 945–952.
- EL-MENSHAWY, M., BENTAHAR, J., WAN, W., AND DSSOULI, R. 2010. Verifying conformance of commitment protocols via symbolic model checking. In *Proceedings of the International Workshop on Agent Communication*. Toronto, 53–72.
- FAGIN, R., HALPERN, J. Y., MOSES, Y., AND VARDI, M. Y. 1995. *Reasoning about Knowledge*. MIT Press, Cambridge, MA, USA.
- GABBAY, D. 1987. The declarative past and imperative future: Executable temporal logic for interactive systems. In *Proceedings of the Colloquium on Temporal Logic in Specification*. LNCS, vol. 398. Springer-Verlag, 409–448.
- LOMUSCIO, A., QU, H., AND RAIMONDI, F. 2009. MCMAS: A model checker for the verification of multi-agent systems. In *Proceedings of the International Conference on Computer Aided Verification*. LNCS. 682–688.
- LOMUSCIO, A. AND RAIMONDI, F. 2006. Model checking knowledge, strategies, and games in multi-agent systems. In *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. ACM, Hakodate, Japan, 161–168.
- MALLYA, A. U. AND SINGH, M. P. 2007. An algebra for commitment protocols. *Journal of Autonomous Agents and Multi-Agent Systems* 14, 2 (Apr.), 143–163.
- MALONE, T. W., CROWSTON, K., AND HERMAN, G. A., Eds. 2003. *Organizing Business Knowledge: The MIT Process Handbook*. MIT Press, Cambridge, MA.
- MCBURNEY, P. AND PARSONS, S. 2002. Games that agents play: A formal framework for dialogues between autonomous agents. *Journal of Logic, Language and Information* 11, 315–334.
- NORMAN, T. J. AND REED, C. 2002. Group delegation and responsibility. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems*. AAMAS. ACM, Bologna, 491–498.
- RESCHER, N. 1998. Collective responsibility. *Journal of Social Philosophy* 29, 3 (Dec.), 46–58.
- SINGH, M. P. 1999. An ontology for commitments in multiagent systems: Toward a unification of normative concepts. *Artificial Intelligence and Law* 7, 1 (Mar.), 97–113.
- SINGH, M. P. 2008. Semantical considerations on dialectical and practical commitments. In *Proceedings of the 23rd Conference on Artificial Intelligence*. AAAI Press, 176–181.
- SINGH, M. P. AND CHOPRA, A. K. 2010. Correctness properties for multiagent systems. In *Proceedings of the 6th AAMAS Workshop on Declarative Agent Languages and Technologies (DALT 2009)*. LNAI, vol. 5948. Springer, Budapest, 192–207.
- WINIKOFF, M. 2006. Designing commitment-based agent interactions. In *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology*. IEEE Computer Society, Washington, DC, USA, 363–370.
- WINIKOFF, M. 2007. Implementing commitment-based interactions. In *Proceedings of the 6th International Conference on Autonomous Agents and Multiagent Systems*. ACM, Honolulu, 1–8.
- YOLUM, P. 2007. Design time analysis of multiagent protocols. *Data and Knowledge Engineering Journal* 63, 137–154.
- YOLUM, P. AND SINGH, M. P. 2002. Commitment machines. In *Proceedings of the 8th International Workshop on Agent Theories, Architectures, and Languages (ATAL 2001)*. LNAI, vol. 2333. Springer, Seattle, 235–247.

Received August 2010; revised January 2011, May 2011

DRAFT, Vol. V, No. N, July 2011.

THIS DOCUMENT IS THE ONLINE-ONLY APPENDIX TO:

Formalizing and Verifying Protocol Refinements

Scott N. Gerard and Munindar P. Singh
North Carolina State University

DRAFT, Vol. V, No. N, July 2011, Pages 1–28.

A. SUPPLEMENTARY THEOREMS

The first theorem connects interpreted system models with our definitions.

THEOREM A.1. *Let $P = \langle \mathcal{R}, \mathcal{M}, \mathcal{C}, \mathcal{A}, \mathcal{S}, \mathcal{G} \rangle$ be a protocol and let \mathcal{I} be its Proton model. A run is allowed by a Proton model \mathcal{I} (Definition 4.2) if and only if it is a well-defined run (Definition 4.3).*

PROOF SKETCH. Proton models allow interleaved, but not concurrent, messages. At each step, the environment schedules some role $r \in \mathcal{R}$. Role r chooses some enabled message $m \in Act^r$, and the ISPL joint action Act is equal to m ,

Runs for both ISPL and Definition 4.3 begin in state \emptyset . At every step in a run, a message is enabled in ISPL by local strategy AP^r if and only if that message's guard is enabled. Therefore, a message can be appended to an ISPL run if and only if it can be appended to a Proton run. \square

The next theorem shows embedding from Definition 4.8 is equivalent to the model checker verifying guards with Equation 12. The idea behind this theorem is that it assumes the two protocols are already mapped, so the guards of the superprotocol can be evaluated in the Proton model generated from the subprotocol.

Let π^i denote the path consisting of the first i steps of π , let $|\pi|$ be the length of path π , and let $\pi + \langle a, s \rangle$ be path π extended by action a resulting in new state s .


THEOREM A.2. *Let P and Q be two protocols, and M a mapping between them. Let \mathcal{I}_Q be the Proton model for Q as specified in Definition 4.2. Let $\pi_P = \langle p_0, b_1, p_1, \dots \rangle$, and $\pi_Q = \langle q_0, c_1, q_1, \dots \rangle$. Then, for all runs $\pi_Q \in \mathbf{runs}(Q)$, there exists a run $\pi_P \in \mathbf{runs}(P)$ such that $\mathbf{emb}(M(\pi_Q), \pi_P)$ if and only if $(\forall a_i \in \mathcal{A}_P : \mathcal{I}, g \models \mathbf{AG}(a_i.\text{sub-guard} \rightarrow a_i.\text{super-guard}))$*

PROOF. From Theorem A.1, checking for well-defined runs is the same as checking the interpreted system model. Figure 8 diagrams the relationships between entities in π_P and π_Q .

Let RHS be $\mathbf{emb}(M(\pi_Q), \pi_P)$, and let LHS be $(\forall a_i \in \mathcal{A}_P : \mathcal{I}, g \models \mathbf{AG}(a_i.\text{sub-guard} \rightarrow a_i.\text{super-guard}))$.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2011 ACM xxxx-3785/2011/0700-0001 \$5.00


 Fig. 8. The mapping between entities in π_P and π_Q .

\Rightarrow Assume RHS. Let $LHS^j = (\forall i : \tau(i) \leq j, a_i \in \mathcal{A}^j : \mathcal{I}, g \models \mathbf{AG}(a_i.\text{sub-guard} \rightarrow a_i.\text{super-guard}))$ where set $\mathcal{A}^j = \{b_k \mid j = \tau(i) \wedge b_k \in \pi_P^i\}$. We prove LHS^j by induction on path length j in π_Q .

—Base case:

- Define $\mathcal{A}^0 = \emptyset$.
- Define $\tau(0) = 0$.
- LHS^0 is vacuously true.

—Inductive Step: Assume LHS^j . Consider action c_j .

—Case: There are no more actions c_j . We are at the end of π_Q and setting $j = |\pi_Q|$ gives $LHS^{|\pi_Q|} = LHS$.

—Case: $c_j.\text{actexp} \notin \mathcal{A}_P$. This corresponds to the case $\tau(i) < j < \tau(i+1)$.

- c_j has no effect on P .
- $\widehat{M}(q_{\tau(i)}) = \widehat{M}(q_j)$ by Definition 4.8.
- Define $\mathcal{A}^{j+1} = \mathcal{A}^j$.
- Therefore, $LHS^{j+1} = LHS^j$ is true.

—Case: $c_j \in \mathcal{A}_P$. This corresponds to the case $j = \tau(i+1)$.

- Define $j = \tau(i+1)$.
- $p_i = \widehat{M}(q_{\tau(i+1)-1})$ by Definition 4.8.
- $p_{i+1} = \widehat{M}(q_{\tau(i+1)})$ by RHS.
- Since $q_{\tau(i+1)} = q_{\tau(i+1)-1} \cup c_{\tau(i+1)}.\text{actexp}$ and $p_{i+1} = p_i \cup b_{i+1}.\text{actexp}$, then $c_{\tau(i+1)}.\text{actexp} = b_{i+1}.\text{actexp}$, and $c_{\tau(i+1)} = b_{i+1}$.
- Let $a_{i+1} = c_{\tau(i+1)} = b_{i+1}$ be the name of the action in LHS.
- Define $\mathcal{A}^{j+1} = \mathcal{A}^j \cup b_{i+1}$. b_{i+1} might already be a member of \mathcal{A}^j .
- $\widehat{M}(q_{\tau(i+1)-1}) \models c_{\tau(i+1)}.\text{guard}$, since π_Q is well defined.
- $p_i \models b_{i+1}.\text{guard}$, since π_P is well defined.
- Therefore, $c_{\tau(i+1)}.\text{guard} = b_{i+1}.\text{guard}$ in state p_i .
- $c_{\tau(i+1)}.\text{guard} \rightarrow b_{i+1}.\text{guard}$ in state p_i , by previous step. Induction shows it holds for all states in all $\pi_Q \in \mathbf{runs}(Q)$. Since we do not consider all runs $\pi_P \in \mathbf{runs}(P)$, there may be states where $b_{i+1}.\text{guard}$ is true, but $c_{\tau(i+1)}.\text{guard}$ is not true.
- $a_{i+1}.\text{sub-guard} \rightarrow a_{i+1}.\text{super-guard}$.
- Since there is a $\pi_Q \in \mathbf{runs}(Q)$ for every true guard, it is true for all reachable states (**AG**).
- Therefore, LHS^{j+1} is true.

\Leftarrow Assume LHS. Let $RHS^j = (\forall i : \tau(i) \leq j : \pi_P^i \text{ is well defined} \wedge \mathbf{emb}(M(\pi_Q^j), \pi_P^i))$.

Given any π_Q , we will construct a π_P such that RHS^j by induction on path length j in π_Q .

—Base case:

- $q_0 = M(q_0) = \emptyset = p_0$.
 - Define $\tau(0) = 0$.
 - Define $\pi_P^i = \langle p_0 \rangle$ which is well defined.
 - Since $p_0 = \emptyset = \widehat{M}(q_0)$, then $\mathbf{emb}(M(\pi_Q^0), \pi_P^0) = RHS^0$.
- Inductive Step: Assume RHS^j . Then $\mathbf{emb}(M(\pi_Q^{\tau(i)}), \pi_P^i)$ so that $p_i = \widehat{M}(q_{\tau(i)})$. Consider the next action $c_{j+1} \in \pi_Q$.
- Case: There are no more actions c_{j+1} . We are at the end of π_Q and setting $j = |\pi_Q|$ gives $RHS^{|\pi_Q|} = RHS$.
- Case: Action $c_{j+1} \rightarrow \mathcal{A}_P$.
- c_{j+1} does not change $\widehat{M}(q_j)$ by Definition 4.8.
 - $RHS^{j+1} = RHS^j$ is true.
- Case: c_{j+1} equals some $b' \in \mathcal{A}_P$.
- $c_{j+1} \in \mathcal{A}_P$ and $c_j.\mathit{guard}$ and $c_j.\mathit{actexp}$ are also elements in \mathcal{A}_P .
 - Define $\tau(i+1) = j+1$.
 - Let a_{i+1} in LHS equals $c_{j+1} = c_{\tau(i+1)}$ and $b' = b_{i+1}$ in RHS, so that $a_{i+1} = c_{\tau(i+1)} = b_{i+1}$. And a_{i+1} 's sub-guard is $c_{\tau(i+1)}.\mathit{guard}$ and a_{i+1} 's super-guard is $b_{i+1}.\mathit{guard}$.
 - Define $p_{i+1} = p_i \cup b_{i+1}.\mathit{actexp}$.
 - Define $\pi_P^{i+1} = \pi_P^i + \langle b_{i+1}, p_{i+1} \rangle$.

Show π_P^{i+1} is well defined.

- (1) • $q_{\tau(i+1)-1} \models c_{\tau(i+1)}.\mathit{guard}$, since π_Q is well defined.
 - $\widehat{M}(q_{\tau(i+1)-1}) \models c_{\tau(i+1)}.\mathit{guard}$, since $c_{\tau(i+1)}.\mathit{guard}$ is an element of \mathcal{A}_P .
 - $\widehat{M}(q_{\tau(i)}) \models c_{\tau(i+1)}.\mathit{guard}$ by Definition 4.8.
 - $p_i \models c_{\tau(i+1)}.\mathit{guard}$ by Definition 4.8 and RHS^j .
 - $a_{i+1}.\mathit{sub-guard} \rightarrow a_{i+1}.\mathit{super-guard}$ by LHS.
 - Since $c_{\tau(i+1)} = a_{i+1} = b_{i+1}$ are all the same action, $c_{\tau(i+1)}.\mathit{guard} \rightarrow b_{i+1}.\mathit{guard}$.
 - Therefore, $p_i \models b_{i+1}.\mathit{guard}$.
- (2) $p_{i+1} = p_i \cup b_{i+1}.\mathit{actexp}$ by definition above.

Show $\mathbf{emb}(\pi_Q^{\tau(i+1)}, \pi_P^{i+1})$.

- $p_i = \widehat{M}(q_{\tau(i+1)-1})$ by Definition 4.8 and RHS^j .
- Since $c_{\tau(i+1)} = a_{i+1} = b_{i+1}$ are all the same action, $p_i \cup b_{i+1}.\mathit{actexp} = \widehat{M}(q_{\tau(i+1)-1}) \cup c_{\tau(i+1)}.\mathit{actexp}$.
- This reduces to $p_{i+1} = \widehat{M}(q_{\tau(i+1)})$.

Therefore, RHS^{j+1} .

□

The next definition characterizes that action a whose $a.\mathit{actexp}$ causes $m.\mathit{actexp}$ to become true.

Definition A.3 Decisive. Let P be a protocol, let e, e' be two Boolean expressions, and let s be any state and $s' = s \cup e'$ be the next valid state after s where e' holds. e' is *decisive* for e , at a state s if and only if $s \not\models e'$, and $s' \models e'$ implies $s \not\models e \wedge s' \models e$.

The state s' in the definition is the state in which e becomes true. An expression e' is decisive for expression e in state s if and only if, making e' true also makes e true. A change in e' causes a change in e .

In particular, we say an action a is decisive for message m at state s exactly when expression $a.actexp$ is decisive for expression $m.actexp$ at state s .

The next theorem shows that diffusion and collection properly maintain the guards and action expressions as a message is decomposed from a protocol P to its derived protocol $P' = \mathbf{col}(\mathbf{dif}(\mathbf{means}(M(P))))$. We prove it for mappings that contain individual actions and conjunctions, as well as mapping that contain disjunctions even though disjunction is not required by later proofs. This theorem is used between the superprotocol's super-gMsg and super-gAct, and between the subprotocol's sub-gMsg and sub-gAct in Figure 7(b).

THEOREM A.4 DIFFUSION AND COLLECTION PRESERVE GUARDS. *Let P be any protocol, let M be any mapping function, and let $gs \in \mathcal{G}$ be any guarded statement in P , possibly containing guarded action expressions. If gs_i is any guarded statement derived from gs by diffusion, and if gs_i is decisive for gs at state s then*

$$s \models gs_i.guard \leftrightarrow s \models gs.guard$$

PROOF. Show diffusion preserves guards.

Diffusion breaks one guarded statement gs into a set of guarded statements gs_i . Let $LHS^i = ((gs_i \text{ is decisive for } gs \text{ at } s) \rightarrow (s \models gs_i.guard \leftrightarrow s \models gs.guard))$ where gs_i is derived from gs_{i-1} by diffusion and collection. We prove LHS^i by induction on the structure of $gs_i.actexp$.

—Base case:

- $gs_0 = gs$ and $gs_0.guard = gs.guard$ is trivially true for all states s .
- LHS^0 is true.

—Inductive Step: Assume LHS^i .

—Case: There is no outermost operator in $gs_i.actexp$.

- Then $gs_i.actexp$ is a single guarded action.
- $gs_{i+1}.guard = gs_i.guard$ by Equation 11.
- $gs_{i+1}.guard = gs.guard$ by LHS^i .
- LHS^{i+1} since this holds in all states s .

—Case: The outermost operator of $gs_i.actexp$ is disjunction.

- Equation 9 applied to gs_i creates multiple guarded statements, one for each disjunct. Let gs_{i+1} be any of those disjuncts.
- $gs_{i+1}.guard = gs_i.guard$ by Equation 9,
- Each gs_{i+1} is decisive for gs .
- $gs_{i+1}.guard = gs.guard$ by LHS^i .
- LHS^{i+1} since this holds in all states s .

—Case: The outermost operator of $e.actexp$ is conjunction.

- Equation 10 applied to gs_i creates multiple guarded statements, one for each conjunct. Let gs_{i+1} be any of those conjuncts.
- For gs_{i+1} to be decisive at s , all other conjuncts must be true at s .
- $gs_{i+1}.guard = gs_i.guard$ by Equation 10 because all other conjuncts are true.
- $gs_{i+1}.guard = gs.guard$ by LHS^i .

- LHS^{i+1} since this holds in all states s for which gs_{i+1} is decisive.

By the no overlap constraint of Definition 4.1, in any state s , at most one of the guarded statement combined by collection can be enabled at a time. Therefore, collection also preserves guards. \square

The next three theorems relate runs of a protocol P expressed in terms of messages and the runs of its derived protocol $P' = \mathbf{col}(\mathbf{dif}(M(P)))$ expressed in terms of actions. These theorems relate runs between both (1) the super-gMsg and super-gAct protocols, and (2) the sub-gMsg and sub-gAct protocols as shown in Figure 7. The first theorem proves every run of P embeds a run of P' .

THEOREM A.5. *If Let $P = \langle \mathcal{R}, \mathcal{M}, \mathcal{C}, \mathcal{A}, \mathcal{S}, \mathcal{G} \rangle$ be a protocol, possibly containing guarded action expressions, and let M be any mapping function. Let $P' = \mathbf{col}(\mathbf{dif}(M(P)))$ be the protocol derived from P by mapping, diffusion, and collection. Then*

$$\forall \pi_r \in \mathbf{runs}(P) : (\exists \pi_s \in \mathbf{runs}(P') : \mathbf{emb}(\pi_s, \pi_r))$$

PROOF. Denote $\pi_r = \langle h_0, m_1, h_1, \dots \rangle \in \mathbf{runs}(P)$ with message $m_i \in \mathcal{M}$, and $\pi_s = \langle g_0, a_1, g_1, \dots \rangle \in \mathbf{runs}(P')$ with mapped action $a_i \in M(\mathcal{A})$. Let $LHS^i = (\pi_s^{\mu(i)} \text{ is well defined} \wedge \mathbf{emb}(\pi_s^{\mu(i)}, \pi_r^i))$. We will construct a well-defined run $\pi_s \in \mathbf{runs}(P')$, and show LHS^i by induction on $0 \leq i \leq |\pi_r|$.

—Base case:

- Define $g_0 = \emptyset$.
- Define $\pi_s^0 = \langle g_0 \rangle$ which is well defined.
- Define $\mu(0) = 0$.
- Then $h_0 = \emptyset = \widehat{M}(g_0) = \widehat{M}(g_{\mu(0)})$.
- LHS^0 is true.

—Inductive Step: Assume LHS^i where $h_i = \widehat{M}(g_{\mu(i)})$. Let m_{i+1} be the next message in π_r .

—Case: No such m_{i+1} exists. All messages in π_r have been considered. $\pi_s = \pi_s^{\mu(i)}$ is a well-defined run and μ demonstrates $\mathbf{emb}(\pi_s^{|\pi_s|}, \pi_r^{|\pi_r|}) = \mathbf{emb}(\pi_s, \pi_r)$. $LHS^{|\pi_r|} = LHS$ is true.

—Case: m_{j+1} exists.

- Let $n = |\mathbf{means}(m_{i+1})|$ be the number of actions in m_{i+1} 's meaning.
- Define $g_{k+1} = g_k \cup a_{k+1} \cdot \mathbf{actexp} \forall k : \mu(i) \leq k < \mu(i) + n$.
- Define $\pi_s^{\mu(i+1)} = \pi_s^{\mu(i)} + \sum_{k: \mu(i) \leq k < \mu(i) + n} \langle a_{k+1}, g_{k+1} \rangle$ by appending all the actions a_k in $\mathbf{means}(m_{i+1})$ onto the end, in any order.
- Let a_d be the decisive action for m_{i+1} in π_r where $\mu(i) \leq d \leq \mu(i) + n$.
- Define $\mu(i+1) = d$.

Show $\pi_s^{\mu(i+1)}$ is well defined.

(1) Show $g_k \models a_{k+1} \cdot \mathbf{guard} \forall k : \mu(i) \leq k < \mu(i) + n$

- $h_i \models m_{i+1} \cdot \mathbf{guard}$ because π_r is well defined.
- For each action $a_j \in \mathbf{means}(m_{i+1})$, $m_{i+1} \cdot \mathbf{guard} \rightarrow a_j \cdot \mathbf{guard}$ by Equations 9, 10, and 11. Since collection disjoins guards, each of the actions' guard is true after collection.

- Therefore, $h_i \models a_k.\text{guard} \forall k : \mu(i) \leq k < \mu(i) + n$.
- (2) $g_{k+1} = g_k \cup a_k.\text{actexp} \forall k : \mu(i) \leq k < \mu(i) + n$ by the definition above.

Show $\mathbf{emb}(\pi_s^{\mu(i+1)}, \pi_r^{i+1})$.

- $g_{\mu(i+1)} = g_{\mu(i)} \bigcup_{\mu(i) \leq k < \mu(i)+n} a_k.\text{actexp}$ by definition of $g_{\mu(i+1)}$ above.
- $\widehat{M}(g_{\mu(i+1)}) = \widehat{M}(g_{\mu(i)} \bigcup_{\mu(i) < k \leq \mu(i)+n} a_k.\text{actexp})$.
- For all j , $\widehat{M}(g_j \cup a_{j+1}.\text{actexp}) = \widehat{M}(g_{j+1})$ if a_{j+1} is not decisive for m_{j+1} in π_s .
- For all j , $\widehat{M}(g_j \cup a_{j+1}.\text{actexp}) = \widehat{M}(g_{j+1}) \cup m_{j+1}.\text{actexp}$ if a_{j+1} is decisive for m_{j+1} in π_s .
- $\widehat{M}(g_{\mu(i+1)}) = \widehat{M}(g_{\mu(i)}) \cup m_{j+1}.\text{actexp}$, by applying the previous two reductions for all a_{j+1} .
- $\widehat{M}(g_{\mu(i+1)}) = h_i \cup m_{j+1}.\text{actexp}$, since $h_i = \widehat{M}(g_{\mu(i)})$ by LHS^{i+1} .
- $\widehat{M}(g_{\mu(i+1)}) = h_{i+1}$ by the definition of h_{i+1} .
- Therefore, LHS^{i+1} is true.

□

The next theorem shows the reverse: every run in P' embeds a run in P .

THEOREM A.6. *If Let $P = \langle \mathcal{R}, \mathcal{M}, \mathcal{C}, \mathcal{A}, \mathcal{S}, \mathcal{G} \rangle$ be a protocol, possibly containing guarded action expressions, and let M be any mapping function. Let $P' = \mathbf{col}(\mathbf{dif}(M(P)))$ be the protocol derived from P by mapping, diffusion, and collection. Then*

$$\forall \pi_s \in \mathbf{runs}(P') : (\exists \pi_r \in \mathbf{runs}(P) : \mathbf{emb}(\pi_s, \pi_r))$$

PROOF. Denote $\pi_r = \langle h_0, m_1, h_1, \dots \rangle \in \mathbf{runs}(P)$ with message $m_i \in \mathcal{M}$, and $\pi_s = \langle g_0, a_1, g_1, \dots \rangle \in \mathbf{runs}(P')$ with mapped action $a_i \in M(\mathcal{A})$.

Let $LHS^j = (i = \mathop{\text{argmax}}_k \mu(k) \leq j : \pi_r^i \text{ is well defined} \wedge \mathbf{emb}(\pi_s^j, \pi_r^i))$. We will construct $\pi_r \in \mathbf{runs}(P)$ and show LHS^j by induction on the path length $0 \leq j \leq |\pi_s|$. We allow additional actions in π_s after $\mu(i)$ as long as they have no effect on P .

—Base case:

- Define $h_0 = \emptyset$,
- Define $\pi_r^0 = \langle h_0 \rangle$,
- π_r^0 is well defined.
- Define $\mu(0) = 0$.
- Then $\widehat{M}(g_0) = \widehat{M}(g_{\mu(0)}) = \emptyset = h_0$ implies $\mathbf{emb}(\pi_s^0, \pi_r^0)$ with $\mu(0) \leq 0$.
- LHS^0 is true.

—Inductive Step: Consider the next action $a_{j+1} \in \pi_s$.

—Case: No such a_{j+1} exists. All actions in π_s have been considered. $\pi_r = \pi_r^j$ is a well-defined run, and μ demonstrates $\mathbf{emb}(\pi_s^{|\pi_s|}, \pi_r^{|\pi_r|}) = \mathbf{emb}(\pi_s, \pi_r)$ with $\forall i : \mu(i) \leq j$. $LHS^{|\pi_r|} = LHS$.

—Case: a_j exists but it is *not* decisive for any $m \in \mathcal{A}_P$.

- Leave π_r^j unchanged which is still well defined.
- Leave μ unchanged.

- No additional $h_i = \widehat{M}(g_{\mu(i)})$ conditions are required to established $\mathbf{emb}(\pi_s^{j+1}, \pi_r^i)$ with $\forall i : \mu(i) \leq j$, and the existing conditions are true by the induction hypothesis.

- LHS^{j+1} is true.

—Case: a_{j+1} exists and it is decisive for some $m \in \mathcal{A}_P$.

- There is at most one such message m by the no overlap constraint of Definition 4.1. Denote the message by $m_{i+1} = m$.
- Define $h_{i+1} = h_i \cup m_{i+1}.\mathit{actexp}$.
- Define $\pi_r^{i+1} = \pi_r^i + \langle m_{i+1}, h_{i+1} \rangle$.
- Define $\mu(i+1) = j+1$.

Show π_r^{i+1} is well defined.

(1) Show $h_i \models m_{i+1}.\mathit{guard}$.

- $g_j \models a_{j+1}.\mathit{guard}$ because π_s is well defined.
- $\widehat{M}(g_j) \models \widehat{M}(a_{j+1}.\mathit{guard})$.
- $\widehat{M}(g_j) \models m_{i+1}.\mathit{guard}$ because $a_{j+1}.\mathit{guard} = m_{i+1}.\mathit{guard}$ by Theorem A.4.
- $\widehat{M}(g_{\mu(i+1)-1}) \models m_{i+1}.\mathit{guard}$ by definition of $\mu(i+1)$.
- $\widehat{M}(g_{\mu(i)}) \models m_{i+1}.\mathit{guard}$ by Definition 4.8.
- Therefore, $h_i \models m_{i+1}.\mathit{guard}$ by LHS^j .

(2) $h_{i+1} \models m_{i+1}.\mathit{actexp}$ by definition of h_{i+1} above.

Show $\mathbf{emb}(\pi_s^{j+1}, \pi_r^{i+1})$.

- $g_{j+1} = g_j \cup a_{j+1}.\mathit{actexp}$ since π_s is well defined.
- $\widehat{M}(g_{j+1}) = \widehat{M}(g_j \cup a_{j+1}.\mathit{actexp})$.
- For all j , $\widehat{M}(g_j \cup a_{j+1}.\mathit{actexp}) = \widehat{M}(g_j) \cup m_{j+1}.\mathit{actexp}$ if a_{j+1} is decisive for $m_{j+1}.\mathit{actexp}$ in π_s .
- $\widehat{M}(g_{j+1}) = \widehat{M}(g_j) \cup m_{j+1}.\mathit{actexp}$, by applying the previous reduction.
- $\widehat{M}(g_{\mu(i+1)}) = \widehat{M}(g_{\mu(i+1)-1}) \cup m_{j+1}.\mathit{actexp}$, by definition of $\mu(i+1)$.
- $\widehat{M}(g_{\mu(i+1)}) = \widehat{M}(g_{\mu(i)}) \cup m_{j+1}.\mathit{actexp}$, by Definition 4.8.
- $\widehat{M}(g_{\mu(i+1)}) = h_i \cup m_{j+1}.\mathit{actexp}$, since $h_i = \widehat{M}(g_{\mu(i)})$ by LHS^j .
- $\widehat{M}(g_{\mu(i+1)}) = h_{i+1}$ by the definition of h_{i+1} .

— LHS^{i+1} is true.

□

THEOREM A.7. *If Let $P = \langle \mathcal{R}, \mathcal{M}, \mathcal{C}, \mathcal{A}, \mathcal{S}, \mathcal{G} \rangle$ be a protocol, possibly containing guarded action expressions, and let M be any mapping function. Let $P' = \mathbf{col}(\mathbf{dif}(M(P)))$ be the protocol derived from P by mapping function M , diffusion, and collection. Then*

$$\begin{aligned} \forall \pi_r \in \mathbf{runs}(P) : (\exists \pi_s \in \mathbf{runs}(P') : \mathbf{emb}(\pi_s, \pi_r)) \\ \forall \pi_s \in \mathbf{runs}(P') : (\exists \pi_r \in \mathbf{runs}(P) : \mathbf{emb}(\pi_s, \pi_r)) \end{aligned}$$

PROOF. Follows immediately from Theorems A.5 and A.6. □