Azorus: Commitments over Protocols for BDI Agents

Amit K. Chopra
Lancaster University
Lancaster, UK
amit.chopra@lancaster.ac.uk

Matteo Baldoni
Università degli Studi di Torino
Torino, Italy
baldoni@di.unito.it

Samuel H. Christie V
North Carolina State University
Raleigh, NC, USA
schrist@ncsu.edu

Munindar P. Singh
North Carolina State University
Raleigh, NC, USA
mpsingh@ncsu.edu

ABSTRACT

Commitments support flexible interactions between agents by capturing the meaning of their interactions. However, commitment-based reasoning is not adequately supported in agent programming models. We contribute Azorus, a programming model based on declarative specifications centered on commitments and aligned with information protocols. Azorus supports reasoning about goals and commitments and combines modeling of commitments and protocols, thereby uniting three leading declarative approaches to engineering decentralized multiagent systems. Specifically, we realize Azorus over three existing technology suites: (1) Jason, a popular BDI-based programming model; (2) Cupid, a formal language and query-based model for commitments; and (3) BSPL, a language and its associated tools for information protocols, including Jason programming. We implement Azorus and demonstrate how it enables capturing interesting patterns of business logic.

CCS CONCEPTS

• **Computing methodologies** → **Multi-agent systems**.

KEYWORDS

Agent Programming; Meaning; Asynchrony; Causality

ACM Reference Format:
Amit K. Chopra, Matteo Baldoni, Samuel H. Christie V, and Munindar P. Singh. 2025. Azorus: Commitments over Protocols for BDI Agents. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 10 pages.

1 INTRODUCTION

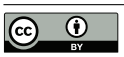
Important domains such as business and healthcare that involve autonomous principals lend themselves to the application of decentralized multiagent systems (MAS). Engineering flexible MAS calls upon programming abstractions for social meaning, operational interactions, and agent reasoning.

Commitments are a high-level abstraction that capture the social *meaning* of a communicative act [31]. For example, an *offer* from a seller to a buyer for some Item and Price may be modeled as a commitment from the seller to the buyer that if *payment* of Price happens, then the *shipment* of Item will happen. Commitments model autonomy by both enabling flexible engagements between agents and yielding a standard for compliance [23, 39, 43]. Several previous works address languages for specifying commitments [2, 10, 15].

However, much of the work on commitments does not address decentralized settings. To support such settings, commitments need to be layered on flexible, decentralized interaction protocols that minimally constrain when agents may perform communicative acts [14]. For example, *refund* without a prior *payment* would be meaningless; and *accept* and *reject* should be mutually exclusive to be meaningful; however, *shipment* and *payment* may happen in any order. Because of their emphasis on message ordering, traditional protocol specification approaches [3, 5, 20, 41] are not suited to specifying flexible protocols. For this purpose, we turn toward information protocols, specifically BSPL [32], a declarative approach for specifying flexible protocols. Indeed a motivation for information protocols was to provide a suitable operational layer for commitments [32, p. 498].

Commitments are not adequately supported in programming models for multiagent systems. Popular approaches such as JADE [6], Jason [8], JaCaMo [7], and SARL [24] provide diverse, useful abstractions for engineering multiagent systems. However, the abstractions for communication in these approaches are either low-level (e.g., messaging in JADE and Jason and *event spaces* in SARL), limited in repertoire, inflexible (support for FIPA Interaction Protocols [22] in JADE), or promote centralization (via *artifacts* in JaCaMo). MOISE (the 'Mo' in JaCaMo) [26] supports a notion of commitments but tightly couples them to agent goals. Baldoni et al. [1] model communicative acts and their effects on commitments via JaCaMo artifacts. Kiko [17], an information protocol-based programming model supports creating flexible, decentralized MAS but does not support commitments.

We contribute *Azorus* (named after the helmsman of Jason's ship, the Argo), a commitment-based programming model that enables implementing flexible MAS via BDI agents. We synthesize, for the first time, three declarative MAS paradigms: commitments, information protocols, and cognitive agents. For the latter, we adopt BDI (belief-desire-intention) agents, which have beliefs and goals, and execute plans in response to changes in beliefs and goals. Jason [8] is a prominent exemplar of the paradigm (and the 'Ja' in JaCaMo). The synthesis makes conceptual sense because, in a multiagent



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

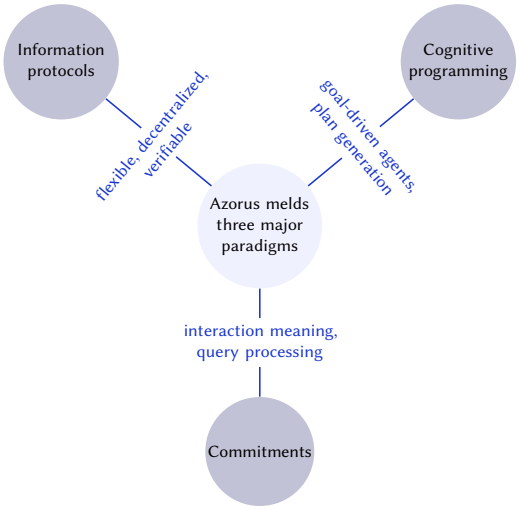


Figure 1: Azorus in a nutshell.

system, agents depend on others for the satisfaction of their goals [30]. Commitments capture such dependencies between agents [25], and, as described above, motivate information protocols. Winikoff [40] notes the lack of support for flexible interactions in agent programming. Accordingly, we contribute:

- A formalization of Cupid [15], an expressive commitment language, as abstract Jason rules. We provide a compiler enabling a declarative, high-level abstraction for commitments in Jason plans.
- A Jason communication adapter that supports an agent's internal reasoning by maintaining the mapping between commitments and enactments of information protocols and providing abstractions for querying and reacting to commitment events and performing valid communicative acts.
- Reasoning patterns for realizing flexible agents in Azorus.

Organization. The rest of the paper is organized as follows. Section 2 provides background on Jason. Section 3 describes how we specify MAS via commitments in Cupid and information protocols. Section 4 introduces the Azorus programming model via its architectural elements, including a semantics in Jason for inferring commitment events from communicative acts. Section 5 demonstrates patterns for implementing flexible agents. Section 6 evaluates our contributions conceptually. Section 7 summarizes our contributions and identifies some future directions.

2 JASON BACKGROUND

Jason extends the AgentSpeak logic-programming language for specifying agent behavior [8]. An agent is modeled as having *beliefs*, which capture the state of the world; *goals*, which capture its objectives; and *plans*, which are methods for realizing its goals. Jason adopts communication primitives based on the Knowledge Query and Manipulation Language, better known as KQML [9].

To illustrate Jason's programming model, especially how it combines communication and reasoning, Listing 1 gives a snippet of how an agent Bob, who plays SELLER in *Ebusiness*, might be implemented in Jason without any special support for protocols.

Listing 1: Jason snippet of a SELLER agent Bob.

```
buyer(alice).
in_stock(figs).
goes_for(figs, 10).

!start.
+!start <-
        ?buyer(Buyer);
        ?goes_for(Item, Price);
        .random(Id);
        .send(Buyer, tell, offer(Id, Item, Price)).

+accept(Id, Item, Price, Decision)[source(Sender)]
 : in_stock(Item) & buyer(Sender)
 <- .send(Sender, tell, shipment(Id, Item, Price,
      done)).
```

The first few lines of Listing 1 assert beliefs that buyer is `alice`, `figs` are in stock, and that they go for the price of `10`. Then the goal *start* is asserted. The following lines show two plans. The first is for the goal *start* and is executed whenever it is asserted. This plan executes two queries to bind variables Buyer, Item, and Price, respectively. It then uses a library function to bind variable Id to a random identifier. Finally, it uses the built-in function for sending an *offer* to Buyer using the KQML speech act *tell*.

Jason asserts beliefs corresponding to received messages. The second plan is for handling a belief corresponding to a received *accept* and is executed whenever the belief is asserted. The plan checks (via guards in the context) that the specific Item is in stock and that the Sender is the buyer and, if so, sends a *shipment* message.

3 MODELING MULTIAGENT SYSTEMS

As explained below, we use information protocols to specify the basic communicative acts in a MAS; meaning-level commitment specifications refer to these acts.

3.1 Information Protocols in BSPL

Information protocols are *declarative* interaction specifications [32, 33]. An interaction is specified as a composition of protocols—a message being an atomic protocol with a special syntax—in terms of the information dependencies between them. The idea is that an agent can emit any message whose information dependencies are satisfied given its *local state*, that is, its communication history. We adopt information protocols because they support flexible and asynchronous multiparty enactments better than traditional message ordering-oriented representations of protocols [12].

We explain BSPL via the protocol *Ebusiness* in Listing 2. It specifies several messages, each with a sender, a receiver, and information parameters. The parameter Id is annotated key, meaning it serves to identify enactments (and correlate messages). Adornments ⌜in⌝, ⌜out⌝, and ⌜nil⌝ for parameters capture information dependencies and are interpreted relative to enactments. A message in some enactment is *viable* (i.e., legal for emission) if the sender's local state includes bindings for all the ⌜in⌝ parameters and none of the ⌜out⌝

or ⌜nil⌝ parameters. Sending the message adds it to the agent's local state (along with the bindings for the ⌜out⌝ parameters, which are computed at that point, thus making them known); the ⌜nil⌝ parameters remain unbound. Receiving a message adds it to the receiver's local state (along with the bindings for all its parameters, thus making them known). Notably, information protocols do not specify message reception order.

Listing 2: An information protocol in BSPL.

```
Ebusiness {
  roles Buyer, Seller, Bank
  parameters out Id key, out Item, out Price, out
      Status

  Seller -> Buyer: offer[out Id key, out Item,
      out Price]
  Buyer  -> Seller: accept[in Id key, in Item, in
      Price, out Decision]
  Buyer  -> Bank: instruct[in Id key, in Price,
      out Details]
  Bank   -> Seller: transfer[in Id key, in Price,
      in Details, out Payment]
  Seller -> Buyer: shipment[in Id key, in Item,
      in Price, out Status]
  Seller -> Bank:  refund[in Id key, in Item, in
      Payment, out Amount, out Status]
}
```

In an enactment of *Ebusiness*, SELLER may send *offer* anytime since all its parameters are ⌜out⌝. Once SELLER has sent *offer*, it would know the bindings for Id, Item, and Price, which means it may send *shipment* provided it does not already know the binding for Status. Analogously, BUYER may send *accept* or *instruct* anytime after receiving *offer*; BANK may send a *transfer* anytime after receiving *instruct*; and SELLER may send *refund* anytime after sending *offer* and receiving *transfer*. And, *shipment* and *refund* are mutually exclusive since they both bind Status (it is ⌜out⌝ in both).

To get a sense of how flexible *Ebusiness* is, consider the fact that it has 658 distinct maximal enactments (each a causally valid permutation of sends and receives of its messages extended until no agent is left with any viable message), including the enactment depicted in Figure 2, which is notable because *accept* and *transfer* are "reordered" in the communication infrastructure and SELLER sends *shipment* even though it has not received *accept*.

3.2 Specifying Commitments

Cupid is an approach for specifying commitments over databases of business events [15].

Listing 3: Commitment specification in Cupid.

```
commitment OfferCom Seller to Buyer
  create offer
  detach transfer[, created OfferCom + 5]
        where "Payment >= Price"
  discharge shipment [, detached OfferCom + 5]

commitment AcceptCom Buyer to Seller
  create accept
```

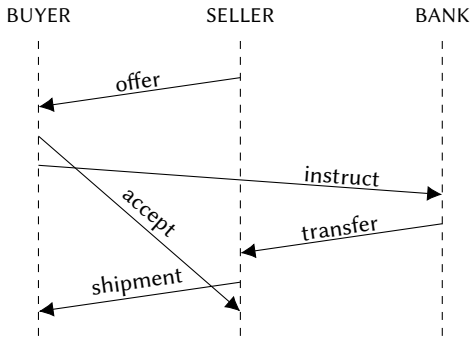


Figure 2: *Ebusiness* enactment in which *shipment* is sent by SELLER even as *accept* was in transit, based on [12, p. 1380].

```
  detach shipment[, created AcceptCom + 5]
  discharge transfer[, detached AcceptCom + 5]
        where "Payment >= Price"

commitment RefundCom Seller to Buyer
  create offer
  detach violated OfferCom
  discharge refund[, detached RefundCom + 2]
      where "Amount >= Payment"

commitment TransferCom Bank to Seller
  create instruct
  discharge transfer[, created TransferCom + 2]
        where "Payment = Price"
```

Listing 3 gives a Cupid specification that gives the meaning of messages in the *Ebusiness* protocol in Listing 2. Specifically, events such as *offer*, *transfer*, and so on refer to the observation of the corresponding message. These events constitute the *base* events for the specification. The attributes of a base event are the parameters of its message plus a unique timestamp attribute.

The commitment `OfferCom` specifies that *offer creates* a commitment (instance) from SELLER to BUYER. This commitment is *detached* if *transfer* happens within 5 time units (for purposes of this paper, seconds) of the creation and Payment in the *transfer* is at least as much as Price in the *offer*. The commitment *expires* (fails to be detached) if either of these conditions is not met. The commitment is *discharged* if *shipment* happens within 5 time units of being detached. The commitment is *violated* (fails to be discharged) if *shipment* fails to occur within the stipulated time.

`AcceptCom` specifies that *accept* creates a commitment from BUYER to SELLER that if *shipment* happens within 5 time units of its creation, then *transfer* will occur within 5 time units of its being detached. `RefundCom` specifies that *offer* creates a commitment from SELLER to BUYER that if `OfferCom` is violated, then *refund* of at least the Amount paid will be done with 2 time units of the violation (else, obviously, the `RefundCom` will be violated). `RefundCom` demonstrates the use of nested commitments, which may be used to capture patterns such as compensation. `TransferCom` captures BANK's commitment to BUYER to do *transfer* upon *instruct*.

Table 1 defines the formal syntax of Cupid, which we include here since we give it a new formal semantics based on Jason. Below, $\mathcal{A}$ and $\mathcal{T}$ are the sets of agent names and time instants, respectively; in particular, $\mathcal{T} = \mathbb{N} \cup \{\infty\}$, where $\mathbb{N}$ is the set of natural numbers and ∞ is an infinitely distant time instant.

ComSpec gives the syntax for a commitment: the debtor and creditor agents, and the create, detach, and discharge clauses. Listing 3 uses a surface syntax for readability. We write and, or, and except for $\sqcap$, $\sqcup$, and $\ominus$ respectively. In time intervals, we omit lower and upper instants when they are 0 and ∞, respectively. An omitted detach clause means the commitment is unconditional. We label commitments to simplify referring to commitment events.

Table 1: Syntax of Cupid [15].

Event	$\longrightarrow$	Base | LifeEvent
LifeEvent	$\longrightarrow$	created($\mathcal{A}$, $\mathcal{A}$, Expr, Expr, Expr) | detached($\mathcal{A}$, $\mathcal{A}$, Expr, Expr, Expr) | discharged($\mathcal{A}$, $\mathcal{A}$, Expr, Expr, Expr) | expired($\mathcal{A}$, $\mathcal{A}$, Expr, Expr, Expr) | violated($\mathcal{A}$, $\mathcal{A}$, Expr, Expr, Expr)
Expr	$\longrightarrow$	Event[Time, Time] | Expr $\sqcap$Expr | Expr $\sqcup$Expr | Expr $\ominus$ Expr | Expr where φ
Time	$\longrightarrow$	Event + $\mathcal{T}$ | $\mathcal{T}$
ComSpec	$\longrightarrow$	commitment($\mathcal{A}$, $\mathcal{A}$, Expr, Expr, Expr)

Cupid specifies five *life events* for every commitment: *created*, *detached*, *expired*, *discharged*, and *violated*. The semantics of Cupid gives a query for each life event for a commitment. The idea is to infer the life events (including their timestamps) from the base events. Time intervals for an event ([Time, Time] in Table 1) are interpreted strictly: the event is required to occur after (including at) the initial moment but before the final moment of the interval.

Chopra and Singh [15] give Cupid's semantics in relational algebra; its existing implementation compiles each life event of a commitment into an SQL query. Azorus provides a new implementation of Cupid into Jason to enable BDI programming using commitments.

4 PROGRAMMING MODEL, ARCHITECTURALLY

Figure 3 describes the Azorus architecture and programming model. A MAS is specified in terms of commitments and an information protocol. The Azorus tooling generates an adapter for the role being played by the agent based on the specifications. The adapter supports implementing agents via programming abstractions for commitments and protocols. The figure shows the computational components of the adapter. Beliefs represent an agent's state.

Each agent sends and receives messages via an *Asynchronous Communication Service*. An agent's *Local State* is the protocol state projected to the messages sent or received by the agent and is represented as a set of beliefs corresponding to the messages. The *Protocol Adapter* captures the protocol constraints relative to the role played by the agent. It sends messages using Jason's communication primitives and adds them to the *Local State*. (Received messages are added to the *Local State* automatically by Jason.) Moreover, the

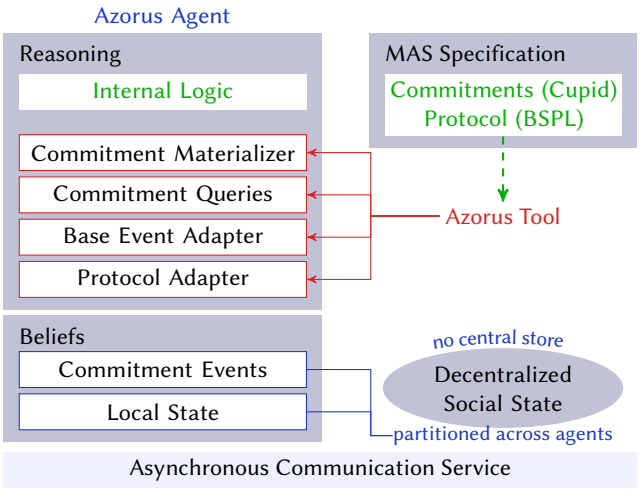


Figure 3: Azorus architecture and programming model.

Protocol Adapter computes the set of *enabled* communicative acts (explained below) from the *Local State*.

As messages are added to the *Local State*, the *Base Event Adapter* asserts the corresponding base events as timestamped beliefs. These events are used by *Commitment Queries* to compute the commitment events. The *Internal Logic* is a set of Jason plans that capture agent behavior (modulo protocol constraints, of course). These plans use *Commitment Queries* and *Protocol Adapter* to reason about commitments and send only enabled communications.

Commitment Queries may be used in the context of a Jason plan. To accommodate a programming style where a Jason plan is triggered by the occurrence of a commitment event, the *Commitment Materializer* asserts commitment events as beliefs as they occur.

In Figure 3, developers provide the MAS specifications and the internal logic of the agents. The value of Azorus arises from generating the *Protocol Adapter*, *Base Event Adapter*, *Commitment Queries*, and *Commitment Materializer* from commitments and protocols and packaging them as the Azorus adapter. Specifically, the agent programmer may focus on writing the *Internal Logic* based on the interface afforded by Azorus adapter: local state (the communicative acts that have occurred), enabled acts (the acts that may be performed), and commitment queries and materialized commitment events (as capturing meaning).

Below we describe each computational component, including how they update the stateful ones.

4.1 Protocol Adapter

Baldoni et al. [4] present Orpheus, a programming model for implementing protocol-based Jason agents. Given an information protocol, the Orpheus protocol adapter enables the implementation of Jason agents that play roles in the protocol. Specifically, an agent's protocol adapter maintains its local state. Based on the state and the protocol specification, it keeps track of information-enabled *forms*. The forms are necessarily partial message instances that would be legal to send if completed. Specifically, a form's $\ulcorner$in$\urcorner$ parameters have bindings from the local state, whereas the $\ulcorner$out$\urcorner$

parameters are unbound because their bindings don't exist in the local state; ⌜nil⌝ parameters are omitted from the form because they are neither bound in the local state nor can be bound.

Listing 4 gives a possible local state for a SELLER agent and Listing 5 shows the forms available to it in that state.

Listing 4: A possible local state for a SELLER agent. It contains instances of messages in the *Ebusiness* protocol.

```
offer(1, fig, 10)
offer(2, jam, 100)
accept(2, jam, 100, yes)
transfer(1, 10, done, 10)
```

Listing 5: Enabled forms, showing parameters to be bound.

```
offer(Id, Item, Price)
shipment(1, fig, 10, Status)
refund(1, fig, 10, Amount, Status)
shipment(2, jam, 100, Status)
```

To write an Orpheus agent, a programmer writes a set of *plans*. Each plan is an event-triggered piece of code that gets some *enabled* forms; *completes* them via some logic; and then *attempts* to send them. If the attempt passes the required integrity checks, the adapter turns the completed forms into messages on the wire and records them in the local state. Listing 6 shows a Jason code snippet (blue for Orpheus constructs; red for what a programmer must implement) that represents a SELLER agent's internal reasoning. The first plan concerns communicating *offers*. If there is an *enabled offer* form, then it *completes* the form by checking if it has something to offer, and then *attempts* to send it. The listing also contains a plan for completing and attempting *shipment* forms. The *enabled* predicate and *attempt* are adapter abstractions. The programmer uses them and also writes the plan for completing the form. Notably, the programmer never writes code to receive messages.

Listing 6: Some Orpheus snippets.

```
@offer_plan[atomic]
+!send_offer
  :  enabled(offer(out, out, out)[receiver(out)])
  <- !complete(offer(Id, Item,
       Price)[receiver(Buyer)]);
     !attempt(offer(Id, Item,
          Price)[receiver(Buyer)]).

@shipment_plan[atomic]
+!send_shipment(Id, Item, Price, Buyer)
  :  enabled(shipment(Id, Item, Price,
       out)[receiver(Buyer)])
  <- !complete(shipment(Id, Item, Price,
       Status)[receiver(Buyer)]);
     !attempt(shipment(Id, Item, Price,
          Status)[receiver(Buyer)]).

+!complete(offer(Id, Item,
     Price)[receiver(Buyer)])
  :  on_offer(Id, Item, Price) & buyer(Buyer)
  <- -on_offer(Id, Item, Price).

+!complete(shipment(Id, Item, Price,
     Status)[receiver(Buyer)])
  :  in_stock(Item) & condition(Status) &
       buyer(Buyer)
  <- -in_stock(Item).
```

Orpheus abstracts away the maintenance of the local state and presents an interface to the programmer that supplies the enabled communicative acts. However, it does not support meaning-based reasoning—the programmer must encode when messages should be sent using low-level reasoning.

4.2 Base Event Adapter

Every time a message m with parameters $\vec{p}$ is sent or received, a belief for the corresponding base event b is asserted with its timestamp t as the current system time. C_1 gives the corresponding rule pattern, whose instance the tooling generates for every message and corresponding base event pair $(m(\vec{p}), b(\vec{p}, t))$. We explain the goal *update* in Section 4.4.

C_1 $+m(\vec{p}) : system_time(Now) \texttt{<-} +b(\vec{p}, Now);\ !update(\vec{k})$.

4.3 Commitment Queries

To support commitment queries, we give abstract Jason rules of the form *head* :- *body*. The rules are substantially more modular than in the previous semantics [15], which facilitates comprehension and enhances confidence that they capture intuitions correctly.

We treat all expressions of type Expr in Table 1, e.g., $X \sqcap Y$, $X \sqcup Y$, and so on, uniformly as events. $[\![X]\!]$ refers to the predicate for event X. For a base event E with attributes $\vec{a}$ and timestamp t, $[\![E]\!]$ is simply $E(\vec{a}, t)$ and its instances are asserted beliefs. For example, the predicate for *offer* is offer(Seller, Buyer, Id, Item, Price, Otime). The rules below lift $[\![]\!]$ to all events.

Below, E, F, and G are base or commitment life events; L is a life event; more generally, X and Y are events; $\vec{a}_X$ and t_X refer to the attributes and timestamp of X, respectively; t_p stands for a globally unique timestamp name in every application of the rules in which it appears. $[\![X]\!]_t^{\vec{a}}$ means that $[\![X]\!]$'s attributes and timestamp are $\vec{a}$ and t, respectively (omitted where obvious from the rule).

C_2 says that an instance of $[\![E[c, \infty]]\!]$ is an instance of E that has occurred at or after c. C_3 is similar.

C_2 $[\![E[c, \infty]]\!]$:- $[\![E]\!]$ & $c \leqslant t_E$.

C_3 $[\![E[0, d]]\!]$:- $[\![E]\!]$ & $t_E < d$.

A compiler uses the abstract Jason to produce actual Jason. Thus, for example, when the compiler encounters the expression offer[0, 5], it will map it to a unique name such as offerPred1 and generate the Jason rule in Listing 7.

Listing 7: Compiler-generated Jason from applying C_3.

```
offerPred1(Seller, Buyer, Id, Item, Price, Otime)
    :- offer(Seller, Buyer, Id, Item, Price,
    Otime) & Otime < 5 .
```

C_4 says that an instance of $X \sqcap Y$ represents correlated instances of X and Y and whose timestamp value is the max of their timestamps. Further, the set of attributes of the instance is the union of the attributes in the X and the Y instances.

C_4 $[\![X \sqcap Y]\!]^{\vec{a}_X \cup \vec{a}_Y}_{t_p}$:- $[\![X]\!]$ & $[\![Y]\!]$ & $.max([t_X, t_Y], t_p)$.

Suppose the compiler encountered the expression offer[0,5] ⊓ accept[0,6]. Listing 8 gives the kind of actual Jason code generated.

Listing 8: Compiler-generated Jason from applying C_4.

```
andPred3(Seller, Buyer, Id, Item, Price, T1) :-
  //offerPred1 as described in Listing 7
  offerPred1(Seller,Buyer,Id,Item,Price,Otime) &
  //Assume a rule for accept[0,6] from applying C3
  acceptPred2(Seller,Buyer,Id,Item,Price,Atime) &
  .max([Otime,Atime],T1).
```

C_6 says that an instance of $E[F+c, \infty]$ is an instance of E that has occurred no earlier than c time units after the correlated F instance. C_8 says that an instance $E[0, G+d]$ is an instance of E such that if the correlated G instance has occurred, then the E should have occurred before d units after the G's occurrence. The rest of the rules in C_5–C_{10} are straightforward applications of C_4.

C_5 $[\![E[c,d]]\!]$:- $[\![E[c,\infty] \sqcap E[0,d]]\!]$.

C_6 $[\![E[F+c,\infty]]\!]^{\vec{a}_E}_{t_E}$:- $[\![E]\!]$ & $[\![F]\!]$ & $t_F + c \leqslant t_E$.

C_7 $[\![E[F+c,d]]\!]$:- $[\![E[F+c,\infty] \sqcap E[0,d]]\!]$.

C_8 $[\![E[0,G+d]]\!]^{\vec{a}_E}_{t_E}$:- $[\![E]\!]$ & $(not\ [\![G]\!] \mid ([\![G]\!]\ \&\ t_E < t_G + d))$.

C_9 $[\![E[c,G+d]]\!]$:- $[\![E[c,\infty] \sqcap E[0,G+d]]\!]$.

C_{10} $[\![E[F+c,G+d]]\!]$:- $[\![E[F+c,\infty] \sqcap E[0,G+d]]\!]$.

C_{11} says that an instance of $X \sqcup Y$ is either an X instance or a Y instance. If correlated X and Y instances have both occurred, then the timestamp is the min of the two. To avoid unbound attributes in the $X \sqcup Y$ instance, the set of its attributes is the intersection of the attributes of the X instance and the Y instance. C_{12} is straightforward.

C_{11} $[\![X \sqcup Y]\!]^{\vec{a}_X \cap \vec{a}_Y}_{t_p}$:- $([\![X]\!]\ \&\ [\![Y]\!]\ \&\ .min([t_X, t_Y], t_p) \mid$
$([\![X]\!]\ \&\ not\ Y\ \&\ t_p = t_X) \mid$
$([\![Y]\!]\ \&\ not\ X\ \&\ t_p = t_Y)$.

C_{12} $[\![X \text{ where } \varphi]\!]$:- $[\![X]\!]$ & φ.

Let commitment(x, y, c, r, u) be a specification with debtor x, creditor y, and create, detach, and discharge expressions c, r, and u, respectively. Below, we write commitment(c, r, u) since the debtor and creditor are the same throughout.

C_{13}–C_{15} give the rules for some of the commitment life events of interest. For commitment(c, r, u), the created instances are the c instances; detached instances represent correlated created and r instances; and discharged instances represent correlated created and u instances. Notice that a commitment may be detached even if it has been discharged. In coming up with the rules, we are guided by flexibility and simplicity.

C_{13} $[\![created(c,r,u)]\!]$:- $[\![c]\!]$.

C_{14} $[\![detached(c,r,u)]\!]$:- $[\![created(c,r,u) \sqcap r]\!]$.

C_{15} $[\![discharged(c,r,u)]\!]$:- $[\![created(c,r,u) \sqcap u]\!]$.

Formulating rules in Jason for computing expired and violated instances of commitments require the notion of failed events. C_{16} says that an instance of E fails to occur at or after c if it occurs before c. C_{17} says that an instance of E fails to occur before d either if it occurs at or after d or it does not occur at all. In both cases, the timestamp of failure is d. C_{21} says that an instance of E fails to occur before $t_G + d$ if either E occurs at or after $t_G + d$ or E does not occur at all. In both cases, the timestamp of failure is $t_G + d$. The rest of the rules in C_{16}–C_{23} are straightforward.

C_{16} $[\![\overline{E[c,\infty]}]\!]$:- $[\![E[0,c]]\!]$.

C_{17} $[\![\overline{E[0,d]}]\!]_{t_p}$:- $[\![E[d,\infty]]\!] \mid not\ [\![E]\!])\ \&\ t_p = d$.

C_{18} $[\![\overline{E[c,d]}]\!]$:- $[\![\overline{E[c,\infty]} \sqcup \overline{E[0,d]}]\!]$.

C_{19} $[\![\overline{E[F+c,\infty]}]\!]$:- $[\![E[0,F+c]]\!]$.

C_{20} $[\![\overline{E[F+c,d]}]\!]$:- $[\![\overline{E[F+c,\infty]} \sqcup \overline{E[0,d]}]\!]$.

C_{21} $[\![\overline{E[0,G+d]}]\!]^{\vec{a}_E}_{t_p}$:- $[\![G]\!]\ \&\ ([\![E[G+d,\infty]]\!] \mid not\ [\![E]\!])\ \&$
$t_p = t_G + d$.

C_{22} $[\![\overline{E[c,G+d]}]\!]$:- $[\![\overline{E[c,\infty]} \sqcup \overline{E[0,G+d]}]\!]$.

C_{23} $[\![\overline{E[F+c,G+d]}]\!]$:- $[\![\overline{E[F+c,\infty]} \sqcup \overline{E[0,G+d]}]\!]$.

C_{24}—C_{26} apply De Morgan's laws to extend failure.

C_{24} $[\![\overline{X \sqcap Y}]\!]$:- $[\![\overline{X} \sqcup \overline{Y}]\!]$.

C_{25} $[\![\overline{X \sqcup Y}]\!]$:- $[\![\overline{X} \sqcap \overline{Y}]\!]$.

C_{26} $[\![\overline{X \text{ where } \varphi}]\!]$:- $[\![\overline{X} \sqcup (X \text{ where } not\ \varphi)]\!]$.

C_{27} says that an instance of $X \ominus Y$ is an instance of X such that the correlated Y has failed to occur. Its timestamp is the max of the two. C_{28} says that an instance of the failure of $X \ominus Y$ is either an instance of the failure of X or an instance of Y.

C_{27} $[\![X \ominus Y]\!]^{\vec{a}_X}_{t_p}$:- $[\![X]\!]\ \&\ [\![\overline{Y}]\!]\ \&\ .max([t_X, t_{\overline{Y}}], t_p)$.

C_{28} $[\![\overline{X \ominus Y}]\!]$:- $[\![\overline{X} \sqcup Y]\!]$.

C_{29}–C_{30} compute expired (failed to detach) and violated (failed to discharge) instances.

C_{29} $[\![expired(c,r,u)]\!]$:- $[\![created(c,r,u) \ominus r]\!]$.

C_{30} $[\![violated(c,r,u)]\!]$:- $[\![detached(c,r,u) \ominus u]\!]$.

Often, we are interested in life events that have occurred, that is, their timestamp is no later than the current time, as C_{31} captures.

C_{31} $[\![nowL]\!]$:- $[\![L]\!]\ \&\ t_L \leqslant Now\ \&\ system_time(Now)$.

4.4 Commitment Materializer

To materialize commitment events as beliefs, we assert an *update* commitment events goal every time an agent asserts a base event (as described above). Any base event affects commitments that are relevant to some subset of enactments, as identified by the bindings of its key attributes. Therefore, for efficiency, the update goal is parameterized by key attributes $\vec{k}$ that are common to all base events and are therefore guaranteed to occur in every life event predicate. C_1 triggers the update ($\vec{k}$ is the set of key attributes common to all base events, therefore $\vec{k} \subseteq \vec{p}$ in C_1).

C_{32} gives the abstract Jason plan for materializing commitment events; $[\![ev_nowL]\!]$ is a predicate with the same attributes and timestamp as $[\![nowL]\!]$. The plan for the update goal consists of asserting a belief corresponding to a life event if it is an instance of the life event predicate but not yet asserted. Assume that the life event predicates are $[\![L_1]\!]$,…, $[\![L_n]\!]$.

C_{32} $+!update(\vec{k})$ <- if $([\![nowL_1]\!]\ \&\ not\ [\![ev_nowL_1]\!])$
$\{ +[\![ev_nowL_1]\!]; \}$

...
if ([[$nowL_n$]] & *not* [[ev_nowL_n]])
{ +[[ev_nowL_n]]; }.

Agent programmers do not need to know either the abstract Jason rules (C_1–C_{32}) or the generated Jason rules. Their API consists of the predicates [[L]], [[$nowL$]], and [[ev_nowL]], where L is a lifecycle event.

5 IMPLEMENTING FLEXIBLE AGENTS

We now give examples of how Azorus agents can reason about commitments to flexibly enact protocols.

5.1 With Commitments as Queries

Azorus offers a set of queries for each commitment as a module (see Figure 3). These queries can be used for driving the choices of the enabled messages computed by the protocol adapter module.

Listing 9: Commitments as queries in Azorus.

```
+!handle_form([shipment(Id, Item, Price,
    out)[receiver(Buyer)]|_])
 :  in_stock(Item) &
     now_detached_OfferCom(Seller, Buyer, Id,
         Item, Price, Bank, Payment, Timestamp)
 <- !send_shipment(Id, Item, Price, Buyer).

+!handle_form([shipment(Id, Item, Price,
    out)[receiver(Buyer)]|_])
 :  not in_stock(Item) &
     now_detached_OfferCom(Seller, Buyer, Id,
         Item, Price, Bank, Payment, Timestamp)
 <- !send_refund(Id, Item, Payment, Bank).

+!handle_form([refund(Id, Item, Payment, out,
    out)[receiver(Bank)]|_])
 : now_detached_RefundCom(Seller, Buyer, Id,
     Item, Price, Bank, Payment, Timestamp)
 <- !send_refund(Id, Item, Payment, Bank).
```

A common reasoning pattern is for an agent to discharge a commitment if it is detached. The first plan in Listing 9 embodies this pattern. The seller executes the goal send_shipment if the Item is in stock and the commitment `OfferCom` is detached, that is, the *shipment* occurs if the *transfer* has been done in a timely manner.

Otherwise, by the second plan, if the Item is not in stock but `OfferCom` is detached, the goal send_refund is executed. The plan for send_shipment is as in Listing 6 and the plan for send_refund is analogous. The last plan is for when the commitment `OfferCom` is violated (because shipping does not occur by the deadline); again, the goal send_refund is executed. Both plans intend *refund*; however, the second does it simply on the basis of the detachment of `OfferCom` whereas the last plan does it upon its violation.

5.2 With Commitments as Events

Besides the set of queries for each commitment, an agent program can exploit the commitment materializer (see Figure 3), which asserts beliefs corresponding to the occurrence of commitment events. These events can be exploited to support reasoning.

Listing 10: Commitments as events in Azorus.

```
+!offer : on_offer(Id, Item, Price)
 <- !send_offer.
+ev_now_detached_OfferCom(Seller, Buyer, Id,
    Item, Price, Bank, Payment, Timestamp)
 : in_stock(Item)
 <- !send_shipment(Id, Item, Price, Buyer).
+ev_now_detached_RefundCom(Seller, Buyer, Id,
    Item, Price, Bank, Payment, Timestamp)
  <- !send_refund(Id, Item, Payment, Bank).
```

For example, in Listing 10, the agent SELLER sends an offer to a potential BUYER. Upon a timely *transfer*, the commitment `OfferCom` is detached and, by exploiting the rule C_{32}, the event +ev_now_detached_OfferCom is produced by adding the corresponding belief to the SELLER agent's belief base. This triggers the plan for dealing with such an event: the agent performs the *shipment*. Analogously, in the case the event +ev_now_detached_RefundCom is generated (the *shipment* does not occur within the deadline) the agent performs the *refund*.

5.3 Timestamp-Based Reasoning

Recall that for a life event L, an instance of [[$nowL$]] is an [[L]] instance that has actually occurred (that is, with current time as the reference point). In general, any time instant, in the past or the future, could be the point of reference.

Suppose the SELLER agent, as a matter of managing its commitments, wanted to discharge the `OfferCom` commitments that will be violated within 10 time units from now (unless, of course, *shipment* is sent). Listing 11 shows how to accomplish this using a future time instant as the point of reference.

Listing 11: Deadline-based reasoning.

```
+!handle_form([shipment(Id, Item, Price,
    out)[receiver(Buyer)]|_])
 :  in_stock(Item) & violated_OfferCom(Id,..., T)
     & system_time(Now) & T <= Now + 10
 <- !send_shipment(Id, Item, Price, Buyer).
```

6 CONCEPTUAL EVALUATION

Let's summarize what must be manually specified or coded and what Azorus provides as abstractions. The commitment specification, the protocol, and an agent's internal reasoning must be manually specified. Azorus supports the coding of internal reasoning by providing abstractions that enable reasoning about commitments and performing communicative acts that are legal from the standpoint of the protocol.

In virtually any multiparty application, commitments and protocols are domain concepts; there is no avoiding reasoning about them. Specifying them cleanly opens up the possibility of building a tool-supported methodology around them, including verification [19, 34, 36, 42] and programming abstractions (as we do in Azorus), and other productivity tools such as IDEs. Not specifying them means architects and programmers must figure out the possible enactments and encode the reasoning using low-level abstractions. Naturally, such code is likely to be ad hoc, complex, error-prone,

and hard-to-maintain even for simple MAS involving rigid interactions between two parties, let alone MAS with more than two parties and flexible engagements (such as the *Ebusiness* protocol which has 658 enactments).

Even with protocol support, as Orpheus provides, the programmer would still have to encode reasoning about commitments manually. Consider Listing 12, which shows a SELLER's code snippet. It says that the agent sends an enabled *shipment* if *transfer* has occurred. Since *transfer* is required for the detach of `OfferCom`, this seems to capture the intent behind the first plan in Listing 9. It does not though because it misses the time-related reasoning. That is, *transfer* could have happened late enough that `OfferCom` would have expired, in which case the agent may not want to send *shipment*.

Listing 12: No support for commitment reasoning can lead to errors by underspecification.

```
+!handle_form([shipment(Id, Item, Price,
    out)[receiver(Buyer)]|_])
 :  in_stock(Item) &
    transfer(Id, Price, _, Payment)
 <- !send_shipment(Id, Item, Price, Buyer).
```

Commitments without protocol support can also go wrong. In Listing 13, *shipment* and *refund* (which should be mutually exclusive) are triggered solely by their respective commitment detachments. If *transfer* takes too long,`OfferCom` will be violated and `RefundCom` detached. Since protocol constraints are not enforced, both *shipment* and *refund* could be sent, violating mutual exclusion.

Listing 13: No support for protocols can lead to erroneous communication.

```
+transfer(Id, Price, Payment)
 :  in_stock(Item) &
     now_detached_OfferCom(Seller, Buyer, Id,
          Item, Price, Bank, Payment, Timestamp)
 <- .send(Buyer, tell, shipment(Id, Item, Price,
     done)).

+ev_now_detached_RefundCom(Seller, Buyer, Id,
    Item, Price, Bank, Payment, Timestamp)
 : Amount=Payment
 <- .send(Bank, tell, refund(Id, Item, Payment,
     Amount, done)).
```

Without protocol support, in Jason, programmers typically use *tell* for every message. We might as well drop KQML support (and FIPA ACL [21] support from JADE) and instead consider the protocol messages themselves as first-class communicative acts and express their meaning via social abstractions such as commitments (see Singh's essay in [11]), as Azorus does.

7 DISCUSSION

Azorus' novelty is twofold. One, it shows how protocols as operational abstractions and commitments as high-level abstractions can be leveraged in a multiagent programming model. Two, it introduces higher-level communication abstractions to Jason, a popular BDI-based programming model. Azorus exploits practical, expressive languages for commitments and protocols and the Azorus adapter is the first careful working out of the interplay between protocol enactment and commitment reasoning. Its significance is also two-fold. One, Azorus simplifies the engineering of flexible, decentralized MAS. Two, it brings goals, commitments, and protocols—all of which represent autonomy—into a single programming model. Below, we discuss concerns that require further investigation.

Specifying Commitments. Different commitment specifications could be overlaid on the same protocol. The specification in Listing 3 is "direct" in that it gives the meaning of both *offer* and *accept* as an exchange of *shipment* and *transfer*. An alternative commitment specification could have a "waterfall" flavor: *offer* means that if *accept*, then *shipment*, and *accept* means that if *shipment*, then *transfer*. The possibility of alternative commitment specifications motivates characterizing the specifications in terms of properties and stakeholder requirements that they satisfy.

Implementing Agents. Consider buyer and seller agents implemented such that the seller waited for the buyer to detach `OfferCom` by effecting *transfer* and the buyer waited for the seller to detach `AcceptCom` by doing *shipment*. Naturally, in every enactment, the agents end up deadlocked (even though the *Ebusiness* protocol itself is live). Such deadlocked enactments are not necessarily problematic: they arise from agents exercising their autonomy by not sending messages.

Notions such as trust and other business requirements can facilitate progress. For example, if a buyer trusts the seller or if the monetary amount involved is small, the buyer may be willing to detach `OfferCom` from the seller, effectively moving first in the exchange. What we need are novel methodologies for implementing agents that take into account the various contextual assumptions and business requirements.

There has been some work on methodologies for specifying commitments and implementing agents. Winikoff [38] and Yolum [42] give methods for designing and checking specifications for properties related to progress, consistency, and flexibility. Marengo et al. [28] and Günay et al. [25] relate commitments to notions of safety and control. Some work has studied relationships between goals (as representation of requirements) and commitments [13, 29, 37]. Yolum and Singh study commitments from the point of view of concession (taking a risk by moving first) [44]. To ensure monitorability of commitments, Azorus could be combined with either organizations and shared artifacts [18] or use alignment-producing techniques [16, 27]. Finally, Langshaw [35] is an even higher-level protocol language than BSPL and may simplify combining commitments and protocols. These works can serve as a starting point for methodologies for building flexible, decentralized MAS, a direction that should yield rich dividends.

8 REPRODUCIBILITY

The entire Azorus codebase and examples as well as other related tools are available online at https://gitlab.com/masr.

ACKNOWLEDGMENTS

MPS thanks the NSF (grant IIS-1908374) for partial support.

REFERENCES

[1] Matteo Baldoni, Cristina Baroglio, Federico Capuzzimati, and Roberto Micalizio. 2018. Commitment-based Agent Interaction in JaCaMo+. *Fundamenta Informaticae* 159, 1-2 (2018), 1–33. https://doi.org/10.3233/FI-2018-1656

[2] Matteo Baldoni, Cristina Baroglio, Elisa Marengo, and Viviana Patti. 2013. Constitutive and Regulative Specifications of Commitment Protocols: A Decoupled Approach. *ACM Transactions on Intelligent Systems and Technology (TIST)* 4, 2, Article 22 (March 2013), 25 pages. https://doi.org/10.1145/2438653.2438657

[3] Matteo Baldoni, Cristina Baroglio, Alberto Martelli, and Viviana Patti. 2006. A Priori Conformance Verification for Guaranteeing Interoperability in Open Environments. In *Proceedings of the 4th International Conference on Service-Oriented Computing (ICSOC) (Lecture Notes in Computer Science, Vol. 4294)*. Springer, Chicago, 339–351. https://doi.org/10.1007/11948148_28

[4] Matteo Baldoni, Samuel H. Christie V, Munindar P. Singh, and Amit K. Chopra. 2025. Orpheus: Engineering Multiagent Systems via Communicating Agents. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI)*. AAAI, Philadelphia, 1–9.

[5] Bernhard Bauer, Jörg P. Müller, and James Odell. 2000. An Extension of UML by Protocols for Multiagent Interaction an existing Multi-Agent Planning System. In *Proceedings of the 4th International Conference on Multiagent Systems (ICMAS)*. IEEE Computer Society, Boston, 207–214. https://doi.org/10.1109/ICMAS.2000.858455

[6] Fabio Bellifemine, Giovanni Caire, and Dominic Greenwood. 2007. *Developing Multi-Agent Systems with JADE*. Wiley, Chichester, UK. https://doi.org/10.1002/9780470058411

[7] Olivier Boissier, Rafael H. Bordini, Jomi Fred Hübner, Alessandro Ricci, and Andrea Santi. 2013. Multi-agent oriented programming with JaCaMo. *Science of Computer Programming* 78, 6 (June 2013), 747–761. https://doi.org/10.1016/j.scico.2011.10.004

[8] Rafael H. Bordini and Jomi Fred Hübner. 2010. Semantics for the Jason Variant of AgentSpeak (Plan Failure and some Internal Actions). In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI) (Frontiers in Artificial Intelligence and Applications, Vol. 215)*. IOS Press, Lisbon, 635–640. https://doi.org/10.3233/978-1-60750-606-5-635

[9] Hans Chalupsky, Tim Finin, Rich Fritzson, Don McKay, Stu Shapiro, and Gio Wiederhold. 1992. *An Overview of KQML: A Knowledge Query and Manipulation Language*. TR. University of Maryland Computer Science Department, Baltimore.

[10] Federico Chesani, Paola Mello, Marco Montali, and Paolo Torroni. 2013. Representing and Monitoring Social Commitments using the Event Calculus. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)* 27, 1 (July 2013), 85–130. https://doi.org/10.1007/s10458-012-9202-0

[11] Amit K. Chopra, Alexander Artikis, Jamal Bentahar, Marco Colombetti, Frank Dignum, Nicoletta Fornara, Andrew J. I. Jones, Munindar P. Singh, and Pınar Yolum. 2013. Research Directions in Agent Communication. *ACM Transactions on Intelligent Systems and Technology (TIST)* 42, 2, Article 20 (March 2013), 23 pages. https://doi.org/10.1145/2438653.2438655

[12] Amit K. Chopra, Samuel H. Christie V, and Munindar P. Singh. 2020. An Evaluation of Communication Protocol Languages for Engineering Multiagent Systems. *Journal of Artificial Intelligence Research (JAIR)* 69 (Dec. 2020), 1351–1393. https://doi.org/10.1613/jair.1.12212

[13] Amit K. Chopra, Fabiano Dalpiaz, F. Başak Aydemir, Paolo Giorgini, John Mylopoulos, and Munindar P. Singh. 2014. Protos: Foundations for Engineering Innovative Sociotechnical Systems. In *Proceedings of the 22nd IEEE International Requirements Engineering Conference (RE)*. IEEE Computer Society, Karlskrona, Sweden, 53–62. https://doi.org/10.1109/RE.2014.6912247

[14] Amit K. Chopra and Munindar P. Singh. 2006. Contextualizing Commitment Protocols. In *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems*. ACM Press, Hakodate, Japan, 1345–1352. https://doi.org/10.1145/1160633.1160884

[15] Amit K. Chopra and Munindar P. Singh. 2015. Cupid: Commitments in Relational Algebra. In *Proceedings of the 29th Conference on Artificial Intelligence (AAAI)*. AAAI Press, Austin, Texas, 2052–2059. https://doi.org/10.1609/aaai.v29i1.9443

[16] Amit K. Chopra and Munindar P. Singh. 2015. Generalized Commitment Alignment. In *Proceedings of the 14th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. IFAAMAS, Istanbul, 453–461. https://doi.org/10.5555/2772879.2772938

[17] Samuel H. Christie V, Munindar P. Singh, and Amit K. Chopra. 2023. Kiko: Programming Agents to Enact Interaction Protocols. In *Proceedings of the 22nd International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. IFAAMAS, London, 1154–1163. https://doi.org/10.5555/3545946.3598758

[18] Mehdi Dastani, Leendert W. N. van der Torre, and Neil Yorke-Smith. 2017. Commitments and Interaction Norms in Organisations. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)* 31, 2 (March 2017), 207–249. https://doi.org/10.1007/s10458-015-9321-5

[19] Mohamed El Menshawy, Jamal Bentahar, Hongyang Qu, and Rachida Dssouli. 2011. On the Verification of Social Commitments and Time. In *Proceedings of the 10th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. IFAAMAS, Taipei, 483–490. https://doi.org/10.5555/2031678.2031687

[20] Angelo Ferrando, Michael Winikoff, Stephen Cranefield, Frank Dignum, and Viviana Mascardi. 2019. On Enactability of Agent Interaction Protocols: Towards a Unified Approach. In *Proceedings of the 7th International Workshop on Engineering Multi-Agent Systems (EMAS) (Lecture Notes in Computer Science, Vol. 12058)*. Springer, Montréal, 43–64. https://doi.org/10.1007/978-3-030-51417-4_3

[21] FIPA. 2002. FIPA Agent Communication Language Specifications. FIPA: The Foundation for Intelligent Physical Agents, http://www.fipa.org/repository/aclspecs.html.

[22] FIPA. 2003. FIPA Interaction Protocol Specifications. http://www.fipa.org/repository/ips.html. FIPA: The Foundation for Intelligent Physical Agents. Accessed 2024-11-24.

[23] Nicoletta Fornara and Marco Colombetti. 2002. Operational Specification of a Commitment-Based Agent Communication Language. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. ACM Press, Melbourne, 535–542. https://doi.org/10.1145/544862.544868

[24] Stéphane Galland, Sebastian Rodriguez, and Nicolas Gaud. 2020. Run-time Environment for the SARL Agent-Programming Language: The Example of the Janus platform. *Future Generation Computer Systems* 107 (June 2020), 1105–1115. https://doi.org/10.1016/j.future.2017.10.020

[25] Akın Günay, Michael Winikoff, and Pınar Yolum. 2015. Dynamically Generated Commitment Protocols in Open Systems. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)* 29, 2 (March 2015), 192–229. https://doi.org/10.1007/s10458-014-9251-7

[26] Jomi Fred Hübner, Jaime Simão Sichman, and Olivier Boissier. 2007. Developing Organised Multiagent Systems using the MOISE$^+$ Model: Programming Issues at the System and Agent Levels. *International Journal of Agent-Oriented Software Engineering* 1, 3/4 (2007), 370–395. https://doi.org/10.1504/IJAOSE.2007.016266

[27] Thomas Christopher King, Akın Günay, Amit K. Chopra, and Munindar P. Singh. 2017. Tosca: Operationalizing Commitments over Information Protocols. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. IJCAI, Melbourne, 256–264. https://doi.org/10.24963/ijcai.2017/37

[28] Elisa Marengo, Matteo Baldoni, Amit K. Chopra, Cristina Baroglio, Viviana Patti, and Munindar P. Singh. 2011. Commitments with Regulations: Reasoning about Safety and Control in Regula. In *Proceedings of the 10th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. IFAAMAS, Taipei, 467–474. https://doi.org/10.5555/2031678.2031684

[29] Felipe Meneguzzi, Mauricio C. Magnaguagno, Munindar P. Singh, Pankaj R. Telang, and Neil Yorke-Smith. 2018. GoCo: Planning Expressive Commitment Protocols. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)* 32, 4 (July 2018), 459–502. https://doi.org/10.1007/s10458-018-9385-0

[30] Jaime Simão Sichman, Rosaria Conte, Yves Demazeau, and Cristiano Castelfranchi. 1994. A Social Reasoning Mechanism Based on Dependence Networks. In *Proceedings of the 11th European Conference on Artificial Intelligence*. John Wiley and Sons, Amsterdam, 188–192.

[31] Munindar P. Singh. 1998. Agent Communication Languages: Rethinking the Principles. *IEEE Computer* 31, 12 (Dec. 1998), 40–47. https://doi.org/10.1109/2.735849

[32] Munindar P. Singh. 2011. Information-Driven Interaction-Oriented Programming: BSPL, the Blindingly Simple Protocol Language. In *Proceedings of the 10th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. IFAAMAS, Taipei, 491–498. https://doi.org/10.5555/2031678.2031687

[33] Munindar P. Singh. 2012. Semantics and Verification of Information-Based Protocols. In *Proceedings of the 11th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. IFAAMAS, Valencia, Spain, 1149–1156. https://doi.org/10.5555/2343776.2343861

[34] Munindar P. Singh and Samuel H. Christie V. 2021. Tango: Declarative Semantics for Multiagent Communication Protocols. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*. IJCAI, Online, 391–397. https://doi.org/10.24963/ijcai.2021/55

[35] Munindar P. Singh, Samuel H. Christie V, and Amit K. Chopra. 2024. Langshaw: Declarative Interaction Protocols Based on Sayso and Conflict. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*. IJCAI, Jeju, Korea, 202–210. https://doi.org/10.24963/ijcai.2024/23

[36] Pankaj R. Telang and Munindar P. Singh. 2012. Specifying and Verifying Cross-Organizational Business Models: An Agent-Oriented Approach. *IEEE Transactions on Services Computing (TSC)* 5, 3 (July 2012), 305–318. https://doi.org/10.1109/TSC.2011.4 Appendix pages 1–5.

[37] Pankaj R. Telang, Munindar P. Singh, and Neil Yorke-Smith. 2019. A Coupled Operational Semantics for Goals and Commitments. *Journal of Artificial Intelligence Research (JAIR)* 65 (May 2019), 31–85. https://doi.org/10.1613/jair.1.11494

[38] Michael Winikoff. 2006. Designing Commitment-Based Agent Interactions. In *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology*. IEEE Computer Society, Hong Kong, 363–370. https://doi.org/10.1109/IAT.2006.53

[39] Michael Winikoff. 2007. Implementing Commitment-Based Interactions. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. IFAAMAS, Honolulu, 868–875. https://doi.org/10.1145/1329125.1329283

[40] Michael Winikoff. 2012. Challenges and Directions for Engineering Multi-Agent Systems. *CoRR* abs/1209.1428 (2012), 12 pages.

[41] Michael Winikoff, Nitin Yadav, and Lin Padgham. 2018. A New Hierarchical Agent Protocol Notation. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)* 32, 1 (Jan. 2018), 59–133. https://doi.org/10.1007/s10458-017-9373-9

[42] Pınar Yolum. 2007. Design Time Analysis of Multiagent Protocols. *Data and Knowledge Engineering* 63, 1 (Oct. 2007), 137–154. https://doi.org/10.1016/j.datak.2006.12.001

[43] Pınar Yolum and Munindar P. Singh. 2002. Flexible Protocol Specification and Execution: Applying Event Calculus Planning using Commitments. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. ACM Press, Bologna, 527–534. https://doi.org/10.1145/544862.544867

[44] Pınar Yolum and Munindar P. Singh. 2007. Enacting Protocols by Commitment Concession. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. IFAAMAS, Honolulu, 116–123. https://doi.org/10.1145/1329125.1329158