Stellar: A Programming Model for Developing Protocol-Compliant Agents

Akın Günay and Amit K. Chopra

Lancaster University
Lancaster, LA1 4WA, UK
{a.gunay, amit.chopra}@lancaster.ac.uk

Abstract. An interaction protocol captures the rules of encounter in a multiagent system. Development of agents that comply with protocols is a central challenge of multiagent systems. Our contribution in this chapter is a programming model, Stellar, that simplifies development of agents compliant with information protocols specified in BSPL. A significant distinction of Stellar from similar approaches is that it does not rely upon extracting control flow structures from protocol specifications to ensure compliance. Instead, Stellar provides a set of fundamental operations to programmers for producing viable messages according to the correct flow of information between agents as specified by a protocol, enabling flexible design and implementation of protocol-compliant agents. Our main contributions are: (1) identification of a set of programming errors that commonly occur when developing agents for protocol-based multiagent system, (2) definition of Stellar's operations and a simple yet effective pattern to develop protocol-compliant agents that avoid the identified errors, and (3) demonstration of Stellar's effectiveness by presenting concrete agents in e-commerce and insurance policy domains.

1 Introduction

Interaction protocols capture the rules of encounter in multiagent systems by defining operational constraints on the occurrence and ordering of messages between agents. Effective interaction of agents in a multiagent system depends on their compliance with the system's protocol. However, development of protocol-compliant agents is challenging in practical settings where communication is asynchronous.

There are several approaches to specify and implement interaction protocols, such as HAPN [14], Scribble [5,15], BPMN in conjunction with BPEL [10], and business artifacts [6,9]. These approaches mainly use procedural control flow structures (e.g., sequencing, branching, etc.) to specify interactions of agents, whose implementations reflect the protocol's control flow to ensure compliance. This is mostly achieved by developing agents on top of rigid code skeletons that are extracted from the protocol specifications. As a result, protocol specifications and implementations of agents who enact them become tightly coupled. An imminent drawback of this approach is the lack of flexibility in agent design,

which is a critical limitation particularly in open multiagent systems, where independent parties implement their own agents according to their private business requirements and logic. Another technical drawback of this approach is the need for synchronization between agents to ensure correct ordering of messages, which is hard to achieve in asynchronous decentralized environments.

Several information-based protocol languages [11,13,3] have been proposed in the recent years to overcome limitations of the procedural protocol specification approaches. These languages specify protocols in a declarative way with respect to the correct flow of information between the agents, rather than specifying rigid messages sequences. Hence, information-based languages do not impose a control flow for implementing protocol-compliant agents. As a result, independent parties can design their own agents as they see fit according to their own requirements, as long as their agents emit messages complying with the protocol's flow of information. Consequently, information-based languages do not rely on synchronization and inherently support asynchronous and decentralized communication.

In this chapter we focus particularly on BSPL [11] which constitutes the base for all later information-based languages. Although BSPL provides a rich protocol specification language, it does not define a systematic methodology for developing protocol-compliant agents. Our contribution, namely Stellar, addresses this issue with a simple yet effective programming model. To this end, Stellar defines a set of fundamental operations and a software pattern over these operations that enables developers to build compliant agents. Hence, developers can focus on the business logic of their agents without worrying about compliance with protocols. Thanks to BSPL's declarative approach, Stellar does not rely on control flow structures (e.g., no code skeleton is created), which enables maximum flexibility when designing and implementing agents. Our main contributions are as follows. One, we identify common pitfalls of protocol-compliant agent development in decentralized multiagent systems. Two, we develop Stellar's programming model, describe its programming pattern, and define its operations. Three, we demonstrate Stellar's effectiveness by developing agents in e-commerce and insurance policy domains.

2 BSPL

In this section we provide an overview of BSPL to establish the necessary background. BSPL [11] is an information-based protocol specification language. The main difference of BSPL from procedural protocol specification approaches is its way of characterizing operational constraints with respect to causality and flow of information between agents. We explain BSPL's main features using an example purchase protocol that we present in Listing 1.

Listing 1: A BSPL protocol for purchase.

```
Purchase {
 roles B, S // buyer, seller
```

```
parameters out pID key, out item, out price, out result

B ↦ S: rfq [out pID, out item]
S ↦ B: quote [in pID, in item, out price]
B ↦ S: accept [in pID, in item, in price, out result]
B ↦ S: reject [in pID, in item, in price, out result]
}
```

A BSPL protocol is composed of a name, a set of roles, a set of public parameters, and a set of message schemas. BSPL is a declarative language and hence the ordering of message schemas in a protocol specification is irrelevant. The name of the protocol in Listing 1 is Purchase. It includes two roles, B and S corresponding to a buyer and a seller, respectively. Purchase has four public parameters pID, item, price, and result, which describe the protocol's interface, intuitively corresponding to the identifier of the protocol, an item to purchase, price of the item, and the outcome of the interaction, respectively. A protocol's enactment is complete when all of its public parameters are bound. BSPL protocols can be composed using their interfaces to build complex interactions. However, we do not consider composite protocols in this chapter for brevity. Each message schema in the form of $s \mapsto r : m[P]$ has a sender s and a receiver r role, a message name m, and a set of parameters P.

For instance, the name of the first message schema in Listing 1 is rfq, corresponding to a request for a quote, in which the sender is B, the receiver is S, and the parameters are pID and item. Instances of message schemas are relational tuples that represent the bindings of message parameters. For instance, Table 1 shows three instance of quote. In the rest of the chapter we use "message" to refer to both a message schema and message instance when there is no ambiguity.

Table 1: Instances of quote message.

pID	item	price
1	*book*	5
2	*bike*	10
3	*phone*	20

An *enactment* of a protocol corresponds to the set of messages that are exchanged between the agents with respect to a unique key. Each unique enactment of a protocol is identified by one or more key parameters. In our example the only key parameter is pID. Hence, each distinct enactment of Purchase must have a unique binding for pID. The uniqueness constraints of typical relational models apply to the bindings of keys in each message instance (i.e., no two instances of a message can have the same binding for a key), and each parameter across the messages (i.e., a parameter has the same binding for the same key in the instances of different messages).

Agents can enact multiple instances of a protocol concurrently. To this end, each agent keeps its own *local history*, which is the set of sent and received messages by the agent in all enactments. Table 2 shows an example local history of an agent who enacts the buyer role. In Table 2, there are four enactments of Purchase (i.e., one for each distinct binding of pID). Note that only the enactments in which pID is bound to 1 and 2 are complete. That is, all the public parameters of Purchase are bound in these two enactments. The local history of an agent is sufficient for the agent to carry out its interactions with other agents complying with a protocol. In other words, an agent does not need any information about the states of the other agents to interact with them. Hence, BSPL protocols can be enacted by agents in a fully decentralized way without referring to any global state.

Table 2: Local history of an agent enacting the buyer role.

(a) rfq

pID	item
1	*book*
2	*bike*
3	*phone*
4	*pen*

(b) quote

pID	item	price
1	*book*	5
2	*bike*	10
3	*phone*	20

(c) accept

pID	item	price	result
1	*book*	5	*OK*

(d) reject

pID	item	price	result
2	*bike*	10	*NOK*

Given an agent's local history, we say that a parameter's binding is *known* to the agent in an enactment of a protocol, if the agent's local history includes a message with a binding of the parameter for that particular enactment. Otherwise, we say that the parameter's binding is *unknown* to the agent in that particular enactment. For instance, according to the local history of the buyer in Table 2, binding of price is known (as 5) to the buyer for the enactment of Purchase where pID is bound to 1. This is due to the quote message that is received by the buyer for this enactment. However, the binding of price is unknown to the buyer for the enactment where pID is bound to 4, since there is no message in the buyer's local history with a binding of price for that enactment.

As we have stated earlier, the key idea of BSPL is to specify operational constraints of a protocol in terms of correct flow of information among agents, instead of using procedural control structures. BSPL models the flow of information in a protocol by adorning parameters with ⌜in⌝, ⌜out⌝, or ⌜nil⌝.

Parameters that are adorned ⌜in⌝ in a message correspond conceptually to the inputs of the message, whose bindings must be known to the sender before sending the message. For instance, the seller must know the bindings of pID

and item before sending a quote. Parameters that are adorned ⌜out⌝ correspond conceptually to the outputs of a message, whose bindings are produced by the sender when sending the message. For instance, the seller must produce the binding of price when sending a quote. Agents cannot violate integrity of information when sending messages. That is, a sender cannot change the known binding of a parameter when sending a message. As a consequence, if two or more message share the same ⌜out⌝ adorned parameter, only one of these messages can be sent in an enactment to ensure integrity, meaning that the messages that share ⌜out⌝ adorned parameters are mutually exclusive (e.g., accept and reject due to ⌜out⌝ adorned parameter outcome). Lastly, if a parameter is adorned ⌜nil⌝ in a message, the sender must not know the binding of the parameter and also must not produce a binding for the parameter when sending the message.

BSPL formalizes the correct flow of information in a protocol by defining *viability* of messages in an enactment. A message is viable for a sender in an enactment, if and only if (1) the sender knows the bindings of all the ⌜in⌝ adorned parameters of the message, and (2) there is no earlier message in the sender's local history that already binds any of ⌜out⌝ or ⌜nil⌝ adorned parameters of the message. Agents comply with a protocol, if they exchange only viable messages in an enactment of the protocol. For instance, considering the local history of the buyer in Table 2, the instance $(3, phone, 20, OK)$ of accept is viable for the buyer, since the bindings of all its ⌜in⌝ adorned parameters are known (due to the earlier rfq and quote messages that are exchanged in the enactment) and the binding of the ⌜out⌝ adorned result is unknown to the buyer for the enactment where pID is 3. Hence, the buyer can send this message by producing the binding of result, which is OK in our case. Similarly, the reject message instance $(3, phone, 10, NOK)$ is also viable.

On the other hand, there is no viable accept message for the enactment where pID is bound to 2, since the ⌜out⌝ adorned result is already bound to NOK in this enactment as a result of the prior reject message (i.e., the value of an ⌜out⌝ adorned parameter is known). Similarly, there is no viable reject message for the enactment where pID is bound to 1, since the ⌜out⌝ adorned result is already bound to OK in this enactment because of the prior accept message. For the enactment where pID is bound to 4, the buyer does not know the binding of price, which is adorned ⌜in⌝ in accept and reject messages. Hence, there are no viable accept or reject messages for this enactment.

3 Pitfalls of Developing Protocol-Compliant Agents

Development of a protocol-compliant agent for an information-based protocol is a challenging task due to factors such as concurrent enactments of the protocol and asynchronous communication between agents. Without a well-defined methodology, developers may easily fail to identify subtle details of a protocol and implement non-compliant agents. In this section we identify such potential pitfalls of agent development for information-based protocols using our Purchase example from the previous section. Although our example is specified in BSPL,

the issues that we discuss here are general and occur when developing agents for protocols that are specified in any language.

Let us start by examining some interactions between a protocol-compliant buyer and seller for our Purchase protocol. Figure 1 shows two such interactions. In both cases, the buyer first sends an rfq in which pID and item are bound to 1 and *book*, respectively. Then, the seller replies with a quote that binds price to 5. Finally, the buyer either sends an accept as in Figure 1(a) or a reject as in Figure 1(b) in response to the received quote. Now we identify several issues that induce non-compliant implementation of agents.


Fig. 1: Compliant interactions between the buyer (B) and seller (S).

Information Integrity: Protocol-compliant agents must ensure integrity of the exchanged information when interacting according to an information-based protocol. An agent may easily violate information integrity (maliciously or accidentally) by either creating information that does not exist or by altering known information. Figure 2(a) shows an interaction that corresponds to the former case, where the buyer sends an accept message to the seller without receiving a quote message by creating the binding of price as 3 even tough price is adorned ⌜in⌝ for accept. Figure 2(b) shows an interaction that corresponds to the latter case, where the buyer alters the binding of the item to *bike* when sending the accept message, which should actually be *book* as in the prior rfq message that she sent to the seller earlier.

In both cases, the integrity of the exchanged information is violated by the buyer leading to a non-compliant enactment of the protocol. These kind of mistakes occur especially when an agent is implemented for concurrently enacting multiple instances of a protocols. For instance, consider a buyer that interacts concurrently with multiple sellers to purchase the cheapest copy of a book. This normally can be achieved by executing multiple instances of the buyer agent concurrently (e.g., in separate threads), each handling a separate enactment of Purchase with a different seller. If each concurrent instance of the buyer code can be executed in complete isolation, the errors we identify cannot happen. However, in most realistic applications, instances of an agent cannot be fully isolated since they must access to some shared data. In our example, the instances of

buyer agent must share the price information they receive from different buyers. A developer may make a mistake when developing the buyer agent, which may cause the price that is received in one enactment to be used in another enactment as we demonstrate in Figure 2.


(a) Value Fabrication　　　　(b) Value Alteration

Fig. 2: Violation of information integrity.

Mutual Exclusion: Realistic information-based protocols usually involve mutually exclusive messages. For instance, when the buyer receives a quote message from the seller, she must either send an accept or reject message, but not both. Figure 3(a) shows violation of the mutual exclusion by the buyer, who sends first an accept message and then a reject message after receiving a quote message. Note that, in this example mutual exclusion is local to the buyer. That is, emission of the accept and reject messages are local choices of the buyer. Hence, violation of mutual exclusion can be avoided by ensuring the buyer's compliance with the protocol.

However, mutual exclusion may also be non-local [7]. Suppose that in an extended version of our purchase protocol the seller may cancel its quote by sending a cancel message, which binds result, between the quote and accept messages. Therefore, if there is a cancel message, there should not be an accept message and vice versa. However, these message are emitted by different agents (i.e., mutual exclusion is non-local) and violation of mutual exclusion may occur as Figure 3(b) shows even though the agents are protocol-compliant. In general, non-local mutual exclusion cannot be satisfied unless behaviors of agents are synchronized, which is costly and hard to achieve (if not impossible) in realistic systems. A protocol is called *safe* if it does have any enactments where two agents may bind the same parameter, and in this chapter we only consider safe protocols.

Concurrency: As we have discussed in information integrity issues, in many practical multiagent systems, agents concurrently enact multiple instances of protocols. For instance, in our purchase scenario, the buyer may concurrently send multiple quote requests to the seller for different items in different enactments. Besides, the buyer (and also seller) can interact with multiple sellers (buyers) concurrently in different enactments of the purchase protocol. In such

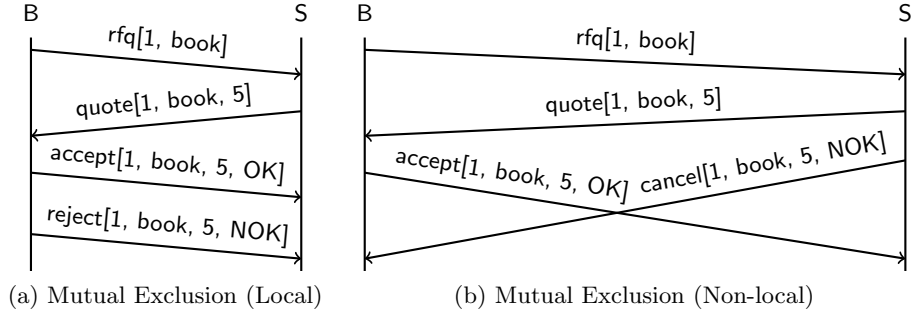(a) Mutual Exclusion (Local) (b) Mutual Exclusion (Non-local)

Fig. 3: Violation of mutual exclusion.

situations, in addition to the integrity issues that we have discussed, developers should also deal with interleaved asynchronous emission and reception of messages in different enactments. To this end, developers normally use multithreading mechanisms (i.e., each concurrent protocol instance is executed in a separate thread). However, this requires the use of complex synchronization mechanisms between the threads to properly handle interleaving messages from different enactments. For instance, the the seller must check if there are sufficient quantity of goods when making concurrent quotes. Achieving such synchronization is an error-prone task and if not done correctly may easily cause agents to act in a non-compliant manner or event stop operating due to deadlock issues.

4 Stellar

Stellar[1] is a programming model to develop protocol-compliant agents for BSPL protocols. Stellar eliminates the pitfalls of protocol-compliant agent development that we have discussed in Section 3. To this end, Stellar provides a a set of well-defined operations around a software pattern for developing agents to enact roles in a protocol. If an agent's interactions are implemented using the operations of Stellar following its software pattern, the developed agent is guaranteed to be protocol-compliant.

We implement Stellar as a Java framework. The implementation provides, (1) a code generation tool, which, given a BSPL protocol, automatically generates a protocol-specific code library that includes classes to represent the roles, messages, and parameters of the given protocol, and (2) a static core library that provides the operations to apply Stellar's software patterns.

The workflow for developing a protocol-compliant agent using Stellar is as follows. First, a BSPL protocol for which an agent is intended to be developed is specified. Second, the protocol is provided to Stellar's code generation tool, which automatically generates a library of classes corresponding to the roles, messages,

---

[1] Stellar is available on https://github.com/akingunay/stellar

and parameters of the given protocol. Third, programmers develop their agents using the static core library and the automatically generated protocol-specific library and by following Stellar's software pattern.

Before explaining Stellar's details, we first highlight its key features for developing protocol-compliant agents using the following Java snippet, which shows a possible implementation of the seller agent in Purchase protocol to handle the reception of an rfq message and respond with the corresponding quote message.

Listing 2: Handling of a received rfq message by a seller.

```
1   public void handleRfq(Rfq rfq)
2   {
3     // create a query to define a criteria for message retrieval
4     Query query = new Query("pID", Query.EQ, rfq.get(Rfq.pID));
5
6     // use adapter to retrieve an enabled message according to the criteria
7     Quote quote = adapter.retrieveEnabled(Quote.class, query).getFirst();
8
9     // seller's business logic to determine the requested item's price
10    String price = priceMap.get(quote.get(Quote.ITEM));
11
12    // send the enabled message by binding necessary parameters
13    quote.send(price);
14  }
```

Stellar follows BSPL's declarative approach. Hence, it does not impose a control flow for developing protocol-compliant agents. Instead, Stellar uses an event-driven model, where viable messages are created and sent according to the local history of an agent when certain events happen. In this regard, the above code snippet shows an event handler for the reception of an rfq message. A key class of Stellar is a role adapter, which provides operations to retrieve and exchange viable messages during enactment of protocols. In our code snippet the seller's adapter is referred via the variable adapter, which is created during the initialization of the seller agent as we will demonstrate later.

A fundamental feature of an adapter is to provide operations for retrieving *enabled* messages from an agent's local history. In an enabled message, all ⌈in⌉ adorned parameters are bound according to the local history of the agent, and all ⌈out⌉ and ⌈nil⌉ parameters are unbound. Hence, the programmer can easily create a viable message from an enabled message, which is retrieved from its local history using the role adapter, simply by producing bindings for all unbound ⌈out⌉ parameters according to the business logic of the agent. In this way Stellar ensures that agents send only viable messages and accordingly guarantees compliance of an agent's implementation with a protocol.

To exemplify, in Line 7 of Listing 2, adapter object's retrieveEnabled method retrieves an enabled quote object of Quote class, which corresponds to a quote message of Purchase protocol, to create a viable response to the received rfq message. The retrieval operation takes a query to determine which particular enabled message(s) it should retrieve. In our example, a single object corresponding to the enabled quote message is retrieved as a response to the received rfq, using the identifier of the received message in the query that is provided to retrieveEnabled. Finally, in Line 13, the send method of the retrieved Quote object is used to send the actual message to its recipient (i.e., the buyer who sent the received

rfq message), which is automatically set by the adapter when retrieving the Quote object from the sender's local history. Note that, in order to make the corresponding message viable, the send method takes a price argument, whose value is determined by the seller's business logic.
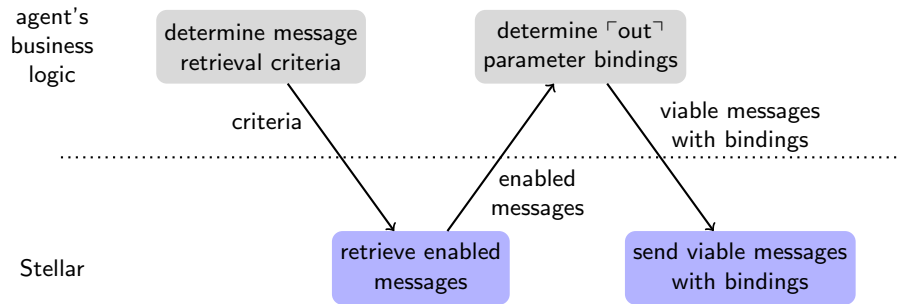

Fig. 4: Pattern to implement protocol-compliant agents using Stellar.

Figure 4 shows an abstract representation of Stellar's software pattern that we use in the above code snippet in Listing 2 to ensure compliance of the seller agent. First, the agent's business logic determines the criteria to retrieve a certain type of enabled message(s) from its local history. Then, the agent uses Stellar to retrieve the enabled message(s) that satisfy its criteria. Next, the agent's business logic determines the bindings of the ⌜out⌝ adorned parameters of the retrieved enabled message(s), and provide them to Stellar. Finally, Stellar compiles viable messages using the enabled messages and the provided bindings for the ⌜out⌝ adorned parameters, and sends them to their recipients.

4.1 Developing Agents using Stellar

In this section we present details of Stellar using an example in which we implement a buyer agent for the Purchase protocol in Listing 1.

Structure of an Agent: Listing 3 shows the overall structure of the Buyer class, which we use to implement the buyer agent. The object adapter of class BAdapter is the buyer's role adapter which is generated by Stellar from the specification of the Purchase protocol. Buyer class implements QuoteHandler interface, which is also generated by Stellar, with a single method handleQuote, which is called by adapter when the buyer receives a Quote message. Note that QuoteHandler could also be implemented by a separate class to enhance modularity. Buyer class may have other variables and methods as usual to represent the buyer's business logic. Also note that Buyer is a programmer defined class and it is not generated by Stellar. That is, Stellar does not require developers to inherit a certain base class or use a certain code skeleton when implementing their agents.

Listing 3: Structure of the buyer's agent.

```
1   public class Buyer implements QuoteHandler
2   {
3
4       private BAdapter adapter;
5
6       // class variables and methods to represent buyer's business logic
7       ...
8
9       // object initialization
10      ...
11
12      public void handleQuote(Quote quote)
13      {
14          ...
15      }
16
17  }
```

Initialization of Agent: Listing 4 shows the constructor of Buyer class. The object adapter of class BAdapter is initialized using the factory method newAdapter according to a Configuration object, which includes information about the buyer's deployment such as its own and other agents network addresses. We discuss these concepts in detail in Section 4.2. Next, the created Buyer object registers itself as the handler for the received Quote messages.

Listing 4: Initialization of the buyer's agent.

```
1   public Buyer(Configuration configuration)
2   {
3       adapter = BAdapter.newAdapter(configuration);
4       adapter.registerQuoteHandler(this);
5
6       // initialization of other class variables
7       ...
8   }
```

Initialization of Interaction: Listing 5 shows how the buyer agent initiates its interaction with a seller agent. The code first retrieves an enabled rfq message object calling the method retrieveEnabled of adapter. Remember that in the Purchase protocol the rfq message does not have any ⌜in⌝ adorned parameters. Hence, this message corresponds to an entry point for a new enactment of the protocol. Therefore, the buyer can send this message at any time to initiate a new enactment by setting its ⌜out⌝ adorned parameters. In other words, an rfq message is always viable. When retrieveEnabled method is called for such a message, Stellar automatically assigns values to the key parameter(s) of the message to create a unique key for initiating a new enactment ensuring information integrity. Hence, the only thing the buyer should do is to determine the item, for which it intends to request a quote, and call the send method of the retrieved rfq object.

Listing 5: Initialization of interaction by the buyer's agent.

```
1   Rfq rfq = adapter.retrieveEnabled(Rfq.class);
2   String item = ... // set by buyer's business logic
3   rfq.send(item);
```

This code snippet can be part of any programmer-defined method that captures the buyer agent's business logic. For instance, if the buyer agent is provided a

list of items to buy, it can iterate over the list and execute the snippet for each item in effect imitating a new (concurrent) enactment of the Purchase protocol. Note that, although the buyer can enact multiple protocols concurrently, it is executed as a single thread, avoiding pitfalls of concurrency. Agents that are developed using Stellar process a single message at a time (similar to actors model). Hence, an agent is always implemented as a single thread even if it enacts multiple protocols concurrently.

Handling of Messages: The next two code snippets show handling of received quotes. Listing 6 shows the interface that is generated by Stellar from the specification of Purchase and Listing 7 shows the implementation of the interface by the buyer's agent.

Listing 6: Specification of QuoteHandler interface.

```
1   public interface QuoteHandler
2   {
3     public void handleQuote(Quote quote);
4   }
```

For simplicity, suppose that the business logic of the buyer is to accept quotes below 50 and reject others. The code in Listing 7 first creates a Query object to represent the buyer's acceptance criteria for the received quotes. The first part of the query calls the get method to determine the identifier of the enactment for which the quote is received (Line 3), and then defines the buyer's criterion for the acceptable value of the price (Line 4). Next, the code calls the retrieveEnabled method to retrieve an enabled Accept message object that matches the given query.

Listing 7: Implementation of QuoteHandler interface by buyer's agent.

```
1    public void handleQuote(Quote quote)
2    {
3      Condition c1 = new Condition("pID", Query.EQ, quote.get(Quote.pID));
4      Condition c2 = new Condition("price", Query.LT, 50));
5      Query aQuery = new Query(new AndCondition(c1, c2));
6
7      Accept accept = adapter.retrieveEnabled(Accept.class,
             aQuery).getFirst();
8
9      if (accept != null) {
10       accept.send("OK");
11     } else {
12       Query rQuery = new Query("pID", Query.EQ, quote.get(Quote.pID));
13       Reject reject = adapter.retrieveEnabled(Reject.class,
             rQuery).getFirst();
14       reject.send("NOK");
15     }
16   }
```

Note that there can only be one enabled accept message for every enactment with a particular binding of pID. However, retrieveEnabled returns a MessageSet object, which implements the Set interface with additional convenience methods. In Line 7, getFirst is one of these convenience methods that retrieves a single message if the set is a singleton and null otherwise.

If there is an enabled accept message that matches the query (i.e., the quoted price is below 50), the code sends the retrieved Accept message using its send

method, providing "OK" as the binding of result parameter (Lines 9–10). Otherwise (i.e., the quoted price is above 50 and hence accept is null), the code sends a Reject message, which is retrieved by calling the retrieveEnabled method with the corresponding query (Lines 11–14).

Remember that the buyer should either send an accept or a reject message for a received quote to comply with Purchase (i.e., accept or reject are mutually exclusive). Let us explain how this is guaranteed by Stellar. Suppose that the programmer of buyer agent made a mistake an wrote the following code to handle received quote messages instead of the code in Listing 7.

```
1   public void handleQuote(Quote quote)
2   {
3     ...
4     Accept accept = adapter.retrieveEnabled(Accept.class,
            aQuery).getFirst();
5     if (accept != null) {
6       accept.send("OK");
7     }
8     Query rQuery = new Query("pID", Query.EQ, quote.get(Quote.pID));
9     Reject reject = adapter.retrieveEnabled(Reject.class,
            rQuery).getFirst();
10    if (reject != null) {
11      reject.send("NOK");
12    }
13    ...
```

This piece of code tries to send first an accept and then a reject message for the same pID binding. However, when an accept message for a pID binding is sent (Line 6), the parameter result is bound for the pID binding, which makes the reject message for the same pID binding disabled, since result is adorned ⌜out⌝ in reject. Accordingly, retrieveEnabled always returns null when it is called to retrieve a Reject object for a binding of pID for which an accept message is already sent (Line 9). Hence, the agent still complies with the Purchase protocol, event though its business logic is not correctly implemented.

4.2   Implementation of Stellar

Management of Local Histories: Stellar stores local history of an agent in a local relational database. Stellar hides the details of the particular database system from programmers. In fact, the programmer should not access the local history of the agent directly. Instead, the programmer should use only the retrieve and send methods provided by Stellar. Our implementation currently uses MySQL to store local histories of agents, however any relational database system that supports fundamental relational operations can be easily adopted.

Emission and Reception of Viable Messages: Stellar uses asynchronous message passing for agent communication, which we implemented using UDP. The messages that are exchanged over the network are serialized into parameter-value pairs and represented in JSON format. Emission of messages is enabled only via the send methods of the generated message classes. The aim of these methods is to ensure that agents send only viable messages by binding all of the necessary parameters. Otherwise, these methods throw exception. Reception of messages and their insertion into an agent's local history is handled by the

adapters of Stellar. Hence, Stellar does not provide any method to programmers for manual message reception. Instead, the programmers should implement handlers of the messages that they want to react, which are automatically called by the adapters of Stellar when a message is received. This simplifies programming of agents in asynchronous settings.

Retrieval of Enabled Messages: Here we provide Stellar's algorithm for the retrieval of enabled messages from an agent's local history. Below, we use, $\mathbb{P}$ for a BSPL protocol, p for individual parameters, P, Q, K for lists (or sets if their ordering is not important) of parameters. We use calligraphic capital letters for relations in the local history of an agent, and apply standard relational algebra operators Π for projection, σ for selection, $\bowtie$ for natural join, $\bowtie$ for full outer join, and $\bowtie$ for left outer join. We also use the utility methods `allParams`, `keyParams`, `inParams`, `nilParams`, and `outParams` with a relation and protocol argument to access the set of all, key, ⌜in⌝, ⌜nil⌝, and ⌜out⌝ adorned parameters of the relation, respectively.

Algorithm 1: retrieveEnabled($\mathcal{M}, \phi, \mathbb{P}$)

1 $P_{in} \leftarrow$ `inParams`$(\mathcal{M}, \mathbb{P})$ // ⌜in⌝ adorned parameters of $\mathcal{M}$
2 $P_{nil} \leftarrow$ `nilParams`$(\mathcal{M}, \mathbb{P})$ // ⌜nil⌝ adorned parameters of $\mathcal{M}$
3 $P_{out} \leftarrow$ `outParams`$(\mathcal{M}, \mathbb{P})$ // ⌜out⌝ adorned parameters of $\mathcal{M}$
4 $K \leftarrow$ `keyParams`$(\mathcal{M}, \mathbb{P})$ // key parameters of $\mathcal{M}$
5 $\mathcal{W}_I \leftarrow \emptyset$
6 **if** P_{in} *is* $\emptyset$ **then**
7 | **return** $\{()\}$
8 $\mathcal{W}_I \leftarrow \bigcup_{\mathcal{N} \in \mathbb{P}} \Pi_K(\mathcal{N})$
9 **foreach** $p \in P_{in}$ **do**
10 | $Q \leftarrow K \cup \{p\}$
11 | $\mathcal{W}_p \leftarrow$ `createRelation`(Q)
12 | **foreach** $\mathcal{N} \in \mathbb{P}$ *such that* $p \in$ `allParams`$(\mathcal{N}, \mathbb{P})$ **do**
13 | | $\mathcal{W}_p \leftarrow \Pi_Q(\mathcal{N}) \cup \mathcal{W}_p$
14 | $\mathcal{W}_I \leftarrow \mathcal{W}_I \bowtie_K \mathcal{W}_p$
15 $\mathcal{W}_E \leftarrow \bigcup_{\mathcal{N} \in \mathbb{P}} \Pi_K(\mathcal{N})$
16 **foreach** $p \in P_{out} \cup P_{nil}$ **do**
17 | $Q \leftarrow K \cup \{p\}$
18 | $\mathcal{W}_p \leftarrow$ `createRelation`(Q)
19 | **foreach** $\mathcal{N} \in \mathbb{P}$ *such that* $p \in$ `allParams`$(\mathcal{N}, \mathbb{P})$ **do**
20 | | $\mathcal{W}_p \leftarrow \Pi_Q(\mathcal{N}) \cup \mathcal{W}_p$
21 | $\mathcal{W}_E \leftarrow \mathcal{W}_E \bowtie_K \mathcal{W}_p$
22 $\mathcal{W} \leftarrow \sigma_{P_{out}=null \land P_{nil}=null}(\mathcal{W}_I \bowtie_K \mathcal{W}_E)$
23 **return** $\sigma_\phi(\mathcal{W})$

Algorithm 1 defines retrieval of enabled messages, given the relation $\mathcal{M}$ that corresponds to a message schema (e.g., a quote message schema), the user defined

query ϕ, and the protocol specification $\mathbb{P}$. Note that the algorithm returns a relation (not actual message objects in Java), where each tuple of the relation corresponds to the parameter bindings of an enabled instance of the message schema $\mathcal{M}$. A Stellar adapter uses this relation to create the corresponding message objects and return them as the result of a retrieveEnabled method call as we demonstrated in earlier examples.

Algorithm 1 can be divided into three phases. In the first phase (from lines 8 to 14), the algorithm builds the relation $\mathcal{W}_I$ where all $\ulcorner$in$\urcorner$ adorned parameters of $\mathcal{M}$ are bound. In the second phase (from lines 15 to 21), the algorithm builds another relation $\mathcal{W}_E$ where one or more $\ulcorner$out$\urcorner$ or $\ulcorner$nil$\urcorner$ adorned parameters of $\mathcal{M}$ are bound. In the last phase (lines 22–23), the algorithm removes the tuples from $\mathcal{W}_I$ for which there is a matching tuple (i.e., identified by the same key) in $\mathcal{W}_E$. Hence, each tuple of the resulting relation corresponds to an enabled message (i.e., all $\ulcorner$in$\urcorner$ parameters are bound and all $\ulcorner$out$\urcorner$ and $\ulcorner$nil$\urcorner$ parameters are unbound). Note that this operation removes the tuples that would cause violation of mutual exclusion. If a message $\mathcal{M}'$ that is mutually exclusive to $\mathcal{M}$ is already emitted in an enactment, then there are tuples in $\mathcal{W}_E$ with bindings of the parameters that are adorned $\ulcorner$out$\urcorner$ in both $\mathcal{M}$ and $\mathcal{M}'$, and accordingly the corresponding tuples in $\mathcal{W}_I$ are removed. Hence, the messages that can cause to violation of mutual exclusion are not enabled. The algorithm applies the given query ϕ to the resulting relation $\mathcal{W}$ to filter the tuples.

4.3 Revisiting Pitfalls

Stellar's retrieval and emission operations ensure causality and information integrity. Specifically, retrieval operations ensure integrity of $\ulcorner$in$\urcorner$ adorned parameters by binding them permanently to the corresponding values according to the agent's local history. Hence, a programmer cannot fabricate or alter $\ulcorner$in$\urcorner$ adorned parameters of a message. The send operations ensure integrity of $\ulcorner$out$\urcorner$ and $\ulcorner$nil$\urcorner$ adorned parameters by enforcing the programmer to assign values to only $\ulcorner$out$\urcorner$ adorned parameters when needed. Hence, the programmer cannot omit assignment of mandatory parameters and thus break integrity. Further, Stellar handles creation of bindings for key parameters ensuring their uniqueness, which prevents key related integrity issues.

Stellar's retrieval operations prevent emission of mutually exclusive messages. That is, if two (or more) messages are mutually exclusive in a protocol and an agent has already sent one of these messages in an enactment of the protocol, the retrieval operation does not consider the other mutually exclusive message(s) as enabled in the same enactment. Hence, the agent cannot retrieve and send mutually exclusive messages. Stellar does not directly handle non-local mutual exclusion. However, safety of a BSPL protocol, which means that the protocol is free from non-local mutual exclusion, can be verified automatically [12] at design time to avoid non-local mutual exclusion issues.

Communication in Stellar is fully asynchronous. Hence a single-threaded agent can easily enact multiple protocols at the same time using Stellar. That is, an agent's execution is never blocked when sending or receiving messages.

Further, Stellar's programming model handles one incoming message at a time. Hence, developers can implement their agents without any thread synchronization that deals with interleaving reception of multiple messages in different enactments. This feature of Stellar substantially simplifies design of an agent, as our case study in Section 5 demonstrates. Note that UDP, which is used in our implementation, is an unreliable protocol (i.e., it does not guarantee delivery of emitted messages), which may compromise liveness of an interaction. This issue can be avoided using a reliable alternative, such as RUDP, TCP, or message queues. However, these alternatives provide features (e.g., ordered message delivery), which are not needed by Stellar. We chose UDP for our implementation to show that lack of such features do not affect protocol-compliance of agents. We will address liveness of interactions in our future work.

5 Case Study

To demonstrate the use of Stellar in a more comprehensive case, where an agent should consider multiple messages for decision making, we use a claim handling scenario from insurance domain. We list the protocol that represents this scenario in Listing 8. In this scenario there is a policy subscriber and an insurer. The subscriber can make multiple claims (claim message) by sending an incident's details and the claimed amount to the insurer. The insurer either approves (approve message) or rejects a claim (reject message). In case of approval, the insurer pays the claimed amount to the subscriber. The insurer can pay its balance immediately for each claim or as lump sum for several claims (pay message). For brevity, we omit some policy aspects such as premium payments.

Listing 8: An insurance policy claim protocol.

```
Insurance {
 roles I, S //insurer, subscriber

 parameters out sID key, out cID key, out pID key, out subscriber, out
     period,
     out type, out date, out incident, out cAmount, out outcome, out pAmount

 S ↦ I: subscribe[out sID, out subscriber, out period, out type]
 I ↦ S: register[in sID, in subscriber, in period, in type, out date]
 S ↦ I: claim[in sID, out cID, out incident, out cAmount]
 I ↦ S: approve[in sID, in cID, in incident, in cAmount, out outcome]
 I ↦ S: reject[in sID, in cID, in incident, in cAmount, out outcome]
 I ↦ S: pay[in sID, out pID, out pAmount]
}
```

Listing 9 shows the implementation of ClaimHandler interface by the insurer agent to handle claim messages when enacting Insurance. As we explained earlier, ClaimHandler interface is generated by Stellar from the specification of Insurance and consists of a single method handleClaim, which is used to define the insurer's business logic for handling claims. Suppose that the insurer handles a claim in two steps. In the first step, the insurer decides whether the received claim is valid or not. In the second step, if the claim is valid and the insurer's policy balance exceeds a minimum payable amount, the insurer pays its balance. Otherwise, the insurer does not make any immediate payment.

The method processClaim (Lines 8–21) captures the first step. The method first decides whether the received claim is valid by calling isValidClaim (Line 12), which returns true for valid and false for invalid claims. We do not present the details of isValidClaim since they are not relevant to our demonstration. Depending on the validity of the claim, processClaim retrieves and sends either the enabled Approve (Lines 13–15) or Reject (Lines 17–19) message. Finally, processClaim returns true or false depending on the validity of the claim.

If the claim is approved (Line 3), handleClaim calls payBalance (Line 4), which captures the second step (Lines 23–30). The method payBalance first computes the insurer's total balance for the policy using approvedClaimAmount and paidClaimAmount methods (Line 24). We describe these methods later in Listing 10. If the insurer's balance is more that the minimum payable amount, processClaim retrieves and sends the enabled Pay message to pay the insurer's balance (Lines 25–29).

Listing 9: Implementation of ClaimHandler interface by the insurance agent.

```
1   public void handleClaim(Claim claim)
2   {
3     boolean isApproved = processClaim(claim);
4     if (isApproved) {
5       payBalance(claim.get(Claim.sID));
6     }
7   }
8
9   private void processClaim(Claim claim)
10  {
11    Condition c1 = new Condition("sID", Query.EQ, claim.get(Claim.sID));
12    Condition c2 = new Condition("cID", Query.EQ, quote.get(Claim.sID));
13    Query query = new Query(new AndCondition(c1, c2));
14
15    if (isValidClaim(claim)) {
16      Approve msg = adapter.retrieveEnabled(Accept.class, query).getFirst();
17      msg.send("APPROVED");
18      return true;
19    } else {
20      Reject msg = adapter.retrieveEnabled(Reject.class, query).getFirst();
21      msg.send("REJECTED");
22      return false;
23    }
24  }
25
26  private void payBalance(String sId)
27  {
28    int balance = payableClaimedAmount(sId) - totalPaidAmount(sId);
29
30    if(MIN_PAYABLE_AMOUNT <= balance) {
31      Query query = new Query("sID", Query.EQ, sId);
32      Pay msg = adapter.retrieveEnabled(Pay.class, query);
33      msg.send(balance);
34    }
35  }
```

Computation of the insurer's balance for a policy requires consideration of multiple messages. That is, we should first compute the total payable claimed amount for the policy according to the approved claims. Then we should compute the total paid amount for the policy according to the previous payments, and subtract it from the total payable claimed amount. The method payableClaimedAmount in Listing 10 (Lines 1–9) computes the total payable claimed

amount. It first retrieves all the sent Approve messages for the policy, which is identified by sId, from the insurer's message history calling retrieveMessage method (Line 3). Note that this is a different method than retrieveEnabled. Next, the total payable amount is computed by iterating over all the retrieved Approve messages and summing up the claimed amount of each message. The method totalPaidAmount (Lines 11–19) repeats the same process to compute the total paid amount for the policy using Pay messages (instead of Approve messages) and corresponding paid amounts in those messages.

Listing 10: Computation of total claimed and paid amounts.

```
1   private int payableClaimedAmount(String sId)
2   {
3     Query query = new Query("sID", Query.EQ, sId);
4     MessageSet<Approve> msgs = adapter.retrieveMessage(Approve.class,
          query);
5
6     int sum = 0;
7     for (Approve msg : msgs) {
8       sum += (int) msg.get(Approve.cAmount);
9     }
10    return sum;
11  }
12
13  private int totalPaidAmount(String sId)
14  {
15    Query query = new Query("sID", Query.EQ, sId);
16    MessageSet<Pay> msgs = adapter.retrieveMessage(Pay.class, query);
17
18    int sum = 0;
19    for (Pay msg : msgs) {
20      sum += (int) msg.get(Pay.pAmount);
21    }
22    return sum;
23  }
```

6 Discussion

We summarize our contributions, relate them to the literature, and discuss directions for future work.

6.1 Summary

This chapter presented Stellar, a programming model for developing protocol-compliant agents for BSPL. Stellar's main idea is to ensure compliance of agents by allowing exchange of only viable messages between them. To this end, Stellar provides a simple yet effective software pattern for retrieving enabled messages from an agent's local history and for sending them ensuring their viability. Communication in Stellar is fully asynchronous. Further, Stellar implicitly ensures information integrity of interaction, prevents emission of locally mutually exclusive messages, and enables concurrent enactment of protocols without relying on multithreading mechanisms. Accordingly, Stellar simplifies agent development in decentralized settings ensuring their compliance. Stellar is different from programming models that ensure agent compliance using control flow structures

(i.e., code skeletons), and more distantly related to distributed programming models without interaction protocols.

6.2 Related Work

Scribble [15] enforces design-time adherence to protocols that specify typed message signatures and messaging constraints with explicit control flow. Scribble [5] extracts state machines from protocol specifications to generate APIs for endpoints. Sending or receiving a message returns a protocol state object and each protocol state object provides message sending and receiving operations that comply with correct subsequent state transitions. Stellar provides more flexibility to developers, since it does not impose a control flow for ensuring compliance when implementing agents. Developers are free to design their agents as they see fit, without being distracted about the state of their interactions. Stellar's retrieve and send pattern ensures exchange of only viable messages and accordingly ensures compliance of agent, which is independent from a control flow.

Business-oriented approaches for web services propose the use of high-level processes according to which correct interactions are enforced in code. Business Process Modeling Notation (BPMN) has been used to specify processes that are then translated into the Business Process Execution Language [10] (BPEL), an executable language for externally invoking web services and their interactions based on event occurrences. The BPMN-BPEL approach is inherently process-oriented and depends on the correct realization of workflows. Stellar is information-oriented and uses a declarative approach without imposing any workflow. Business artifacts [9] are high-level representations of both processes, interaction, and the relational data that they operate on. Artifact interoperation hubs [6] enforce correct messaging with business processes by acting as central communication points between web services. Stellar only uses local information and does not rely on any centralized communication artifacts to ensure compliance.

Programming models for distributed systems generally do not consider protocols, e.g., functional reactive programming [2], the Sunny Event-driven programming model [8], and the Actor programming model [4] implemented in Akka [1]. These programming models support interaction derived from internal system or actor code, without a protocol specification against which correct implementations must comply. In Stellar, we focus on supporting independent agent development against protocol specifications, where programmers are protected from violating protocol compliance and integrity of information.

6.3 Future Work

Stellar is a first step toward declarative agent programming based on declarative information-based protocols. A fuller exploration of agent programming would need to consider abstractions for agent policies and how they fit with Stellar. Another interesting direction would be to explore how normative abstractions (e.g.,

commitments) may be used alongside Stellar, especially since norms represent the meaning of information communicated via protocols.

References

1. Akka: 2.5.6 (2017), http://akka.io
2. Bainomugisha, E., Carreton, A.L., van Cutsem, T., Mostinckx, S., de Meuter, W.: A survey on reactive programming. ACM Computing Surveys **45**(4), 52:1–52:34 (2013)
3. Chopra, A.K., Christie V., S.H., Singh, M.P.: Splee: A declarative information-based language for multiagent interaction protocols. In: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems. pp. 1054–1063 (2017)
4. Hewitt, C., Bishop, P., Steiger, R.: A universal modular actor formalism for artificial intelligence. In: Proceedings of the 3rd International Joint Conference on Artificial Intelligence. pp. 235–245. IJCAI'73, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1973)
5. Hu, R., Yoshida, N.: Hybrid session verification through endpoint api generation. In: Proceedings of the 19th International Conference on Fundamental Approaches to Software Engineering - Volume 9633. pp. 401–418. Springer-Verlag, New York, NY, USA (2016)
6. Hull, R., Narendra, N.C., Nigam, A.: Facilitating workflow interoperation using artifact-centric hubs. In: Proceedings of the 7th International Joint Conference on Service-Oriented Computing. pp. 1–18. ICSOC-ServiceWave '09, Springer-Verlag, Berlin, Heidelberg (2009)
7. Ladkin, P.B., Leue, S.: Interpreting message flow graphs. Formal Aspects of Computing **7**(5), 473–509 (1995)
8. Milicevic, A., Jackson, D., Gligoric, M., Marinov, D.: Model-based, event-driven programming paradigm for interactive web applications. In: Proceedings of the 2013 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming & Software. pp. 17–36. Onward! 2013, ACM, New York, NY, USA (2013)
9. Nigam, A., Caswell, N.S.: Business artifacts: An approach to operational specification. IBM Systems Journal **42**(3), 428–445 (2003)
10. Ouyang, C., Dumas, M., van der Aalst, W.M.P., Hofstede, t.A.H.M., Mendling, J.: From business process models to process-oriented software systems. ACM Transactions on Software Engineering and Methodology **19**(1), 2:1–2:37 (2009)
11. Singh, M.P.: Information-driven interaction-oriented programming: BSPL, the Blindingly Simple Protocol Language. In: Proceedings of the 10th International Conference on Autonomous Agents and MultiAgent Systems. pp. 491–498 (2011)
12. Singh, M.P.: Semantics and verification of information-based protocols. In: Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems. pp. 1149–1156 (2012)
13. Singh, M.P.: Bliss: Specifying declarative service protocols. In: Proceedings of the 2014 IEEE International Conference on Services Computing. pp. 235–242 (2014)
14. Winikoff, M., Yadav, N., Padgham, L.: A new hierarchical agent protocol notation. Autonomous Agents and Multi-Agent Systems **32**(1), 59–133 (2018)

15. Yoshida, N., Hu, R., Neykova, R., Ng, N.: The scribble protocol language. In: 8th International Symposium on Trustworthy Global Computing - Volume 8358. pp. 22–41. TGC 2013, Springer-Verlag New York, Inc., New York, NY, USA (2014)