

NRC Publications Archive Archives des publications du CNRC

Rule responder agents framework and instantiations

Boley, Harold; Paschke, Adrian

For the publisher's version, please access the DOI link below./ Pour consulter la version de l'éditeur, utilisez le lien DOI ci-dessous.

Publisher's version / Version de l'éditeur:

https://doi.org/10.1007/978-3-642-18308-9_1

Semantic Agent Systems: Foundations and Applications, Studies in Computational Intelligence; Volume 344, pp. 3-23, 2011

NRC Publications Archive Record / Notice des Archives des publications du CNRC :

<https://nrc-publications.canada.ca/eng/view/object/?id=0c2ad733-a7ed-4a66-a843-0ee30643bc69>

<https://publications-cnrc.canada.ca/fra/voir/objet/?id=0c2ad733-a7ed-4a66-a843-0ee30643bc69>

Access and use of this website and the material on it are subject to the Terms and Conditions set forth at

<https://nrc-publications.canada.ca/eng/copyright>

READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THIS WEBSITE.

L'accès à ce site Web et l'utilisation de son contenu sont assujettis aux conditions présentées dans le site

<https://publications-cnrc.canada.ca/fra/droits>

LISEZ CES CONDITIONS ATTENTIVEMENT AVANT D'UTILISER CE SITE WEB.

Questions? Contact the NRC Publications Archive team at

PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the first page of the publication for their contact information.

Vous avez des questions? Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n'arrivez pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.

Rule Responder Agents

Framework and Instantiations

Harold Boley¹, Adrian Paschke²

¹ Institute for Information Technology, National Research Council Canada

Fredericton, NB, Canada

harold.bole AT nrc.gc.ca

² Freie Universitaet Berlin, Germany

paschke AT mi.fu-berlin.de

Abstract This chapter introduces Rule Responder and its applications. Rule Responder is a framework for specifying virtual organizations as semantic multi-agent systems that support collaborative teams. It provides the infrastructure for rule-based collaboration between the distributed members of such a virtual organization. Human members of an organization are assisted by (semi-)autonomous rule-based agents, which use Semantic Web rules to describe aspects of their owners' derivation and reaction logic.

To implement different distributed system/agent topologies with their negotiation/coordination mechanisms Rule Responder instantiations employ three core classes of agents - Organizational Agents (OA), Personal Agents (PAs), and External Agents (EAs). The OA represents goals and strategies shared by its virtual organization as a whole, using a rulebase that describes its policies, regulations, opportunities, etc. Each PA assists a group or person of the organization, semi-autonomously acting on their behalf by using a local knowledge base of rules defined by the entity. EAs can communicate with the virtual organization by sending messages to the public interfaces of the OA. EAs can be human users using, e.g., Web forms or can be automated services/tools sending messages via the multitude of transport protocols of the underlying enterprise service bus (ESB) middleware. The agents employ ontologies in their knowledge bases to represent semantic domain vocabularies, normative pragmatics and pragmatic context of conversations and actions, as well as the organizational semiotics.

1 Introduction

Rule Responder¹ extends the Semantic Web towards a Pragmatic Web infrastructure for collaborative rule-based agent networks realizing distributed inference services, where independent agents engage in conversations by exchanging messages and cooperate to achieve (collaborative) goals. Rule Responder can be characterized on three levels, from general to specific.

- It models a virtual organization of agents recursively as again being a single agent, forming what has been called [15] a hierarchy of *holons* (or, a *holarchy*).
- It supports different *interaction/coordination models*, where information is interchanged within a *pragmatic context* (e.g. language action speech acts, deontic norms, etc.).
- It provides a technical Web-based *multi-agent architecture* which supports different distribution models (distributed agent system topologies).

A virtual organization as a whole is represented by an Organizational Agent (OA), which uses ontologies and rules to assign and delegate incoming tasks (e.g., queries) to responsible Personal Agents (PAs). Rule Responder agents communicate in conversations that allow implementing different agent coordination and negotiation protocols. The interaction and interpretation is driven by the organizational semiotics which details how the information flow works within and between organizations. For instance, an OA can use a responsibility assignment matrix, represented as an ontology, to find an appropriate PA in its organization. The OA can then send a message (e.g., a query) to that PA and receive results (e.g., answers), typically using reaction rules. By means of pragmatic primitives, such as speech acts, deontic norms, etc., which are represented as ontologies, Rule Responder attaches the semantic and pragmatic context, e.g. organizational norms, purposes or goals and values, to the interchanged messages.

In its multi-agent architecture Rule Responder utilizes messaging reaction rules from Reaction RuleML² for communication between the distributed agent inference services. The Rule Responder middleware is based on modern enterprise service technologies and Semantic Web technologies for implementing intelligent agent services that access data and ontologies, receive and detect events (e.g., for complex event processing in event processing agent networks), and make rule-based inferences and (semi-)autonomous pro-active decisions for reactions based on these representations.

The core of a Rule Responder agent is a rule engine, such as Prova³, OO jDREW, DR-Device (initially in Emerald), Euler, or Drools, which implements the decision and behavioral reaction logic of the agents' roles. An agent can employ vocabularies defined as Semantic Web ontologies (e.g., based on RDFS or OWL)

¹<http://responder.ruleml.org>

²<http://reaction.ruleml.org>

³<http://prova.ws>

to give its rules a domain-specific meaning. The vocabularies can be used within the conversation with other agents to enable a semantic and pragmatic interpretation of the messages. For the deployment of agents on the Web and for the communication in agent networks, Rule Responder uses the Mule-based enterprise service bus middleware, which supports a multitude of synchronous and asynchronous transport protocols (> 40) -- such as MS, SMTP, JDBC, TCP, HTTP, XMPP, Jade -- to transport rulebases, queries and answers between the agents. Reaction RuleML, the de facto standard for XML-serialized reaction rules, is used as a platform-independent rule interchange format for agent conversation.

In summary, Rule Responder can be seen to support a *digital ecosystem*, evolving from the Semantic Web [4] to the Pragmatic Web, which consists of all the semantic agents in one or more virtual organizations, as well as all the other components of this environment with which the agents interact, such as other services, tools, the ESB middleware, etc.

Several instantiations of Rule Responder have been developed, including the eScience infrastructure for Health Care and Life Sciences [11], the Rule-based IT Service Level Management and Semantic BPM system [12, 13], multiple versions of the deployed SymposiumPlanner system [9], two versions of the WellnessRules prototype [5], and the PatientSupporter prototype.⁴

The rest of the chapter is organized as follows. Section 2 discusses the agent architecture and used technologies of the Rule Responder framework. Section 3 explains a typical distributed agent topology for virtual organizations and the types agents used to implement it. Section 4 focuses on interchange between the semantic agents which communicate by using (Reaction) RuleML as common rule interchange format. Section 5 demonstrates some application use cases of Rule Responder by means of selected Rule Responder instantiations. Section 6 concludes the paper.

2 The Rule Responder Framework

Three interconnected architectural layers constitute the Rule Responder framework, listed here from top to bottom:

- Computationally independent user interfaces such as template-based Web forms or controlled English rule interfaces.
- Reaction RuleML as the common platform-independent rule interchange format to interchange rules, events, queries, and data between Rule Responder agents and other agents (e.g., Semantic Web services or humans via Web forms).
- A highly scalable and efficient enterprise service bus (ESB) as agent/service-broker and communication middleware on which platform-specific rule engines are deployed as distributed agent nodes (resp. semantic inference services).

⁴<http://ruleml.org/PatientSupporter>

These engines manage and execute the logic of Rule Responder's semantic agents in terms of declarative rules which have access to semantic ontologies.

In the following, the Rule Responder framework will be refined, and explained from bottom to top.

2.1 Mule Enterprise Service Bus

To seamlessly handle message-based interactions between the Rule Responder agents/services and other agents/services using disparate complex event processing (CEP) technologies, transports, and protocols, an enterprise service bus (ESB) -- the Mule open-source ESB ⁵ -- is used in Rule Responder as the communication middleware. This ESB allows deploying the rule-based agents as highly distributed rule inference services installed as Web-based endpoints on the Mule object broker and supports the communication in this rule-based agent processing network via a multitude of transport protocols (see Figure 1). That is, the ESB provides a highly scalable and flexible application messaging framework to communicate synchronously or asynchronously amongst the ESB-local agents and with agents/services on the Web.




Fig. 1. Distributed Rule Responder Agent Services

Mule is a messaging platform based on principles of ESB architectures, but goes beyond the typical definition of an ESB as a transit system for carrying data between applications by providing a distributable object broker to manage all sorts of service components such as the Rule Responder agent services. The three processing modes of Mule are:

- **Asynchronous:** many events (messages) can be processed by the same component at a time in various threads. When the Mule server is running asynchronously instances of a component run in various threads all accepting incoming events, though an event will only be processed by one instance of the component.
- **Synchronous:** when a component receives an event message, in this mode the whole request is executed in a single thread.

⁵www.mulesoft.org

- Request-Response: this allows for a component to make a specific request for an event and wait for a specified time to get a response back.

The object broker follows the Staged Event Driven Architecture (SEDA) pattern [20]. The basic approach of SEDA is to decomposes a complex, event-driven application into a set of stages connected by queues. This design decouples event and thread scheduling from application logic and avoids the high overhead associated with thread-based concurrency models. That is, SEDA supports massive concurrency demands on Web-based services and provides a highly scalable approach for asynchronous communication.

Figure 2 shows a simplified breakdown of the integration of Mule into the Rule Responders framework.

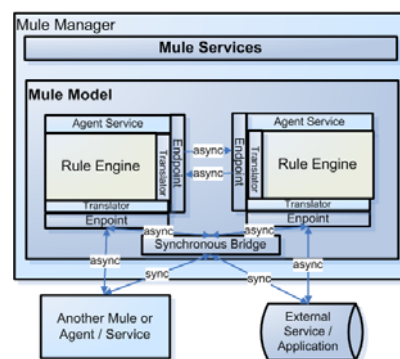


Fig. 2 Layering of Rule Responder on Mule ESB

Distributed agent services (see Figure 1), which at their core run a rule engine, are deployed as Mule components which listen at configured endpoints, e.g., JMS message endpoints, HTTP ports, SOAP server/client addresses or JDBC database interfaces, etc. Reaction RuleML is used as a common platform-independent rule interchange format between the agents (and possible other rule execution/inference services). Translator services are used to translate inbound and outbound messages from platform-independent Reaction RuleML into the platform-specific execution syntaxes of rule engines, and vice versa. XSLT and ANTLR based translator services are provided as Web forms, HTTP services and SOAP Web services on the Reaction RuleML Web page.

The large variety of transport protocols provided by Mule can be used to transport the messages to the registered endpoints or external applications/tools. Usually, JMS is used for the internal communication between distributed agent instances, while HTTP and SOAP is used to access external Web services. The usual processing style is asynchronous using SEDA event queues. However, sometimes synchronous communication is needed. For instance, to handle communication with external synchronous HTTP clients such as Web browsers where requests, e.g. by a Web form, are sent through a synchronous channel. In this case, a synchronous bridge component dispatches the requests into the asynchronous messaging framework and collects all answers from the internal service nodes,

while keeping the synchronous channel with the external service open. After all asynchronous answers have been collected, they are sent back to the still connected external service via the HTTP-synchronous channel.

2.2 Selected Platform-Specific Rule Engines for Rule Responder Agents

The core of a Rule Responder agent, which is deployed as a service component on the Rule Responder ESB, is a platform-specific rule engine. These engines might differ, e.g., in their supported rule types, state representation, rule evaluation mechanism, conflict resolution and truth maintenance. Hence, depending on their expressiveness and functionalities, these rule engines might be capable of implementing agents in the strong sense of cognitive architectures for intelligent agents with goal/task-based, utility-based and learning-based functionalities, or in the weak sense of inference agent services with simple reflexive functionalities for, e.g., deductive query-answering capabilities. Following the general consensus defined by the strong notion of agency in [21], a Rule Responder agent, in addition to being (semi-)autonomous, should be capable of reactive, proactive, and communicative behavior. Additionally, it is often important that certain mentalistic notions⁶ can be used in the rule language for describing the agent behavior in an abstract and intuitive way, e.g. in the interactions between agents to communicate the pragmatics of the interchanged information.

In the following, the interplay between our most often used rule engines Prova, OO jDREW, Euler will be discussed, although there are other engines such as DR-Device and Drools supported by Rule Responder.

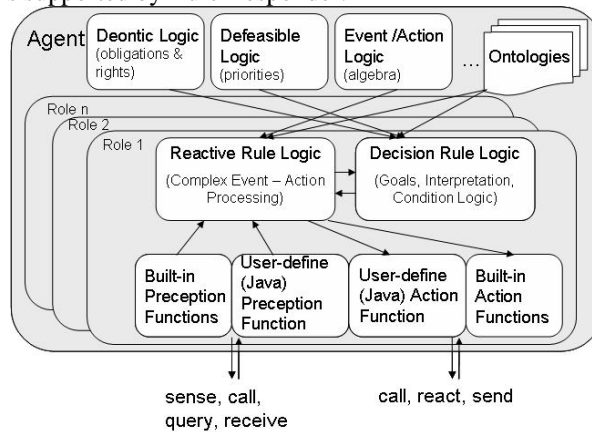


Fig. 3 Rule Responder Agent

⁶The term *mentalistic notions* aka *mental attitudes* refers to human-like properties such as beliefs, goals, etc. when transferred to describing machine agents.

Figure 3 shows the architecture of an intelligent cognitive Rule Responder agent which is implemented in Prova. Prova is an enterprise-strength, highly expressive distributed Semantic Web logic programming (LP) rule engine. The Prova rule engine supports different rule types:

- Derivation rules to describe the agent's decision logic
- Integrity rules to describe constraints and potential conflicts
- Normative rules to represent the agent's permissions, prohibitions and obligation policies
- Global ECA-style reaction rules to define global reaction logic which are triggered on the basis of detected (complex) events
- Messaging reaction rules to define conversation-based workflow reaction and behavioral logic based on complex event processing

Prova follows the spirit and design of the W3C Semantic Web initiative and combines declarative rules, ontologies and inference with dynamic object-oriented programming and access to external data sources via built-in query languages such as SQL, SPARQL, and XQuery.

File Input / Output

```
..., fopen(File,Reader), ...
```

XML (DOM)

```
document(DomTree,DocumentReader) :- XML(DocumentReader),...
```

SQL

```
...,sql_select(DB,cla,[pdb_id,"1alx"],[px,Domain]).
```

RDF

```
...,rdf(http://..., "rdfs",Subject,"rdf_type", "gene1_Gene"),...
```

XQuery

```
..., XQuery = 'for $name in StatisticsURL//Author[0]@name/text()
return $name', xquery_select(XQuery,name(ExpertName)),...
```

SPARQL

```
...,sparql_select(SparqlQuery,...
```

One of the key advantages of Prova is its elegant separation of logic, data access, and computation as well as its tight integration of Java, Semantic Web technologies, with service-oriented computing and complex event processing. In particular, Prova supports external type systems such as, e.g., Java class hierarchies or Semantic Web ontologies (RDFS, OWL) via its typed order-sorted logic [18]. For instance, in the following example all agents from an external OWL ontology responsibility assignment matrix (RAM) are assigned to the typed variable *Agent* of type *Organizing_Committee* (with the namespace *ruleml2010*), where *Organizing_Committee* is a concept defined in the RAM ontology. The query then selects all agent individuals of type *ProgramChair* which is a subtype of *Organizing_Committee*, i.e. the query selects a subset with appropriate subtype from the bound variable.


```

% import external ontology representing responsibility assignment
matrix (RAM)
import("http://2010.ruleml.org/RuleML-2010.owl").
% bind all agent instances of type "Organizing_Committee" from the
RAM to the variable Agent
agent(Agent:ruleml2010_Organizing_Committee).
% query for all agents of type "ProgramChair"
:- solve(agent(Agent:ruleml2010_ProgramChair)

```

Prova can be run in a plain Java environment as stand alone application or rule inference service on the Rule Responder ESB, or as an OSGI component. Prova has a modular knowledge base to implement several different roles an agent might play in the same agent instance. Each role has its own set of reaction rules to autonomously react (potentially proactive) on detected situations (complex events) and its own set of decision rules to interpret goals and derive decisions according to conditional proofs. For instance, it is possible to consult (load) distributed rulebases from local files, a Web address, or from incoming messages transporting a rulebase.

```

%load from a local file
:- eval(consult("organization2009.prova")).
% import from a Web address
:- eval(consult("http://ruleml.org/organization2010.prova")).

```

The rulebases are managed as modules in the knowledge base. Their module label can be used for asserting or retracting complete modules from the knowledge base and for scoping queries/goals to a particular module, i.e. the query only applies to the particular scoped module. In the following example the subgoal *agent(Agent)* applies on the modules *organization2009.prova* and not on the module *http://ruleml.org/organization2010.prova*.

```

responsible(Agent, Task) :-
    @src("organization2009.prova") agent(Agent),
    ...

```

To sense the environment and trigger actions, query data from external sources such as databases, call external procedural code such as Enterprise Java Beans, and receive/send messages from/to other agents or external services, Prova provides a set of built-in functions and additionally can dynamically instantiate any Java object and call its API methods at runtime. For instance, the following simple rule creates a response sentence with the name using Java string computations and displays it via to the Java system out console.

```

hello(Name) :-
    S = java.lang.String("Hello, your name is "),
    S.append (Name),
    java.lang.System.out.println (S).

```

Additional libraries can be imported, e.g. to represent rights and obligations of agents, implement conflict handling rules, or describe complex events and actions. In its cognitive cycle a Prova agent follows the sense-reason-act pattern. However,

Prova does not define one particular cognitive cycle, but allows configuring an agent with user-defined conversation-based negotiation and coordination protocols or workflow patterns. Via constructs for asynchronously sending and receiving event messages within rules, an agent interacts with the environment. The main language constructs of messaging reaction rules are: *sendMsg* predicates to send messages, reaction *rcvMsg* rules which react to inbound messages, and *rcvMsg* or *rcvMult* inline reactions in the body of messaging reaction rules to receive one or more context-dependent multiple inbound event messages:

```
sendMsg(XID,Protocol,Agent,Performative,Payload |Context)
rcvMsg(XID,Protocol,From,Performative,Payload|Context)
rcvMult(XID,Protocol,From,Performative,Payload|Context)
```

Here, *XID* is the conversation identifier (conversation-id) of the conversation to which the message will belong. *Protocol* defines the communication protocol. *Agent* denotes the target party of the message. *Performative* describes the pragmatic envelope for the message content. A standard nomenclature of performatives is, e.g., the FIPA Agents Communication Language (ACL). *Payload* represents the message content sent in the message envelope. It can be a specific query or answer or a complex interchanged rule base (set of rules and facts). For instance, the following rule snippet shows how a query is sent to an agent via the ESB and then an answer is received from this agent.

```
...
sendMsg(Sub_CID,esb,Agent,acl_query-ref, Query),
rcvMsg(Sub_CID,esb,Agent,acl_inform-ref, Answer),
...
```

Prova does not define a specific set of mentalistic notions as first-class programming constructs. Instead, interchanged messages besides the conversation's metadata and payload also carry the pragmatic context of the conversation such as communicative situations/acts, mentalistic notions, organizational and individual norms, purposes or individual goals and values. The payload of incoming event messages is interpreted with respect to the local conversation state, which is denoted by the conversation id, and the pragmatic context, which is given by a pragmatic performative. For instance, a standard nomenclature of pragmatic performatives, which can be integrated as external (semantic) vocabulary/ontology, is e.g., defined by the Knowledge Query Manipulation Language (KQML) (Finin et al. 1993), by the FIPA Agent Communication Language (ACL), which gives several speech act theory-based communicative acts, or by the Standard Deontic Logic (SDL) with its normative concepts for obligations, permissions, and prohibitions. Depending on the pragmatic context, the message payload is used, e.g. to update the internal knowledge of the agent (e.g., add new facts or rulebases), add new tasks (goals), or detect a complex event pattern (from event-instance sequences).

Several expressive logic formalisms are supported by Prova [17], e.g., for updating the knowledge base (transactional update logic), defining and detecting complex events (complex event algebra), handling situations/states (event calculus), as well as for reasoning (e.g., deontic logic for normative reasoning on permissions, prohibitions, obligations) and planning (abductive reasoning on plans

and goals).

In summary, Prova agents can interchange event information, rules (tasks), and queries/answers in agent conversations, including information about the semantics and pragmatics of the interchanged information.

Besides Prova, Rule Responder supports rule engines such as OO jDREW, Euler, DR-Device, and Drools for implementing such query answering agents as inference services in Rule Responder.

3 Rule Responder Agents

With the support of Prova's agent conversations, various distributed coordination topologies can be implemented, from centralized orchestration, executed in star-like agent nodes, to decentralized ad-hoc choreography within the Rule Responder agent network. In the following, we describe a common hierarchical agent topology which represents a centralized star-like structure for virtual organizations (and many orchestrated distributed systems). Organizational Agents (OAs) act as central orchestration nodes which control and disseminate the information flow from and to their internal Personal Agents (PAs) and the External Agents/Services (EAs).

3.1 Organizational Agent

An Organizational Agent (OA) represents its virtual organization as a whole. An OA manages its local Personal Agents (PAs), providing control of their life cycle and ensuring overall goals and policies of the organization and its semiotic structures. OAs can act as a single point of entry to the managed sets of local PAs to which requests by EAs are disseminated. This allows for efficient implementation of various mechanisms of making sure the PAs functionalities are not abused (security mechanisms) and making sure privacy of entities, personal data, and computation resources is respected (privacy & information hiding mechanisms). For instance, an OA can disclose information about the organization to authorized external parties without revealing private information and local data of the PAs, although this data might have been used in the PAs to compute the resulting answers to the external requester.

OAs, which require high levels of expressiveness to represent the logic of cognitive agents, are implemented using the Prova Semantic Web rule engine. In the following we will discuss some of the expressive language constructs of Prova that are required to implement the Rule Responder framework.

For implementing the Rule Responder communication flows in the OAs, Prova messaging reaction rules are used. A typical coordination pattern implemented in a Rule Responder OA is the following messaging reaction rule (Prova variables start with an upper-case letter), which waits for an incoming query from an EA and delegates this query to an internal responsible PA.

```

% receive query and delegate it to another party
rcvMsg(CID,esb, Requester, acl_query-ref, Query) :-
    responsibleRole(Agent, Query),
    sendMsg(Sub-CID,esb,Agent,acl_query-ref, Query),
    rcvMsg(Sub-CID,esb,Agent,acl_inform-ref, Answer),
    ... (other goals)...
    sendMsg(CID,esb,Requester,acl_inform-ref,Answer).

```

When activated by an incoming request from an EA, e.g. an HTTP request coming from a Web form, this messaging reaction rule first selects the responsible role for the query. Then the rule sends the query in a new sub-conversation to the selected party and waits for the answer to the query. That is, the rule execution waits until an answer event message is received in the inlined sub-conversation, which activates the process flow again, e.g. to prove further 'standard' goals, e.g. with information from the received answer, which is assigned to variables in the normal logic programming way, including also backtracking to other variable assignments. Finally, in this example, the rule sends back the answer to the original requesting EA.

The selection logic for the dissemination of queries to PAs is, e.g., implemented by a standard derivation rule which, e.g., accesses, via a Prova SPARQL query built-in, an external responsibility assignment matrix (RAM) (see section 3.4). The following rule selects responsible agents with a SPARQL query on a triple store Web interface, where the responsibility assignment matrix is stored.

```

% receive query and delegate it to another party
rcvMsg(CID,esb, Requester, acl_query-ref, Query) :-
    responsibleRole(Agent, Query),
    sendMsg(Sub-CID,esb,Agent,acl_query-ref, Query),
    rcvMsg(Sub-CID,esb,Agent,acl_inform-ref, Answer),
    ... (other goals)...
    sendMsg(CID,esb,Requester,acl_inform-ref,Answer).

```

RAMs (RACI matrices, Linear Responsibility Charts, etc.) are often in project management, when responsibilities are clearly defined for each role. It should be noted that Prova OAs can also implement other well-known agent coordination and negotiation mechanisms: for instance, a Contract Net coordination protocol, where PAs bid for the task offered by the OA and the OA selects the best PA according to the received bids, or a publish-subscribe protocol, where PAs are selected according to their subscriptions with the OA.

3.2 *Personal Agents*

Personal Agents (PAs) assist the local entities of a virtual organization. Often these are human roles in the organization. But, it might be also services or applications in, e.g. a service oriented architecture. A PA runs a rule engine which accesses different sources of local data and computes answers according to the local rule-based decision logic of the PA. Depending on the required

expressiveness to represent the PAs rule logic arbitrary rule engines can be used as long as they provide an interface to ask queries and receive answers which are translated into the common Reaction RuleML interchange format in order to communicate with other agents.

Importantly, the PAs might have local autonomy and might support privacy and security implementations. In particular, local information used in the PA rules becomes only accessible by authorized access of the OA via the public interfaces of the PA which act as an abstraction layer supporting security and information hiding. A typical coordination protocol is that all communication to EAs is via the OA, but the OA might also reveal the direct contact address of a PA to authorized external agents which can then start an ad-hoc conversation directly with the PA [6]. A PA itself might act as a nested suborganization, i.e. containing itself an OA providing access to a suborganization within the main virtual organization. This can be useful to represent nested organizational structures such as departments, project teams, and service networks.

3.3 External Agents

External Agents (EAs) constitute the points-of-contact that allow an external user or service to query the Organizational Agent (OA) of a virtual organization. An EA is based, e.g., on a Web (HTTP) interface that allows such an enquiry user to pose queries, employing a menu-based Web form, which gets translated to an equivalent RuleML/XML message. An external agent -- from the point of view of a Rule Responder agent organization -- can be an external human agent, a service/tool, or another external Rule Responder organization, thus leading to cross-organizational Rule Responder communication.

3.4 Responsibility Assignment Matrix

As one possible way for coordination in a virtual organization the Rule Responder framework uses a 'pluggable' Responsibility Assignment Matrix (RAM) to support the OA in its selection of a PA and its optional participating profiles underneath. A RAM describes the responsibility of agent roles in completing certain tasks or deliverables in a virtual organization. A standard RAM is a RAI matrix, with

- **Responsible** -- agents who do the work to achieve the task. Typically, the PAs are the responsible roles.
- **Accountable** (also Approver or final Approving authority) -- agent who is ultimately accountable for the correct and thorough completion of the deliverable or task, and the one to whom Responsible is accountable. Typically, this is the OA which receives the answer from the PA and further processes it

before forwarding it to the EA.

- Informed -- the agent who is kept up-to-date on progress, often only on completion of the task or deliverable; and with whom there is just one-way communication. Typically, this is the EA who is informed about the result by the OA.

In a simple star-like Rule Responder agent topology, a single RAI matrix can be used in the OA to map an incoming query to the PA whose local knowledge base is deemed to be best suited for answering it. The RAI matrix is represented as an OWL ontology (OWL Lite) and can be used by a Rule Responder agent via querying it with the Semantic Web built-ins of Prova, binding the respective roles and their responsibilities to typed variables in the agent's rule logic. Many variants of the RAM with different role distinctions are possible such as RACI (with Consulted agents), RASCI (with Supporting agents) etc. - see, e.g., table 1.

Table 1 Responsibility Assignment Matrix

	General Chair	Program Chair	Publicity Chair
Symposium	responsible	consulted	supportive
Website	accountable	responsible	
Sponsoring	informed, signs	verifies	responsible
Submission	informed	responsible	
...

For instance, the RAM has been split so that role responsibility assignment is done on the 'higher' level of a Group Responsibility Matrix (GRM) in the OA and on the 'lower' level of a Profile Responsibility Matrix (PRM) in the PAs. Figure 2 shows these two central matrices in the larger context of the Rule Responder architecture used in the WellnessRules instantiations (cf. Section 5.2), which has been further evolved for the PatientSupporter instantiation (cf. Section 5.3).

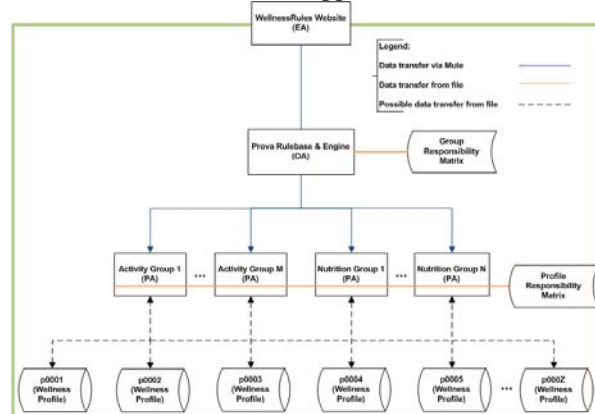


Fig. 2 Rule Responder architecture instantiated to WellnessRules


```

    <quantification> <!-- e.g. variable bindings-->          </quantification>
    <on>             <!-- event part -->                    </on>
    <if>            <!-- condition part -->                 </if>
    <then>         <!-- (logical) conclusion part -->      </then>
    <do>          <!-- action part -->                     </do>
    <after>       <!-- postcondition part after action, e.g.
                  to check effects -->                    </after>
</Rule>

```

Depending on which parts of this general rule syntax are used different types of reaction rules can be expressed, e.g. if-then (derivation rules), if-do (production rules), on-do (trigger rules), on-if-do (ECA rules). For communication between distributed rule-based (agent) systems Reaction RuleML provides a general message syntax:

```

<Message>
  <oid>           <!-- conversation ID-->                </oid>
  <protocol>     <!-- used protocol -->                  </protocol>
  <agent>       <!-- sender/receiver agent/service --> </agent>
  <directive> <!-- pragmatic primitive, i.e. context --> </directive>
  <content>     <!-- message payload -->                 </content>
</Message>

```

Using these messages agents can interchange events (e.g., queries and answers) as well as complete rule bases (rule set modules), e.g. for remote parallel task processing. Agents can be engaged in long running possibly asynchronous conversations and nested sub-conversations using the conversation id to manage the conversation state. The protocol is used to defines the message passing and coordination protocol. The directive attribute corresponds to the pragmatic instruction, i.e. the pragmatic characterization of the message context broadly characterizing the meaning of the message.

The Reaction RuleML translator services are configured in the transport channels of the inbound and outbound links of the deployed rule engines on the ESB. Incoming Reaction RuleML messages (receive) are translated into platform-specific rulebases which can be executed by the rule engine, e.g. Prova, and outgoing rulebases (send) are translated into Reaction RuleML in the outbound channels before they are transferred via a selected transport protocol.

The semantic agent architecture in Rule Responder supports privacy and security implementations. In particular, local information used in the PAs becomes only accessible by authorized access via the public interfaces of the OAs which act as an abstraction layer supporting security and information hiding. To achieve this, Prova supports an interface definition language (Reaction RuleML IDL) which allows descriptions of the signatures of publicly accessibly rule functions together with their mode and type declarations. *Modes* are states of instantiation of the predicate described by mode declarations, i.e. declarations of the intended input-output constellations of the predicate terms with the following semantics:

- "+" The term is intended to be input
- "-" The term is intended to be output
- "?" The term is undefined/arbitrary (input or output)

For instance, the interface definition for the function $add(Arg1, Arg2, Result)$ is $interface(add(-,+,+))$, i.e. the function is a public interface which expect two input arguments and returns one output argument. $add(X,1,1)$ would be a valid query to this public function.

External agents can access the virtual organization only via these public interfaces, which often only reveal abstracted information to authorized users and hence hide local information of the organization and its PAs.

5 Rule Responder Instantiations

Early instantiations of Rule Responder include the Health Care and Life Sciences eScience infrastructure [11], the Rule-based IT Service Level Management, and Semantic BPM system [12, 13]. Recent instantiations include multiple versions of the deployed SymposiumPlanner system [9], two versions of the WellnessRules prototype [5], PatientSupporter, a reputation management system, and a SCEP agent network. We will here highlight the principles of Rule Responder instantiations with an emphasis on the recent ones.

5.1 SymposiumPlanner

SymposiumPlanner is a series of deployed applications created with Rule Responder for the Q&A parts of the official websites of the RuleML Symposia. Rule Responder started to support the organizing committee of the RuleML Symposium [8] and was further developed to assist the yearly RuleML Symposia since 2007. These applications embody responsibility assignment, automated first-level contacts for information regarding the symposium, helping the publicity chair with sponsoring correspondence, helping the panel chair with managing panel participants, and the liason chair with coordinating organization partners. SymposiumPlanner utilizes a single organizational agent to handle the filtering and delegation of incoming queries. Each committee chair has a personal agent that acts in a rule-governed manner on behalf of the committee member. Each agent manages personal information, such as a FOAF-like profile containing a layer of facts about the committee member as well as FOAF-extending rules. These rules allow the PA to automatically respond to requests concerning the RuleML Symposium. Task responsibility for the organization is currently managed through a responsibility matrix, which defines the tasks committee members are responsible for. The matrix and the roles assigned within the virtual organization are defined by an OWL (Ontology Web Language) Lite Ontology. Request users and personal agents can communicate by sending messages that transport queries, answers, or complete rulebases through the public EA interface of the OA (typically, an EA uses an HTTP port to which post and get requests are

sent from a Web form). The Rule Responder instantiations to SymposiumPlanner are published and deployed online.⁷

5.2 *WellnessRules*

This is a Web 3.0 case study, where ontology-structured rules (including facts) about wellness opportunities are created by participants in rule languages such as Prolog and N3, and translated for interchange within a wellness community using RuleML/XML. The wellness rules are centered around participants, as profiles, encoding knowledge about their activities, nutrition, etc. conditional on the season, the time-of-day, the weather, etc. This distributed knowledge base extends fact-only FOAF profiles with a vocabulary and rules about wellness group networking.

The communication between participants is organized through Rule Responder, permitting translator-based reuse of wellness profiles and their distributed querying across engines. WellnessRules interoperates between rules and queries in the relational (Datalog) paradigm of the pure-Prolog subset of POSL and in the frame (F-logic) paradigm of N3. These derivation rule languages are implemented in the engines OO jDREW and Euler, and connected via Rule Responder to support wellness communities.

WellnessRules is a system supporting the management of wellness practices within a community based on rules plus ontologies. The idea is the following. As in Friend of a Friend (FOAF)⁸, people can choose a (community-unique) nickname and create semantic profiles about themselves, here about their wellness practices, for their own planning and to network with other people supported by a system that 'understands' those profiles. As in FindXpRT [10], such FOAF-like fact-only profiles are extended with rules to capture conditional person-centered knowledge such as each person's wellness activity depending on the season, the time-of-day, the weather, etc. People can use rules of various refinement levels and rule languages ranging from pure Prolog to N3, which will be interoperated through RuleML/XML [3].

Interoperating with translators, WellnessRules thus frees participants from using any single rule language. In particular, it bridges between Prolog as the main Logic Programming rule paradigm and N3 as the main Semantic Web rule paradigm. The distributed nature of Rule Responder profiles, each queried by its own (copy of an) engine, permits scalable knowledge representation and processing.

WellnessRules has recently been developed to WellnessRules2, using a fourth kind of agent, the Computing Agent (CA), for accessing Google weather data. From the point of an OA, a CA can be queried similarly to a PA. However, while a PA is a personal assistant to a human owner, a CA is a pure machine agent, in

⁷<http://ruleml.org/SymposiumPlanner>

⁸<http://www.foaf-project.org/>

WellnessRules2 acting as a wrapper for a Google service. The Rule Responder instantiations to WellnessRules are further described and demoed online.⁹

5.3 *PatientSupporter*

Patients are increasingly seeking interaction in support groups, which provide shared information and experience about diagnoses, treatment, etc. PatientSupporter is an instantiation of Rule Responder that will permit a patient to query other patients' profiles for finding or initiating a matching group.

Rule Responder's External Agent (EA) is a Web-based patient-organization interface that passes queries to the Organizational Agent (OA). The OA represents the common knowledge of the virtual patient organization, delegates queries to relevant Personal Agents (PAs), and hands validated PA answers back to the EA. Each PA represents the medical subarea of primary interest to a corresponding patient group. The PA assists its patients by advertising their interest profiles employing rules about diagnoses and treatments as well as interaction constraints such as time, location, age range, gender, and number of participants.

PAs can be distributed across different rule engines using different rule languages (e.g., Prolog and N3), where rules, queries, and answers are interchanged via translation to and from RuleML/XML. The current implementation of PatientSupporter applies to a use case where the PA's medical subareas are defined through sports injuries structured by a partonomy of affected body parts.

PatientSupporter uses ontologies and rules for organizing geographically distributed patients -- here, suffering from sports injuries -- into virtual support groups around classes of an ontology of injuries -- here, a sports-injury partonomy. The prototype is designed to help patients with a similar sports injury to interact with a virtual support group having that common interest. Patients in an online PatientSupporter virtual organization create their semantic profile referring to classes in a disease ontology -- here a partonomy of body parts affected by sports injuries. Profiles contain rules about diagnoses and treatments as well as interaction constraints such as time, location, age range, gender, and number of participants. A patient can pose queries against the semantic profiles of other patients in his or her virtual organization to find or initiate a matching group.

PatientSupporter allows patients to have their profiles expressed in either Pure Prolog (Logic Programming rules) or N3 [2] (Semantic Web rules). Providing these quite different rule language paradigms permit patients to choose the language that best suits them. Rule Responder handles the interoperation between the rule languages of different patients using translators to and from RuleML/XML as the interchange format [7, 3].

Patients using the PatientSupporter Social Semantic Web portal are able to initiate the virtual support group about their sports injury on a global scale. They also

⁹<http://ruleml.org/WellnessRules> and <http://ruleml.org/WellnessRules2>

benefit from PatientSupporter's interoperation facility in the background -- to transform patient profiles between Pure Prolog and N3 through RuleML/XML. The system employs a partonomy of sports-injury-affected body parts (a 'body partonomy'), which makes it easy for patients to navigate hierarchically up or down to increase recall or precision, respectively. A patient's queries invoke other patients' interaction rules, allowing him or her to narrow down the search in a step-wise fashion. All of this saves a patient from browsing through a large set of irrelevant patient profiles and permits him or her to efficiently converge on a first Skype call.

The Rule Responder instantiation to PatientSupporter is being described and demoed online.¹⁰

5.4 Reputation Management System

The Rule Responder reputation management system [1] is based on distributed Rule Responder rule agents, which use rules for implementing the reputation management functionalities as rule agents, and which use Semantic Web ontologies for representing simple or complex multi-dimensional reputation objects. This Semantic Web reputation ontology model enables reputation portability, eases the management of reputation data, mitigates risks in open environments, and enhances the decision making process in the reputation processing agents. The reputation management system computes, manages, and provides reputation about entities which act on the Web. It is implemented as a *Reputation Processing Network (RPN)* consisting of *Reputation Processing Agents (RPAs)* that have two different roles:

1. *Reputation Authority Agents (RAAs)*: Act as reputation scoring services for the repute entities whose Reputation Objects (ROs) are being considered or calculated in the agents' rule-based Reputation Computation Services (RCSs). An RCS runs a rule engine which accesses different sources of reputation (input) data from the reputors about an entity and evaluates an RO based on its declarative rule-based computational algorithms and contextual information available at the time of computation.
2. *Reputation Management Agents (RMAs)*: Act as a reputation trust center offering reputation management functionalities. An RMA manages the local RAAs providing control of their life-cycle in particular, and also ensuring goals such as fairness. It might act as a Reputation Service Provider (RSP) which aggregates reputations from the reputation scores of local RAAs. Based on the final calculated reputation, it might also perform actions, e.g. compute trust-worthiness, make automated decisions, or trigger reactions. It also manages the communication with the reputors, collecting data about entities from them, generates reputation data inputs for the reputation scoring, and distributes the

¹⁰<http://ruleml.org/PatientSupporter>

data to the RAAs. It might also act as central point of communication for the real repute entities (e.g., persons) giving them legitimate control over their reputation and allowing entities the governance of their reputations.

The agent-based approach to online reputation management ensures efficient automation, semantic interpretability and interaction, openness in ownership, fine-grained privacy and security protection, and easy management of semantic reputation data on the Web.

5.5 Semantic Complex Event Processing Agent Network

The Event Processing Network (EPN) [16] consists of Semantic Event Processing Agents (EPA) implemented as distributed Prova inference services which detect complex events using Prova's rule-based Semantic Complex Event Processing (SCEP) logic. [19]. The multi-agent approach allows for a highly-available distributed implementation with redundant Event-Calculus based state processing where events are processed concurrently in the EPN.

6 Conclusion

Rule Responder is a framework for specifying virtual organizations as semantic multi-agent systems. Characteristics of Rule Responder include

- the coverage of the distributed processing spectrum from Web Services to agents in one framework
- the recursive (holonic) modeling of a virtual organization of services and agents as a single agent,
- the use of ESBs, especially Mule, as a foundation for the Semantic and Pragmatic Web infrastructure,
- the use of Semantic-Pragmatic Web rules as the main knowledge representation, complemented by ontologies,
- the introduction of PAs as human-assisting agents into a virtual organization, besides the traditional computation-performing agents (CAs),
- the design of a 'pluggable' agent-finding mechanism from role assignment to Semantic Service discovery.

The Rule Responder framework, with its increasing number of users and engines (Prova, OO jDREW, DR-Device, Euler, and Drools), is thus being proposed as a reference architecture for distributed knowledge representation and processing.

Acknowledgments The international Rule Responder initiative has greatly helped us with work leading to this chapter. In particular, we want to thank Alexander Kozlenkov, Benjamin Craig, Taylor Osmun, Derek Smith, Omair

Shafiq, Mahsa Kiani, Kia Teymourian, Rehab Alnemr, Irfan ul Haq, Nick Bassiliades, Stratos Kontopoulos, and Kalliopi Kravari.

References

- [1] Rehab Alnemr and Adrian Paschke and Christoph Meinel. Enabling Reputation Interoperability through Semantic Technologies. *ACM International Conference on Semantic Systems*, 2010. ACM.
- [2] Tim Berners-Lee and Dan Connolly and Lalana Kagal and Yosi Scharf and Jim Hendler. N3Logic: A Logical Framework For the World Wide Web. *Theory and Practice of Logic Programming (TPLP)*, 8(3), 2008.
- [3] Harold Boley. Are Your Rules Online? Four Web Rule Essentials. In A. Paschke and Y. Biletskiy, editors, *Proc. Advances in Rule Interchange and Applications, International Symposium (RuleML-2007), Orlando, Florida in LNCS*, pages 7--24, 2007. Springer.
- [4] Harold Boley and Elizabeth Chang. Digital Ecosystems: Principles and Semantics. *Proc. IEEE Intl. Conf. Digital Ecosystems and Technologies, Cairns, Australia*, 2007.
- [5] Harold Boley and Taylor Michael Osmun and Benjamin Larry Craig. Social Semantic Rule Sharing and Querying in Wellness Communities. In Asunción Gómez-Pérez and Yong Yu and Ying Ding, editors, *The Semantic Web, Fourth Asian Conference, ASWC 2009, Shanghai, China, December 6-9, 2009. Proceedings in Lecture Notes in Computer Science*, pages 347--361, 2009. Springer.
- [6] Harold Boley and Adrian Paschke. Expert Querying and Redirection with Rule Responder. In Anna V. Zhdanova and Lyndon J. B. Nixon and Malgorzata Mochol and John G. Breslin, editors, *Proceedings of the 2nd International ISWC+ASWC Workshop on Finding Experts on the Web with Semantics, Busan, Korea, November 12, 2007 in CEUR Workshop Proceedings*, pages 9--22, 2007. CEUR-WS.org.
- [7] Harold Boley and Said Tabet and Gerd Wagner. Design Rationale of RuleML: A Markup Language for Semantic Web Rules. *Proc. Semantic Web Working Symposium (SWWS'01)*, pages 381-401, 2001. Stanford University.
- [8] Benjamin Craig. The OO jDREW Engine of Rule Responder: Naf Hornlog RuleML Query Answering. In Adrian Paschke and Yevgen Biletskiy, editors, *Advances in Rule Interchange and Applications - International Symposium, RuleML 2007, Orlando, Florida, October 25-26, 2007, Proceedings in Lecture Notes in Computer Science*, 2007. Springer.
- [9] Benjamin Larry Craig and Harold Boley. Personal Agents in the Rule Responder Architecture. In Nick Bassiliades and Guido Governatori and Adrian Paschke, editors, *Rule Representation, Interchange and Reasoning on the Web, International Symposium, RuleML 2008, Orlando, FL, USA, October 30-31, 2008. Proceedings in Lecture Notes in Computer Science*, pages 150--165, 2008. Springer.

- [10] Jie Li and Harold Boley and Virendrakumar C. Bhavsar and Jing Mei. Expert Finding for eCollaboration Using FOAF with RuleML Rules. *Montreal Conference of eTechnologies 2006*, pages 53-65, 2006.
- [11] Paschke, Adrian. Rule responder HCLS eScience infrastructure. *ICPW '08: Proceedings of the 3rd International Conference on the Pragmatic Web*, pages 59--67, New York, NY, USA, 2008. ACM.
- [12] Paschke, Adrian and Bichler, Martin. Knowledge representation concepts for automated SLA management. *Decis. Support Syst.*, 46(1):187--205, 2008.
- [13] Adrian Paschke and Alexander Kozlenkov. A Rule-based Middleware for Business Process Execution. *Multikonferenz Wirtschaftsinformatik*, 2008.
- [14] Geoff Sutcliffe and Randy Goebel, editors. *Proceedings of the Nineteenth International Florida Artificial Intelligence Research Society Conference, Melbourne Beach, Florida, USA, May 11-13, 2006*, 2006. AAAI Press.
- [15] A. Koestler. *The Ghost in the Machine*. Hutchinson & Co, London, 1967.
- [16] Alexander Kozlenkov and David Jeffery and Adrian Paschke. State management and concurrency in event processing. *DEBS*, 2009.
- [17] Paschke, A. *Rule-Based Service Level Agreements - Knowledge Representation for Automated e-Contract, SLA and Policy Management*. Idea Verlag GmbH, Munich, 2007.
- [18] Adrian Paschke. A Typed Hybrid Description Logic Programming Language with Polymorphic Order-Sorted DL-Typed Unification for Semantic Web Type Systems. *CoRR*, abs/cs/0610006, 2006.
- [19] Kia Teymourian and Adrian Paschke. Towards semantic event processing. *DEBS*, 2009.
- [20] Welsh, M. and Culler, D. and Brewer, E. SEDA: An Architecture for Well Conditioned, Scalable Internet Services. *Proceedings of Eighteenth Symposium on Operating Systems (SOSP-18)*, Chateau Lake Louise, Canada, 2001.
- [21] M. Wooldridge. *An Introduction to MultiAgent Systems*. John Wiley & Sons, 2001.