



I5-D6

Reactive rule ontology: RDF/OWL level

Project title:	Reasoning on the Web with Rules and Semantics
Project acronym:	REWERSE
Project number:	IST-2004-506779
Project instrument:	EU FP6 Network of Excellence (NoE)
Project thematic priority:	Priority 2: Information Society Technologies (IST)
Document type:	D (deliverable)
Nature of document:	R (report)
Dissemination level:	PU (public)
Document number:	IST506779/Lisbon and Göttingen/I5-D6/D/PU/a1
Responsible editors:	José Júlio Alferes and Wolfgang May
Reviewers:	n/a
Contributing participants:	Göttingen, Lisbon
Contributing workpackages:	I5
Contractual date of deliverable:	28 February 2007
Actual submission date:	11 April 2007

Abstract

This deliverable is concerned with the definition of a concrete ontology for reactive rules at the RDF/OWL level, in line with the general framework defined in deliverable I5-D4 – Specification of a Model, Language and Architecture for Reactivity and Evolution.

Keyword List

ECA rules, Reactivity, OWL ontology, Language and data heterogeneity

Reactive rule ontology: RDF/OWL level

José Júlio Alferes¹, Ricardo Amador¹, Erik Behrends², Oliver Fritzen², Tobias Knabke², Wolfgang May², Franz Schenk², Daniel Schubert²

¹ Centro de Inteligência Artificial - CENTRIA, Universidade Nova de Lisboa

² Institut für Informatik, Universität Göttingen

11 April 2007

Abstract

This deliverable is concerned with the definition of a concrete ontology for reactive rules at the RDF/OWL level, in line with the general framework defined in deliverable I5-D4 – Specification of a Model, Language and Architecture for Reactivity and Evolution.

Keyword List

ECA rules, Reactivity, OWL ontology, Language and data heterogeneity

Contents

1	Introduction	1
I	MARS: Modular Active Rules in the Semantic Web	3
2	Introduction	5
3	Ontology of Rule-Based Behavior	7
3.1	Requirements Analysis	7
3.1.1	Web vs. Semantic Web	8
3.1.2	Abstraction Levels	8
3.1.3	Domain Ontologies including Dynamic Aspects	10
3.1.4	Events	11
3.1.5	Types of Rules	15
3.2	Simple ECA Rules: Data Model Triggers	18
3.2.1	Triggers on XML Data	18
3.2.2	Triggers on (Plain) RDF Data	20
3.2.3	Triggers on RDFS and OWL Data	20
3.2.4	Triggers vs. ECA Rules	20
3.3	ECA Language Structure	21
3.3.1	Language Heterogeneity and Structure: Rules, Rule Components and Languages	22
3.3.2	Components and Languages of ECA Rules	23
3.3.3	Markup Proposal: ECA-ML	24
3.3.4	Hierarchical Structure of Languages	25
3.3.5	Common Structure and Aspects of E, C, T and A Sublanguages	25
3.3.6	Language Information	28
3.3.7	Opaque Rules and Opaque Components	29
3.3.7.1	Opaque Rules	29
3.3.7.2	Opaque Components	30
4	Abstract Semantics: Rule Level	35
4.1	Abstract Declarative Semantics of Rule Execution	35
4.1.1	Rule Semantics	35
4.1.2	Logical Variables	36
4.1.3	Horizontal Communication: Logical Semantics	37
4.1.4	Vertical Communication	38
4.1.5	Communication Modes and Declaration of Variables	39
4.2	Logical Variables: Markup for Communication	44
4.2.1	Basic Interchange of Variable Bindings	44
4.2.2	Downward Communication: Variable Bindings	44
4.2.3	Upward Communication: Results and Variable Bindings	45

4.3	Markup: Binding and Using Variables	47
4.3.1	Alternative Syntaxes	47
4.3.2	Discussion: Variable Syntax	48
4.3.3	Variable Bindings by ECA Rules	49
4.4	Operational Aspects of Rule Execution	51
4.4.1	Firing ECA Rules: the Event Component	51
4.4.2	The Query Component	54
4.4.3	The Test Component	56
4.4.4	Summary of Event, Query and Test Semantics	57
4.4.5	The Action Component	57
4.4.6	Transactions	59
4.4.7	Examples	59
5	Abstract Semantics and Communication: Component Services	65
5.1	General Communication Patterns	65
5.2	The Event Component: Structure and Languages	67
5.2.1	Atomic Events	67
5.2.1.1	XML Representation of Atomic Events	67
5.2.2	Atomic Event Descriptions and Formalisms	68
5.2.2.1	Event Specification by XML-QL Style Matching	69
5.2.2.2	Navigation-Based Event Specification	70
5.2.2.3	Event Specification by Opaque XQuery	70
5.2.3	Event Algebras	71
5.2.4	Embedding Algebraic Languages	72
5.2.5	Embedding Atomic Events in Composite Events and Rules	72
5.2.6	Example: SNOOP	74
5.2.7	Related Work and Existing (Sub)languages	75
5.3	Architecture and Communication: ECA, CED, and AEM	76
5.3.1	Abstract Semantics and Markup of Event Detection Communication	77
5.3.2	Communication for Event Components between ECA and CED	78
5.3.3	Communication for Event Matching with AEMs	79
5.3.4	Identification of an AEM for an Atomic Event Specification	81
5.3.4.1	Example	82
5.4	The Query Component	84
5.4.1	Opaque Queries	85
5.4.2	Atomic Queries	85
5.4.3	Composite Queries	85
5.4.4	Result Semantics	86
5.5	The Test Component	86
5.5.1	Test Component Languages	87
5.5.2	Atomic Tests: Predicates	88
5.6	The Action Component	89
5.6.1	Atomic Actions	90
5.6.2	Composite Actions	92
5.6.3	Atomic and Leaf Items	96
5.6.4	Example: CCS	97
5.6.5	Examples in CCS	98
5.6.6	Processing	100
5.6.7	ECA Rules vs. CCS	100
5.7	Summary	101

6	Domain Ontologies in the MARS Framework	103
6.1	Domain Ontologies	103
6.2	Description of Application Services	105
6.3	Rules in Ontologies	105
6.3.1	Derivation Rules	106
6.3.2	ECE Rules	106
6.3.3	ACA Rules	107
6.3.3.1	Definitory ACA Rules on the Ontology Level	107
6.3.3.2	Implementational ACA Rules from the Ontology Level to the Data Model Level	107
6.3.4	Discussion: Comparison with RuleML Proposal	108
7	Domain Application Nodes	109
7.1	External Interface of Domain Nodes	109
7.1.1	Providing Static Data	109
7.1.2	Executing Action Requests	110
7.1.3	Reporting About Behavior: Providing Atomic Events	110
7.1.4	Summary	110
7.2	Domain Node Local Behavior	110
7.3	Domain Node Architecture Proposal	111
7.3.1	Functionality of the Node Core	111
7.3.2	Functionality of the Node Wrapper: Mapping of Actions	112
7.3.3	Installing and Running a Domain Node	113
8	Domain Brokering	115
8.1	Event Brokering	115
8.2	Brokering of Derived Events	116
8.2.1	RSS-based Event Brokering	117
8.3	Query Brokering	117
8.4	Action Brokering	118
8.4.1	Required Information for Action Brokering	118
8.4.2	Handling of Composite Actions by ACA Rules	120
9	The MARS Architecture, Ontology, Language and Service Metadata	123
9.1	Architecture and Processing: Cooperation between Resources	123
9.1.1	Event Detection	123
9.1.2	Query Processing	124
9.1.3	Action Processing	124
9.2	Architectural Variants	125
9.3	Ontology of Languages and Services	126
9.3.1	The Ontology	126
9.3.2	The MARS Subontology of Domain-Related Notions	128
9.3.3	MARS Languages and Services Ontology	129
9.4	Service Interfaces and Functionality	132
9.4.1	ECA Services	132
9.4.1.1	Upper Interface:	132
9.4.1.2	Lower Interface:	133
9.4.2	(Algebraic) Component Languages/Services (General)	133
9.4.3	Domain Brokers	134
9.4.4	Domain Services	134
9.4.5	The Services Ontology	135
9.5	Locating and Contacting Language Services	136
9.5.1	Language & Service Registries	137
9.5.2	Interface Descriptions of Individual Tasks	137

9.5.3	The Languages and Services Registry RDF Model	140
9.5.4	Service Brokering: Generic Request Handlers	145
9.5.5	Using the LSR	148
9.6	Contacting Domain Services	149
9.7	Issue: Redundancy and Duplicates in Communication	149
10	Implementation and Prototype	151
10.1	Simple Setting for an ECA Prototype	151
10.1.1	ECA Module Implementation	152
10.1.2	Handling Opaque Queries	153
10.1.2.1	Communicating with Primitive Services	153
10.1.2.2	Framework-Aware Wrappers	153
10.1.2.3	Raising Events by Opaque Atomic Actions	154
10.1.2.4	Web Service Calls via HTTP/SOAP	155
10.1.2.5	Opaque: Matching Regular Expressions	155
10.1.2.6	Summary	155
10.2	Extending the Prototype with Component Services	155
10.2.1	Composite Event Detection and Atomic Event Matchers	155
10.2.2	Queries and Updates	156
10.2.3	Composite Actions and Processes	157
10.3	Application Domains	157
10.4	Infrastructure	157
10.5	Using the Prototype Demonstrator	158
10.6	Anticipated “Foreign” Modules	158
10.6.1	Wrapped Composite Event Detection Engines	158
10.6.2	Wrapped Query Engines	158
11	The RDF Level: Language Elements and their Instances as Resources	159
11.1	Ontologies vs. Markup Languages	159
11.1.1	Fragment Translation between both Representations	160
11.1.2	Glue in XML and RDF Representations	160
11.1.3	Design of Ontologies, Languages and Namespaces	161
11.2	Ontology of Expressions	162
11.3	Subontology of General Expressions	162
11.4	Subontology of Algebraic Expressions	163
11.5	Subontology of Variable Usage	164
11.6	The ECA-ML Ontology	166
11.6.1	Handling Opaque Stuff	166
11.6.2	Assertions	166
11.6.3	Handling Atomic Stuff	167
11.6.4	The ECA Ontology, OWL File, and ECA-ML Markup	167
11.6.5	The Implicit Event Variable	169
11.6.6	The XML Representation: ECA-ML	170
11.6.7	Some Consequences	170
11.7	The Snoop Ontology	170
11.7.1	The Snoopy OWL File	170
11.7.2	The XML Representation: Snoopy-ML	176
11.8	The CCS Process Ontology	176
11.9	The Composite Formulas and Queries Ontology	179

12 Atomic Event Matching, Queries and Atomic Actions in RDF	181
12.1 Discarded Pattern Approaches	182
12.1.1 The Embedded Pattern Approach	182
12.1.2 Embedded Event and Query Patterns as parseType=XML	184
12.2 Formulas, Queries, Tests, Assertions, Atomic Event Matching, and Atomic Actions in RDF	185
12.2.1 The OWLQ Conjunctive Query Language	186
12.2.2 Specification of Atomic Events by OWLQ Queries	188
12.2.3 Evaluation of OWLQ Conjunctive Query Specifications	190
12.2.4 Specification of Atomic Actions: RDF-CL	190
13 Rules and Rule Components as Resources	193
13.1 Putting the Ontologies Together	193
13.2 Reasoning about Rules	201
13.2.1 Validity	201
13.2.2 Safety and Evaluation	201
13.2.3 Typing	202
13.2.4 Compositionality	202
14 MARS Appendices	203
14.1 Abbreviations	203
14.2 MARS Infrastructure URIs and Filenames	204
14.2.1 RDF Files	204
14.2.2 URI Schemata	205
14.3 Language DTDs	205
14.3.1 DTD for Logical Stuff: Variable Bindings etc.	205
14.3.2 DTD of ECA-ML	205
14.3.3 Snoop DTDs	205
14.3.3.1 The SNOOP Striped DTD	205
14.3.3.2 The SNOOP-ML DTD	206
II r^3: Towards a foundational ontology for reactive rules	209
15 The r^3 ontology	211
15.1 An Ontology for Reactive Rules	211
15.2 Defining Reactive Rule Languages	213
15.3 Defining Reactive Rule Constructions	217
15.4 From ECA-ML to r^3 DF	220
15.5 Related Work	222
15.6 Final remarks	222

Chapter 1

Introduction

The present deliverable includes the ever-growing report on the ECA Framework whose previous versions were regularly attachments to previous reports, starting with Deliverable I5-D4 where the General Model for Evolution and Reactivity in the Web was first defined. It substantiates that the way followed during the project successfully lead to (i) a model of rules acting in the Semantic Web, and (ii) a *Semantic* representation of these rules.

The report substantially progressed in order to lift the results to the level of RDF/OWL, as required by the present deliverable. The main changes from the previous description (attached to I5-D5) were:

- The Framework got a name: “MARS – Modular Active Rules in the Semantic Web”. On the concept level, the Framework describes both the MARS Prototype at Göttingen, and the r^3 Prototype at Lisbon.
- Chapter 9 which contains the “runtime” metadata, i.e., Languages, Services, etc. is now fully represented in RDF.
- Chapters 11, 12, and 13 describe how the rules themselves are represented in RDF. Chapter 11 gives the ontologies of the higher levels, ECA-ML, general expressions, and the sample languages SNOOP and CCS that have already been discussed in the previous deliverables in their XML representation. Chapter 12 describes how atomic expressions are *described* in RDF, and Chapter 13 concludes this to an overall RDF/OWL description of the rules.

Besides the inclusion of the report on the current status of the general framework, this deliverable also presents results of the r^3 project towards defining an (OWL-DL) foundational ontology for reactive rules, that serves as basis for the r^3 implementation described in Deliverable I5-D7 – Completion of the prototype scenario. This proposal is at a low (structural) abstraction level; the extension of this proposal towards characterizing higher abstraction level concepts, like domain/application specific languages (vs. algebraic and general-purpose languages) is under consideration taking into account the results described in the MARS part.

Part I

MARS: Modular Active Rules in the Semantic Web

Chapter 2

Introduction

The static aspect of the Semantic Web deals with providing computer-understandable semantic information of Web data. An important task for this is to provide homogeneous languages and formats throughout the Web as an extension to today's *portals* – at least as views over heterogeneous data sources. With URIs, XML, RDF and recently OWL, there is a clear perspective what the Semantic Web will look like.

In contrast to the current Web, the *Semantic Web* should be able not only to support querying, but also to propagate knowledge and changes in a semantic way. This *evolution* and *behavior* depends on the cooperation of nodes. In the same way as the goal of the *Semantic Web* is to bridge the heterogeneity of data formats, schemas, languages, and ontologies used in the Web to provide semantics-enabled unified view(s) on the Web, the heterogeneity of concepts for expressing behavior requires for an appropriate handling on the semantic level. When considering dynamic issues, the concepts for describing and implementing behavior will surely be diverse, due to different needs, and it is unlikely that there will be a unique language for this throughout the Web. Since the Web nodes are prospectively based on different concepts such as data models and languages, it is important that *frameworks* for the Semantic Web are modular, and that the *concepts* and the actual *languages* are independent. As such, besides having a concrete language for dealing with evolution and reactivity, the Semantic Web calls for the existence of a framework able to deal with this heterogeneity of languages.

In this respect, *reactivity* and its formalization as *Event-Condition-Action (ECA) rules*, provide a suitable common model because they provide a modularization into clean concepts with a well-defined information flow. ECA rules can be used for providing a generic uniform framework for specifying and implementing communication, local evolution, policies and strategies, and altogether global evolution in the Semantic Web.

In this paper, we propose an ontology-based approach for describing (reactive) behavior in the Web and evolution of the Web that follows the ECA paradigm: MARS is a modular framework for *composing* languages for events, conditions, and actions by separating the ECA semantics from the underlying semantics of events, conditions and actions. This modularity allows for high flexibility wrt. the heterogeneity of the potential sublanguages, while exploiting and supporting their meta-level *homogeneity* on the way to the Semantic Web.

Another important aspect when considering the *Semantic Web* is that of abstraction levels of behaviour. As it will be detailed below, it is our opinion that evolution and reactivity will appear at several levels in the Semantic Web: there will be local basic events and actions, similar to local database triggers; events on local XML data; global rules on the XML level that are able to react on views that include remote data; and application-level events and rules referring to the terms of the ontology of an application. With the MARS framework we propose to deal with different levels of behaviour.

Moreover, the ECA rules do not only operate on the Semantic Web, but are themselves also part of it. In general, ECA rules (and their components) must be communicated between different nodes, and may themselves be subject to being updated; also reasoning about evolution might

be desired. For that, the ECA rules themselves must be represented as objects of the (Semantic) Web, adhering to an own ontology of rules, and marked up in an (XML) Markup Language of ECA Rules. A markup proposal for active rules can be found already in RuleML [66], but it does not tackle the complexity and language heterogeneity of events, actions, and the generality of rules, as described here.

Structure of the Presentation. The next chapter, Chapter 3, develops the MARS ontology for behavior in the Semantic Web. The abstract semantics from the point of the view of the rule level, abstracting from the actual semantics of the components is described in Chapter 4. In Chapter 5 we proceed to the component level and investigate the abstract semantics and communication structure of the Event/Query/Test/Action component languages, including some (sample) concrete languages. Here we focus on the event component by describing how the event algebra similar to that of SNOOP [24] is mapped into the MARS framework. Then, three chapters deal with domain aspects: Chapter 6 analyzed the notion of domain ontologies wrt. the behavioral notions dealt with in the MARS Framework. The architecture of actual domain nodes is then described in Chapter 7. Domain Brokering is then discussed in Chapter 8. The Web architecture and the formal ontology and metadata of the MARS framework is described in Chapter 9. The language and service metadata contained in this ontology is also used for establishing actual communication between the component services. Chapter 10 gives an overview of the prototype on the XML level.

Chapters 11 – 13 continue the approach to the RDF level: Chapters 11 lifts the ECA language and the component languages to ontologies. Chapter 12 investigates how to describe atomic event specification and atomic action patterns on the RDF level. Finally, Chapter 13 puts these things together to complete rules on the RDF level.

Publications from this report. A preliminary version of the MARS framework described in this chapter has been published in [3], the ontology of rules, rule components and languages, and the service-oriented architecture proposal have been published in [53], and the languages and their markup, communication and rule execution model can be found in [51].

A first version of the ECA engine has been described in [11]. The proposal to use CCS in the action part is discussed in [8]; a detailed description of the algebraic handling of composite events and actions is given in [10]. The development of the domain node and domain broker architecture is described in [55] and [7].

Chapter 3

Ontology of Rule-Based Behavior

We start by summarizing some requirements posed by evolution and reactivity on the *Semantic Web*. Here, after some brief points on the important differences between dealing with the level of the Web versus the level of the Semantic Web, we elaborate on the various abstraction levels of behaviour that need to be considered in the semantic level. This leads a discussion of how to extend the domain ontologies with actions and events, and the various types of events that need to be considered in a general framework. Then, we summarize what kinds of rules, according to their tasks in the MARS framework, have to be covered. Section 3.2 analyzes ECA rules on the lowest level, namely triggers on the data model level.

We then discuss the main issue of dealing with the above-mentioned heterogeneity of languages. For this we start, in Section 3.3, by proposing a general structure, rule level ontology and corresponding markup, for (ECA) reactive rules. In it, basically each rule consists of:

- An *event part*, in which there must be a description of something that, if it happens, fires the rule;
- A *condition part*, which, depending on the detected event, may collect some further (static) data and test conditions on both the event and the collected information to check that an action has actually to be executed. Accordingly, this condition part can be further decomposed into one or more *query parts* for collecting data, and one *test part* for checking the (mainly boolean) condition;
- An *action part*, describing what to do when the event is detected and the condition test succeeds.

The components of a rule use different sublanguages for expressing events, queries, tests or actions. Not only the sublanguages of each family share some properties, but there are common properties of all kind of such sublanguages. We analyse the structure of these languages from the semantical and structural aspects and summarize their common aspects: they are algebraic languages, consisting of nested expressions.

3.1 Requirements Analysis

This section analyses the requirements for ECA-based evolution and reactivity in the Semantic Web and sets some working hypotheses. We investigate the abstraction levels used in the Semantic Web and its infrastructure, the structure of domain ontologies when dealing with dynamic aspects, then have a more detailed look on the kinds of events that have to be modeled, and classify several kinds of rules that are then actually needed. Finally we “initialize” the way towards our approach by discussing *triggers* as very simple ECA rules and illustrate why this approach provides a good base, but that a comprehensive framework must extend this idea in many aspects.

3.1.1 Web vs. Semantic Web

Whereas in the conventional XML/HTML Web, ECA models and languages that operate on the data level (as often in the literature, we distinguish between the *data level* and the *information or knowledge level*, which includes an additional knowledge base and reasoning) and on explicit events are sufficient, the situation for a Semantic Web framework is much more complex:

- Model (RDF): With RDF, the same resource can be described at different physical locations, using its URI. Thus, changes in the description of “something” are not necessarily located at a given node.
- Model (OWL): While some research has already been done in the area of queries and static reasoning on the OWL level, the extension to events and actions is still completely open. Domain ontologies must define their (derived) atomic events in terms of changes to the underlying data, and, in case in addition to just execution of rules, *reasoning* is intended, also actions must be described in terms of their effects on the data.

Developing an approach and case-studies for this has been identified as an important task for understanding what functionality and expressiveness should be provided by languages for describing behavior in the Semantic Web.

- Model and Languages: Rules in the Semantic Web exist on different abstraction levels (see Section 3.1.2) and should “cover” existing approaches. For this, composite events, conditions involving several nodes, and complex actions must be supported. Here, the existing and expected future heterogeneity has to be taken into account.
- Model, Languages and Architecture: Rules are themselves part of the Semantic Web. For this, they have to be seen as resources. On a smaller granularity, rule components and smaller identifiable parts like individual event descriptions are also resources. Rules and rule components have to be described not only syntactically in terms of a *programming language*, but on the (*rule*) *ontology level* which is then *translated* to actual, executable specifications in one or more programming languages.
- Languages and Architecture: Languages are resources (that have to be described by a suitable ontology). From this point of view, semantics and processors that implement this semantics are also resources and have to be described and correlated on the ontology level.

3.1.2 Abstraction Levels

Data Model Abstraction Levels. As described above, the *Semantic Web* can be seen as a network of autonomous (and autonomously evolving) nodes. Each node holds a *local* state consisting of extensional data (facts), metadata (schema, ontology information), optionally a knowledge base (intensional data), and, again optional, a behavior base. In our case, the latter is given by the ECA rules under discussion that specify which actions are to be taken upon which events under which conditions.

According to the Semantic Web Tower, there are already from the static point of view several abstraction levels.

In classical database systems, the *physical level/model/schema*, the *logical level/model/schema* (i.e., an *abstract data type* that can have different implementations/physical models, and as a database model comes with a generic *database query language*), and sometimes the *export schema* are distinguished.

In the early *network data model*, there was only the physical model where the query language constructs were also directly based on. For *relational databases*, the physical model includes the tables (i.e., ordered sets) and storage data structures (including indexes etc.), the logical model is the relational/SQL schema, and the export schema is in most cases also a relational schema, including views.

With *object-oriented* databases, there came different physical models, including relational and “native” storage. The logical model is the object-oriented model with an ODL/OQL schema; the export model is also the object-oriented one. *Object-relational* architectures are those where the physical model is the relational one, the logical model is split into two levels, the low-level one is the relational model, and the high-level is the object-oriented model, which also serves as export model.

With *XML*, the relationships became even more complex. There are several physical models that can serve for XML data; one of them is again the relational one. XML data can also be stored in object-oriented structures, as done in the early products Tamino [69] (based on Adabas) and Excelon [29] (based on Object Store). Several products claimed to use a “native” XML data model. The “lowest” well-specified XML-related notion is then the *Document Object Model (DOM)* [27] as an *abstract datatype*. Since this model is only on the level of an abstract datatype and does not support any query language, it is not a logical data model. The *logical data model* then is XML, with query languages like XQuery. On the other hand, XML serves as an *export data model* when XML views are defined over relational data [28].

In the research community, models like OEM [61] or F-Logic [41] (for which several internal physical models can be used, e.g., a frame-based one, or a relational one together with Datalog) came up that are used as logical models, or as export models for integrating data from other models.

Adding more abstraction with *RDF*, RDF is seen as the *logical data model*. Then, between it and –numerous– physical models, there can be a layer that uses the XML data model or the relational data model. Using “native” RDF databases, the physical data model can be any data structure that stores triples, or a frame-based structure like F-Logic. When exporting or integrating relational or XML data in RDF, RDF serves as export model.

Considering *OWL*, it qualifies as an *export model* since –by its reasoning– it defines views over RDF data as *logical model* that are queried by the user. From the point of view of the OWL/RDF user, XML then is not a logical model, but “below” this.

In general, the user works on the *export level* (which uses in general a data model which is the same or closely related to the *logical level*). Queries against the export level are mapped down to the logical level.

Abstraction Levels in the Conventional Web and in the Semantic Web. For the conventional (XML) Web and for the Semantic Web, there are different “towers” of data models. In the conventional Web, there are two levels; the upper of which is XML:

- Data level. Files, SQL databases, XML databases etc. Here *local* behavior of the *databases* (e.g., integrity maintenance) is located.
- Logical Level: XML. Here *local* behavior of the *nodes* (e.g., local application behavior) is located. Remote actions between tightly coupled nodes (i.e., that use common XML Schemas etc.) are also possible on this level. Interfaces for Web-Services like SOAP are also on this (syntactical) level.

In the Semantic Web, the structure of the levels has to be seen from a local and from a global aspect:

- Data level. Files, SQL databases, XML databases etc. Here again *local* behavior of the *databases* (e.g., integrity maintenance) is located.
- Local Logical Level: this level is provided by an XML or relational model, sometimes omitted (for RDF databases). Here *local behavior* of the *nodes* (e.g., local application behavior) is located. Remote actions between tightly coupled nodes (i.e., that use common XML Schemas etc.) are also possible on this level.

- Global Logical/Integrated Level: RDF. Here, *integrated* behavior (i.e., simple push/pull communication) will be located; messages between loosely coupled nodes that communicate in an application domain will be exchanged on this level.
- Semantic Level: OWL. Here, *intelligent integrated behavior* will be located, i.e., business rules, policies and strategies that often use derived data (and derived events).

Abstraction Levels of Behavior. In the same way as there are different levels from the static point of view, behavior can be distinguished wrt. these levels (programming language/data structure level, logical level, integrated level, and semantic level), and with different scope (local or global).

On all these levels, there is behavior. The user and the applications rely on the behavior on the *export model level* (OWL), whereas the actual, persistent changes to the database take place on the *physical level* (SQL, XML). On this lowest level, several proposals for “triggers” for XML data analogous to the SQL triggers exist, where with the distributed environment of XML data on the Web now, two types can be distinguished:

- Local triggers where event, condition, and action components use only the local database (like for SQL triggers),
- “Web-Level triggers” whose event component is based on data-level events in the local database, the condition component uses the local database and possibly also remote ones, and the action component can include arbitrary actions on the Web level (sending messages, interactions via HTTP, SOAP),

Such trigger concepts for XML and RDF as simple ECA rules will be investigated in more detail in Section 3.2.

For realizing behavior in the Semantic Web, also vertical transmission between the levels is required, including the XML and RDF models. A classification of types of rules wrt. their functionality and roles in the whole framework will be given in Section 3.1.5.

3.1.3 Domain Ontologies including Dynamic Aspects

The coverage of *domain ontologies* differs already in the classical data models (and conceptual models): in the relational model and in the Entity-Relationship model, a domain ontology consists only of the static notions, expressed by relations and attributes, or entity types with attributes and relationships (similar in first-order logic). In the object-oriented model and in UML, the static issues are described by classes, properties and relationships, and the dynamic issues are described by actions; in UML also their effects can partly be described.

A *complete* ontology of an application domain requires to describe not only the static part, but also the dynamic part, including actions and events (cf. Figure 3.1):

- describing actions in terms of agents, preconditions, and effects/postconditions,
- describing events, i.e., correlating actions and the resulting events, and specifying composite events, and
- describing composite actions (processes),
- in fact, business rules themselves can also be seen as parts of the ontology of an application.

For designing a Semantic Web application or service, in general ontologies of several domains interfere:

- Application domain ontologies define the static and dynamic notions of the application domain (banking, traveling, etc.), i.e., predicates or literals (for queries and conditions), and events and actions (e.g. events of train schedule changes, actions of reserving tickets).

- Application-independent domains that *talk about* an application; mostly related to classes of services (messaging, transactions, calendars, generic data manipulation). They can be generically used in combination with arbitrary application domains. They also provide static and dynamic notions.

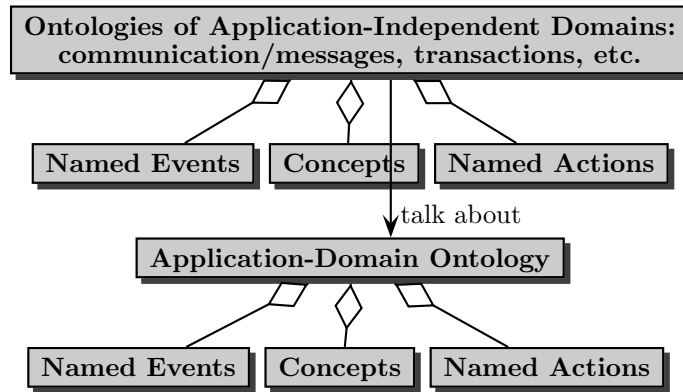


Figure 3.1: Kinds and Components of Ontologies

Figure 3.2 further describes more detailed the structure of an ontology and the interferences between its components. Concepts can be divided into classes, relationships and individuals (note that by reification, classes and relationships can also be seen as individuals). The actual languages talk about these things by literals. Actions influence classes, relationships, and individuals and raise events.




Figure 3.2: Structure and Interference in Ontologies

3.1.4 Events

An important aspect is the analysis of types of events that have to be considered. The ontology of events has to consider the abstraction levels and the application-dependent and application-independent domain ontologies. There are different kinds of (atomic) events (see explanations below and Figure 3.3):

- events of a given application domain (e.g., in banking, travel organizing, administration); such atomic events are described in terms of the ontologies of the application domain,
- generic parametric events that are not from a specific application domain but that instantiate generic event patterns, e.g., communication (“receive a message about ...”) and transactional events that *talk about* application domains.

Data level events are also a special kind of such generic (= generic to the data model) events.

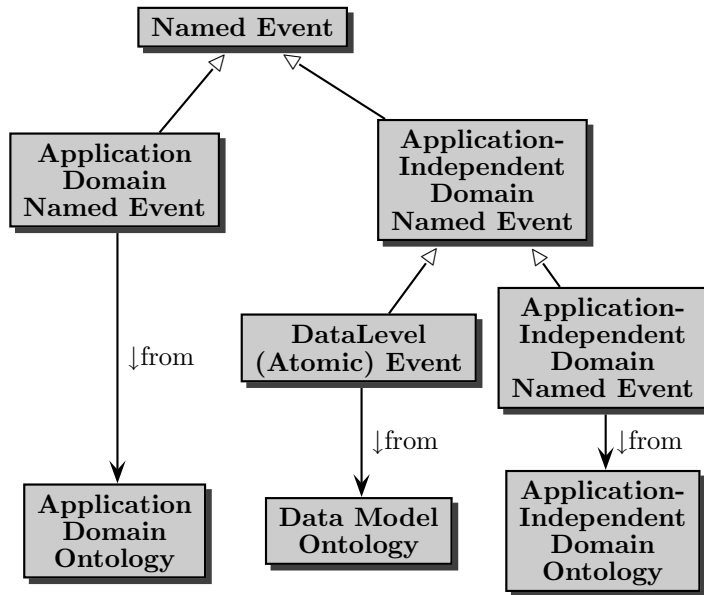


Figure 3.3: Ontology of Named Events

We start with the conceptually simpler events in application-independent domains.

Events in Application-Independent Domains. The application-independent domains provide *patterns* of atomic events that are ontologically independent from the actual application, but *talk about* an application; mostly related to classes of services, e.g., messaging or operations in a data model (cf. Figure 3.4). In general, such events are associated with a certain node.

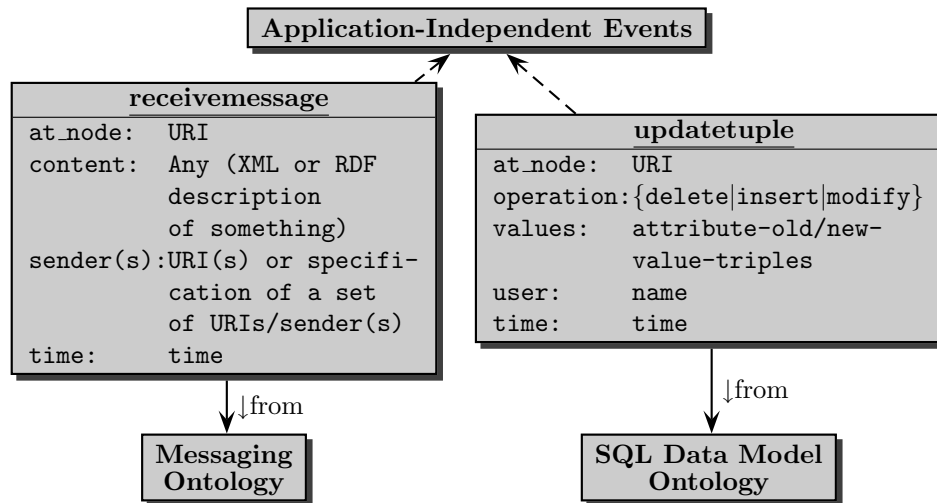


Figure 3.4: Example Events of Application-Independent Domains

Data Model Events. In the same way as for SQL data, there are atomic generic data model-level (i.e., XML or RDF) events (as will be discussed in more detail Section 3.2 for triggers). The actions that raise such events are operations of the underlying data model. Thus, they are *generic*

in the sense that they apply to the schema level of a given data model and only their parameters, i.e., names of classes and relationships and actual data, are taken from the application.

Other Generic Events. In the same way as the above data model events, there are *generic* events that are not raised by data model-level updates but belong to high-level application-independent ontologies that often *deal* with application-specific information:

- Communication: receiving or sending messages. Such events are parametric with e.g., sender and receiver (as URIs) and application-specific data (content of the message, possibly the roles of sender/receiver in this communication).
- Transactions: start, commit, and rollback of transactions. E.g., the “decision” or the message from another node that a transaction cannot be executed successfully raises such an event.

In the following, we also call such events *generic parametric events*.

Application Domain Events. Atomic application domain events are the visible happenings in the application domain. High-level rules, e.g., *business rules* use application specific events (e.g., `professor_hired($object, $subject, $university)`). Such events must be described by the ontology of an application.

Actions vs. Events. In contrast to simple data-level events on XML or RDF data, on the application-level there is an important difference between *actions* and *events*: an event is a visible, possibly indirect or derived, consequence of an action. For instance, the action is to “debit 200E from Alice’s bank account”, and visible events are “a debit of 200E from Alice’s bank account”, “a change of Alice’s bank account” (that is also immediately detectable from the update operation), or “the balance of Alice’s bank account becomes below zero” (which has to be derived from an update). Another example is the action “book person *P* on flight LH123 on 10.8.2005” which results in the internal action “book person *P* for seat 42C of flight LH123 on 10.8.2005” and the events “a person has been booked for a seat behind the kitchen”, “flight LH123 on 10.8.2005 is fully booked”, “there are no more flights on 10.8.2005 from *X* to *Y*”, or “person *P* has now more than 10.000 bonus miles”. All these events can be used for formulating (business) rules.

Reaction to Events. For actually reacting on events, a node must be *aware* of them. Here, some issues have to be considered in the Semantic Web since events can be derived ones, and that they are not necessarily located at a certain node.

- Explicit events are events that have a *direct* relationship with an action (and, from the implementation point of view, with a code fragment where they can be “caught”). The database-level events are obviously explicit ones: an update of the database is an action *and* is seen as an event. The receipt of a message is also an explicit event (although originally resulting from a different (sending) action, it can be caught on the socket level of a node).
- Implicit and derived events: an event is derived if it is defined based on other events, e.g. “flight LH123 on 10.8.2005 is fully booked” or “there are no more flights on 10.8.2005 from *X* to *Y*” are derived from “person *P* is booked for seat 5C of flight LH123 on 10.8.2005” under certain conditions. Note that there can be multiple derivations for a derived event (for the second one above, also e.g. “flight AF678 on 10.8.2005 is canceled”).

Raising and Derivation of Events. The task of becoming aware of implicit events is not the task of the rule execution, but of application-level reasoning, based on the application ontology. Thus, there are *derivation rules* also for events (that can be seen as event-condition-action rules where the action consists of raising an event); see Section 3.1.5.

Localization of Events. Orthogonal to being derived or not, application-level atomic events can be associated with a certain node or can describe happenings on the Web-wide level:

- local events: these happen locally at the node. E.g., data model events are in most cases used locally (by triggers, which then can raise higher-level events or trigger a remote action)
- remote events: From the point of view of a rule, a *remote* event is an event that is local (and can be localized) at another node, e.g., “if Amazon offers a new book on *X*”. Here, event detection can be done e.g. by monitoring the node (*continuous querying*) or by a *publish-subscribe-service*.
- global events: global events happen “somewhere in the Web”, e.g., “a new book on the Semantic Web is announced”, or “election of a new chancellor of Germany”. *Global* events are (mostly) application-level events where it is not explicitly specified where they actually occur.

In these cases, event detection is even more complicated since it must also be searched and derived *where* and *how* the event can be detected. Rules using global events require appropriate communication and notification mechanisms by the Semantic Web infrastructure (that can in turn also be based on ECA rules). For dealing with global (not located or locatable at a certain node) implicit (derived), the Semantic Web must provide *event broker* functionality.

Temporal Delay and Event Wrapping. Event brokering leads to the effect that the time-points of actual events and the event detection may differ. The actual infrastructure will probably also use messaging functionality, i.e., not raising an event, but packing it in a message or event “I became aware that ...”.

Example 1 Consider a customer *C* who wants to buy a Christmas tree, *C* can state the following rules:

- “if some *X* announces to sell Christmas trees, go there”. The rule is correct, but not “complete”: Probably, *C* will not become aware of the event of announcing, so he will never get a Christmas tree.
- “if I become aware [by a message] that some *X* announced to sell Christmas trees, go there”. This rule is more complete, but, if *C* becomes aware too late (e.g., after New Year) he will also go there.
- The correct rule is thus “if I become aware [by a message] [before Christmas] that some *X* announced to sell Christmas trees, go there”.
- Nevertheless, the real-world formulation of the rule will be of the style “if some *X* sells Christmas trees, I should go there” – which is formally based on an event that *X* sells a tree to some *Y* (which will eventually happen, but *C* will most probably not be informed about it).

Thus, in a future stage of Semantic Web rules, the interpretation of rules should provide a reasoning-based rewriting of rules to get their intended semantics.

Composite Events. Composite events are subject of heterogeneity in that there are multiple formalisms and languages for describing them. As already mentioned, most of them use *event algebras*. The target framework for the Semantic Web must support this heterogeneity.

3.1.5 Types of Rules

In our ECA-based approach, the behavior of domains *described* in an axiomatic way by (OWL) ontologies is specified and implemented by ECA rules. In a first step, these descriptions have to allow to run applications by ECA rules. In the next step, the ontology has to be extended to preconditions and effects as a base for *reasoning*.

There are several types of rules that are used for actually specifying the ontology and the behavior of an application:

- Rules that axiomatize the *ontology*, i.e., mandatory relationships between actions, objects, and events that are inherent to the domain. The correctness of the rules must be proven against the ontology.
- Rules that specify a given *application* on this domain, e.g., business rules. Changing such rules result in a different behavior of the application.

ECA Rules. From the external user’s point of view, *ECA-Business Rules* specify the actual behavior: “when something happens and some conditions are satisfied, something has to be done”. Here, events and actions refer to a very high and abstract level of the ontology.

- Such rules are “actual”, user-defined ECA rules since they trigger an action upon an event “to keep the application running”. Such rules exist on different abstraction levels and granularity, designed to the notions of the application domain. Changing them changes the behavior of the application.
- Internally, such rules are also used for implementing mechanisms for detection of derived and composite events on the respective level.

ECE Event Derivation Rules: Providing High-Level Events. For implementing high-level rules, it is necessary that these high-level events are provided somehow: They must be *derived*.

- horizontal ECE rules: Here, the event is derived from another high-level event under certain conditions, e.g., “*when a booking for a flight is done, and this is the last seat, then the plane is completely booked*”. The rule is an E-C-E (event-condition-event) rule, e.g., the “action” consists in deriving/raising an event. The events are logically related and inherent in terms of the application. Changing such rules would invalidate the application wrt. its ontology.
- upward vertical ECE rules: an abstract event is derived from changes in the underlying database, e.g., “*when the arrival time in a database of a flight of today is changed, this is actually a delayed flight*”. The rule is again an E-C-E (event-condition-event) rule. The events are not logically related and inherent in terms of the application, but are related due to the physical implementation of the application (i.e., since an explicit message “flight *F* is delayed” is missing, and only visible due to a modification of the database”. Changing such rules would invalidate the application wrt. its ontology.

As another example, “`professor_hired($object, $subject)`” is (locally) derived at a node from an insertion of a fact into an SQL, XML, or RDF database.

These rules correspond to the *bottom-up* semantics of derivation rules: Given the body, do the head. While “classical” ECA rules are *active* rules, the above enumeration presented several kinds of *derivation* rules. The main difference of these wrt. classical *derivation* is that the latter define continuously existing *views* and are used for *querying*. In contrast, the *event derivation rules* “fire” only once when an event is detected and another event must be raised. Thus, these *ECE* rules are more similar to ECA rules than to derivation rules.

The derivation of events by such rules can be done similar to views:

- bottom-up style: they can be “materialized” by raising them explicitly when (and where) they occur (even if probably nobody is actually interested in them), or
- top-down style: when an application uses a derived event, it runs the rule locally.

High-level events can also be raised as side-effects of high-level actions, see below. When designing rules, it must be cared that such effects do not cause any behavior twice.

ACA/ACE Rules: Talking about High-Level Actions. High-level actions like “(at a travel agency) book a travel by plane from Hanover to Lisbon” cannot be executed directly, but there must be another rule that says how this is implemented (by searching for connections, possibly via Madrid). Such rules are *reduction* rules that reduce an abstract action to actions on a lower level.

On the other hand, there may be another business rule that should be executed whenever somebody does a plane travel from Germany to Portugal, putting this person on a list for sending them advertisements about (questionable) tax saving tricks by investing in resorts in the Algarve. The latter rule should not be defined on the basis of “if there are bookings for a person via some places that lead from Germany to Portugal” (which would e.g. also fire if a Tyrolian flies from Innsbruck to Munich and then to Lisbon; the German tax tricks do probably not apply to him), but could –most declaratively– directly use the *abstract action* “book a travel by plane from a German airport to a Portuguese airport” for firing the rule.

Thus, there are rules that regard (abstract) actions as events – or in the above case more exactly, use the event of *committing an abstract action*. Such rules can be expressed based on transactional events, or on messages (“if we get the message that such an abstract action should be executed”), but in both cases this blurs the declarativity that the *action* actually is the reason to react.

Thus, there are several kinds of ACA rules:

- horizontal reduction ACA rules, e.g., “*the action of transferring 200E from account A to B is implemented by debiting 200E from account A and depositing 200E on account B*”. This rule is a kind of an A-C-A rule that explains a composite action by its components, both still in terms of the application domain.
- vertical reduction ACA rules, e.g., “*the action of debiting 200E from account A is realized by reading the account value, adding 200 and writing it*”. This rule is also a kind of an A-C-A rule that reduces a composite action into its components on a lower level.
- horizontal non-reduction ACA rules see an action (that has to be executed for itself) as an event that should trigger another action, known also as *rule chaining*.

This kind of ACA rules is more directly related to ECA rules. Changing such rules would not invalidate the application wrt. the ontology, but just change its behavior.

The above reduction ACA rules correspond to SQL’s *INSTEAD* triggers: in SQL, *INSTEAD*-triggers are used for specifying what updates should actually be done instead of inserting something into a view (which is not possible), whereas here, simpler actions are specified instead of an abstract one. Such rules are closely related to the *top-down* semantics of derivation rules: To obtain the head, realize the body (cf. Transaction Logic Programming [19]). These rules describe actions that are logically related and inherent in terms of the application. Some of these rules are inherent to the ontology of the underlying domain, others specify only the behavior of a given application (including e.g. policies that are not inherent to the domain).

Since high-level events can also be seen as observations, it is also reasonable to raise them when an appropriate action is executed. In most cases, this amounts to a simple mapping from high-level actions to raise high-level events: the action “*hire_professor(\$object, \$subject)*” directly raises the event “*professor_hired(\$object, \$subject)*” at the same node (and is internally executed by inserting a fact into an SQL, XML, or RDF database; see downward ACA vertical rules above).

The execution of both, ACA and ACE rules, must usually be located at the node where the action is executed (which makes completely sense because ACA is actually the algorithm to execute an action, and ACE is the communication of its effect on a high level).

Low-Level Rules. The base is provided by update actions on the database level (to which all abstract actions must eventually be reduced to actually change the state of any node) and low-level ECA rules, e.g., database triggers for referential integrity or bookkeeping.

Here, neither the event nor the action is part of the application ontology, but both exist and are related only due to the physical implementation of the application. Such rules guarantee –together with the RDF views on the database– the integrity of the model; thus, when verifying a process they also have to be taken into account.

Summary and Interference. The resulting information flow between events and actions is depicted in Figure 3.5.

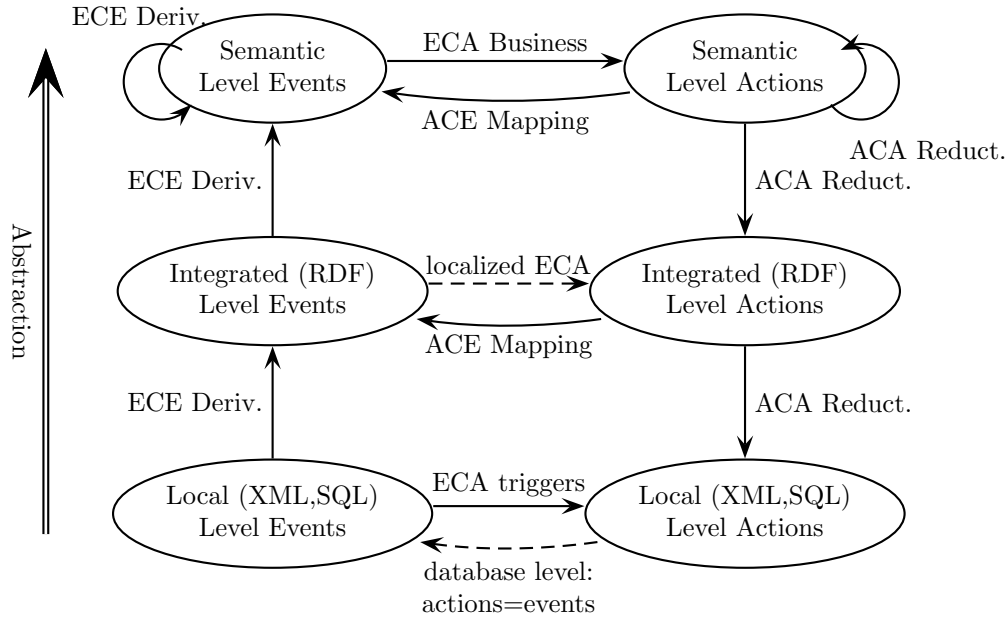


Figure 3.5: Types of Rules

Example 2 (Actions and Events) Consider the following scenario: A –rule-driven– review process leads to the acceptance of a paper. Here, “acceptance of a paper” is an action which is also an event on the Web level – “the paper P of A has been accepted for conference C ”. This event is communicated inside the program committee by mail (push communication). An internal rule of the conference of the form “when a paper is accepted, send a message to the author and list it on the conference Web page”. By the latter, the event of “the paper P of A has been accepted” becomes actually accessible for the public. Communication in the Semantic Web then should lead to firing other business rules, e.g., at the DBLP publications server and at citeseer, where lots of “business rules” (later) react upon.

Such rules in the Semantic Web are not formulated on the level of messaging, but assume the notification about events as given (which is, on a lower level, done by messaging).

The author of an accepted paper probably has a rule “when a paper is accepted (event), then book a travel to the conference”. This action is submitted as a message “Person P wants to book a travel from X to Y on date D ” to a travel agency. The travel agency reacts on incoming messages by searching for connections and book an available one, possibly by flights AA123 from X to Z and BB456 from Z to Y (horizontal rule for decomposing an action into its constituents). These

actual bookings of flights are then submitted to the airline, seats are assigned, and the booking actually takes place by modifications of the database content (vertical rules). Assume that this booking reserves the last seat in AA123.

On the other hand, there are several other rules that should fire in this process, e.g., a rule that removes all completely booked flights from some list, and that raises the price for all flights between two destinations *M* and *N* in case that more than 50% of the total capacity on this connection for that day is booked.

Both can be done by vertical rules, reacting on database events for booking actual seats, or on the higher level, e.g. “on any booking between *M* and *N* (event), check all flights between two destinations *M* and *N*, and in case that more than 50% of the total capacity on this connection for that day is booked, raise the price for the remaining ones by 10E”. In the latter case, the booking action is immediately also seen as an event “somebody books ...”.

Moreover, there can be a business rule “whenever a person *P* books a travel from country *C* to country *D*, do ...”. If *X* is located in *C* and *Y* is located in *D*, but *Z* is located in a different country, then, this event cannot be detected from the actual, independent bookings of the individual flights. Instead, the action “book a travel from *X* to *Y* on date *D* for person *P*” should be considered as a high-level event for immediately firing appropriate rules.

3.2 Simple ECA Rules: Data Model Triggers

A simple form of active rules that is often provided by database systems, are *triggers*. As shown in Figure 3.5, these triggers on the database level form the lowest level of rules. Reacting directly to changes in the database, they provide the basic level of behavior. Triggers are simple rules on the (*database*) *programming language and data structure level*. They are associated with the logical level (i.e., not referring to the implementation of the logical data model, but acting on its notions). They follow a simple ECA pattern where the conditions are given in the database query language and the action component is given in a simple, operational programming language. In SQL, triggers are of the form

```
ON database-update WHEN condition BEGIN pl/sql-fragment END .
```

In the Semantic Web, this base level is assumed to be in XML or RDF format. While the SQL triggers in relational databases are only able to react on changes of a given tuple or an attribute of a tuple, the XML and RDF models call for more expressive event specifications according to the (tree or graph) structure.

3.2.1 Triggers on XML Data

Work on triggers for XQuery has e.g. been described in [16] with *Active XQuery* (using the same syntax and switches as SQL, with XQuery in the action component) and in [6, 62], emulating the trigger definition and execution model of the SQL3 standard that specifies a syntax and execution model for ECA rules in relational databases. The following proposal refines our previous one developed in [4].

We propose to use *simple-expressions* that consist only of the downward axes *child* and *descendant*, and the final step is allowed to be an *attribute* step for basically locating events of interest¹.

- ON {DELETE|INSERT|UPDATE} OF *simple-expr*: if a node matching the *simple-expr* is deleted, inserted, or updated,
- ON MODIFICATION OF *simple-expr*: if anything in the subtree rooted in a node matching the *simple-expr* is modified,

¹Alternative possibilities are: allow arbitrary XPath expressions, or allow even for specifications like ON UPDATE OF *subexpr* IN *xpath-expr*.

- ON INSERT INTO *simple-expr*: if a node is inserted (directly) into a node matching the *simple-expr*,
- ON {DELETE|INSERT|UPDATE} [SIBLING [IMMEDIATELY]] {BEFORE|AFTER} *simple-expr*: if a node (optionally: only sibling nodes) is modified (immediately) before or after a node matching the *simple-expr*.

All triggers should make relevant values accessible, e.g., OLD AS ... and NEW AS ... (like in SQL), both referencing the node to which the event happened, additionally INSERTED AS, DELETED AS referencing the inserted or deleted node.

Similar to the SQL STATEMENT and ROW triggers, the granularity has to be specified for each trigger:

- FOR EACH STATEMENT (as in SQL),
- FOR EACH NODE: for each node in the *simple-expr*, the rule is triggered only at most once (cumulative, if the node is actually concerned by several matching events) per transaction,
- FOR EACH MODIFICATION: each individual modification (possibly for some nodes in the *simple-expr* more than one) triggers the rule.

The implementation of such triggers in XML repositories can e.g. be based on the *DOM Level 2/3 Events* or on the triggers of relational storage of XML data. Events/triggers on this logical level are local (and internal) to the database that provides an RDF view to the outside. Usually, the actions are local updates of the database (that then effect the RDF view indirectly), or they raise events on the RDF level (it is see also below), but it is also possible to send XUpdate or HTTP requests or SOAP messages to other nodes, or to state remote XML updates explicitly:

```
ON INSERT OF department/professor
let $prof:= :NEW/@rdf-uri, $dept:= :NEW/parent::department/@rdf-uri
RAISE RDF_EVENT(INSERT OF has_professor OF department)
  with $subject:= $dept, $property:=has_professor, $object:=$prof;
RAISE RDF_EVENT(CREATE OF professor)
  with $class=professor, $resource:=$prof;
```

XML: Local and Global Rules in the “Conventional” XML Web. The above events occur always local in a node and can be detected at this node.

Rules on the XML level of the Web can either be local to a certain node, or they can include Web data, e.g., reacting on events on views that include remote data, or raising actions on the Web. For that reason, we call them *Web Level Triggers* (note that these can already be applied to the conventional non-semantic Web; whereas for integration reasons in the *Semantic Web*, the level of RDF events is preferable).

Actual rules on this level usually are not only based on atomic data-level events, but use own event languages that are based on a set of atomic events (that are not necessarily just simple update operations) and that usually also allow for composite events. Their event detection mechanism is not necessarily located in the database, but can be based on the above triggers.

Such rules require knowledge of the actual XML schema of the corresponding nodes. Provided a mapping between rules on the XML level and those of the RDF view level, the implementation can (more efficiently) be kept on the XML level, whereas reasoning about their behavior can be lifted to the Semantic Web level. A Semantic Web framework should also support this kind of rules.

From the *Semantic Web* point of view, events on the XML level should usually not be communicated to other nodes (except very close coupling with nodes using the same schema); instead semantic events should be derived from them.

3.2.2 Triggers on (Plain) RDF Data

RDF triples describe properties of a resource. In contrast to XML, there is no subtree structure (which makes it impossible to express “deep” modifications in a simple event), but there is some metadata². A proposal for RDF events can be found in RDFTL [62, 63]. The following proposal refines our previous one developed in [4]:

- ON {INSERT|UPDATE|DELETE} OF *property* [OF *class*] is raised if a property is added to, updated, or deleted from a resource (optionally: of the specified class).
- ON {CREATE|UPDATE|DELETE} OF *type* is raised if a resource of a given class is created, updated or deleted.
- ON NEW TYPE is raised if a new type is introduced,
- ON NEW INSTANCE [OF *type*] is raised if a new instance of a type is introduced,
- ON NEW PROPERTY OF INSTANCE [OF *type*] is raised, if a new property is added to an instance (optionally: to a specified class). This extends ON INSERT OF *property* to properties that cannot be named (are unknown) during the rule design.

Besides the OLD and NEW values mentioned for XML, these events bind variables Subject, Property, Object, Class, Resource, referring to the modified items (as URIs), respectively. Trigger granularity is FOR EACH STATEMENT or FOR EACH TRIPLE.

Note that for plain RDF data, there is no *derived* information (subproperties, subclasses, domains, etc.). Thus, events *immediately* correspond to operations (inserting, deletion or modification of corresponding triples).

3.2.3 Triggers on RDFS and OWL Data

For the Semantic Web, triggers on RDFS and OWL data are more relevant. In these cases, *reactivity* must be combined with *reasoning*. A proposal for triggers in such a scenario has been developed and implemented in [55, 73] based on the Jena Framework [38]; see also Section 10.3.

Application-level events (that must be characterized appropriately in the application ontology) can then be raised by such rules, e.g.,

```
ON INSERT OF has_professor OF department
  % (comes with parameters $subject=dept, $property:=has_professor,
  %   and $object=prof)
  % $university is a constant defined in the (local) database
  RAISE EVENT (professor_hired($object, $subject, $university))
```

which is then actually an event (e.g., `professor_hired(prof, dept, univ)`) of the application ontology on which a “business rule” of a publisher could react that says, if a new professor is hired at a university, then the appropriate list of textbooks should be sent to him. Note that in the above trigger, this event is only raised – the complete issues of communicating it and detecting it by the node that actually processes the business rule have to be dealt with separately.

3.2.4 Triggers vs. ECA Rules

As shown in Figure 3.5, triggers on the database level form the lowest level of rules. They deal with the data model level instead of the application level. Thus, their “home” is inside the database, using notions of the database model, and their implementation often even depends on the availability of information from the physical level of the database. Especially, their “events” coincide directly

²We consider here only `rdf:type` as a special property. RDFS vocabulary, including RDFS reasoning is considered in the next section.

with the update operations of the database, which are also the *actions* on that level. Triggers are not necessarily subject of the modularization of our model. In contrast, often triggers, although following the ECA paradigm, are only very restricted. They are usually closely intertwined with the database (e.g., in relational databases), although very efficient external implementations exist, cf. TriggerMan [34].

In case that triggers are implemented external to a database, the data items that are concerned by the triggering event must usually be identifiable:

- SQL: OLD and NEW are the modified tuples; all related data can be identified externally by foreign keys. Furthermore, usually, the ROWID is used that identifies a tuple internally.
- XML: OLD and NEW should at least allow to access the subtrees, but also parent or ancestor nodes may be of interest. For this, an internal identifier must be accessible. Note that there may also be triggers that do a modification inside the database relative to the modified node, e.g., in the subtree.
- RDF: OLD and NEW are the current nodes. Since they in general have a URI, they can be identified without any problem.

These triggers are usually defined in a homogeneous way on the programming language level (i.e., data model events of the data model's update language, queries in the data model's query language, and actions are given as data model updates or as a program segment in a programming language that includes the data model's query language). On the higher level, ECA rules make use of more abstract languages.

Rules on the Ontology Level. In rules on the level of an application, the events, conditions and actions refer to the domain ontology. Even more, often the local knowledge of the node is not sufficient, but in general, OWL events refer to the distributed scenario. Thus, it is not always appropriate to locate them in a database like triggers.

Local, active rules working on this level are e.g. available in the *Oracle Rule Manager*: Events are similar to tuples of a view (but volatile) and can e.g. be raised by triggers. Then, active rules can be defined that react on such events, including a restricted form of composite events. We give here just an example to complete the “upward” transmission of events:

```
ON (professor_hired($prof, $dept, $univ))
WHEN $Books := select relevant books for people at this dept
BEGIN do something END
```

More complex rules also use composite events and queries against the Web. Composite events in general consist of subevents at different locations. Additionally, higher-level events are in general not explicitly raised. In the above example, both was simple: the source explicitly raised `professor_hired($object, $subject, $university)`, and the publisher can e.g. register at all universities to be notified about such events. In general, events like “when a publication p becomes known that deals with ...” cannot be detected in this simple way, but must be derived and obtained from other, more general information. Here, Semantic Web reasoning is required even for detecting atomic events “somewhere in the Web”. The timepoints of actual events and the event detection may differ.

3.3 ECA Language Structure

After having discussed and analyzed the requirements for specifying and implementing behavior in the Semantic Web by active, ECA-style rules, we develop now the structure of such rules and the required languages.

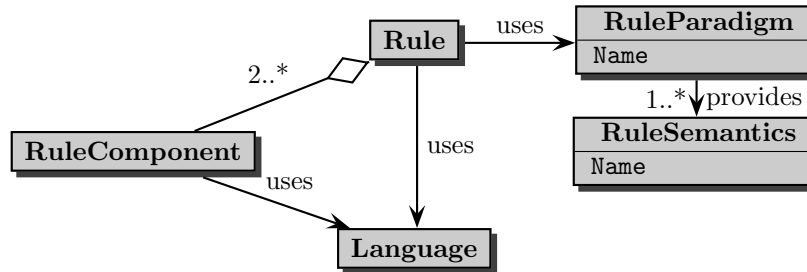
3.3.1 Language Heterogeneity and Structure: Rules, Rule Components and Languages

Rules in general consist of several components. For instance, deductive rules like in Prolog/Datalog consist of a *rule head* and a *rule body*; similar e.g. for F-Logic or Transaction Logic rules. These languages are *rule-based languages* – their head and body are both expressions in this language (but note that e.g. negation may in most cases only occur in rule bodies). In the ECA paradigm, rules are not homogeneous: they consist of an *event component*, a *condition component* (that together roughly correspond to the “body” since they describe the situations where the rule is applicable), and an *action component* (that roughly corresponds to the “head”) – and all these components use different languages. For example, for database triggers, events are database events like “on update” (that are raised by update operations), where-clause conditions are expressed in SQL, and actions are SQL programs or updates (that again may raise events). Another example is the language XChange discussed in [22], where the event component consists of an event query, the condition component is a test by evaluating an Xcerpt [23] query, and the action component may contain update actions, transactions, or raising of an event. Yet another example is ruleCore [13], where the (different) languages for each component are described. The ECA languages mentioned in the introduction ([17, 16, 6, 63]) also use different languages in each of the components. In these cases, the event, condition, and action language are usually closely related, but this is not necessarily always the case. The more complex a scenario is, the more specialized are the used component languages.

Thus, the semantics of a rule is determined by two constituents:

- the rule semantics or “rule paradigm” that characterizes the interplay between the components, and the
- language(s) used in its components.

For instance, *deductive rules* have several common semantics, either top-down, bottom-up as fixpoints, or well-founded or stable semantics, independent what underlying language (first-order logic, F-Logic, Transaction Logic etc.) is used. In the same way, ECA rules have a fixed semantics, independent what languages are used in the E, C, and A components. An important common feature here is that the communication both for derivation rules and for ECA rules is done by *logical variables*; we will discuss this in detail in Chapter 4.



Constraints:

RuleParadigm determines number of RuleComponents

RuleComponents ordered or named; using appropriate languages

Figure 3.6: Rules, Rule Components and Languages

An XML markup of rules must cover the *structure* of rules and rule component languages. Here, the RuleML language [66] provides *general* guidelines that must then be specialized for each paradigm.

In the following, we will investigate general ECA rules. The analysis of the languages will be continued in two aspects: semantics/ontology and syntax (i.e., algebraic, variables etc.).

3.3.2 Components and Languages of ECA Rules

In usual Active Databases in the 1990s, an ECA language consisted of an event language, a condition language, and an action language. For use in the Semantic Web, the ECA concept needs to be more flexible and adapted to the “global” environment of a world-wide living organism where nodes “speaking different languages” should be able to interoperate. So, different “local” languages, for expressing events, queries and conditions, and actions have to be integrated in a common framework.

The target of the development and definition of languages for (ECA) rules and their components should be a semantic approach, i.e., based on an (extendible) ontology for these notions that allows for *interoperability* and also turns the instances of these concepts into objects of the Semantic Web itself. The upper level of this ontology is shown as an UML-style diagram in Figure 3.8, which will be explained below.

In contrast to previous ECA languages from the database area, we aim at a more succinct, conceptual separation between the event, condition, and action components, which are (i) possibly given in separate languages, and (ii) possibly evaluated/executed in different places. Each of the components is described in an appropriate language, and ECA rules can use and combine such languages flexibly.

Analysis of Rule Components. A basic form of ECA/active rules are the well-known *database triggers*, e.g., as already shown above, in SQL, of the form

ON *database-update* WHEN *condition* BEGIN *pl/sql-fragment* END.

For them, *condition* can only use very restricted information about the immediate *database update*. In case that an action should only be executed under certain conditions which involve a (local) database query, this is done in a procedural way in the *pl/sql-fragment*. This has the drawback of not being declarative: reasoning about the actual effects would require to analyze the program code of the *pl/sql-fragment*. Additionally, in the distributed environment of the Web, the query is probably (i) not local, and (ii) heterogeneous in the language – queries against different nodes may be expressed in different languages. For the MARS framework, we prefer a *declarative* approach with a *clean, declarative* design as a “Normal Form”: Detecting just the dynamic part of a situation (event), then check *if* something has to be done by probably obtaining additional information by a query and then evaluating a *boolean* test, and, if “yes”, then actually *do* something – as shown in Figure 3.7.




Figure 3.7: Components and Phases of Evaluating an ECA Rule

With this further separation of tasks, we obtain the following structure:

- every rule uses an event language, one or more query languages, a test language, and one or more action languages for the respective components (for that, we allow several action components in different languages that *all* have to be executed),
- each of these languages and their constructs are described by metadata and an ontology, e.g., associating them with a processor,
- there is a well-defined *interface* for communication between the E, Q&T, and A components by variables (e.g., bound to XML or RDF fragments).

This model can be readily extended by adding a fifth optional component – the post-condition (another *test*) – resulting in a variation usually called ECAP rules. In most cases, this post-condition can be omitted by allowing the action language to test for conditions inside the action rules. But it may have particular relevance when considered together with transactional rules, and for reasoning about the effects of sets of rules.

For applying such rules in the Semantic Web, a uniform handling of the event, query, test, and action sublanguages is required. For this, rules, their components, and the languages must be objects of the Semantic Web, i.e., described in XML or RDF/OWL in a generic *rule ontology* that contains all required information as shown in the UML model in Figure 3.8.

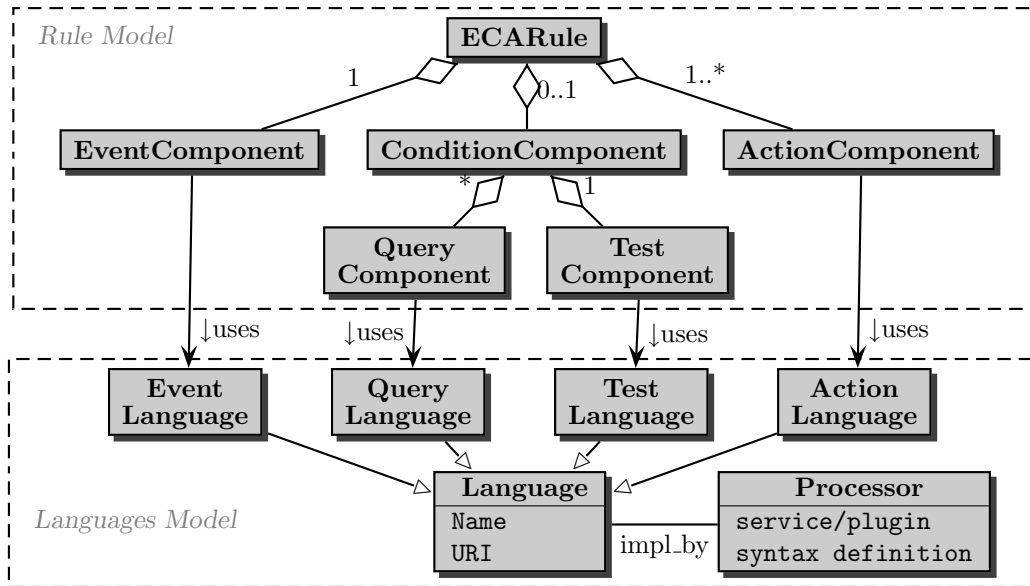


Figure 3.8: ECA Rule Components and Corresponding Languages

3.3.3 Markup Proposal: ECA-ML

The model is accompanied by an XML ECA rule (markup) language, ECA-ML. The relationship between the rule components and languages is provided by identifying the languages with namespaces (from the RDF point of view: resources), which in turn are associated with information about the specific language (e.g. an XML Schema, an ontology of its constructs, a URL where an interpreter is available). The latter issues are discussed in Section 9; here we investigate the languages and the markup itself.

For an XML representation of ECA rules as shown in Figure 3.8, we propose the below (basic) markup (ECA-ML). Later, in Section 11, we develop an *ontology* of rules that defines and relates these concepts. ECA-ML will turn out to be a *canonic* projection of an *also canonic* serialization in RDF/XML of the RDF graph of ECA rules.

```

<!ELEMENT rule (%variable-decl,event,query*,test?,action+)>
  <!-- %variable-decl is not yet specified in detail here-->
<eca:Rule rule-specific attributes>
  rule-specific content, e.g., declaration of logical variables
  <eca:Event identification of the language >
    event specification, probably binding variables; see Section 4.4.1
  </eca:Event>
  <eca:Query identification of the language > <!-- there may be several queries -->

```

```

    query specification; using variables, binding others; see Section 4.4.2
</eca:Query>
<eca:Test identification of the language >
    condition specification, using variables; see Section 4.4.3
</eca:Test>
<eca:Action identification of the language > <!-- there may be several actions -->
    action specification, using variables, probably binding local ones; see Section 4.4.5
</eca:Action>
</eca:Rule>

```

The actual languages of the components (and on deeper nested levels) are usually identified by namespaces, sometimes also by explicit `language` attributes.

A similar markup for ECA rules (without separating the query and test components) has been used in [17] with *fixed* languages (a basic language for atomic events on XML data, XQuery as query+test language and SOAP in the action component). This fixed approach falls short wrt. the language heterogeneity, and especially the use and integration of languages for composite events. The same structure is also followed by the XChange [22] transaction rules above, there again without any intention to deal with heterogeneity of languages: fixed languages are used for specifying the event, condition (without separating query/test), and action component. In contrast, here we generalise the approach to allow for using arbitrary languages. Thus, these other proposals are just possible configurations. Our approach even allows to mix components of both these proposals.

3.3.4 Hierarchical Structure of Languages

The approach defines a hierarchical structure of language families (wrt. the embedding of language expressions) which relates the different kinds of ontologies (application-dependent and application-independent) and their components as already described above in Section 3.1.3 and Figure 3.1 to the languages of rules and rule components as shown in Figure 3.9 [here directly associating ontologies with the corresponding languages over the same alphabet]: As described until now, there is an ECA language (that is already described by the above markup), and there are (heterogeneous) event, query, test, and action languages. Rules will combine one (or more) language of each of the families. In general, each such language consists of its own, application-independent syntax and semantics (e.g., event algebras, query languages, boolean tests, process algebras or programming languages) that is then applied to a domain (e.g. traveling, banking, universities, etc.). The domain ontologies define the static and dynamic notions of the application domain, i.e., predicates or literals (for queries and conditions), and events and actions (e.g. events of train schedule changes, actions of reserving tickets, ...). Additionally, there are domain-independent languages that provide primitives (with arguments), like general communication, e.g. `received_message(M)` (where M in turn contains domain-specific content).

In the next section, we discuss common aspects of the languages on the “middle” level (that immediately lead to the tree-style markup of the respective components – thus, here the XML markup is straightforward). The semantics of rule execution and communication issues between the rule components are discussed in Chapter 4. Samples of component languages will be discussed in Chapter 5; a short account on domain languages including events and actions has been given in Section 3.1.3.

3.3.5 Common Structure and Aspects of E, C, T and A Sublanguages

The four types of rule components use corresponding types of languages that share the same algebraic language structure, although dealing with different notions:

- event languages: every expression gives a description of a (possibly composite) event. Expressions are built by composers of an event algebra, and the leaves here are atomic events of the underlying application domain or an application-independent domain;

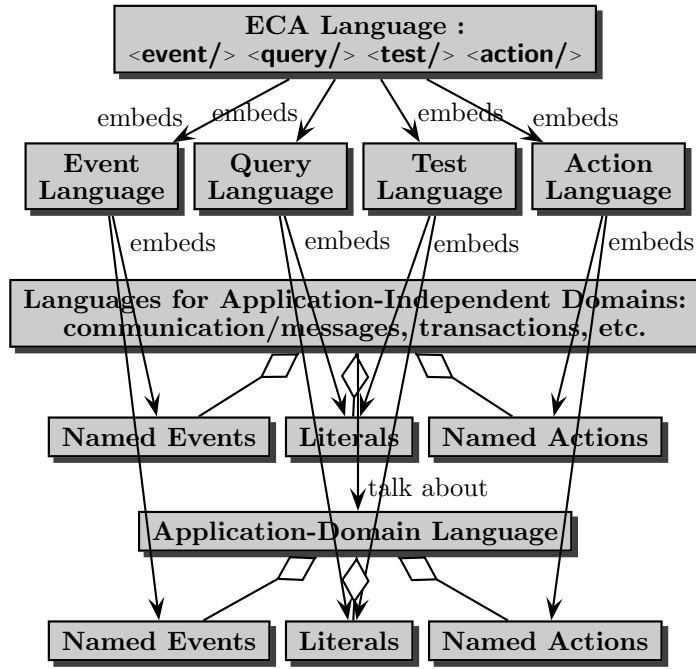


Figure 3.9: Hierarchy of Languages

- query languages: expressions of an algebraic query language, embedding literals from the domains;
- test languages: they are in fact formulas of some logic over literals of that logic in the underlying domains (that determine the predicate and function symbols, or class symbols etc., depending on the logic);
- action languages: every expression describes an action. Here, algebraic languages (like process algebras) or “classical” programming languages (that nevertheless consist of expressions) can be used. Again, the atomic items are actions of the underlying domains.

Algebraic Languages.

As shown in Figure 3.10, all components have in common that the component languages consist of an algebraic language defining a set of *composers*, and embedding *atomic* elements (events, literals, actions) that are contributed by the *domain languages*. Expressions of the language are then (i) atomic expressions, or (ii) composite expressions recursively obtained by applying composers to expressions. Due to their structure, these languages are called *algebraic languages*, e.g. used in *event algebras*, *algebraic query languages*, and *process algebras*. Each composer has a given *cardinality* that denotes the number of expressions (of the same type of language, e.g., events) it can compose, and (optionally) a sequence of parameters (that come from another ontology, e.g., time intervals) that determines its *arity* (see Figures 3.10 and 3.11).

For instance, “ E_1 followed by E_2 within t ” is a binary composer to recognize the occurrence of two events (atomic or not) in a particular order within a time interval, where t is a parameter. Event languages define different sets of composers, such as XChange for its composite event queries [22], the ruleCore detectors [13], or the SNOOP event algebra of [24]. Similar composers are used in process algebras, or also –but in general syntactically covered– in algebra-based query languages. The boolean algebra with its composers is well-known.

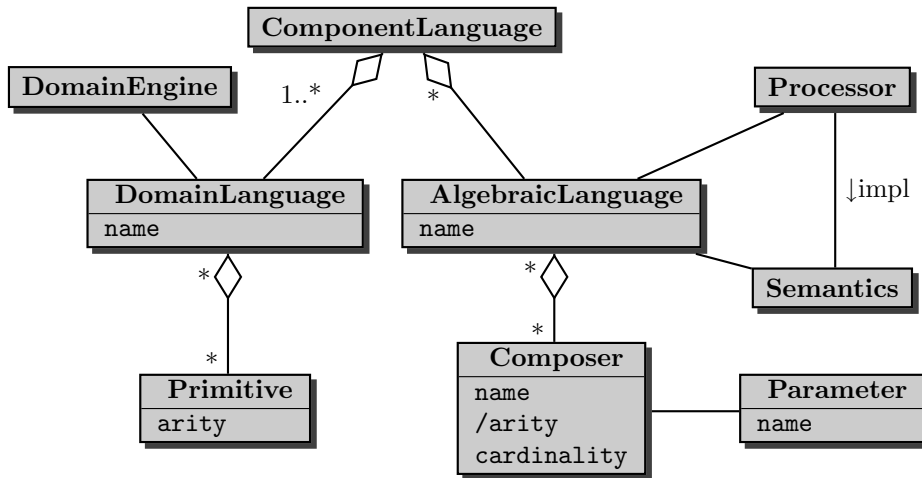


Figure 3.10: Notions of an Algebraic Language

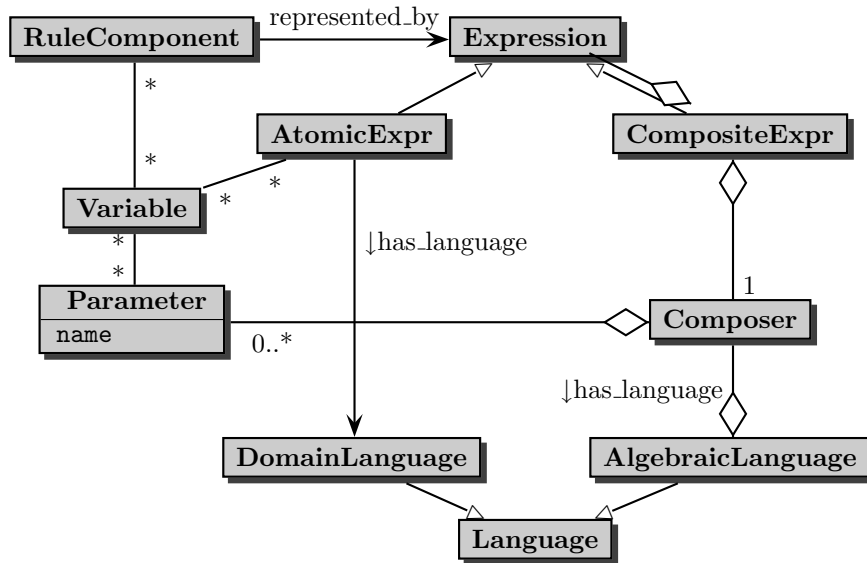


Figure 3.11: Syntactical Structure of Expressions of an Algebraic Language

Semantics of Algebraic Languages.

Every algebra expression is assigned a semantics, i.e., from evaluating it (in case of queries: in a given state). Starting with the semantics of atomic expressions, the semantics of composite expressions is determined by the composer. The semantics of the different types of algebraic languages are as follows:

- event languages: in most cases, the event instance (or sequence) that matches the given expression pattern,
- query languages: a query result, e.g., a relation or an XML structure,
- test languages: a truth value of a logic (i.e., for classical logics, true or false)
- action languages: the formal semantics of terms of e.g. process algebras are denotational

or operational semantics as state transformers; nevertheless, here we are only interested in their side effects on the underlying data.

Composition of Algebraic Languages.

For each type of such algebraic languages (i.e., event, query, test, and action languages), the expressions define and combine entities of the same kind, i.e., again, events, queries, tests, or actions, possibly with appropriate parameters and logical variables (see below) and using events, literals, or actions from the respective part of the domain language. Thus, from the ontology point of view, entities of the same kind described in different languages can be combined (e.g., a conjunction of a (sub)formula in first-order logic with one in description logic, or a sequence (or, more generic, any binary composer of any event language) of an event specified in language EL_1 and one specified in EL_2).

The global language concepts and the markup will support this, and we will also explain how evaluation also smoothly crosses these language borders. Note that Figure 3.11 does not associate the whole rule components with a language, but each *expression* is associated to a language.

Tree Markup of Algebraic Languages.

Thus, language expressions are in fact trees which are marked up accordingly. The markup elements are provided by the definition of the individual languages, “residing” in and distinguished by the appropriate namespaces. As described above, it is also possible to nest composers and expressions from different languages of the same kind, distinguishing them in the markup by the namespaces they use. Thus, languages are in fact not only associated once on the *rule component* level, but this can also be done on the *expression* level.

Note that leaves of expressions can be either *atomic* expressions (events, predicates, actions) that are defined in domain languages, or composite expressions in a suitable embedded algebraic language.

3.3.6 Language Information

As stated above, languages are associated on the *expression* level. We first investigate this from the point of view of the ECA engine (which is the first module to be developed), where the expressions of interest are the rule components. Here, the functionality of selecting services according to the language information is located (having nested subexpressions in different languages requires to have this functionality also in the processors of component languages, e.g., for having a subevent specification in a different event algebra).

As shown in Figure 3.11, every expression (i.e., the rule components and their subexpressions) can be associated with a language: an expression is either

- an atomic one (atomic event, literal, action) that belongs to a domain language (either application-dependent or application-independent), or
- a composite expression that consists of a composer (that belongs to a language) and several subexpressions (where each recursively also belongs to a language – in many cases, the same as the composer).

Language Binding for Components and Expressions. The language binding is made explicit by the namespace that is used in the root node of an expression (and declared there or in one of its ancestor elements; e.g. in the `<eca:Rule>` element). The namespace declaration always yields a URI.

The markup of an ECA rule with language information has the following form:

```
<eca:Rule>  
  <eca:Event xmlns:evns="ev-lang-uri" >
```

```

    <evns:element-name> ... </evns:element-name>
  </eca:Event>
  <eca:Query xmlns:qns="q-lang-uri" >
    <qns:element-name> ... </qns:element-name>
  </eca:Query>
  :
</eca:Rule>

```

Here, *ev-lang-uri* and *q-lang-uri* are URIs associated with the namespaces of an event language and a query language.

The meaning of the URIs will be discussed in Chapter 9 (since there is not yet a standardization what is “behind” the namespace URI, we propose an intermediate solution that is sufficient for the infrastructure in our approach). Each of these languages (i.e., their URIs) has an associated engine that captures the semantics of the (composers of its) language. The engines provide the (expected) interfaces for communication, must keep their own state information, including at least the current variable bindings. Specific tasks of the engines then include e.g. the evaluation of composite events (for the event languages), or the execution of transactions (for the action engines). Thus, the MARS framework itself does not have to deal with actual event detection or transaction execution, but only with employing suitable services (provided by the “owners” of these sublanguages) on the Web.

The leaves of the markup trees are then atomic events, literals, or actions, contributed by the underlying domains (and residing in the domain’s ontology and namespace).

Special markup elements are provided for using and binding *variables* inside the expressions and on the rule level (e.g., results of the event detection or of functional queries); see Section 4.3.1.

3.3.7 Opaque Rules and Opaque Components

Rules and components can also be given in an “opaque” form. This means that they are given as program code in some already existing language and have an operational semantics as an ECA rule or component. Opaque rules and components include e.g. the following:

- SQL Triggers (opaque rules)
- SQL/SQLX, XPath or XQuery queries
- matching regular expressions
- program code for actions
- Web Service calls via HTTP

Especially during the development of a prototype, such functionality will be used. Wrappers are employed to integrate them into the MARS framework (cf. Section 10.1.2).

3.3.7.1 Opaque Rules

There are the existing trigger-style languages that handle specific, simple database events, simple conditions and actions, with their own syntax as discussed in Section 3.2 above. Since these triggers work on the logical level and are in general (very efficiently) implemented based directly on the physical database level, they are not necessarily marked up in ECA-ML. Often, they are even not subject to the “Semantic Web”, since they are just used to locally *implement* something that is *specified* in a completely different way (e.g., integrity constraints), or for raising RDF-level events based on an SQL or XML storage. In our ontology, we embed this as *opaque* rules as shown in Figure 3.12.

The ECA-ML language integrates opaque rules by the following markup: An `<eca:opaque>` element with text content (program code of some rule language) and an attribute `language` that indicates

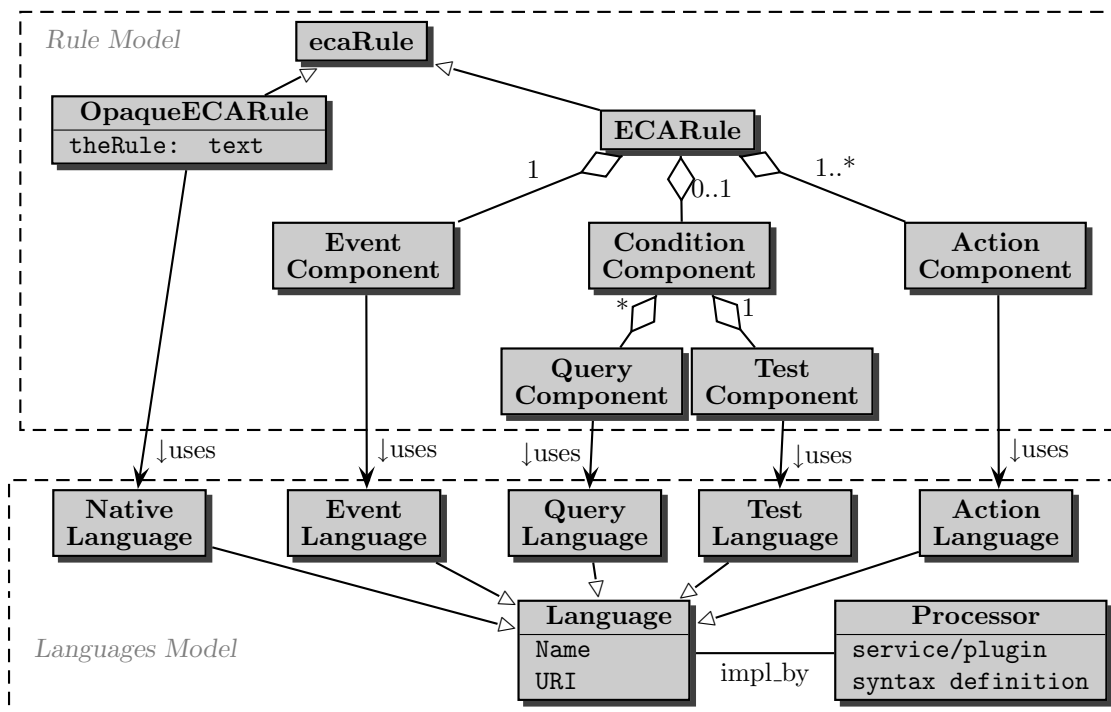


Figure 3.12: ECA Rule Components and Corresponding Languages II

the name of a language or the URI where an interpreter is found, similar to the namespace. Names can also be allowed for language if it is designed similar to XML's NOTATION concept that associates resources with a program URI.

```

<eca:Rule>
  <eca:opaque language="uri of the trigger language" >
    ON database-update WHEN condition BEGIN action END
  </eca:opaque>
</eca:Rule>

```

Since opaque rules are ontologically "atomic" objects, their event, condition, and action components cannot be accessed by Semantic Web concepts.

Note that there are canonic mappings between such triggers and their components and the ECA ontology.

3.3.7.2 Opaque Components

Analogous to *opaque* rules, we allow for opaque rule components. An opaque component consists of a code fragment (not in XML markup) of some event/query/logic/action language, e.g., conditions that are not expressed in a markup language, but in XQuery, or actions expressed in Java or Perl. Opaque components extend the possible syntax of expressions as shown in Figure 3.13. Opaque components play an important role especially during the development of the MARS framework when not yet any dedicated sublanguages are available. In such cases, again `<eca:opaque>` elements are used that refer to a URI for that language.

Opaque fragments are marked up as shown below and do not reside in a namespace:

```

<eca:{event|query|test|action}>
  <eca:opaque attributes>

```

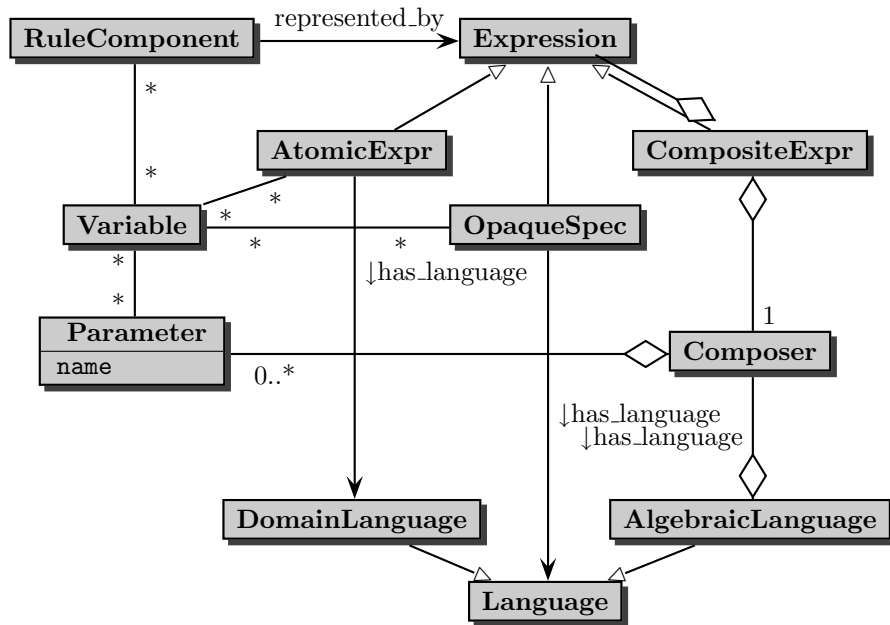



Figure 3.13: Syntactical Structure of Expressions (Algebraic and Opaque)

```

opaque code
</eca:opaque>
</eca:{event|query|test|action}>

```

For being able to execute such opaque fragments, the following cases must be considered (note that the examples do not deal with the handling of the result; this will be discussed in the next section):

- Opaque component language expressions against generic (language) services: *opaque code* is e.g. an XQuery query against the Web or a Web source given as `document("...")` that can be executed by any XQuery service (e.g., based on saxon). In this case, the `eca:opaque` element has a `language` attribute whose value may either be a URI, or a namespace prefix that is declared before (and must be resolved to a URI by the respective application).

```

<eca:Query>
  <eca:opaque language="http://www.w3.org/XQuery">
    for $c in document("http://dbis.uni-goe.de/mondial/mondial.xml")//country
    return $c/name
  </eca:opaque>
</eca:Query>

```

Note that an important instance of such a generic service is a *local* service for evaluating XPath expressions against (small) instances bound to variables (e.g., for analyzing XML fragments returned by event or query components; see Section 4.3.1).

- Specific Domain Services: e.g., XML repositories, SQL databases, or Web Services can be queried by HTTP GET. In all cases, there is a URL to which a certain query is submitted. In this case, the `eca:opaque` element has a `uri` attribute that indicates the service. Additional attributes can be used to give further specifications.

For HTTP GET there are two possible syntaxes:

- append the query part to the base URI and having an empty content:

```
<eca:Query>
  <eca:opaque uri="http://exist.dbis.uni-goe.de/exist/servlet/db/mondial.xml
    ?_query=//country[name='Germany']&_wrap=no"
    method="get"/>
</eca:Query>
```

- having the base URI as the uri attribute, and the local part as content:

```
<eca:Query>
  <eca:opaque uri="http://exist.dbis.uni-goe.de/exist/servlet/db/mondial.xml" >
    method="get" >
    ?_query=//country[name='Germany']&_wrap=no
  </eca:opaque>
</eca:Query>
```

The latter has the advantage that it separates the service part from the query part:

- * having the service part separate allows to use information about the service,
- * the query part can contain variables whose values must be inserted,
- * and it is easier accessible to analysis and reasoning.

This holds especially when the service is not a database query service, but a form-like “lookup” Web Service:

```
<eca:Query>
  <eca:opaque uri="http://simple-service.dbis.uni-goe.de" method="get" >
    ?name='Germany'
  </eca:opaque>
</eca:Query>
```

Discussion (including definitions of Section 4.1.5):

- when input variables are declared, the GET syntax can be derived when the URI is known: $uri?var_name_1=value_1\&\dots\&var_name_n=value_n$
- in case that a service with multiple “methods” is used, the service should be given in the URI part (since the service description according to Chapter 9 is associated with it) with this part, all other information in the contents.
- More advanced HTTP methods (POST) and HTTP-based protocols (like SOAP) that require pre-processing before sending the request are also indicated as attributes of the `eca:opaque` element. The content of the opaque element is then sent as message content; possibly modified for communication variable bindings; cf. Section 4.2.2).

```
<eca:Query>
  <eca:opaque uri="http://soap-service.dbis.uni-goe.de/xpath-servlet" >
    for $c in doc("/db/mondial.xml")//country
    return $c/name
  </eca:opaque>
</eca:Query>
```

In case that service descriptions according to Section 9.5 are used, it should *not* be necessary to indicate the used communication protocol explicitly.

If a component is given as opaque code, the communication is chosen according to the language indicators:

service URI	protocol	method	used protocol
http://...	-	-	native (framework)
http://...	-	get/post	http with indicated method
ip-address	tcp	-	tcp
ip-address	-	-	native (framework)

Note that it is possible and useful to use opaque code with framework-aware services: This is the exact characterization of wrappers around non-framework-aware services and provides a good means for rapid prototyping of an ECA engine.

Note that so far we did not deal with submitting values of variables to the services (either as values for comparison in predicates, or as “documents” that have to be queried); this will be discussed in Sections 4.1.2, 4.2.2, and 4.2.3.

Chapter 4

Abstract Semantics: Rule Level

This chapter develops the MARS Framework from the point of view of the rule level. For this, a rough intuitive idea what the event, query, test and action components do is sufficient. They will be considered in more detail in Chapter 5.

4.1 Abstract Declarative Semantics of Rule Execution

This section deals with the overall semantics of ECA rules and the abstract semantics for each of the components, and with the actual communication. Although the languages are heterogeneous wrt. the components, there is a common requirement: to support language heterogeneity at the rule component level, there must be a precise convention between all languages how the different components of a rule can exchange information and interact with each other.

There are two kinds of communication in the rules:

- Horizontal communication, following the data flow in the rule according to Figure 3.7 from the event component to the query component and so on.

For this, we propose to use *logical variables* in the same way as in Logic Programming (cf. Section 4.1.2).

- Vertical communication between the ECA engine and the engines that are responsible for processing the component languages.

We align this way of communication with the concepts used for logical variables by regarding every component as a mapping that takes as input a set of tuples of variable bindings and returns another set of tuples of variable bindings.

We introduce the notion of variable bindings that are used for both kinds of communication and provide a logic-based declarative semantics on the rule level and for data exchange with component services.

4.1.1 Rule Semantics

Comparison: Firing Deductive Rules. For deductive rules (that do not have an event component) in *bottom-up* evaluation, the body is evaluated and produces a set of tuples of variable bindings (Datalog, F-Logic, Transaction Logic, Statelog, XML-QL, XPathLog [49], Xcerpt [23]; in some sense also the basic form of XQuery). Then, the rule head is “executed” by *iterating* over all bindings, for *each* binding instantiating the structure described in the head (in some languages also executing actions in the head).

The semantics of ECA rules should be as close as possible to this semantics, adapted to the temporal aspect of an event:

ON event AND additional knowledge, IF condition then DO something.

We transfer the concept of variable bindings to ECA rules in the following examples:

Example 3 (Cancelled Flights) Consider a rule that should do the following: Whenever a flight is canceled, send a message to the destination airport that the flight will not take place. Assume that events are of the following form

`<travel: canceled-flight code="LH1234" reason="weather" />`

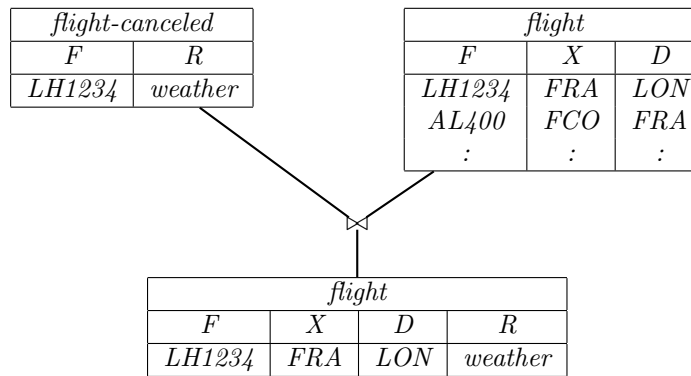
indicating the code F of the canceled flight and the reason R of the cancellation. Then, a query against a database is evaluated that yields the destination airport D . Finally, a mail is sent to the target airport D , telling that and why the flight F is canceled.

A logic programming-style rule (using a relation `flight(flight_no, from, to)`) would look like this:

`travel:canceled-flight(F , R) \wedge flight(F , $_$, D) \rightarrow send_mail(D , F , "canceled" , R).`

In LP terminology, the event component `travel:canceled-flight(F , R)` binds variables F and R . The query component `flight(F , $_$, D)` alone binds variables F and D . Both results are joined and the result is then propagated to the action component (which corresponds to the rule head in LP) where they are then used.

Formally, the rule semantics is join-based:



An alternative, algebraically equivalent, evaluation strategy (which corresponds to evaluating a join based on an index) is to use the (only) binding of F from the event part to lookup the destination and to just to extend the variable bindings.

4.1.2 Logical Variables

We propose to use *logical variables* in the same way as in Logic Programming that can be bound to several things: values/literals, references (URIs), XML or RDF fragments, or events. The representation of the bound items must be in ASCII, e.g., URIs, serialized XML, or XML-serialized RDF. The binding of a variable to an event (or a sequence of events) e.g. occurs in the SNOOP language [24] in *cumulative aperiodic events*; such variables can then be used for extracting values from these events. Variables can be bound by the rule (as constants upon registration) or by the components and used in later components.

For logical variables used in LP rules, there are several definitions that carry over to EQTA rules:

- Similar to deductive rules, variables used for communication occur *free* in the components, their scope is the rule,

- While in deductive rules, variables must be bound by a positive literal in the body¹ to serve as join variables in the body (adhering to safety requirements!) and to be used in the head, in ECA rules we have four components that induce an information flow according to Figures 3.7 and 3.8.
- Positive occurrences are defined analogously to deductive rules based on the term/formula structure (must be done with the semantics of each individual such language).
- Positive occurrences can be used to bind variables to a value. In case that a variable occurs positively several times, it acts then as a join variable, i.e., the values must coincide; this e.g. allows for an event component that in some cases binds a variable which is then used as a join variable in the condition, and in other cases is only bound by the latter.
- Negative occurrences of a variable *use* the value the variable has been bound to *before*.
- Thus, during execution of a rule, any variable occurring negatively must be bound to a value earlier on the rule level (e.g., with the rule’s initialization, or by deriving its value from another variable) or in an “earlier” component (E<Q<T<A) or at least “earlier” in the same component as where the negative occurrence is. This leads to the usual definition of safety of rules.
- Expressions can also use local variables, e.g., in first-order logic conditions. In this case, the scope of a variable is local, e.g., by a quantifier.
- Variables in the action component: Using variables as parameters to an action in the action component counts as negative occurrences (in case that a language used in the action component does not define positive or negative occurrences). Thus, for languages used in the action component that do not define the notions of positive/negative occurrences, all occurrences are negative. Note that this allows for binding a variable in the action component, e.g., by allowing for evaluation of queries in that component, like in Transaction Logic [18] that defines its own notions of positive occurrences.

The relationships between rules, rule components and variables on this abstract level are shown in Figure 4.1:

- for each rule R , $R.\text{scopes}$ (the set of all logical variables whose scope is the rule) is the union of $R.\text{occurs_positive}$ (before or between evaluating components) and $C.\text{free}$ for all its components C ;
- for each component C of a rule R , and also for each expression inside any component, $C.\text{positive}$, $C.\text{negative}$, $C.\text{free}$, $C.\text{bound}$, denote the sets of variables that occur positively, negatively, free or bound.

Free variables occur either positively or negatively:

$$C.\text{free} = C.\text{positive} \cup C.\text{negative}.$$

4.1.3 Horizontal Communication: Logical Semantics

So far, seeing components as abstract predicates yields a preliminary declarative semantics of ECA rules which is based on the LP join semantics. The *horizontal* communication during processing ECA rules is actually based on propagating sets of tuples of variable bindings. Several issues are dealt with in the subsequent sections:

¹Note that there are different terminologies in the literature about “positive” and “negative” literals: In Logic Programming, these notions are defined wrt. the rule body, whereas in works based on resolution of disjunctive clauses, they are defined wrt. those literals (as in Xcerpt/XChange). Since a (Horn) clause $p(x) \vee \neg q(x)$ corresponds to a LP rule $p(x) \leftarrow q(x)$, in the first case, $q(x)$ is a negative binding whereas in the second case, it is a positive binding.





Figure 4.1: Logical Rules and Variables

- replacing the abstract predicates by actual event, query, and action components that are specified by sublanguages and evaluated by appropriate services. This requires vertical communication;
- special semantics of events since they are not like usual predicates; especially, the semantics of event algebras (event sequences, cumulative events) must be integrated accordingly;
- special semantics of Web Services since they use specific notions of input and output variables;
- appropriate adaptations of the notions of positive and negative occurrences of variables and their influence on the evaluation of joins;
- procedural aspects like event detection and at least basic transactional functionality.


4.1.4 Vertical Communication

In the same way, the communication with the engines that process the sublanguages is done by sets of tuples of variable bindings. The abstract predicates are now replaced by the actual evaluation. Here, the interplay between the variable bindings on the ECA level and the component services and their results can be designed in different ways:

- **Join-Style:** at a given state of evaluation, the next query is evaluated independent from the already obtained variable bindings, and the result is joined with the already existing variable bindings. Here, the functionality of the ECA engine consists of invoking a service and joining results:



- Similar to *sideways information passing strategies* in algebraic evaluation that use variables that are already bound as constraint when evaluating the next predicate or subquery for minimizing results as soon as possible (cf. internal evaluation of Florid [31]), existing variable bindings can be communicated downwards when calling a component language engine. In this case, the join semantics is realized as a restriction in the component language engine:



- The actual evaluation is a mixture for several reasons:
 - only actual Datalog services would implement the pure LP-style strategy;
 - negated variables must be bound before; here the join acts as a set difference,
 - variables that are not used at all do not need to be communicated downwards and upwards again (this also reduces the number of actual tuples of variable that have to be communicated);
 - some services (especially, many query languages like SQL and XQuery) show a functional behavior that does not bind variables at all;
 - functionality of the services (e.g. Web Services that require a single input value and return a single output value);

In the next sections, downward and upward communication are investigated.

4.1.5 Communication Modes and Declaration of Variables

Usually, languages based on logical variables do not use explicit variable declarations. Nevertheless, for controlling variable exchange, it must be known to the ECA engine which variable must or should be exchanged with the component services (especially during the development phase where opaque components and non-semantic services are employed). This also allows to extend the binding mechanism with a type system.

Consider the abstract structure of a rule over a set x_1, \dots, x_n of variables (without loss of generality, we consider only one query and one action component):

$$\text{event}(x_{e_1}, \dots, x_{e_{n_e}}) \text{ query}(x_{q_1}, \dots, x_{q_{n_q}}) \text{ test}(x_{t_1}, \dots, x_{t_{n_t}}) \text{ action}(x_{a_1}, \dots, x_{a_{n_a}})$$

where $1 \leq n_e, n_q, n_t, n_a \leq n$ and $1 \leq e_1, \dots, e_{n_e} \leq n$ are pairwise disjoint, same for $q_1, \dots, q_{n_q}, t_1, \dots, t_{n_t}$, and a_1, \dots, a_{n_a} .

From a logical point of view, these variables are the *free* ones in the event, query, test and action part. Actually, the components may have an operational semantics that uses some of the x_i as input, and generates some others as output or “result”. From an abstract view, they all can be considered as predicates.

In the following, we will clarify four kinds of variables. Considering the following example query (event, test, and action components are specialized cases):

Example 4 (Variables) *Consider a car rental company that holds a database that holds for each location a list of all models with a classification and prices available. Seen as a predicate, the service has the following extension as a predicate available:*

```

answer("Frankfurt", "Golf", "B", "40")
answer("Hanover", "Golf", "B", "50")
answer("Hanover", "Passat", "C", "120")
answer("Hanover", "S500", "D", "250")
answer("Paris", "C4", "B", "50")
answer("Paris", "Golf", "B", "65")
answer("Paris", "C6", "D", "150")
answer("Bucharest", "Logan", "B", "25")
:

```

Consider now a travel agent service that has some personalized knowledge about its clients, e.g., what car they drive at home. The service has now a rule that states

“when a client books a flight to *place* with an airline *airline*, offer him a list of cars equivalent to his own to rent”.

The variable *place* is bound by the “flight booking” event, the second relevant variable is *class* that can be bound by a query against a local car database. The third step is then to collect offers at the given place, which is the part under discussion now. Logically, the rule reads as

```
offer(Model,Price) :-booking(Person,_From,To), owns(Person,OwnCar),
                      class(OwnCar,Class), available(To,Model,Class,Price).
```

Consider the owner of a VW Golf traveling to Paris, i.e., *To*=“Paris” and *OwnCar*=“Golf” are known from the first line.

- *Bottom-up evaluation*: take the extensions of *class* and *available*, join them with the known values for *To* and *OwnCar* and return the result.
- *Service (1)*: call the service for “Paris”, look up the own car’s class in a local database, and join the results,
- *Service (2)*: look up the own car’s class in a local database, call the service for (“Paris”, “B”), and return the results.

The actual choice in the rule (i) depends on the functionality of the data source, and (ii) results in different costs of query evaluation and size of the transferred result. While (i) depends on the design (and the rule designer must adhere to it), (ii) can be dealt with by an optimizer [to be developed].

Consider now a service that requires the location as an input; giving the class is optional.

- *used variables*: the service only uses the variables *To* and *Class*. The value of *Airline* is irrelevant.
- *input variables*: *To* is an input variable that must be given.
- *result variables*: variables that are only returned; if they are submitted to the service, they are ignored.

As said above, it is then possible to call the service only by submitting the *To* location and doing the join of the answer tuples with the known *Class* afterwards, or call the service with both *To* and *Class*, i.e., the service already applies the join condition. Probably, the second case would be more efficient.

Consider an extension where a traveler who owns several cars is offered all models that are equivalent to some of his own ones. Now, for people owning several cars, the first case would be more efficient in most cases. Here, it would be useful to have a service that can be called with multiple tuples.

Thus, having an environment where variables are not only logical ones, but also are used as arguments and results, the augmented relationships between rules, rule components and variables are shown in Figures 4.2 (adapting Figure 4.1) and Figure 4.3. The *abstract, declarative* semantics is still logical, only the *operational* semantics needs a more detailed look.

Used variables; free variables. In logical languages, these are the variables that occur free. In languages that have a term markup, they can be derived by analyzing the term structure (simply by `//variable/@name`). Inside an algebraic evaluation, it is usually preferable to communicate all used variables that are already bound to use them as constraints and to minimize results as soon as possible (cf. internal evaluation of *Flord*).

Input variables. In logical languages, these are the variables that occur free, but negatively, e.g., the z in $p(x, y) \wedge \neg q(y, z)$: when stating this query against any source, values for z must be provided to be *safe*.

In logic-based languages, they are defined inductively on the structure of the query. In the Semantic Web environment, they must either be indicated by the rule designer (based on knowledge about the services), or they must be distinguished by the service description.

Output variables. Output variables do not exist in logical languages. Nevertheless, in the Semantic Web, they do exist. Output variables are variables that are bound by a service, independent if they have been bound before or not. To provide join variable semantics, they *must* be considered by equi-joining the result with the values bound before.

Returned variables. When considering to join the answer set of a predicate with a set of already existing variable bindings, it must be taken into consideration whether an answer actually contains all variables that are relevant for a join:

- Datalog answer-style: a query, e.g., `?-capital('Germany',X)` results only in the variable binding `X/'Berlin'` .

- When considering a given set of variable bindings as a constraint, a query, e.g.,

`?-capital(C,X)[{(C/'Germany'),(C/'France')}]`

results in extended variable bindings

`{(C/'Germany',X/'Berlin'), (C/'France',X/'Paris')}` .

- functional “lookup” services: the input variables are only communicated downwards, and only the result is communicated upwards. Such services are e.g. provided by HTTP forms:

`http://capitalsoftheworld.com?country='Germany'`

would result only in the value 'Berlin', which must then be bound to a variable. Such lookup services can also return frame-like data; e.g.

`http://countriesoftheworld.com?country='Germany'`

could return a structure like

```
<code>D</code>
<area>356910</area>
<population>83536115</population>
```

which could be interpreted as binding three variables.




Figure 4.2: Use of Variables in Components of ECA Rules

The evaluation of the event component (i.e., the successful detection of a (composite) event) binds variables to values that are then extended in the query component, possibly constrained in the test component, and propagated to the action component.

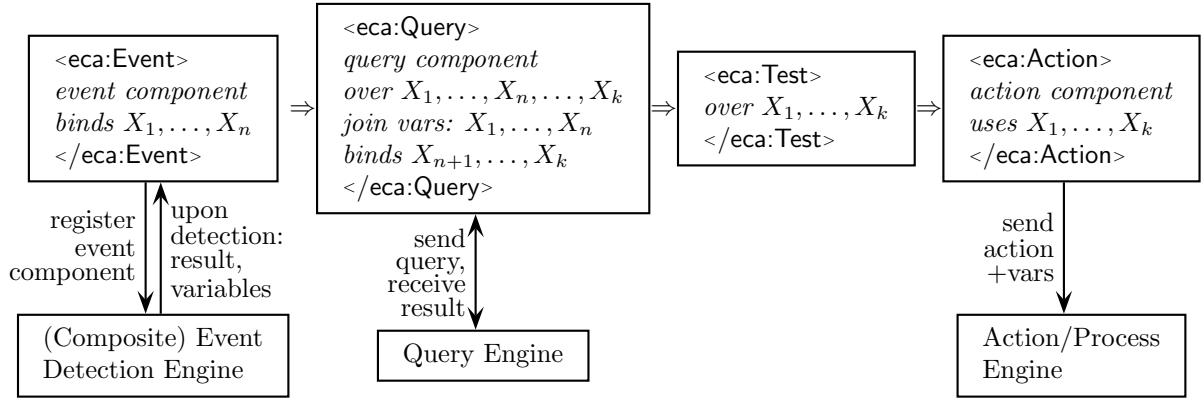


Figure 4.3: Use of Variables in an ECA Rule

Variables and kinds of components. Depending on the type of the component (event, query, test, and action), the following kinds of variables can occur:

	input	used and either input nor output	output	returns
event	(X)	-	X	X
query	X	X	X	X
test	(neg)	(pos)	-	X
action	X	-	-	-

(X) means that this can be only variables bound to constants at the time of registering the rule. Since a test is a logical *condition*, the usual definitions of positive and negative occurrences of variables apply. The above table shows that the query part, by *obtaining* information and correlating it to already present information needs the most complex communication interface.

Declarations of Variable Use. For a rule component C (and an expression in general, since the usage of variable is defined recursively), $\text{used-vars}(C)$, $\text{input-vars}(C)$, $\text{output-vars}(C)$ and $\text{returned-vars}(C)$ denote the above sets of variables.

For that reason, at least when evaluating an ECA rule according to the specification that will be given in Section 4.1, knowledge of the above-mentioned sets of variables is useful:

- Downward communication:
 - If for a component C , $\text{used-vars}(C)$, $\text{input-vars}(C)$ and $\text{output-vars}(C)$ are known, check if a value for each of the $\text{input-vars}(C)$ is contained in each tuple (otherwise, abort). Then, any set $\text{input-vars}(C) \subseteq V \subseteq \text{used-vars}(C) - \text{output-vars}(C)$ is useful to be communicated. Note that $\text{used-vars}(C) - V$ must be considered afterwards in a join.
 - In case that the set $\text{used-vars}(\text{component})$ of variables that are used in a component is known (which, as mentioned can be done relatively simple if the component is marked up), it is safe (i) to communicate all existing bindings downwards *and* perform a join afterwards.
 - Otherwise, all existing bindings must be communicated downwards.
- Upward communication:
 - if $\text{used-vars}(C) = \text{returned-vars}(C)$, then the result can be joined immediately.
 - if $\text{used-vars}(C) \supsetneq \text{returned-vars}(C)$, then the component must be evaluated separately for each bindings of $\text{used-vars}(C) - \text{returned-vars}(C)$ and the missing variables must be taken into consideration.

- In reality, there are mainly two kinds of services:
 - * Logic Programming semantics: $\text{outputvars}(C) = \emptyset$ and $\text{returned-vars}(C) = \text{used-vars}(C)$, and
 - * lookup services: $\text{outputvars}(C) \neq \emptyset$ and $\text{returned-vars}(C) = \text{output-vars}(C)$.
- Default assumptions: if no `returned-vars`, then the above cases are assumed.

The markup of components may thus *optionally* contain explicit declarations of the usage of variables.

However, the goal is that this information can be derived from the markup and the semantic information about the component languages. For this, every service that “offers” a language should provide the following functionality (see also Chapter 9).

- Algebraic component languages: given an XML or RDF fragment of an instance of the language: Validation of the fragment, list of all variables that are used, and all variables that occur positively (i.e., can be bound by this fragment).
- For every Domain Language and Domain Web Service the input and result variables should be defined in a service description.

Explicit declarations are mainly important for handling opaque fragments and non-Semantic Web Services. For the ECA prototype, such services are frequently used.

Example 5 (Rental Cars (Revisited)) *Consider again the car rental example from above. Since we focus on the handling of variables, we give the components as (atomic) Datalog predicates with variable communication mode declarations. Assume that the Datalog service is local and wrapped appropriately with a framework-aware wrapper (see Section 10.1.2.2).*

```
<eca:Rule xmlns:eca="http://www.semwebtech.org/languages/2006/eca-ml#"
  xmlns:xpath="http://www.w3.org/XPath"
  xmlns:pseudocode="http://www.pseudocode-actions.nop">
  <eca:Event>
    <eca:opaque language="datalog-match" output-variables="Person To">
      booking(Person,_From,To)
    </eca:opaque>
  </eca:Event>
  <eca:Query>
    <eca:opaque language="datalog-match" used-variables="Person OwnCar">
      owns(Person,OwnCar)
    </eca:opaque>
  </eca:Query>
  <eca:Query>
    <eca:opaque language="datalog-match" used-variables="OwnCar Class">
      class(OwnCar,Class)
    </eca:opaque>
  </eca:Query>
  <eca:Query>
    <eca:opaque uri='http://localhost/lookup-cars' input-variables="To" output-variables="Class Model Price">
      ?_place=To
    </eca:opaque>
  </eca:Query>
  <eca:Action>
    <eca:opaque uri='localhost' input-variables="Model Price">
      show all model-price tuples
    </eca:opaque>
  </eca:Action>
</eca:Rule>
```

4.2 Logical Variables: Markup for Communication

4.2.1 Basic Interchange of Variable Bindings

We propose the following representation for variable bindings that will be used for interchange (see also Section 4.2.3):

```
<!ELEMENT variable-bindings (tuple+)>
<!ELEMENT tuple (variable+)>
<!ELEMENT variable ANY>
<!ATTLIST variable name CDATA #REQUIRED
            ref URI #IMPLIED> <!-- variable has either ref or content-->
  <logvars:variable-bindings>
    <logvars:tuple>
      <logvars:variable name=" name" ref=" URI" />
      <logvars:variable name=" name" >
        any value
      </logvars:variable>
    :
  </logvars:tuple>
  <logvars:tuple> ... </logvars:tuple>
  <logvars:tuple> ... </logvars:tuple>
  :
  <logvars:tuple> ... </logvars:tuple>
</logvars:variable-bindings>
```

Note that such data exchange is required not only for ECA rules, but for all kinds of services that are based on logical variables (rules, queries, reasoners). For this, we propose to use a separate namespace, here called `logvars`, referring to the URI <http://www.semwebtech.org/lang/2006/logic#>. The complete `logvars` DTD can be found in Appendix 13.2.4.

In the next sections, the exchange of variable bindings with the component services is discussed, extending this basic structure and markup.

4.2.2 Downward Communication: Variable Bindings

The actual handling of downward communication depends on the interface provided by the respective services.

Downward communication occurs in the following cases (note that the handling of the event component differs significantly from that of queries, tests, and actions):

- When a rule is *registered*, the event component is submitted to an event detection engine (that understands the event language). Optionally, variables that are already bound on the rule level (when handling *rule patterns*) can also be contained in the message. The message must contain the following information (see Section 5.1 for the actual structure and syntax):
 1. administrative
 - where the answer should go, and
 - an identification to be used in the answer.
 2. contents:
 - the (event) component or a reference to it, and
 - optional: the current variable bindings (in the format described in Section 4.1.2). Note that only those variables need to be communicated that are actually used in the component according to the above considerations.

For events, there is necessarily an asynchronous communication of requests and answers:

- downward communication by `register` at *registration time*,
 - upward communication (see Section 4.2.3) by `logvars:answers` or `logvars:answer` at *rule evaluation time* when the event is actually detected (see Section 5.3).
- Queries, tests, and actions are submitted at rule evaluation time for certain variable bindings. The information contained in the message is similar to that for events.

Note that it is also possible to *register* a query or an action with free variables at registration time, and to *request* answers or execution (for given variable bindings) at *rule evaluation time*.

4.2.3 Upward Communication: Results and Variable Bindings

Upward communication is concerned with the exchange of results and variable bindings, i.e., returning results and bound variables from evaluating an expression. For this, the structure and markup shown in Section 4.2.1 is extended.

Upward Communication: (Functional) Results and (Logical) Variables. There are several possibilities what the “result” of evaluating a rule component can be:

- Logic-Programming-style languages that bind variables by matching free variables (e.g. query languages like Datalog, F-Logic [41], XPathLog [49]). Here, the matches can be *literals* (Datalog) or literals and structures (e.g., in F-Logic, XPathLog, Xcerpt). Similar techniques can also be applied to design languages for the event component.
- Functional-style languages that are designed as functions over a database or an event stream and a set of input/environment variables:
 - query languages that return a set of data items (e.g., SQL, OQL) that can be interpreted as producing a set of variable bindings (attribute names as variables; probably obtained by the “renaming” operator of the relational algebra),
 - query languages that return a data fragment (e.g. XQuery, Xcerpt [23] – this is only possible since “schema-free” data like XML exists). Here, the result is not bound to an obvious variable name. Note that this should result in a set/sequence of variable bindings if a set/sequence of nodes is created.
 - for event languages, the “result” of an expression can be considered the sequence of detected events that “matched” the event expression in an event stream (e.g., XChange).

In this case, the languages can also *use* variables that are bound before. Thus, “downward” communication is explicit, whereas the upward communication is implicit (and the result must be bound to a variable by the *surrounding* language). Note that the answer can even be empty which *sometimes* is interpreted as “false”.

- Both forms can be mixed (F-Logic, cumulative operators in event languages).

To cover all of them, we propose a structure as e.g. used in the Florid system [31] where with *each* result, a set of variable bindings is associated (cf. [46, Section 2.2.1]). We propose the following representation for interchange of results and variable bindings that extends the markup already used in the `logvars` namespace (the complete `logvars` DTD can be found in Appendix 13.2.4).

```
<!ELEMENT answers (answer*)>
<!ELEMENT answer (result—variable-bindings—(result,variable-bindings))>
<!ELEMENT result ANY>
<!ELEMENT variable-bindings (tuple+)>
```

```

<!ELEMENT tuple (variable+)>
<!ELEMENT variable ANY>
<!ATTLIST variable name CDATA #REQUIRED
          ref URI #IMPLIED>  <!-- variable has either ref or content-->
<logvars:answers>
  <logvars:answer>
    <logvars:result>
      any result structure
    </logvars:result>
    <logvars:variable-bindings>
      <logvars:tuple>
        <logvars:variable name=" name" ref=" URI" />
        <logvars:variable name=" name" >
          any value
        </logvars:variable>
        :
      </logvars:tuple>
    </logvars:variable-bindings>
  </logvars:answer>
  <logvars:answer>
    :
  </logvars:answer>
</logvars:answers>

```

A set of answers consists of multiple answers, where each answer consists of a result value and/or a set of tuples of variable bindings. A variable binding can either be given inline as serialized XML, or as a URI reference (e.g., to a Web page, or an RDF URI). The answers can optionally contain a reference to an ID that in case of asynchronous communication indicate to what they are an answer (this is needed for event detection; query handling can be synchronous or asynchronous). Note the following:

- In cases where only one single answer is produced (which is often the case for event detection, or when calling a “functional” Web Service), the outer `<logvars:answers>` may be omitted, returning only one `<logvars:answer>` structure.
- for services that return no functional result (e.g., a Datalog query service) or no variable bindings (e.g., an XQuery service), each `<logvars:answer>` structure contains only the relevant subelement.
- for services that return only a single functional result (an event sequence, or an answer to an XQuery or SQLX query), it is allowed not to mark it up at all. It is then treated as `<result>` element of a single answer and can be bound to an ECA-level variable as described below.

Upward communication occurs in the following cases (see Section 5.1 for the actual structure and syntax):

- Event detection: an event has been detected. Usually, the result consists of the sequence of relevant events (as functional result) and variable bindings. The message must contain the following information:
 1. administrative: the identification of the event that has been detected
 2. contents: the result (in the format described above for `<answers>`).
- Query answering. Here the result also consists of the a set of answers as above. There is the possibility that a query answer is sent in several parts.

1. administrative: the identification of the query that is answered, and whether the answer is complete (default: yes);
2. contents: the result (in the format described above for <answers>).

4.3 Markup: Binding and Using Variables

Variables can be bound to many kinds of data: XML nodes like elements and attribute nodes, text contents, strings, numbers, and –in an RDF setting– also RDF URIs. Thus, the *mechanisms* for dealing with variables must be generic. On the other hand, since variable binding and using happens inside the scope of individual languages, also the in general depends on the individual language.

The MARS framework only defines how <ecans:variable> can be used on the ECA level (and immediately below it), and in the next paragraphs we discuss some design alternatives how to integrate the handling of variables in to component languages.

4.3.1 Alternative Syntaxes

While the semantics of the ECA rules provides the infrastructure for these variables, the markup of specific languages must provide the actual handling of variables (mainly: binding variables) in its expressions. We propose to use a uniform handling of variables in the ECA language (see also Section 4.3.3), and in the E, Q, T, and A component languages.

Variables: Syntax. We propose to use the following constructs for binding variables, where we borrow from several language designs (note that in contrast to some other languages, binding and using variables is here the same):

1. (syntax for use borrowed from XQuery, XSLT and XML-QL; syntax and semantics for binding borrowed from XML-QL [25]): (use) variables by `{var-name}`:

```
<travel:canceled-flight number=$flight/>
```

or

```
<travel:canceled-flight number="{$flight}" />
```

matches an event (e.g., <travel:canceled-flight number="LH0815" />) and binds `$flight` to the number of the flight. Note that the second variant is valid XML where the `"{$name}"` is embraced in quotes and parentheses. This *atomic event specification language (AESL)* will be considered again in Section 5.2.2.1.

2. (borrowed from XSLT): (bind) variables by <variable name="..."> elements:

```
<foo:variable name="name">
  content
</foo:variable>
```

where *content* is any expression (e.g., an XML fragment in a match-style query language, an event specification or a query) that returns some value. The variable is then bound to this value (see also Examples 14 and 9 below). Note that the `foo:variable` elements reside in the component language's namespace (the variable is part of the component language tree).

3. (reminiscent of F-Logic and XPathLog): alternative way of binding results of subexpressions to variables:

```
<foo:bla foo:variable="name" >
  content
</foo:bla>
```

where `<foo:bla>` is any expression (e.g., an XML fragment in a match-style query language, an event specification or a query) that returns some value. The variable *name* is then bound to this value.

This syntax corresponds to $o[m \rightarrow V[a \rightarrow b; \dots]]$ (V bound to the results of the property m of o) in F-Logic and even more $e[b \rightarrow V[c \rightarrow d; \dots]]$ in XPathLog (V bound to the b subelements of e). Xcerpt uses a similar construct for bindings variables in its query terms.

4. XLink-out-of-line or RDF style: Variable bindings can be seen as references into a term (XLink) or as edges between a variable name and a term. Then, they can be expressed by elements of the form

```
<foo:variable name="name" select="$component/xpath-expression" />
```

where *xpath-expression* references into the term structure of a component expression (should be relative to the rule, the component, or self). This can mainly be applied for *matching* purposes e.g. for (atomic) events or (XML) pattern queries.

We propose to use similar constructs also in the component languages, but the actual decision is up to the language designers.

4.3.2 Discussion: Variable Syntax

The above syntaxes have their advantages and disadvantages:

- all three are “used” to people from certain communities.
- case (1) is the most simple and “clear”, both for people from Logic Programming (matching of variables) and from programming languages like XQuery. We currently do not see any problems with it. But it is only usable for simple cases (attributes or whole element contents).
- case (2): this interferes with the term structure. This is not relevant for *executing* a rule, but when rules and their components are seen as resources and have to be addressed. The navigation expressions for addressing subterms in a pattern depends on where such variables are mentioned.
- case (3): here, the use of variables is indicated *inside* of elements (which violates their DTD). This is problematic when submitting components to services (in this case, the attribute must be removed; it is intended to be used by the service that processes the *outside*).
- case (4): here, the variables do not interfere with the term structure. On the other hand, since they are decoupled from the structure, reasoning requires to parse the language used in the `select` attribute.

The path expressions should be restricted to the same level or component where the `<variable>` is located (e.g., on rule level binding the result of a query, but not binding the second subexpression of an event expression) since otherwise communication is required that is not provided by the loosely coupled language modules.

It is obvious that this must be considered carefully. When using RDF instead of XML markup, things become easier since these things can be dealt with in separate triples (in any namespace). Variables are then like “annotations” to a structure. Note that the above case (4) is already in RDF style.

Next, we discuss this issue for ECA-ML. Another discussion can be found in Section 5.2.6 when a markup for a composite event language closely related to the SNOOP event algebra [24] is presented.

4.3.3 Variable Bindings by ECA Rules

For making the functional result part of a component (e.g., i.e., the of the event component, or of a query component) accessible in the ECA rule, it must immediately be bound to a variable on the rule level. For this ECA-ML provides several equivalent mechanisms according to the above mechanisms:

- `eca:variable` elements may occur on the ECA rule level. Then they contain an expression whose results are bound to that variable:

```
<eca:Rule... >
  <eca:variable name="result-var" >
    <eca:Event ... >
      :
    </eca:Event
  </eca:variable>
  :
</eca:Rule>
```

(analogous for queries)

- `eca:variable` elements may occur on the upper level *inside* a component. The following is equivalent to the above:

```
<eca:Rule... >
  <eca:Event... >
    <eca:variable name="result-var" >
      contents
    </eca:variable>
  </eca:Event>
  :
</eca:Rule>
```

Here, the use of a variable is indicated *inside* of the event element – although actually intended to be used by the service that processes the *outside*. Only the contents of the `eca:variable` element must be sent to the event component service. (Note that if the component language, say, `foo` also provides a variable syntax as above, also `<foo:variable` would have the same semantics.)

- the following syntax according to (3) is also equivalent:

```
<eca:Rule... >
  <eca:Event [eca:]variable="result-var" >
    contents
  </eca:Event>
  :
</eca:Rule>
```

Again, the use of a variable is indicated *in* the event element – although actually intended to be used by the service that processes the *outside*. The `eca:variable` attribute must not be sent to the event component service.

- `<eca:variable>` elements of the form (4)

```
<eca:variable name="name" language="xpath" select="expr" />
```

can be used for binding a new variable based on already bound ones in *expr*, e.g.,

```
<eca:Rule... >
  something that binds variable x
  <eca:variable name="y" language="xpath" select="$x/foo/@bar" />
  :
</eca:Rule>
```

These expressions can be in any language (e.g. XPath; the `language` attribute can be omitted if the service uses a default language) that an ECA service implements for such simple *local* evaluations. The above is a shorthand for

```
<eca:variable name="name" >
  <eca:Query>
    <eca:opaque language="xpath" > expr </eca:opaque>
  </eca:Query>
</eca:variable>
```

- a similar style is used for “forward” declarations for binding a variable to a functional result of a rule component by

```
<eca:variable name="name" language="xpath" bind-to="expr" />
```

(note: `bind-to` instead of `select`) where *expr* begins with `$rule`, e.g.,

```
<eca:Rule... >
  <eca:variable name="queryresult2" language="xpath" bind-to="$rule/query[2]" />
  :
</eca:Rule>
```

Such variables are bound as follows:

- If the result from evaluating a component is one or more `<logvars:answer>`, then for each `<logvars:answer>`, every `<logvars:tuple>` in the `<logvars:variable-bindings>` part is extended with the variable *result-var* which is bound to the `<logvars:result>` part of the `<logvars:answer>` (see Examples 14 and 9 and below).
- If the result is other XML (e.g., from evaluating an XPath query or analogously for RDF), *result-var* is simply bound to it in the same way as in XSLT (this often saves writing wrappers to the above exchange format).

Note that most currently existing tools (e.g. query interfaces) do not return their data in such a format. In such cases, wrappers (that can be provided locally by the ECA service) can be used. Note that the XQuery `return` clause (and similar constructs like XML generating functions in SQLX) can be used to return this format directly (this functionality is especially used in the prototype for rapid prototyping using existing XQuery services; see Section 10.1.2).

Usually, event detection returns the relevant event sequence in this way, we illustrate this situation in Section 4.4.1.

Requirements.

- if a request contains multiple tuples over variables X_1, \dots, X_n , then *each* resulting tuple must bind a superset of X_1, \dots, X_n (to allow an unambiguous correlation). Such services incorporate a full join semantics (cf. declaration of `returned-vars` in Section 4.1.5).

- if a service allows only for requests that contain a single tuple of variable bindings X_1, \dots, X_n , the ECA engine must invoke it for each tuple of bindings X_1, \dots, X_n that occurs in the current variable bindings. It must then correlate each of the results to the correct original tuple(s). Such services are in general only lookup services.
- ⇒ it should be part of the service description, to indicate whether the resulting variable bindings are a superset of the input/used variables (cf. Section 4.1.5 and Section 9.5).

4.4 Operational Aspects of Rule Execution

The previous sections presented logical variables used for the declarative semantics of ECA rules, together with explicit markup for representation and communication of variable bindings

We did not yet discuss the component languages and their engines, or the domain languages that provide the atomic notions. So far, it is sufficient to be able to express them in an illustrative way by opaque expressions, and assume the actual component evaluation to be done by “demons” that understand an agreed format for downward communication and that return answers in an agreed format for upward communication. Component services will be discussed later in Chapter 5. We first clarify the canonic operational semantics of the ECA engine.

The operational semantics of ECA rules differs then from that of logical rules in that the rule body is not just a query, but consists of a triggering event, and then of the evaluation of queries.

Example 6 Consider again the situation from Example 3: For each “canceled flight” event, the rule is “fired”. “Fired” means that the event component produces one “answer”. The next “canceled flight” message fires another rule instance that will be completely independent.

For that “answer” to the event component, the (one and only) destination is selected, bound to the *Destination* variable, and for this pair (*Flight*, *Destination*), the action is triggered.

A different (and more complex) situation occurs, if the query component produces several answers, e.g. when a query selects all customers who have a reservation for this flight and sends each of them an e-mail.

4.4.1 Firing ECA Rules: the Event Component

An event is something that occurs (or, that is detected – in contrast to local databases that represent a closed world, the occurrence of an event somewhere in the Web does not necessarily mean that it is actually detected anywhere where it is relevant). Formally, detection of an event results in an occurrence indication, together with information that has been collected (consider here the data exchange format discussed in Section 4.1.4). The ECA engine must then execute the rule accordingly, using the obtained variable bindings. Note that the cardinality of answers must be considered in this case:

Example 7 (Exam Registration) Consider the following scenario: for an exam, first the (online) registration is opened, then students register, and at a given timepoint, the registration closes. Assume the events to be marked up as e.g.

```
<uni:reg_open subject="Databases" />
<uni:register subject="Databases" name="John Doe" />
<uni:reg_close subject="Databases" />
```

A can e.g. describe an action to be taken “if registration for an exam E is opened, students X_1, \dots, X_n register, and registration for E closes”. The composite event is then formulated as “registration to an exam is closed after (it had been opened and) students s_1, \dots, s_n registered” and is reported at the timepoint when the registration closes. With its occurrence indication, information about the subject E and registered students X_1, \dots, X_n is given. Such a cumulative semantics is provided by appropriate event operators, e.g. by SNOOP [24].

The rule must then be fired for this one event.

Example 8 Consider the following rule: “If flight *F* is delayed for more than 25 minutes, do ...”. Information about delayed flights is available all 10 minutes as a message including a report of the form

```
<msg:receive-message sender="service@raport.com" >
  <msg:content xmlns:travel="http://www.semwebtech.org/domains/2006/travel#" >
    <travel:delayed-flight flight="LH1234" minutes="30" />
    <travel:delayed-flight flight="AF0815" minutes="90" />
    <travel:delayed-flight flight="CY42" minutes="60" />
    :
    <travel:canceled-flight flight="AL4711" />
    :
  </msg:content>
</msg:receive-message>
```

There are two ways how to maintain this rule:

- XML level: the rule designer knows that such messages come in, and formulates a suitable atomic event pattern (cf. Section 5.2) as

```
<eca:Rule>
  <eca:Event xmlns:travel="http://www.semwebtech.org/domains/2006/travel#" xmlns:mail="..." >
    <eca:atomic-event>
      <mail:receive-message sender="service@raport.com" >
        <mail:content>
          <travel:delayed flight="$flight" minutes= "$minutes" />
        </mail:content>
      </mail:receive-message sender="service@raport.com" >
    </eca:atomic-event>
  </eca:Event>
  :
</eca:Rule>
```

which binds variables *flight* and *minutes*.

- Formulating the rule purely in the application domain as

```
<eca:Rule>
  <eca:Event xmlns:travel="http://www.semwebtech.org/domains/2006/travel#" >
    <eca:atomic-event>
      <travel:delayed flight="$flight" minutes= "$minutes" />
    </eca:atomic-event>
  </eca:Event>
</eca:Rule>
```

and using an ECE (see Section 3.1.5) event derivation rule “if there is a message whose content matches event *E*, then consider *E* to be detected”.

Note that use of `<eca:atomic-event>` is optional, and everything inside is here a non-normative example. Details of embedding atomic events into surrounding languages are discussed in Section 5.2.5.

The answer from the event detection module can result in

- one answer, containing several tuples, or

- several answers, containing one tuple.

Note that later, the same event pattern will be detected again, which again fires the rule.

In case that an occurrence indication contains *multiple* tuples of variable bindings, the semantics must be carefully considered: The tuples must be regarded as semantically independent since they –although “just by chance” detected at the same time– represent independent events. For that reason, a correct (but not always most efficient) semantics would be that the ECA engine immediately separates them and fires independent instances of the rule (see Section 4.4.4).

Thus, if a final event should report about something set-like, this must be contained in the semantics of the event language, not of the ECA language – then it is practically too late. As stated above, an explicitly cumulative semantics is e.g. supported by the SNOOP event algebra.

In many approaches, the “result” of event detection is the sequence of the events that “materialized” the event pattern to be detected. In this case, an appropriate way is to bind this result to a variable as shown above and afterwards the values of other variables can be extracted from this one.

Example 9 Consider the following situation from Example 7 and an event specification (using XPathLog and regular expression syntax in an obvious way):

```
<eca:Rule ... >
  <eca:variable name="Subj" />
  <eca:variable name="regseq" >
    <eca:Event>
      <eca:opaque language='xpathlog-events'>
        <eca:bind-variable name="Subj" />
        reg_open[@subject→Subj], register[@subject→Subj]*, reg_close[@subject→Subj]
      </eca:opaque>
    </eca:Event>
  </eca:variable>
  :
</eca:Rule>
```

The returned information from the event detection service in the markup proposed in Section 4.2.3 looks as follows, returning the relevant event sequence. The variable *Subj* has also been bound in the event component:

```
<logvars:answer component="event" ref="identifier" >
  <logvars:result>
    <uni:reg_open subject="Databases" />
    <uni:register subject="Databases" name="John Doe" />
    <uni:register subject="Databases" name="Scott Tiger" />
    :
    <uni:reg_close subject="Databases" />
  </logvars:result>
  <logvars:variable-bindings>
    <logvars:tuple>
      <logvars:variable name="Subj" >Databases</logvars:variable>
    </logvars:tuple>
  </logvars:variable-bindings>
</logvars:answer>
```

Next, the variable *regseq* is bound to the *<result>* part. The variable bindings after completely evaluating the event component look as follows:

```

<logvars:variable-bindings>
  <logvars:tuple>
    <logvars:variable name="regseq" >
      <uni:reg_open subject="Databases" />
      <uni:register subject="Databases" name="John Doe" />
      <uni:register subject="Databases" name="Scott Tiger" />
      :
      <uni:reg_close subject="Databases" />
    </logvars:variable>
    <logvars:variable name="Subj" >Databases</logvars:variable>
  </logvars:tuple>
</logvars:variable-bindings>

```

4.4.2 The Query Component

This second component is concerned with *static* information that is obtained and restructured from two areas:

- analyzing the data that has been collected by the event component (in the variable bindings), and
- based on this data, stating queries against databases and the Web.

Whereas the event component of a rule may be “answered” by detecting occurrences of the event pattern several times, the query component returns all answers at the same time. The query component is very similar to the evaluation of database queries and rule bodies in Logic Programming: in general, it results in a set of tuples of variable bindings (that are possible answers to a query).

Grouping: Set-Valued vs. Multi-Valued. An important issue here is to deal with sets (e.g., in the above examples, all customers who booked a flight that has been canceled, or all students that registered for an exam):

- bind a variable to a collection, e.g.,
 $\beta = \{\text{Subj} \rightarrow \text{'Databases'}, \text{student} \rightarrow \{\text{'John Doe'}, \text{'Scott Tiger'}, \dots\}\}$, or
- produce separate tuples of variable bindings:
 $\beta_1 = \{\text{Subj} \rightarrow \text{'Databases'}, \text{student} \rightarrow \text{'John Doe'}\}$,
 $\beta_2 = \{\text{Subj} \rightarrow \text{'Databases'}, \text{student} \rightarrow \text{'Scott Tiger'}\}$.

We follow again the Logic Programming specification that every answer produces a variable binding. For variable binding by matching (as in Datalog, F-Logic, XPathLog, Xcerpt etc.), this is obvious. Since we also allow variable bindings in the functional XSLT style, the semantics is adapted accordingly:

- each answer node of an XPath expression yields a variable binding;
- each node that is *returned* by an XQuery query yields a variable binding; if the XQuery query is of the form
 $\langle \text{name} \rangle \{ \text{for } \dots \text{ where } \dots \text{ return } \dots \} \langle / \text{name} \rangle$,
then the whole result yields a single variable binding.

Example 10 Consider again Example 9 where the resulting event contained several registrations of students. For doing anything useful, their names have to be extracted.

1. as multiple string-valued variables:


```

<eca:Rule ... >
  :
  same as above, binding variables "Subj" and "regseq"
  :
  <eca:variable name="Student" >
    <eca:Query>
      <eca:opaque language='xpath'>
        $regseq//uni:register[@subject=$Subj]/@name/string()
      </eca:opaque>
    </eca:Query>
  </eca:variable>
  :
</eca:Rule>

```

The above query generates the extended variable bindings
 $\beta_1 = \{ Subj \rightarrow 'Databases', regseq \rightarrow (as\ above), Student \rightarrow 'John\ Doe' \}$,
 $\beta_2 = \{ Subj \rightarrow 'Databases', regseq \rightarrow (as\ above), Student \rightarrow 'Scott\ Tiger' \}$.

2. or as a single variable:

```

<eca:Rule ... >
  :
  same as above, binding variables "Subj" and "regseq"
  :
  <eca:variable name="Students" >
    <eca:Query>
      <eca:opaque language='xquery'>
        <students>
          for $s in $regseq//uni:register[@subject=$Subj]/@name/string()
          return <name> { $s } </name>
        </students>
      </eca:opaque>
    </eca:Query>
  </eca:variable>
  :
</eca:Rule>

```

This query generates the extended variable binding
 $\beta = \{ Subj \rightarrow 'Databases', regseq \rightarrow (as\ above), Students \rightarrow \langle students \rangle \langle name \rangle John\ Doe \langle /name \rangle \langle name \rangle Scott\ Tiger \langle /name \rangle \langle /students \rangle \}$

The above query showed how data from the variable bindings obtained from the event detection is extracted. Note that this query is very similar to the *event queries* mentioned in XChange (instead of the above opaque query, also an Xcerpt/XChange query could have been used).

The next example shows a query against an XML repository in the Web. Again, the answer of the query component is represented in a single XML variable. This example then leads immediately to questions about handling the action component.

Example 11 Consider an ECA rule with opaque components (using different languages) that, whenever a flight is canceled, notifies every customer who has a reservation for this flight (e.g., by SMS), and sends a message to the airport hotel with the names of all customers to make a pre-reservation for this night.

```

<eca:Rule xmlns:eca="http://www.semwebtech.org/languages/2006/eca-ml#"

```

```

xmlns:xpath="http://www.w3.org/XPath"
<eca:variable name="Bookings" > http://localhost/schedule.xml </eca:variable>
<eca:variable name="Flight" />
<eca:Event>
  <eca:opaque language='datalog-match'>
    <eca:bind-variable name="Flight" />
    flight_cancellation(Flight) <!-- matches Flight against received message -->
  </eca:opaque>
</eca:Event>
<eca:variable name="Customers" >
  <eca:Query>
    <eca:opaque language="xquery">
      <eca:use-variable name="Flight" />
      <eca:use-variable name="Bookings" />
      return
        <customers>
          { for $c in document($Bookings)//flight[@id=$Flight]/reservation/customer
            return $c }
        </customers>
    </eca:opaque>
  </eca:Query>
</eca:variable>
<!-- evaluates XPath expression and binds each of the results to the variable 'Customers' -->
<eca:Test>
  <eca:opaque language='xpath'>
    $Customers/customer
  </eca:opaque>
</eca:Test>
<eca:Action>
  <eca:opaque language='pseudocode'>
    <eca:use-variable name="Customers" />
    <eca:use-variable name="Flight" />
    send one message with all Customers/customer/name to the hotel,
    then
    for each N in Customers/customer/@phonenr do
      notify_cancellation(Flight, sms:N)
  </eca:opaque>
</eca:Action>
</eca:Rule>

```

In the first case, the action has a multi-valued semantics, whereas in the second case it has a set-valued semantics. This simple case could be solved by binding the *set* of customers to a variable. The above “solution” moves the solution inside the query component, and is sufficient in this case, but not in general (and less declarative).

Note that there can be query services that return the query answer stepwise (which makes sense if the receiver is then able to continue already for some of the answers). This should be considered in the service descriptions (of the ECA service and of the query service) and in the message answer (default: complete).

4.4.3 The Test Component

The test component is still concerned with the information obtained so far. It evaluates a condition which is a mapping from a knowledge base to true/false. In general, the evaluation of conditions is based on a logic over literals with boolean combinators and quantifiers. A Markup Language exists with FOL-RuleML [15]. Since first-order logic is in general undecidable, it is recommended to use suitable fragments. Instead of first-order atoms, also “atoms” of other data models can be used. Additionally, we envisage to allow to use simple expressions like XPath of a language that

is locally supported in the ECA engine. Note that XPath expressions are also literals that result in a true/false (true if the result set is non-empty) value (as in the above example). Variables are communicated to the test in the same way as above. The test component returns then the set of tuples that satisfy the condition (for further propagation to the action component).

4.4.4 Summary of Event, Query and Test Semantics

Given the variable bindings resulting from the event component, the semantics of evaluating queries is based on the usual answer & join semantics of Logic Programming: for the cases where queries only contain positive occurrences of variables, the resulting variable bindings of the event and query components are just the join of the individual sets of variable bindings. As long as only positive queries are used, the semantics of the query component is commutative. In case of negative occurrences, the usual safety constraints apply, variables must be bound positively before they can be used in a negated occurrence. Then, negation is interpreted as set difference. The test component acts as a selection that removes all variable bindings that do not satisfy the test.

The normal form induces a sequential operational semantics; other evaluation and execution strategies are possible based on equivalence transformations .

In each stage, the variable bindings are considered as a set of tuples that is represented in XML as above, and communicated with the event, query and test services. Note that this allows for aggregation and grouping constructs in queries and tests.

Multiple Occurrences. The “plain” semantics of an ECA rule assumes that a rule is fired for each individual occurrence of the event pattern given in its event component. As shown above, that multiple independent instances of an event pattern are detected at the same time, and reported as a whole. In such cases, it is possible to fire the rule for all these occurrences together (since we do not bind program variables, but use the notion of logical variables and sets of tuples of variable bindings, the underlying declarative semantics allows for this). Semantically, there is a only a difference if transactional issues in the action component are considered. In this case it is enough to separate the instances when executing the action component.

4.4.5 The Action Component

The action component is the one where *actually* something is done in the ECA rule: for each variable binding, the action component is executed. Actions do not have a result. Thus, there is only a set of input variables (usually called parameters in operational semantics) that is submitted in the same way as above.

The action component may consist of several `<eca:Action>` elements which can use different action languages. The semantics is that all actions are executed. Note that actions are not allowed to bind variables, thus they are independent on this level. (Note that sequential and conjunctive execution of actions can also be specified on the level of action languages *inside* the `<eca:Action>` element.)

Grouping and de-grouping execution steps. Here, the possibility of grouping and de-grouping is required: it makes a strong difference if a set is represented by one tuple of variables, or the whole set is bound as a set in one tuple.

Example 12 Consider again Example 9 where the names of students that have registered for an exam have been collected:

1. In the first case, we had the following variable bindings:

$$\beta_1 = \{ \text{Subj} \rightarrow \text{'Databases'}, \text{regseq} \rightarrow (\text{as above}), \text{Student} \rightarrow \text{'John Doe'} \} ,$$

$$\beta_2 = \{ \text{Subj} \rightarrow \text{'Databases'}, \text{regseq} \rightarrow (\text{as above}), \text{Student} \rightarrow \text{'Scott Tiger'} \} ,$$
 i.e., a set of two tuples.

If the action now is “(for each tuple), send the lecturer a mail with the value of the variable *Student*”, the lecturer will get two mails, each one with one student.

Instead, we want to send the lecturer one mail with the names of all registered students.

- For such cases, a functionality for “group by Subject” (which results in only one tuple) would be useful, e.g. a way to group by *Subj*, collecting all names in a list (and forgetting about *regseq*, resulting in

$$\beta_{grp} = \{ \text{Subj} \rightarrow \text{'Databases'}, \text{Student} \rightarrow \{ \text{'John Doe'}, \text{'Scott Tiger'} \} \},$$

for which the action then can be called.

- In the second case, we have only one variable binding (very similar to β_{grp} above,

$$\beta = \{ \{ \text{Subj} \rightarrow \text{'Databases'}, \text{regseq} \rightarrow (\text{as above}), \text{Students} \rightarrow \langle \text{students} \rangle \langle \text{name} \rangle \text{John Doe} \langle \text{/name} \rangle \langle \text{name} \rangle \text{Scott Tiger} \langle \text{/name} \rangle \langle \text{/students} \rangle \}$$

for which the action can be called immediately (and e.g. submit the XML structure of *Students* as the message content).

On the other hand, if the task is “to send a message to each of the registered students”, the first case is immediately suitable, whereas the second one needs either an iteration inside the action component (“for each name in *Students* do ...”) or, – more declaratively on the ECA level – an ungrouping, resulting in the variable bindings of β_1 and β_2 .

Such cases are very frequent. A rule can even contain actions of both kinds, e.g., in

For the above E&Q components, “send the lecturer an e-mail with the names of all students *and* send a confirmation message to each of the registered students”.

In such cases, the number of variable bindings must be changed by grouping or de-grouping, dependent on what the result of the E&Q components are. Grouping and ungrouping on the ECA level is allowed *before* each `<eca:Action>` element. We propose to inherit these elements also to the action languages:

- all tuples that coincide in *all* named variables are grouped together; the other variables are aggregated as lists, sum, avg etc., or omitted (a default can be specified; *aggr-op* = list|sum|avg|omit|...):

```
<eca:group-by aggr="aggr-op" >
  <eca:group-variable name="name" />
  :
  <eca:group-variable name="name" />
  <eca:aggr-variable op="aggr-op" name="name" aggr-name="name" />
  :
  <eca:aggr-variable op="aggr-op" name="name" aggr-name="name" />
</eca:group-by>
```

- Flattening a list or sequence: The new variable binding is obtained by splitting a list, or applying a query to the variable (XPath, or any default language configured by the ECA service or the rule). Note that this can also be done by a sequence of `<eca:Query>` elements.

```
<eca:ungroup-by>
  <eca:ungroup-variable>variable-name
  [<eca:ungroup-query> ...</eca:ungroup-query>]
</eca:ungroup-variable>
:
  <eca:ungroup-variable>variable-name
  [<eca:ungroup-query> ...</eca:ungroup-query>]
</eca:ungroup-variable>
</eca:ungroup-by>
```

4.4.6 Transactions

Although the issue of transactions does not directly have to do with the semantics of ECA rules, some issues should be raised here. Transactional issues are only concerned with the action component (events can neither fail nor be rolled back, queries and tests can also not “fail” and there is nothing to be rolled back). Transactional functionality can be offered independently by the action languages *inside* the `<eca:Action>` elements.

Since we stated above that the semantics for execution of the action component is the same as for executing the head of a deductive rule, i.e., handling it separately for each tuple of variable bindings, transactions that should cover the whole group must explicitly be expressed on the ECA level.

For this, we propose an

```
<eca:transaction attributes> ... </eca:transaction>
```

element that can occur around or anywhere in an `<eca:Action>` element. With this, e.g., actions for *all* tuples can be grouped as a transaction by

```
<eca:Rule ...>
  <eca:Event> ...</eca:Event>
  <eca:Query> ...</eca:Query>
  <eca:Test> ... </eca:Test>
  <eca:transaction>
    <eca:Action> ... </eca:Action>
  </eca:transaction>
</eca:Rule>
```

A further element allows to take a group of tuples together for execution of a transaction (note that in contrast to `<eca:group-by>`, the number of tuples does not change):

```
<eca:transaction-group-by/>
  <eca:group-variable name="name" />
  :
  <eca:group-variable name="name" />
</eca:transaction-group-by>
```

(e.g., if an event contains a list of delayed flights, the query component returns a pair of variable bindings for each customer, and the appropriate actions should be separate transactions for each flight).

4.4.7 Examples

Example 13 (Rental Cars (Revisited)) *Execution is then as follows: after the event part, the variable bindings look as follows, containing only one tuple:*

```
<logvars:variable-bindings>
  <logvars:tuple>
    <logvars:variable name="Person">John Doe</logvars:variable>
    <logvars:variable name="To">Paris</logvars:variable>
  </logvars:tuple>
</logvars:variable-bindings>
```

The first query for looking up what car(s) the person owns declares `Person` and `OwnCar` to be used variables. Since up to now, only `Person` is bound, `owns(Person,OwnCar){(Person/"John Doe")}` is evaluated by the local Datalog database. Consider now the case that John Doe owns two cars. It results in the bindings

```

<logvars:variable-bindings>
  <logvars:tuple>
    <logvars:variable name=" Person" >John Doe</logvars:variable>
    <logvars:variable name=" OwnCar" >Golf</logvars:variable>
  </logvars:tuple>
  <logvars:tuple>
    <logvars:variable name=" Person" >John Doe</logvars:variable>
    <logvars:variable name=" OwnCar" >Passat</logvars:variable>
  </logvars:tuple>
</logvars:variable-bindings>

```

These bindings are joined with the first ones, resulting in

```

<logvars:variable-bindings>
  <logvars:tuple>
    <logvars:variable name=" Person" >John Doe</logvars:variable>
    <logvars:variable name=" To" >Paris</logvars:variable>
    <logvars:variable name=" OwnCar" >Golf</logvars:variable>
  </logvars:tuple>
  <logvars:tuple>
    <logvars:variable name=" Person" >John Doe</logvars:variable>
    <logvars:variable name=" To" >Paris</logvars:variable>
    <logvars:variable name=" OwnCar" >Passat</logvars:variable>
  </logvars:tuple>
</logvars:variable-bindings>

```

The next query against the local Datalog database for obtaining the class of the respective car(s) declares OwnCar and Class to be used. Up to now, only OwnCar is bound, thus variable, i.e., the query `class(OwnCar,Class) [{(OwnCar="Golf"),(OwnCar="Passat")}]` is evaluated against the local Datalog database. It results in the bindings

```

<logvars:variable-bindings>
  <logvars:tuple>
    <logvars:variable name=" OwnCar" >Golf</logvars:variable>
    <logvars:variable name=" Class" >B</logvars:variable>
  </logvars:tuple>
  <logvars:tuple>
    <logvars:variable name=" OwnCar" >Passat</logvars:variable>
    <logvars:variable name=" Class" >C</logvars:variable>
  </logvars:tuple>
</logvars:variable-bindings>

```

which is again joined with the existing ones:

```

<logvars:variable-bindings>
  <logvars:tuple>
    <logvars:variable name=" Person" >John Doe</logvars:variable>
    <logvars:variable name=" To" >Paris</logvars:variable>
    <logvars:variable name=" OwnCar" >Golf</logvars:variable>
    <logvars:variable name=" Class" >B</logvars:variable>
  </logvars:tuple>
  <logvars:tuple>
    <logvars:variable name=" Person" >John Doe</logvars:variable>
    <logvars:variable name=" To" >Paris</logvars:variable>
    <logvars:variable name=" OwnCar" >Passat</logvars:variable>
  </logvars:tuple>

```

```

    <logvars:variable name="Class">C</logvars:variable>
  </logvars:tuple>
</logvars:variable-bindings>

```

The next query for looking up the available cars at the destination declares *To* to be an input variable, i.e., the query `http://localhost/lookup-cars?.place='Paris'` is evaluated, resulting in the bindings

```

<logvars:variable-bindings>
  <logvars:tuple>
    <logvars:variable name="Class">B</logvars:variable>
    <logvars:variable name="Model">C4</logvars:variable>
    <logvars:variable name="Price">50</logvars:variable>
  </logvars:tuple>
  <logvars:tuple>
    <logvars:variable name="Class">B</logvars:variable>
    <logvars:variable name="Model">Golf</logvars:variable>
    <logvars:variable name="Price">65</logvars:variable>
  </logvars:tuple>
  <logvars:tuple>
    <logvars:variable name="Class">D</logvars:variable>
    <logvars:variable name="Model">C6</logvars:variable>
    <logvars:variable name="Price">150</logvars:variable>
  </logvars:tuple>
</logvars:variable-bindings>

```

Note that the ECA engine must keep the knowledge that all these tuples refer to *To*/*"Paris"*. Joining the result removes the tuple dealing with the class *"C"* since no such cars are available in Paris and results in

```

<logvars:variable-bindings>
  <logvars:tuple>
    <logvars:variable name="Person">John Doe</logvars:variable>
    <logvars:variable name="To">Paris</logvars:variable>
    <logvars:variable name="OwnCar">Golf</logvars:variable>
    <logvars:variable name="Class">B</logvars:variable>
    <logvars:variable name="Model">C4</logvars:variable>
    <logvars:variable name="Price">50</logvars:variable>
  </logvars:tuple>
  <logvars:tuple>
    <logvars:variable name="Person">John Doe</logvars:variable>
    <logvars:variable name="To">Paris</logvars:variable>
    <logvars:variable name="OwnCar">Golf</logvars:variable>
    <logvars:variable name="Class">B</logvars:variable>
    <logvars:variable name="Model">Golf</logvars:variable>
    <logvars:variable name="Price">65</logvars:variable>
  </logvars:tuple>
</logvars:variable-bindings>

```

The final action declares *Model* and *Price* as input. Thus, only the bindings

```

<logvars:variable-bindings>
  <logvars:tuple>
    <logvars:variable name="Model">C4</logvars:variable>
    <logvars:variable name="Price">50</logvars:variable>
  </logvars:tuple>

```

```

</logvars:tuple>
<logvars:tuple>
  <logvars:variable name=" Model" >Golf</logvars:variable>
  <logvars:variable name=" Price" >65</logvars:variable>
</logvars:tuple>
</logvars:variable-bindings>

```

are submitted to it.

Example 14 Consider an ECA rule with opaque components (using different languages) that, whenever a flight is canceled, sends a message to the destination airport that the flight will not take place:

```

<eca:Rule xmlns:eca="http://www.semwebtech.org/languages/2006/eca-ml#"
  xmlns:xpath="http://www.w3.org/XPath"
  <eca:variable name=" Schedule" >http://localhost/schedule.xml</eca:variable>
  <eca:variable name=" Flight" />
  <eca:variable name=" Destination" />
  <eca:variable name=" event" >
    <eca:Event>
      <eca:atomic-event>
        <travel:canceled-flight/>
        <!-- matches any travel:canceled-flight event, e.g. <travel:canceled-flight code="LH1234" /> -->
      </eca:atomic-event>
    </eca:Event>
  </eca:variable>
  <!-- the matched event is now bound to the variable $event -->
  <eca:variable name=" Flight" language="xpath" select="$event/canceled-flight/string(@code)" />
  <eca:variable name=" Destination" >
    <eca:Query>
      <eca:opaque language='xpath'>
        <eca:input-variable name=" Flight" use="$Flight" />
        <eca:input-variable name=" Schedule" use="$Schedule" />
        string(document($Schedule)//flight[@id=$Flight]/@to)
      </eca:opaque>
    </eca:Query>
  </eca:variable>
  <!-- evaluates XPath expression and binds the result to the variable 'Destination' -->
  <eca:Test>
    <eca:opaque language='boolean'>true</eca:opaque>
  </eca:Test>
  <eca:Action>
    <eca:opaque language='pseudocode'>
      <eca:input-variable name=" Flight" use="$Flight" />
      <eca:input-variable name=" Destination" />
      send "Flight $Flight has been canceled today" to the destination airport ...
    </eca:opaque>
  </eca:Action>
</eca:Rule>

```

The ECA engine proceeds as follows: It binds the variable *Schedule* as a constant to the given value and allocates variables *Flight* and *Destination*. The event component consists only of an atomic event. If such an event, e.g. `<travel:canceled-flight code="LH1234"/>` is detected (by matching), it is bound to the variable *event* (this semantics coincides with most event detection semantics that return the relevant event sequence as the result). In the next step, the variable *Flight* is bound by evaluating an XPath expression against the value of *event*, yielding the binding *Flight*/"LH1234".

Next, the ECA engine submits the query where (*\$Schedule* is replaced with the constant URL, and *\$Flight* replaced with the actual binding "LH1234") to the XPath engine, that evaluates the

*expression and returns its result, i.e., the identifier of the destination airport (e.g., "FRA"). The ECA engine binds the returned result to its variable **Destination**. The condition is then empty (every flight has a destination). Next, the pseudocode fragment in the action component is equipped with the flight number and the destination airport and a message is sent.*

Chapter 5

Abstract Semantics and Communication: Component Services

Component languages/services and domain languages/services can be provided by everybody somewhere in the (Semantic) Web. We do not propose a central registry of languages. A language becomes known to an ECA engine by being referenced in some rule (with its namespace URI). The algebraic components provide the glue between atomic notions and the ECA level. Usually, the “result” of an algebra expression is of the same type as the elements of its (formal) domain. Thus, different algebras of the same type can usually be nested (in some sense forming a bigger algebra that provides the union of the operators).

Communication issues from the point of view of the *rule semantics*, i.e., transmission of variable bindings has already been discussed in Chapter 4; especially the formats for downward and upward communication in Section 4.2. In this chapter, this is extended from the point of view of the component services.

We first discuss the general communication patterns. Then, we continue in detail with the structure, syntax, and semantics of the event component and its communication issues. The abstract semantics and markup of the query and test components are simpler since they are in general an embedding of query languages in the MARS Framework. The discussion of the action component completes the discussion of the components. A short summary concludes the chapter.

5.1 General Communication Patterns

For the communication between services, messages (in XML) have to be exchanged. Several communication patterns are used:

- request-response in asynchronous and synchronous way,
- some requests do not have a direct response (e.g., submitting composite event specification) but later on, one or more “answers” referring to the request will be sent,
- answers always refer to some request,
- additionally, events are communicated that do not refer to a request, but are (implicitly) “answers” on a registration.

Service URLs. The MARS Framework is flexible wrt. the actual URLs where the services expect the requests. A service can e.g. receive requests for all tasks at a single URL, or provide

a separate method URL for each method. This information is managed by *Language and Service Registries (LSRs)*, see Section 9.5.1.

Communication Protocols. The communication can be implemented by HTTP or SOAP. The prototype uses the HTTP POST method. Nevertheless, it is recommended to plan to allow for both.

Message Contents.

- component expressions (i.e., composite events, atomic event specifications, queries, actions): in plain XML representation as elements in the respective namespaces (note: without the `eca:Event` etc. elements around them – they must only contain markup of the service’s namespace),
- answers to requests: as `<logvars:answers>` or `<logvars:answer>` XML elements.

The design of messages follows the ideas of the simple SMTP protocol. Any message that awaits an answer contains the following components:

MESSAGES/TASKS THAT WILL BE ANSWERED:	
sender	by URL of the sending process
Reply-To	where the answer should go (URL)
Subject	an identification that allows for uniquely identifying the answer (e.g., this can be the URI of an event component which is submitted for detection - the URI will then be used in the answer).
contents	(application-dependent)

Tasks given as XML Fragments. A request uses the same parameters as a request in real life, or an e-mail. Most tasks are specified by XML fragments (registering a rule, registering a composite event, registering an atomic event description, registering or answering a query, registering or executing an action). In these cases, the XML fragment is submitted to a suitable service. Optionally, variables that are already bound on the rule level (when handling *rule patterns*) can also be contained in the message. The message must contain the following information:

MESSAGES/TASKS ON XML FRAGMENTS:	
Sender, Reply-To	as above
Subject	URI of the component. The answer will refer to this subject (in case of asynchronous communication).
content(1)	the component or a reference to it
content(2)	optionally: variable bindings (cf. Section 4.2.2).

Answers from Evaluation. Answers are expressed as `<logvars:answers>` or `<logvars:answer>` XML elements. They go to the url that has been specified as `Reply-To` by the requester. The message must contain the following information:

MESSAGES/TASKS ON XML FRAGMENTS:	
Subject	the URI sent as <code>Subject</code> by the requester for identifying the answer.
Sender, Reply-To	empty, optional.
content(1)	in the <code><logvars:answers></code> format described in Section 4.2.3
content(2)	optional “comments” like “answer is complete” or “more answers follow”, or error messages.

Communication of Events. As already discussed, events are plain XML fragments:

- Reply-To and Subject: none
- contents/task: event in XML markup.

Actual Markup. The Sender is always submitted in the HTTP header. The Reply-To and Subject can alternatively be submitted in the HTTP header (according to the convention that private/non-standard properties start with “X-”, they are called X-Reply-To and X-Subject then). If no variable bindings are submitted. Optionally, all information sent in the body can be wrapped into an XML element hull to be not only a sequence of elements, but a single element node (Such things have to be specified in the Service Descriptions; see Section 9.5.) Additional comments can also be added to answers, e.g., for incomplete query answers in a stream-like strategy.

First examples will be given in Section 5.3.1.

5.2 The Event Component: Structure and Languages

The semantics of the event component and its services can be distinguished into two levels. The event component language is embedded in the language hierarchy as shown in Figure 3.9. It is based on the event ontology as discussed in Section 3.1.4. Composite events consist of certain combination of atomic events. Thus, the event component of a rule, which is a *specification* of a composite event (using a *Composite Event Language (CEL)*), consists of the specification of the combination, and of specifications of the contributing atomic events. Thus, languages of two types are needed:

Event Algebras: They are used to define composite events, including their specification as algebra terms as shown in Figure 3.11 and their detection algorithms. Since the usual semantics of evaluating an event algebra expression is to return the matching event sequence, subexpressions from different algebras can be combined easily.

Atomic Event Description Formalisms: They are used to specify which atomic events are considered relevant and how these are matched (at runtime), and what is the result of detecting them. They have to allow to deal with events according to the ontology of atomic events given in Figure 3.3

Thus, the event component uses a combination of one or more event algebras, using atomic events of one or more applications, and possibly atomic data-level events from several data models, and atomic events from application-independent services.

5.2.1 Atomic Events

Events occur as atomic events on the Web. For atomic events, two issues have to be considered:

- the atomic event as an occurrence on the Web (which is communicated somehow in some representation), and
- specification of atomic events to react upon in the event components of ECA rules.

5.2.1.1 XML Representation of Atomic Events

We assume that events are available as XML or RDF fragments.

Example 15 (Atomic Events) *Events are data fragments that are available in XML markup or as RDF fragments.*

- *The event*

```
<travel:canceled-flight flight="LH1234">
  <travel:reason>bad weather</travel:reason>
</travel:canceled-flight>
```

is an event in the traveling domain that means that the flight “LH1234” is canceled and also the reason is given. Note that more information (e.g., that this concerns today’s flight that should depart in one hour) must then be accessible from the context. Complete information would be available if the event is of the following form `<travel:canceled-flight flight=“LH1234” date=“10062005” />`.

An example for a delayed flight is

```
<travel:delayed-flight flight=“LH1234” minutes=“30” />
```

- The following events have already been used above in an example:

```
<uni:reg_open subject=“Databases” />
```

```
<uni:register subject=“Databases” name=“John Doe” />
```

- in RDF, events are resources of a type “event” that also have a name and are connected to other resources as parameters (cf. Chapter 13).

5.2.2 Atomic Event Descriptions and Formalisms

The atomic event specifications (AESs) form the leaves of the event component tree. In general, the AES specifies which occurring events are relevant to take a reaction, e.g., the name of the event, or also its contents (i.e., its parameters). Thus, there are always *two* languages involved as shown in Figure 5.1:

- a domain language (associated with the namespace of the event),
- and an atomic event description/matching/query language/formalism for *describing* what events should actually be matched. Since the event is seen as an XML or RDF fragment, these conditions can be stated as patterns, or as queries against this fragment.

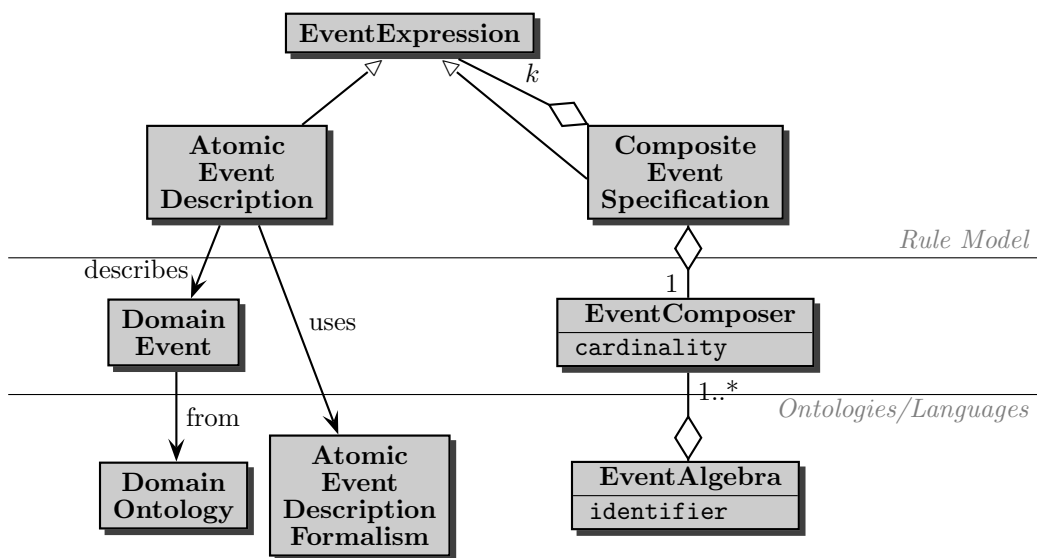


Figure 5.1: Event Expressions: Languages

Since the event is seen as an XML or RDF fragment, the AES formalism is actually a query language that states these conditions against an XML or RDF fragment. The result of the evaluation should be yes/no (as “the specified event occurred”), optionally (but recommended) the formalisms also should support for using and binding variables. The implementation of AES formalisms is provided by *Atomic Event Matcher (AEM)* services. An *Atomic Event Matcher* that

implements such a formalism should usually return the event as functional result (since usually expected by the semantics of the surrounding event algebra), together with the variable bindings. From the point of view of the abstract semantics, the AES must provide the following information:

- **Mandatory:** information about the URI of the service that is responsible for processing the AES must be accessible to the processor of the surrounding language. (Note that this is not available in the above example.)
- Information about the URI of the service(s)/domains that provide the actual events is required to be accessible for the service that then actually processes the AES. This information may be accessible for the surrounding service.
- The content must state (expressed by the AESL) the requirements on the events on which a reaction should be taken, e.g., the name of the event, or also its contents (i.e., its parameters). It is only required that the content is understandable for the service that then actually processes the AES, not for the surrounding service.

Some sample AES formalisms are discussed below. The actual embedding in the rule markup is discussed later in Section 5.2.5.

5.2.2.1 Event Specification by XML-QL Style Matching

Pattern-based (e.g. like XML-QL [25]) event specification continues the tree-like style of event algebras. In this case, the AES specifies the actual atomic event by a pattern and variables can be bound to fragments.

```
<travel:delayed-flight xmlns:travel="http://www.semwebtech.org/domains/2006/travel#"
  foo:language="xmlqlmatch" flight="{ $flight}" time="{ $minutes}" >
  $content
</travel:delayed-flight/>
```

specifies that an event is relevant if it is a `travel:delayed-flight` event. In case that `flight` is already bound, this acts as a (join) condition on the code of the canceled flight, otherwise `flight` is bound by the matching. `$content` is bound to the complete content of the element.

The simple XML-QL matching style does not allow for binding specific elements to variables (only the whole contents as above). As an extension, `variable` elements (in the namespace of the matching formalism!) or variable references of the form `$var-name` can be used inside the pattern to express that a fragment of the event is bound or must match a variable:

```
<travel:canceled-flight xmlns:travel="http://www.semwebtech.org/domains/2006/travel#"
  xmlns:aes-xmlql="http://www.semwebtech.org/languages/2006/aes-xmlql" flight="{ $flight}" >
  <aes-xmlql:variable name="reason" >
    <travel:reason/>
  </aes-xmlql:variable>
</travel:canceled-flight>
```

matches any `travel:canceled-flight` event concerning a given flight and binds the variable `reason` to the `travel:reason` element of the flight. If `flight` is already bound, it acts as a (join) condition on the code of the canceled flight.

Note that the above syntax is valid XML: variables as attribute values are enclosed into quotes; for distinguishing them from strings (e.g., an attribute `price="$30"`), they are put inside braces. A similar syntax has been implemented in [70].

5.2.2.2 Navigation-Based Event Specification

Another alternative uses XPath style matching. Consider that for the matching, the event itself (as an XML fragment) is available as *\$event*. Then,

```
<aed-navig:variable name="var-name" select="$event/relative-expr..." />
```

can be used to access data within the event, and elements of the form

```
<aed-navig:test condition="xpath-expr" />
```

can be used for tests (note that this is the same as *eca:Test* for the test component). Variables can also be addressed by *\$var-name* as in XQuery (for using them as join variable or for binding them to the matched value). This formalism can be designed with a surrounding *domain-ns:name* element, or with a *aed-navig:name* element. Here, the variants look as follows:

```
<travel:delayed-flight xmlns:travel="http://www.semwebtech.org/domains/2006/travel#"
  xmlns:aes-xpath="http://www.semwebtech.org/languages/2006/aes-xpath" >
  <aes-xpath:test condition="$event/@flight=$flight" />
  <aes-xpath:variable name="minutes" select="$event/@minutes" />
</travel:delayed-flight>
```

or

```
<aes-xpath:match xmlns:aes-xpath="http://www.semwebtech.org/languages/2006/aes-xpath"
  xmlns:travel="http://www.semwebtech.org/domains/2006/travel#" >
  <aes-xpath:test condition="$event/name()='travel:delayed-flight'" />
  <aes-xpath:variable name="flight" select="$event/@travel:flight" />
  <aes-xpath:variable name="minutes" select="$event/@travel:minutes" />
</aes-xpath:match>
```

5.2.2.3 Event Specification by Opaque XQuery

Going one step further regarding the event as a small XML document, XQuery can be used to check its properties. For every event that is known to the service, the query is evaluated.

```
<eca:opaque language="xquery-eventchecker" [name="name"]>
  xquery expression over $event
</eca:opaque>
```

- the check is e.g. expressed as *xquery expression over \$event*,

```
let $event := /
where $event/name()="travel:flight-canceled"
and $event/@flight="LH123"
return $event
```
- optional: indicate *event-element-name* to allow early filtering. The XQuery evaluation is only done if the name matches.
- The XQuery syntax can also be used in an experimental service to generate the upwards communication format with bound variables immediately.

Event Specification by Xcerpt. Another possibility is to use the language Xcerpt [23] developed in REVERSE WG I4. Xcerpt query terms can use and bind variables in a declarative way. For “standard” Xcerpt, the opaque mechanism has to be used. Xcerpt’s XML markup would be an interesting example of an XML embedding of a query language.

5.2.3 Event Algebras

Event algebras, well-known from the Active Databases area, serve for specifying *composite* events by defining *terms* formed by nested application of composers over *atomic* events. There are several proposals for event algebras, defining different composers. Each composer has a semantics that specifies what the composite event means.

For dealing with composite events in the context of the ECA rules proposed here, we propose at least the following composers: “ E_1 OR E_2 ”, “ E_1 AND E_2 ” (in arbitrary order), and “ E_1 AND THEN E_2 [AFTER PERIOD {< | >} *time*]” the latter one composing two events and using an additional parameter *time*, indicating the time that has passed between the occurrence of E_1 and E_2 . Detection of a composite event means that its “final” atomic subevent is detected:

- (1) $(E_1 \nabla E_2)(t) \quad :\Leftrightarrow \quad E_1(t) \vee E_2(t) ,$
- (2) $(E_1 \triangle E_2)(t) \quad :\Leftrightarrow \quad \exists t_1 \leq t \leq t_1 : (E_1(t_1) \wedge E_2(t)) \vee (E_2(t_1) \wedge E_1(t)).$
- (3) $(E_1;_{\Delta t} E_2)(t) \quad :\Leftrightarrow \quad \exists t_1 \leq t \leq t_1 + \Delta t : E_1(t_1) \wedge E_2(t).$

Event algebras contain not only the aforementioned straightforward basic connectives, but also additional operators. A bunch of event algebras have been defined that provide also e.g. “negated events” in the style that “when E_1 happened, and then E_3 but not E_2 in between”, “periodic” and “cumulative” events, e.g., in the SNOOP event algebra [24] of the “Sentinel” active database system.

Example 16 (Cumulative Event, [24]) A “cumulative aperiodic event”

$$A^*(E_1, E_2, E_3)(t) \quad :\Leftrightarrow \quad \exists t_1 \leq t : E_1(t_1) \wedge E_3(t)$$

occurs with E_3 and reports the collected occurrences of E_2 in the meantime. Thus, its detection is defined as “if E_1 occurs, then for each occurrence of an instance of E_2 , collect it, and when E_3 occurs, report all collected occurrences (in order to do something)”.

A cumulative periodic event can be used for “after the end of a month, send an account statement with all entries of this month”:

$$E(\text{Acct}) := \\ A^*(\text{first_of_month}(m), (\text{debit}(\text{Acct}, \text{Am}) \nabla \text{deposit}(\text{Acct}, \text{Am})), \text{first_of_month}(m + 1))$$

where the event occurs with `first_of_next_month`. The “result” of the expression is the list of all contributing events.

XML Markup for the Event Component. The `<eca:Event>` elements contain elements according to event algebra languages. Every subexpression is associated by its namespace with the appropriate components of the language – i.e., an event algebra or atomic expressions from underlying domains. In general, if an event algebra supports an XML markup, it will define its own ways for dealing with atomic events and variables.

Example 17 The cumulative event from Example 9 (there given as a regular expression) can be given in SNOOP as

$$A^*(\text{reg_open}(\text{Subj}), \text{register}(\text{Subj}, \text{Stud}), \text{reg_close}(\text{Subj})) .$$

The following markup binds the complete sequence to `regseq` and the subject to `Subj`:

```
<eca:Rule ... >
  <eca:variable name="Subj" />
  <eca:variable name="regseq" >
    <eca:Event xmlns:snoopy="http://www.semwebtech.org/languages/2006/snoopy#" xmlns:uni="..." >
      <snoopy:Cumulative>
        <!-- ignoring the identification of the AES formalism -->
```

```

    <uni:reg_open subject="{Subj}" />
    <uni:register subject="{Subj}" />
    <uni:reg_close subject="{Subj}" />
  </snoopy:Cumulative>
</eca:Event>
</eca:variable>
:
</eca:Rule>

```

(Note that the language could e.g. also provide redundant `<snoopy:Atomic>` elements for being more explicit and providing means for identifying the AES formalism; see Section 5.2.5.)

5.2.4 Embedding Algebraic Languages

The embedding of the event, query, test and action components is straightforwardly represented by a namespace change: outside, there is the `eca:` namespace, inside there is e.g. the SNOOP namespace, or an `<eca:opaque>` element that indicates the embedded language. In the same way, algebraic languages can be embedded into each other (e.g., embedding an expression of one event algebra as subexpression of another).

Example 18 (Embedding of Algebraic Languages) Consider two event algebras, e.g., SNOOP and rCML (RuleCore Markup Language) [13]. An rCML expression can be embedded in a SNOOP expression as follows:

```

<eca:Rule ... >
  <eca:Event xmlns:snoopy="http://www.semwebtech.org/languages/2006/snoopy#" >
    <snoopy:Cumulative>
      <snoopy:Or> ...</snoopy:Or>
      <rcml:times xmlns:rcml="..." > ... </rcml:times>
      <snoopy:Sequence> ...</snoopy:Sequence>
    </snoopy:Cumulative>
  </eca:Event>
  :
</eca:Rule>

```

From the point of view of the ECA engine, the event part resides in the `snoopy` namespace. Thus, when registering this rule, the ECA engine will register the event component at a SNOOP engine. The SNOOP engine prepares the usual algorithm, in this case, for detecting a cumulative event pattern. The `<snoopy:Or>` and `<snoopy:Sequence>` subexpressions are also mapped to the local algorithm. The `<rcml:times>` subevent cannot be detected locally since the semantics is unknown to the SNOOP engine. Instead, it is registered at an rCML engine. Whenever the latter detects this event, it sends an appropriate `<answers>` message as described in Section 4.2.3 to the SNOOP engine. This once more illustrates the use of generic communication schemata.

For embedding atomic events specifications, things are a bit more difficult since an atomic event specification involves a domain language/namespace and an AES formalism.

5.2.5 Embedding Atomic Events in Composite Events and Rules

For the surrounding language (which can be an event algebra or also the ECA-ML language in case of rules that react upon atomic events), an *atomic event specification (AES)* is a leaf element which does *not* belong to the surrounding namespace. The namespace border indicates the language border and initiates the handover between the “responsible” processors.

Since an atomic event specification involves a domain language/namespace and an AES formalism, one of these languages can be represented by the namespace of the root of the AEM XML

structure. In the above examples in Section 5.2.2, this was always the domain namespace, but taking e.g. an XQueryX or Xcerpt XML markup of the subtree, the outer expression would be in the formalism's namespace.

The actual design of the markup is up to the composite event specification language: its interpreter must know where to submit the leaf atomic event description. This can be done either if the AES's namespace identifies the AES formalism, or if an intermediate element contains the necessary information. Additionally, the AES interpreter must be able to detect the domain namespace from the AES.

Note that in an RDF setting, the assignment of languages to an AES is simply done by two triples

```
(aes, has_domain, domain)
(aes, uses_formalism, formalism).
```

Explicit Markup Borders. The surrounding (i.e., composite event) language can specify that the leaves are surrounded by e.g. an explicit `<atomic-event>` element:

```
<cel:atomic-event cel:language="atomic event specification language AESL"
                  cel:domain="domain of the event" >
  atomic event specification in formalism AESL
</cel:atomic-event>
```

Note that depending on the namespace of the inner element, either domain or language can be omitted.

Example 19 A specification of an event that reacts upon a delayed-flight event and extracts some data can be embedded as follows:

```
<eca:Rule>
  <eca:Event>
    :
    <cel:atomic-event
      cel:language="http://www.semwebtech.org/languages/2006/aes-xmlql"
      cel:domain="http://www.semwebtech.org/domains/2006/travel#" >
      <travel:delayed-flight xmlns:travel="http://www.semwebtech.org/domains/2006/travel#"
        flight="{flight}" time="{time}"/>
    </cel:atomic-event>
    :
  </eca:Event>
  :
</eca:Rule>
```

In the same way as the surrounding language can provide redundant explicit elements that carry the language information, the markup of the *embedded* (AESL) formalism can provide this, e.g. by

```
<eca:Rule>
  <eca:Event>
    :
    <aes-xmlql:atomic-event
      xmlns:aes-xmlql="http://www.semwebtech.org/languages/2006/aes-xmlql"
      aes-xmlql:domain="http://www.semwebtech.org/domains/2006/travel#" >
      <travel:delayed-flight xmlns:travel="http://www.semwebtech.org/domains/2006/travel#"
        flight="{flight}" time="{time}"/>
    </aes-xmlql:atomic-event>
    :
  </eca:Event>
  :
</eca:Rule>
```

Note that if the domain namespace can be derived from the inner markup it is not necessary to indicate it explicitly (the Languages-and-Services Registry knows which namespaces are languages and which are domains).

```

<eca:Rule>
  <eca:Event>
    <aes-xmlql:atomic-event domain="http://www.semwebtech.org/domains/2006/travel#"
      xmlns:xmlqlmatchns="http://www.semwebtech.org/languages/2006/aes-xmlql"
      xmlns:travel="http://www.semwebtech.org/domains/2006/travel#" >
      <travel:delayed-flight flight="{ $flight}" time="{ $time}" />
    </aes-xmlql:atomic-event>
  </eca:Event>
  :
</eca:Rule>

```

We recommend the designers of such languages to support such explicit elements as *optional* markup.

No Explicit Element. Such an explicit surrounding element is often not necessary since there is a namespace change, and the language information can be given as an attribute in the subelement.

5.2.6 Example: SNOOP

As an example, a CED service based on the SNOOP [24] event algebra, extended with logical variables has been implemented in the prototype.

A DTD-style syntax description is as described below. It contains the usual simple combinators, additionally “any n out of *alternatives*”, different kinds of periodic and aperiodic, optionally cumulative events. Operands can also be atomic events, or variables (whose contents are again algebra trees). The actual DTD can be found in Appendix 13.2.4.

Language Identification for Nested Subexpressions. As discussed above, the SNOOP interpreter must be able to identify the languages of non-snoop subexpressions (atomic events or nested expressions from other CELs) accordingly. By default, this is done by the namespace used by these elements. Additionally, `snoop:language` (usually, expecting the namespace; but it can be allowed to use short names as “aliases”) and `snoop:domain` can be given if needed. Note that both are required in case of opaque AESs (an XQuery AES in fact is an opaque AES).

Variable Bindings in Algebra Expressions. Variables can be bound to literal values or subtrees as discussed above in Section 5.2.2. Furthermore, variables can be bound to algebraic subexpressions. Our SNOOP-style proposal provides different ways to handle variables on this level.

Example 20 Consider the following event specification that should react if a *b*-event occurs after an *a*-event (by using a join variable *nr*). The shortest form here does not require to use explicit `<snoop:atomic-event>` elements:

```

<eca:Rule... >
  <eca:Event xmlns:dom="anydomain" ... >
    <snoopy:Sequence>
      <dom:a snoopy:language="http://www.semwebtech.org/languages/2006/aes-xmlql" nr="{ $id}" />
      <dom:b snoopy:language="http://www.semwebtech.org/languages/2006/aes-xmlql" nr="{ $id}" />
    </snoopy:Sequence>
  </eca:Event>
  :
</eca:Rule>

```

Assume now that for every such pair, the complete second event, i.e., the complete structure

```
<dom:b nr="x123" attributes> contents </dom:b>
```

should be bound to a variable *var-x*. This can be achieved by putting it inside a `<snoopy:variable>` element:

```
<eca:Rule... >
  <eca:Event xmlns:dom="anydomain" ... >
    <snoopy:Sequence>
      <dom:a snoopy:language="xmlqmatch" nr="{ $id}" />
      <snoopy:variable name="var-x" >
        <dom:b snoopy:language="xmlqmatch" nr="{ $id}" />
      </snoopy:variable>
    </snoopy:Sequence>
  </eca:Event>
  :
</eca:Rule>
```

Then, `/rule/event/seq/*[1]` addresses the first event specification, but the second one is not addressed by `/rule/event/seq/*[1]` but by `/rule/event/seq/*[2]/*`.

Instead, adding the relationship to a variable as an attribute in the `snoopy:` markup would be equivalent and favorable:

```
<eca:Rule... >
  <eca:Event xmlns:dom="anydomain" ... >
    <snoopy:Sequence>
      <dom:a snoopy:language="xmlqmatch" nr="{ $id}" />
      <dom:b snoopy:language="xmlqmatch" nr="{ $id}" snoopy:variable name="var-x" />
    </snoopy:Sequence>
  </eca:Event>
  :
</eca:Rule>
```

Then, `/rule/event/seq/*[1]` and `/rule/event/seq/*[2]` address the atomic event specifications. Both atomic events specifications also include the necessary information about the used formalism (by attribute) and the domain (implicit by namespace).

5.2.7 Related Work and Existing (Sub)languages

Detection of Composite Events.

The internals of the event detection engine are then concerned with implementing the semantics of the event combinators, which can be done in different ways:

- Operators as Classes – (cf. RuleCore [13]),
- Tree and event queries – (cf. SNOOP [24], XChange [22]),
- Automata and Petri Nets – ODE and SAMOS (here, also a representation of automata states in XML and transformations by XSLT can be used),
- RDF: describe how an event description transforms into another upon an atomic event. This semantic solution would require an ontology of event combinators, but allows then for a very high-level specification of rules.

Existing (sub)languages.

Especially, existing tools can be employed in a service-oriented architecture:

- XChange’s event query mechanism [22]. Then, the event component is e.g. marked up by

```
<eca:Rule xmlns:banking="http://www.banking.nop"
  xmlns:temporal="http://www.some.webservice" >
  <eca:Event xmlns:xchange="http://xcerpt.org/xchange" >
    <!-- xchange fragment in opaque or xml-markup form -->
  </eca:Event>
  :
</eca:Rule>
```

- RuleCore [13] is an ECA system that provides an event detection component where new operators can be added by appropriate classes. Then, RuleCore can be used to implement and experiment with arbitrary event algebras.
- Xcerpt [23] can be used as an AESL.

Composability: Embedding Different Event Algebras

Note that with the above design, it is also possible to embed terms (events) of an event algebra into terms of another one: the namespace identifies the subterm as an algebra expression, and the whole subterm is sent to the responsible processor that in turn answers with a “detected event” message as usual.

5.3 Architecture and Communication: ECA, CED, and AEM

Event processing is done in cooperation of an ECA engine, one or more *Composite Event Detection Engine (CED)* that implement the event algebras, and one or more *Atomic Event Matchers (AEM)* that implement the *Atomic Event Specification Languages (AESL)*. The architecture part that is relevant for handling events is shown in Figure 5.2.

The event *detection* then happens when actual events occur. Here, an *independent* and a *dependent* variant are considered:

Independent: here, the owner of the rule does not supply the events (e.g., when a user defines a rule for travel planning). In this case, the AEM is responsible for being informed about relevant events.

Dependent: here, the owner of the rule also supplies all relevant events. This is e.g. the case when an application service (e.g. an airline) defines rules for their own business that use only the events where the airline is aware of.

The independent version works as follows:

- The ECA engine registers the event component at an appropriate CED (composite component) [or AEM (only atomic event)] service,
- the CED registers the AESs at appropriate AEMs that implement the AESLs,
- the AEMs are assumed to be informed about all relevant events in the format given in Section 5.2.1. Details about how this happens are described in Section 8.1.
- The domain nodes or event brokers forward relevant atomic events to the AEMs,
- the AEMs match them against the (more detailed) specification and inform the CEDs,
- the CEDs process them and inform the ECA engine in case that a composite event has been detected.

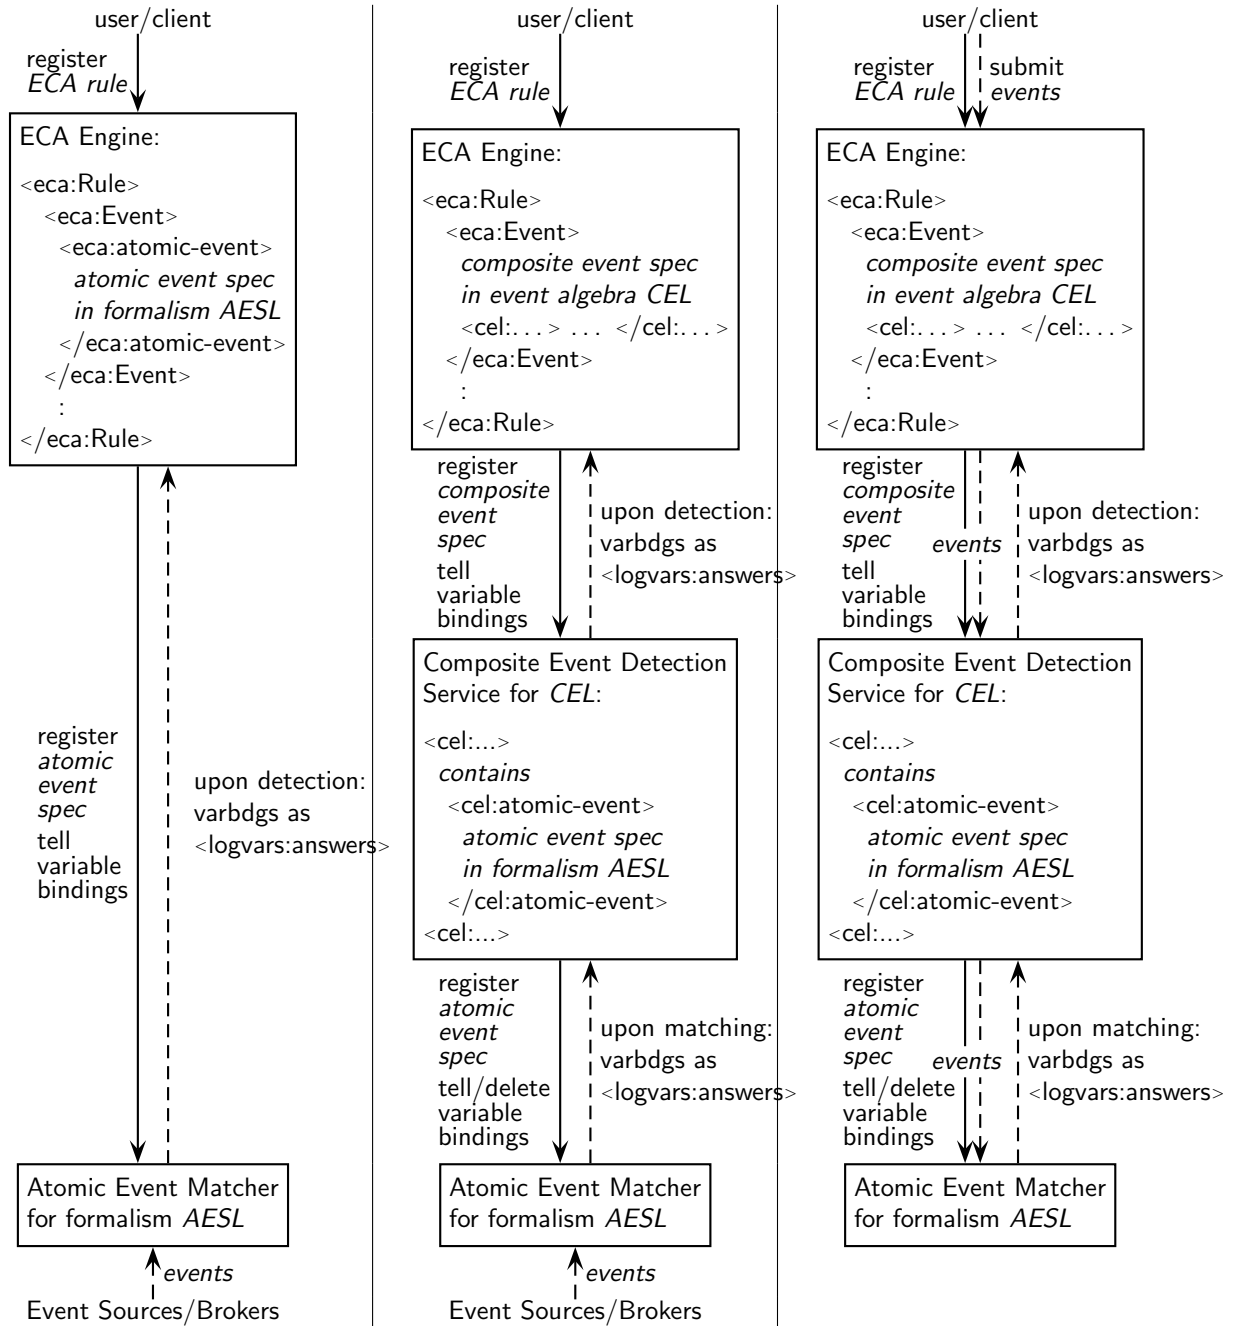


Figure 5.2: Architecture: Processing Event Components and Events (left and middle: independent; right:dependent)

5.3.1 Abstract Semantics and Markup of Event Detection Communication

The above considerations show that the communication with CED services and AEM services uses the same patterns. The all communication $ECA \rightarrow CED$, $ECA \rightarrow AEM$, $CED \rightarrow CED$ (nested) and $CED \rightarrow AEM$ follows patterns described in Section 5.1:

- Downwards: Reply-To, Subject/identifier of request, fragment, optional: variable bindings

(depends whether the service can handle multiple bindings, see Section 9.5).

- Upwards: Subject/identifier, functional result (matched event (sequence)) with optional tuples of variable bindings.

The responses of the event component services must match the format for variable bindings given in Section 4.2.3. Moreover, the answer from detecting an atomic event should have the same format as when detecting a composite event for the following reasons:

- conceptual cleanliness: both are “detected events”,
- architecture: the event component of a rule may be either a composite event or an atomic one. Having different answer formats would require more work. The same holds for allowing embedded event algebra expressions into another event algebra expression.

As usual in the literature about Active Databases, the “semantics” of an event expression is usually not only “yes/no”, but the sequence of atomic events that contributed to the detection. Additionally, event algebras is extended with using free logical variables with join semantics. Thus, here we have a case where the semantics consists of

- a functional result, and
- for each such result (in this case, it is only one), a set of tuples of variable bindings (which actually can be several ones if the sequence allows for several matches).

5.3.2 Communication for Event Components between ECA and CED

When a rule is registered at the ECA engine, the ECA engine registers the event component at an appropriate service. Here we consider only composite event components (atomic ones are dealt with in the following section since their communication is analogous to the one between CED and AEM). Consider the following generic example:

```
<eca:Rule... >
  <eca:Event>
    <evt:foo xmlns:evt="URI of the 'evt' event language" >
      specification of the event component
    </evt:foo>
  </eca:Event>
  :
</eca:Rule>
```

The language used in the event component is identified by the namespace of the event subexpression. How this service identification throughout the Web is done exactly is described in Chapter 9 – the result is a URL where the request has to be sent to and some minor directives how it should be sent (HTTP, SOAP) and wrapped (header, body, hull element).

The above abstract semantics is e.g. communicated in XML as

to: service-URL of the CED

```
<register>
  <Reply-To>URL where the answer is expected (at the ECA engine) </Reply-To>
  <Subject>identifier of the request –
    in most cases the RDF URI of the event component</Subject>
  <!-- next the event component -->
  <evt:foo xmlns:evt="URI of the 'evt' event language" >
    specification of the event component
  </evt:foo>
  <!-- and optionally the existing variable bindings -->
```



```

<logvars:variable-bindings xmlns:logvars="http://www.semwebtech.org/lang/2006/logic#" >
  variable bindings
</logvars:variable-bindings>
</register>

```

The upwards communication is even simpler: the result is sent to the address originally given as Reply-To. The agreed subject is used, and the <answers> or a single <answer> element is submitted:

```

to: Reply-To-address from the request that is answered
<anyname>
  <Subject>identifier of the request</Subject>
  <!-- next the event component -->
  <logvars:answers xmlns:logvars="http://www.semwebtech.org/lang/2006/logic#" >
    answers
  </logvars:answers>
</anyname>

```

If the subject is sent in the HTTP header, also the surrounding <anyname> element can be omitted.

5.3.3 Communication for Event Matching with AEMs

Similar to the ECA-CED communication, it is asynchronous and consists of *registering* relevant CESs (downwards) and reporting occurrences (upwards). In case that the event component of a rule is just atomic, there is a direct communication between ECA and AEM; otherwise the CED communicates with the AEM (note that this induces that the upward communication from the AEM to CED or ECA should be the same as from the CED to the ECA, using the answers format described in Section 4.2.3). In the following we assume a communication from the CED with the AEM, which subsumes the ECA-AEM case.

The proceeding is the following: The CED selects all AESs (=leaf expressions) inside the CES and registers each of them at an appropriate AEM. We here assume that the AEM is aware of all (relevant) events (this will be described in more detail in Section 8.1) and notifies the CED upon occurrence of an event matching a registered AES.

Registration of an AES. If the ECA engine or any composite event detection engine is interested in being informed about occurrences of atomic events wrt. some AES, it sends a message with the following contents to an appropriate atomic event detection service:

- Reply-To, Subject/identifier of request, the AES (as XML fragment or URI reference), optional: variable bindings.

The overall format is the same as when registering a CES at a CED:

```

to: service-URL of the AEM
<register>
  <Reply-To>URL where the answer is expected (at the ECA or CED engine) </Reply-To>
  <Subject>identifier of the request –
    in most cases the RDF URI of the AES wrt. event component</Subject>
  <!-- next the AES -->
  the AES fragment
  <!-- and optionally the existing variable bindings -->
  <logvars:variable-bindings xmlns:logvars="http://www.semwebtech.org/lang/2006/logic#" >
    variable bindings
  </logvars:variable-bindings>
</request>

```

Occurrence Indication of Relevant Events. When such an event is actually detected, it is immediately reported in an answer by the AEM that uses the format described in Section 4.2.3 by `logvars:answers`, `logvars:answer`, `logvars:result` and `logvars:variable-bindings`.

Note that if multiple occurrences are detected together, more than one `logvars:answer` elements can be sent in one `logvars:answers` element (although event notifications should always be sent as soon as possible, it is possible that one event leads to multiple occurrences of a specified atomic event).

Example 21 *Consider that notifications of delayed flights are published all 10 minutes as a list. The domain broker for travel: gets this information and processes it.*

```
<msg:receive-message sender="service@fraport.com" >
  <msg:content>
    <travel:delayed-flight flight="LH1234" minutes="30" />
    <travel:delayed-flight flight="AF0815" minutes="90" />
    :
    <travel:canceled-flight flight="AL4711" />
    :
  </msg:content>
</msg:receive-message>
```

Consider the following rule at an ECA engine:

```
<eca:Rule>
  <eca:Event>
    <eca:atomic-event language="http://www.semwebtech.org/languages/2006/aes-xmlql" >
      <travel:delayed-flight flight="{flight}" minutes="{minutes}" />
    </eca:atomic-event>
  </eca:Event>
  :
</eca:Rule>
```

The ECA engine registers the atomic event at an appropriate atomic event matcher (that in turn will contact the travel service to be informed about relevant events), see Section 5.2.

Later, an answer upon arrival of the above message from the airport will be of the following form:

```
<logvars:answers subject="rule-id/event" >
  <logvars:answer>
    <logvars:result>
      <travel:delayed-flight flight="LH1234" minutes="30" />
    </logvars:result>
    <logvars:variable-bindings>
      <logvars:tuple>
        <logvars:variable name="flight" >LH1234</variable>
        <logvars:variable name="minutes" >30</variable>
      </logvars:tuple>
    </logvars:variable-bindings>
  </logvars:answer>
  <logvars:answer>
    <logvars:result>
      <travel:delayed-flight flight="LH1234" minutes="30" />
    </logvars:result>
    <logvars:variable-bindings>
      <logvars:tuple>
```

```

    <logvars:variable name="flight">AF0815</variable>
    <logvars:variable name="minutes">90</variable>
  </logvars:tuple>
</logvars:variable-bindings>
</logvars:answer>
</logvars:answers>

```

Note that in an RDF environment, there would be a URI reference for LH1234 and AF0815.

Alternative: Sideways Information Passing for Atomic Event Detection

In case that composite events consist of multiple atomic events that share logical *join* variables during evaluation of a composite event, in the same way as applying a *sideways information passing strategy* in query evaluation and rule evaluation, variable bindings obtained by “earlier” atomic events can be used for constraining the relevant event occurrences.

Example 22 (Join Variables in Composite Events) Consider again Example 17 which uses a cumulative SNOOP event with a join variable $\$Subj$:

```

<snoopy:Cumulative xmlns:snoopy="http://www.semwebtech.org/languages/2006/snoopy#"
  xmlns:uni="..." >
  <uni:reg_open snoopy:language="xml-ql-match" subject="{\$Subj}" />
  <uni:register snoopy:language="xml-ql-match" subject="{\$Subj}" />
  <uni:reg_close snoopy:language="xml-ql-match" subject="{\$Subj}" />
</snoopy:Cumulative>

```

Here, when the event component is registered at Snoop, it can do the following:

- it can immediately register all three atomic event patterns at the XML-QL-Matcher service. Assume that then `reg_open("Databases")` is detected. Later, the AEM will report for `register("Scott Tiger", "Databases")`, but also for `register("John Doe", "Algorithmics")` for that rule. The Snoop engine must then apply the join semantics.
- it can register only the first event pattern `reg_open(\$Subject)`. If then, `reg_open("Databases")` is detected and reported, Snoop registers next `register(., "Databases")` which will only report registrations for databases, not for other courses. (The same can be achieved by registering once and disabling/enabling event patterns at the AEM.)

5.3.4 Identification of an AEM for an Atomic Event Specification

As stated above, at the language border to the AES it must be possible for the processor of the surrounding language to identify which service is responsible for processing the AES. Having an element that is suspected to be an atomic event description (i.e., it leaves the namespace from the surrounding language, or it is an **opaque** element), the processor determines the relevant service as follows:

- Check if the namespace of the first element “behind” the language border is an AESL (which can be found out in the Language & Services Registry, see Section 9.5.1). If yes, use it.
- Check if the first element “behind” the language border has a `language` attribute (which is not in the domain namespace of the element). If this language is an AESL, use it. (note that this covers explicit wrapping elements and also allows to add this information just as a shortcut to e.g. the element to be matched.)
- Check if the last element “before” the language border (note that this is “the” language of the currently processing service - so it is directly understood) is a wrapping element for atomic events and contains information about the embedded language (e.g. in a `language` attribute). If yes, use it.

- Determine the domain of the first element “behind” the language border. If the domain broker for that domain supports an AED mechanism, then submit the AED directly to the domain broker.
- If the element is an `<opaque language=“language”>` element, then identify the responsible service (which can be the `language` service or a domain broker which is automatically aware of relevant events). Note that the `opaque` can either belong to the surrounding namespace or to the namespace of the embedded language.
- Otherwise apply suitable heuristics or return an error message.

Note that a language border in a fragment of an event algebra language does not necessarily lead into an AESL, but can also lead to an embedded event algebra language.

5.3.4.1 Example

Example 23 *The following specifies, in an illustrative, non-normative (XML) markup, an event for (very simplified) detection of a late train. It is a composite event specification in the SNOOP (algebraic) language, and uses atomic events from messaging and the domain of train travels. The detection of late trains is made either by being warned by mail from the travel agency, or by the occurrence of a domain-specific event signaling changes in a given (pre-defined) source with expected arrival times:*

```
<eca:Rule xmlns:msg=“http://www.messages.msg/messages”
  xmlns:travel=“http://www.semwebtech.org/domains/2006/travel#” >
  <eca:Event>
    <snoopy:Or xmlns:snoopy=“http://www.semwebtech.org/languages/2006/snoopy#” >
      <snoopy:Atomic language=“xml-ql-match” >
        <msg:receive-message sender=$myTravelAgent” >
          <msg:content>
            <travel:delayed-train train=$myTrain arrivalTime=$newArrival/>
          </msg:content>
        </msg:receive-message>
      </snoopy:Atomic>
      <snoopy:Atomic>
        <aes-xpath:event xmlns:aes-xpath:“http://www.semwebtech.org/languages/2006/aes-xpath” >
          <aes-xpath:test cond=“$event/name()='travel:delayed-train'”/>
          <aes-xpath:test cond=“$event/@train=$myTrain”/>
          <aes-xpath:variable name=“newArrival” select=“$event/@arrivalTime”/>
        </aes-xpath:event>
      </snoopy:Atomic>
    </snoopy:Or>
  </eca:Event>
  <eca:Action> an action specification in any markup
</eca:Action>
</eca:Rule>
```

When a customer registers the rule, the values for the variables `myTrain` and `myTravelAgent` have to be supplied. The ECA rule engine registers the whole event component at the SNOOP service (identified by the URL of the `snoopy` namespace):

- *Reply-To:* ECA engine (exact URL where the response should arrive)
- *Subject/task identifier:* the-rule-id/event
- *the whole event component, i.e., the snoopy:Or element.*

- a `logvars:variable-bindings` element that contains the bindings for `myTrain` and `myTravelAgent`.

Snoop parses the tree as far as it belongs to its namespace. The composite event is an “or” of two atomic events which are not “understandable” to *Snoop*.

For the first one, `language=“xml-ql-match”` is given which can be looked up in the Languages & Services Registry for identifying a service. The whole AES is then sent to `xml-ql-match-service`, together with the bindings of the variables `myTrain` and `myTravelAgent`:

- *Reply-To*: `me`
- *Subject/task identifier*: `the-rule-id/event/disj/*[1]`
- the whole event component, i.e., a `snoopy:Or` element.
- optional (if the `xml-ql-match-service` supports variables): a `logvars:variable-bindings` element that contains the binding for `myTrain` and `myTravelAgent`; otherwise the value must be inserted as string into the code fragment.

The AES describes the receipt of a message (marked-up in XML) with an attribute `sender` which is equal to the value of the variable `myTravelAgent`, and with a `content` with a `delayed` element with an attribute `train` coinciding with that of `myTrain`. After detection, the variable `newArrival` will be bound to the value of the attribute `arrivalTime` of the `delayed` element.

The second AES is identified by its namespace “`navigation-based-formalism-uri`” where also a service URL can be looked up. The whole AES is sent there, this time, only the bindings of the variable `myTrain` are required. The AES specifies a domain-specific event `travel:delayed-train` (that occurs “somewhere in the Web” and has to be detected by Semantic Web mechanisms). The event is implicitly bound to `$event`. The details are then checked by XPath expressions against `$event`: If its attribute `train` equals the value of the variable `myTrain`, then `newArrival` is bound to the value of the `newTime` attribute of the event.

Both AEM services will scan all events that they are aware of for matching, and in case of success they will return an answer (according to the format given in Section 4.2.3).

Assume that the second AEM becomes aware of an event

```
<travel:delayed-train xmlns:travel=“http://www.semwebtech.org/domains/2006/travel#”
  train=“ICE 773” plannedArrival=“11:30” arrivalTime=“13:00” /> .
```

Then, it responds to the *Snoop* service with a message

- *task identifier*: `the-rule-id/event/disj/*[2]`
- *contents*:

```
<logvars:answer>
  <logvars:result>
    <travel:delayed-train xmlns:travel=“http://www.semwebtech.org/domains/2006/travel#”
      train=“ICE 773” plannedArrival=“11:30” arrivalTime=“13:00” />
    </logvars:result>
  <logvars:variable-bindings>
    <logvars:tuple>
      <logvars:variable name=“myTrain”>ICE 773</variable>
      <logvars:variable name=“newArrival”>13:00</variable>
    </logvars:tuple>
  </logvars:variable-bindings>
</logvars:answer>
```

Snoop will then evaluate the semantics of the disjunctive event, which is detected if one alternative has been detected and sends a message to the ECA engine:

- *task identifier*: the-rule-id/event
- *contents*:

```

<logvars:answer>
  <logvars:result>
    <travel:delayed-train xmlns:travel="http://www.semwebtech.org/domains/2006/travel#"
      train="ICE 773" plannedArrival="11:30" arrivalTime="13:00" />
  </logvars:result>
  <logvars:variable-bindings>
    <logvars:tuple>
      <logvars:variable name="myTrain" >ICE 773</variable>
      <logvars:variable name="myTravelAgent" >...</variable>
      <logvars:variable name="newArrival" >13:00</variable>
    </logvars:tuple>
  </logvars:variable-bindings>
</logvars:answer>

```

Note that upon receipt of the message at the ECA service, the event sequence itself is not bound to a variable (since it is not indicated) – all relevant information can be extracted without this.

Example 24 *In Example 17, we illustrated SNOOP’s cumulative event: registration for an exam begins, students register and the registration is closed. The returned message in this case from the Snoop service could e.g. contain the following answer:*

```

<logvars:answer>
  <logvars:result>
    <uni:reg_open subject= "Databases" />
    <uni:register subject= "Databases" name= "John Lennon" />
    <uni:register subject= "Databases" name= "Paul McCartney" />
    <uni:register subject= "Databases" name= "George Harrison" />
    <uni:register subject= "Databases" name= "Ringo Starr" />
    <uni:reg_close subject= "Databases" />
  </logvars:result>
  <logvars:variable-bindings>
    <logvars:tuple>
      <logvars:variable name="Subj" >Databases</variable>
    </logvars:tuple>
  </logvars:variable-bindings>
</logvars:answer>

```

For collecting all names, the event component of the rule is contained in an <eca:variable name="regseq"> element that assigns the returned event sequence to that variable. The names of the registered students can then be extracted by a query.

5.4 The Query Component

There are a lot of query languages for XML or RDF data around, such as XQuery (where an XML markup proposal has been presented in [75]), F-Logic, RDFQL, XPathLog, or Xcerpt. Note that queries that bind (join) variables can already be used for restricting the results.

Queries that are only concerned with the contents a certain database node can use the query language of this node, often in an opaque way. In all cases, the result of evaluating queries is communicated in the usual form described above.

5.4.1 Opaque Queries

For the development of a first XML level prototype of the MARS Framework we rely on *opaque* queries; see Section 10.1.2.

5.4.2 Atomic Queries

Analogously to atomic events or actions, atomic queries refer to atomic notions of the domain ontologies. In algebra-based query languages like SQL or Description Logic queries, the atomic queries are the leaf expressions (table names, concept names, property names) which are combined by operators.

Relational Model/SQL/Datalog Concerning the relational model and the relational algebra or calculus, atomic queries are the relation names (i.e., the leaves in algebra trees), or predicates (in conjunctive queries in the relational calculus and Datalog). The “algebraic” language is then the relational algebra, SQL, or simple conjunctive queries as in Datalog.

RDF. For the RDF model, atomic queries are triples. The “algebraic” language is then provided by SPARQL etc., in the same style as the relational algebra.

XML. For XML, this notion does not exist so clearly. XQuery FLWR expressions are obviously not atomic. Are XPath expressions atomic queries? According to the algebra evaluation, they are not. Atomic here corresponds to decomposing XPath expressions in single steps and expressing an XPath expression as a conjunctive query (as e.g. done when mapping it to F-Logic).

Since there is not yet a favorite Semantic Web query language, there are numerous proposals and interpretations.

5.4.3 Composite Queries

Composite queries can be expressed in two ways:

- by query algebras over atomic query expressions, or
- logic-based.

The result of the evaluation of an expression is returned in the format described in Section 4.2.3.

GG(M)QL: Generic Graph Query (Markup) Language. An XML Markup can be defined based on algebraic operations. For instance, a tree notation of the relational algebra (providing domain-independent operations) with (domain-dependent) relations as atomic expressions provides an example of a tree-structured query component.

Dealing with graph-shaped data, an extended algebra is recommended, e.g., in the style of the one used for the implementation of F-Logic [41, 42] in the Florid system [31] or with some adaptations for XPathLog [49, 50] in the LoPiX system [48].

In addition to the operations known from the relational algebra, i.e., “union”, “intersection”, “difference” (as derived operation), “selection”, “projection”, “join” and “division” (derived), a “navigation” operation is common for graph or tree data. Additionally, “grouping/aggregates” (leading to denormalized data) should be provided.

We follow here the definition of the formal semantics of F-Logic as e.g. given in [46, Section 2.2.1]. The (preliminary) DTD is as follows:

```
<!ENTITY % operand "(... | ANY-from-other-namespace | opaque)">
<!ELEMENT union (%operand;, %operand;)>
<!ELEMENT intersect (%operand;, %operand;)>
<!ELEMENT difference (%operand;, %operand;)>
```

```

<!ELEMENT select (%operand; condition)>
<!ELEMENT project (%operand;)>
  <!ATTLIST project variables NMTOKENS #REQUIRED>
<!ELEMENT join (%operand;, %operand;)>
<!ELEMENT navigation (%operand;, %operand;?)>
  <!ATTLIST navigation
    axis (xml-axes) #IMPLIED
    method-type (f-logic method types) #IMPLIED
    nodetest (xml nodetest) #IMPLIED
    property (nmtoken) #IMPLIED
    variable (variable-name) #IMPLIED >

<!ELEMENT condition (any-test-language-expression)>

```

- navigation covers F-Logic, XML and RDF. The optional parameters correspond to the XML “axis”, the F-Logic “method type”, the XML nodetest and the RDF/F-Logic property name. Optionally, the nodetest/property can be replaced by a second operand for *computed* names (as in F-Logic, XPathLog, or SPARQL). In this case, the answer part of the result of the second operand acts as property name, the variable bindings are joined. The *variable* attribute indicates that the result of the navigation should be bound to a new variable or joined with an existing one.
- condition: use an “external” test language (i.e., a subexpression of another namespace), or an opaque statement. It is recommended that a query language provides a set of predicates such as equal, <, >, substring etc. as in XPath locally.

Markup of an RDF-based Query Language. The markup of queries should support *reasoning* about queries and rules on the RDF level.

5.4.4 Result Semantics

Query languages can have a result of the following types:

- functional: just a set of things
- logical: a set of tuples of variable-bindings
- functional+logical: format as described in Section 4.2.3.

When dealing with the ontology of languages (see Section 9), the type of result semantics must be given.

5.5 The Test Component

According to the distinction between the query component (obtaining information) and the test component (which is only allowed to use existing information), evaluating the test means evaluating a condition. A condition is a formula, its composers are the boolean operators, quantifiers etc. *Atomic* tests are thus only domain-independent tests, i.e., equality, comparisons like “<” and “>” etc. Often, the test part can already be included into the query part if the query language allows for conditions, filters, or even in form of join conditions in conjunctive queries.

Tests can be expressed formulas in any logic, e.g., First-Order logic (where an XML markup is given in [15]), F-Logic, XPath-Logic, or XQuery or Xcerpt (note that in these languages, expressions that have an empty result also yield the truth value “false”).

For tests, similar considerations as for queries apply. In contrast to the query part, tests cannot bind variables and their result is only true or false.

- Join/restriction semantics: tuples of variable bindings are communicated downwards, and the test service returns only those tuples that pass the test. These tuples are then joined with the bindings on the ECA level.
- functional semantics: tuples of variable bindings are communicated downwards, and the test service labels each tuple with true/false.

Test expressions are thus composite expressions over atomic predicates (that can be “simple” classical predicates, or pattern terms as in F-Logic or Xcerpt).

Simple tests (comparison predicates etc.) can usually be evaluated locally at the ECA engine. For this, also the simple boolean algebra operators “and”, “or” and “not” are used.

5.5.1 Test Component Languages

In contrast to queries, tests work on a set of tuples of (input) variable bindings. They cannot bind additional (non-local) variables (local variables may be used in subqueries).

Boolean Logic. The simplest tests only use atomic tests and combine them by boolean operators, i.e., “and”, “or” and “not”. For the markup, see below. The (preliminary) DTD is as follows:

```
<!ENTITY % operand "(... | ANY-from-other-namespace | opaque)">
<!ELEMENT and (%operand;, %operand;)>
<!ELEMENT or (%operand;, %operand;)>
<!ELEMENT not (%operand;)>
<!ELEMENT condition (any-test-language-expression)>
```

The operators can be evaluated according to their usual boolean semantics by verifying them (truth-table style), or relationally: and is actually a join or cartesian product, or is a union (care for same set of bound variables), and not is actually a set difference. Quantifiers are not allowed.

Quantifiers and Grouping and de-grouping. A test can be evaluated to each tuple individually, or for a group of tuples. In the latter case, quantifiers can be applied to each group.

Example 25 Consider the following situation: the input consists of a set of pairs (X, Y) where for each X , multiple pairs are allowed. Those X where all corresponding Y s are > 100 , or at least one Y is > 1000 , should be continued, the other ones are discarded.

```
<group-by variable="X">
  <or>
    <all variable="Y">
      <fn:greater-than>
        <variable>Y</variable>
        100
      </fn:greater-than>
    </for-each>
    <exists variable="Y">
      <fn:greater-than>
        <variable>Y</variable>
        1000
      </fn:greater-than>
    </exists>
  </or>
</group-by>
```

(Note that similar behavior can be expressed also by the action part with alternative branches, grouped by X)

Note that, given a variable X , any test of the form $\forall X : (p(X, Y) \rightarrow q(Y))$ actually involves a query, namely to obtain all Y such that $p(X, Y)$ holds. Thus, this *should* be done in a previous query step. Moreover, “for all” is critical when an open world is assumed.

Often, it is preferable to express “typical” predicates in the domain ontologies as derived properties.

5.5.2 Atomic Tests: Predicates

A test language can support built-in predicates that are evaluated locally. Additionally, domain-specific predicates and terms can be used in the test. The latter have to be evaluated by calling a corresponding domain broker. To save this additional communication overhead, it is recommended to state already the query as strict as possible.

Example 26 Consider the following simple rule: a person P books a travel to city C . The person has a contract with a very attentive car rental company WorldWidePersonalizedMobility (WWPM) that provides a highly personalized service. Customers can register for different profiles, e.g., “ H ” (habit) customers are offered a car of the same size as the person owns at home. For this, the car rental company maintains a database about which cars the person owns, and how much the person is willing to pay. If an appropriate car is available at the target city, the most expensive one is reserved (WWPM wants to make profit, and to offer the customers the “highest” choice).

- Rule “ H ”:
- event: person P books a flight to C .
- initial query: person has profile “ H ”?
- next queries: which cars X are available at P for which price PX ? After that query, all values that are potentially needed for the action are obtained.
- the test could be the following: test whether X is of the same size as any car that P owns, and costs below the maximal amount A that P is willing to pay.
- the action then consists of booking the “best” one.

The test requires to access the underlying database and to evaluate a comparison. In the normal form, the queries and test are as follows (e.g., using XPath)

```
<eca:Rule xmlns:travel="http://www.semwebtech.org/domains/2006/travel#" >
  <eca:Event language="xml-ql-match" >
    <travel:booked-flight name="{P}" to="{City}" />
  </eca:Event>
  <eca:Query variable="$OwnCar" >
    <eca:opaque language="XPath" >
      /customers/customer[@name="$P"]/owncar/@type
    </eca:opaque>
  </eca:Query>
  <eca:Query variable="OwnClass" >
    <eca:opaque language="XPath" >
      /types/type[@name="$OwnCar"]/@class
    </eca:opaque>
  </eca:Query>
  <eca:Query variable="$MaxPrice" >
    <eca:opaque language="XPath" >
```

```

    /customers/customer[@name="$P"]/@maxprice
  </eca:opaque>
</eca:Query>
<eca:Query>
  <eca:use-variable name="$City" use="City" />
  <eca:opaque language="XPathLog">
    /branches/branch[@city→$City]/cars[@type→Type and @price→Price] and
    /types/type[@name→Type and @class→Class]
  </eca:opaque>
</eca:Query>
<eca:Test xmlns:test="http://www.semwebtech.org/mars/2006/test" >
  <test:and>
    <test:equals><test:variable>$OwnClass</test:variable><test:variable>Class</test:variable></test:equals>
    <test:less-equal><test:variable>Price</test:variable><test:variable>$MaxPrice</test:variable></test:less-equal>
  </test:and>
</eca:Test>
<eca:Action>
  <!-- book the cheapest one -->
</eca:Action>
</eca:Rule>

```

5.6 The Action Component

The action component actually enforces the consequences of an ECA rule. The action component is either an atomic action of an application domain, or a specification of a composite action, given in a *Composite Action Language (CAL)* implemented by *Composite Action Execution Engines (CAEs)*. For executing actions, two scenarios can be distinguished:¹

- Rules can implement *global* behavior (general dynamic integrity constraints like policies in an application domain). Here, it is *not* clear which nodes should actually execute the action. In this case, “action brokers” are required that control the communication with the respective domain nodes; see Sections 8.4.
- Many rules are registered by maintainers of domain nodes to implement reactions on events somewhere in the Semantic Web that also use data obtained by queries against the Semantic Web, but whose actions should be executed in the *own* domain node.

For dealing with the first case, we propose to extend the action component with an attribute `eca:execute-at` whose value can be either a pseudo-variable “\$owner” (the owner of the rule) or a hardcoded URL:

```

<eca:Rule >
  <eca:Event> ... </eca:Event>
  :
  <eca:Action execute-at="domain node identifier" >
    action specification
  </eca:Action>
</eca:Rule>

```

We propose that CALs also provide a similar means to add such information to atomic actions [the user can see if this is provided from the ontology metadata of the CAL].

¹Probably there are even more ...

5.6.1 Atomic Actions

Atomic components of composite actions can be the following:

- actions of domain nodes (see Chapter 7) by framework-native mechanisms,
- opaque actions that invoke arbitrary Web Services (by HTTP or SOAP messages),
- implementation of ECE rules: their action component is usually atomic and just raises an event in an application domain.

Invoking Actions at Framework-Aware Nodes. For executing the action for a given rule instance, the request contains the contents of the action component and the variable bindings given in the format described in Section 4.2.2.

The request is sent to an action broker (cf. Section 8.4) that identifies the actual receiver of the request (using the data contained in the request). Note that for a single such message with several variable bindings, *different* application nodes may be finally responsible. The selection is done by the action broker.

Example 27 (Canceling flights due to bad weather) *Consider the following case: at some airport the weather conditions are forecasted to become bad, so that incoming flights cannot land. There is then a rule “if the forecast is ... then for all flights landing in the afternoon, cancel these flights”. All flights that are concerned can easily be selected from the flight schedule. Since these are operated by different airlines, the actions must be executed at different application nodes (that can in an RDF world be determined from the association between flight numbers and the corresponding airlines).*

```
<eca:request>
  <eca:Action>
    <eca:input-variables names="flight reason" />
    <travel:cancel-flight flight=$flight>
      <travel:reason>$reason</travel:reason>
    </travel:cancel-flight>
  </eca:Action>
  <logvars:variable-bindings>
    <logvars:tuple>
      <logvars:variable name="flight">LH123</logvars:variable>
      <logvars:variable name="reason">bad weather</logvars:variable>
    </logvars:tuple>
    <logvars:tuple>
      <logvars:variable name="flight">AL400</logvars:variable>
      <logvars:variable name="reason">bad weather</logvars:variable>
    </logvars:tuple>
  </logvars:variable-bindings>
</eca:request>
```

Opaque Actions. Opaque actions are similar to opaque queries. They contain a code fragment that has to be executed by some service.

- the code fragment is part of an URL to be appended to the service URL:

```
<eca:Rule>
  :
  <eca:Action>
    <eca:opaque uri="domain-service-url" >
```

```

    <eca:input-variables names="flight reason" />
    cancel-flight($flight,$reason)
  </eca:opaque>
</eca:Action>
</eca:Rule>

```

will call e.g.

```
domain-service-url:cancel-flight("LH123", "bad weather")
```

- the code fragment is just an update operation, e.g. an SQL statement DELETE FROM *table* WHERE *condition*. In this case, the URL where the action has to be sent must be given in the opaque element. (The case is similar to submitting an XPath query to a given database.)

```

<eca:Rule>
:
<eca:Action>
  <eca:opaque uri="domain-service-url" >
    <eca:input-variables names="flight reason" />
    INSERT INTO cancellations VALUES($flight,$reason)
  </eca:opaque>
</eca:Action>
</eca:Rule>

```

- the code fragment is a code fragment in some programming language that contains the URLs to be updated. In this case it can be evaluated by a “free” language service (the case is similar to evaluate an XQuery query of the form FOR \$x document(*url*).

```

<eca:Rule>
:
<eca:Action>
  <eca:opaque lang="XQuery+Updates" >
    <eca:input-variables names="flight reason" />
    do insert <cancel>$reason</cancel>
    as last into doc("url")/flights/flight[@nr=$flight]
  </eca:opaque>
</eca:Action>
</eca:Rule>

```

In all cases, the variable occurrences in the expressions are replaced in the same way as for opaque (HTTP) queries. Again, a wrapper can be used that provides a framework-aware interface and iterates over the bindings.

Raise Events. Events are raised in the XML format given in Section 5.2.1. This can be encoded into an opaque action of sending an HTTP message with the raised event to an appropriate target (the event brokers that support the respective domain).

```

<travel:canceled-flight flight="LH123" >
  <travel:reason>bad weather</travel:reason>
</travel:canceled-flight>
<uni:reg_open subject="Databases" />
<uni:register subject="Databases" name="John Doe" />

```

5.6.2 Composite Actions

Composite actions can e.g. be described by *Process Algebras*. Process Algebras describe the semantics of processes in an algebraic way, i.e., by a set of elementary processes (carrier set) and a set of constructors. The semantics can either be given as *denotational semantics*, i.e., by specifying the denotation of every element of the algebra (e.g., CSP – Communicating Sequential Processes, [36]), or as an *operational semantics* by specifying the behavior of every element of the algebra (e.g., CCS – Calculus of Communicating Systems, [56, 57]). Processes defined by Process Algebras can e.g. be used for the specification of *communication*, i.e., for basic protocols, or for defining the behavior of interacting (Semantic) Web Services (note that Process Algebras provide concepts for defining infinite processes).

The structure and markup of composite action languages becomes relevant when reasoning about a system should be done. For the above algebras, model checking is available for verification.

Basic Process Algebra (BPA). For a given set \mathcal{A} of atomic actions,

$$BPA_{\mathcal{A}} = \langle \mathcal{A}, \{\perp, +, \cdot\} \rangle$$

is the basic algebra – i.e., containing the least reasonable set of operators – for constructing processes over \mathcal{A} . \perp is a constant denoting a deadlock, $+$ denotes alternative composition, and \cdot denotes sequential composition: if x and y are processes, then $x + y$ and $x \cdot y$ are processes. These are essentially the processes that can be characterized in Dynamic Logic [35] and Hennessy-Milner-Logic [58].

A term markup of BPA ist straightforward:

```
<!ENTITY % operand "(alt | seq | atomic-action |
  ANY-from-other-namespace | opaque)">
<!ELEMENT alt (%operand;, %operand;, %operand;*)>
<!ELEMENT seq (%operand;, %operand;, %operand;*)>
<!ELEMENT atomic-action ANY>
```

BPA action specifications occur in rules in the usual form as

```
<eca:Rule xmlns:travel="http://www.semwebtech.org/domains/2006/travel#" >
  <eca:Event> ... </eca:Event>
  <eca:Query>... </eca:Query>
  <eca:Action xmlns:bpa="..." >
    contents in bpa namespace
  </eca:Action>
</eca:Rule>
```

We do not give a concrete URL (and also no implementation) since composite actions are more comprehensively handled below with CCS.

CCS and CSP. The more sophisticated Process Algebras *CCS – Calculus of Communicating Systems* [56, 57]), *CSP – Communicating Sequential Processes* [36]), ACP [12], COSY [37], or the *Box Algebra* [14] based on the *Petri Box Calculus*, provide more involved composers, and use several kinds of actions. In addition to the common actions, they include “communication actions”, i.e., sending and receiving messages, and also “reading actions” that access the state of a system. In our ontology, these are modeled as events and queries, or tests. Since Process Algebras allow not only for executing a piece of program code, but also the definition of more complex *processes*, including the definition of independent, communicating processes, the resulting model is also more expressive than current formalisms and languages for active rules.

With this, the action component can be used to specify e.g. the following concepts:

1. a sequence of actions to be executed (as in simple ECA rules),

2. a process that includes “receiving” actions (which are actually events in the standard terminology of ECA rules),
3. guarded (i.e., conditional) execution alternatives,
4. the start of a fixpoint (i.e., iteration or even infinite processes), and
5. a family of *communicating, concurrent processes*.

These patterns can be employed for specifying behavior: (2) can e.g. be used to define a negotiation strategy that communicates with a counterpart. (3) can include different reactions to the answers of the counterpart, (4) extends the behavior even to try again. Note that in these cases, only one side of the communication is specified, whereas the behavior of the counterpart is defined by other rules (with another owner). (5) can be used to specify even more complex behavior of interacting (Semantic) Web Services as a reaction as known from the agent community.

These tasks can also be expressed by (sets of) simple ECA rules, but this leads to a much less intuitive, and hard-to-understand specification. The composition of the ECA and Process Algebra concepts (and ontologies) provides a comprehensive framework for describing behavior in the Semantic Web.

Calculus of Communicating Systems (CCS). CCS [56] extends BPA by more expressive operators. A CCS algebra with a carrier set \mathcal{A} is defined as follows, using a set of process variables:

1. Every $a \in \mathcal{A}$ is a process expression.
2. With X a process variable, X is a process expression.
3. With $a \in \mathcal{A}$ and P a process expression, $a : P$ is a process expression (prefixing; sequential composition).
4. With P and Q process expressions, $P \times Q$ is a process expression (parallel composition).
5. With I a set of indices, $P_i : i \in I$ process expressions, $\sum_{i \in I} P_i$ (binary notation: $P_1 + P_2$) is a process expression (alternative composition).
6. With I a set of indices, X_1, \dots, X_k process variables, and P_1, \dots, P_k process expressions, $\text{fix}_j \vec{X} \vec{P}$ is a process expression (definition of a communicating system of processes). The fix operator binds the process variables X_i , and fix_j is the j th one of the k processes which are defined by this expression.

Process expressions not containing any free process variables are *processes*.

The (operational) semantics of CCS is given by transition rules that immediately induce a naive implementation strategy (note that the semantics of CSP [36] is given as *denotational semantics*). By carrying out an action, a process changes into another process. Considering the modeling as a Labelled Transition System, a process can be regarded as a state or a configuration, which allows to use Model Checking for verifying properties of CCS specifications.

$$\begin{array}{c}
 a : P \xrightarrow{a} P \quad , \quad \frac{P_i \xrightarrow{a} P}{\sum_{i \in I} P_i \xrightarrow{a} P} \text{ (for } i \in I) \\
 \frac{P \xrightarrow{a} P' \quad Q \xrightarrow{b} Q'}{P \times Q \xrightarrow{ab} P' \times Q'} \quad , \quad \frac{P_i \{\text{fix } \vec{X} \vec{P} / \vec{X}\} \xrightarrow{a} P'}{\text{fix}_i \vec{X} \vec{P} \xrightarrow{a} P'}
 \end{array}$$

Additionally, asynchronous CCS allows for delays:

$$\begin{array}{lcl}
 \partial P & := & \text{fix } X(1 : X + P) \text{ , } X \text{ not free in } P \text{ , and} \\
 P_1 | P_2 & := & P \times \partial Q + \partial P \times Q \\
 a.P & := & a : \partial P \text{ .}
 \end{array}$$

with the corresponding transition rules

$$\begin{aligned} \partial P \xrightarrow{1} \partial P \quad , \quad \frac{P \xrightarrow{a} P'}{\partial P \xrightarrow{a} P'} \\ \frac{P \xrightarrow{a} P'}{P|Q \xrightarrow{a} P'|Q} \quad , \quad \frac{Q \xrightarrow{a} Q'}{P|Q \xrightarrow{a} P|Q'} \\ \frac{P \xrightarrow{a} P' \quad Q \xrightarrow{b} Q'}{P|Q \xrightarrow{ab} P'|Q'} \end{aligned}$$

The possibility of Delay is especially important when “waiting” for something to occur, e.g., for synchronization.

Actions with Parameters. Actions are usually parameterized, e.g. “book flight no N on $date$ ”. Communication between the rule components is provided by variable bindings. Accordingly, the specification of the action component uses variables as parameters to the actions.

Example 28 (Process Specifications in CCS)

- A money transfer (from the point of view of the bank) is already a simple process:

transfer(Am, Acc_1, Acc_2) :=
debit(Acc_1, Am) : deposit(Acc_2, Am) .

- a standing order (i.e., a banking order that has to be executed regularly) is defined as a fixpoint process, involving an event. The following process transfers a given amount from one account to another every first of a month (where “first_of_month” is a temporal event):

fix X .(first_of_month : debit(Acc_1, Am) :
deposit(Acc_2, Am) : ∂X)

- A more detailed view could e.g. check if the balance will stay positive, and if not, notify the account holder:

fix X .(first_of_month : send_query($Acc_1 \geq Am?$) :
((∂ : rec_msg(yes) :
debit(Acc_1, Am) : deposit(Acc_2, Am)) +
(∂ : rec_msg(no) : send_msg(\$owner,...))) : ∂X)

(using messaging for queries and message receipt events for answers).

Another way would be to express the same as a complete ECA rule “if the event first_of_month occurs, then do ...” instead of a fixpoint process.

Example 29 Consider the following scenario: if a student fails twice in an exam, he is not allowed to continue his studies. If the second failure is in a written exam, it is required that another oral assessment takes place for deciding upon final passing or failure.

This can be formalized as an ECA rule that reacts upon an event failed(\$Subject,\$StudNo) and then in a further query checks whether this is the second failure of \$StudNo in \$Subject, and whether the exam was a written one. The action component of the rule should then specify the process of (organizing) the additional assessment: as an action, the responsible lecturer will be asked for a date and time (send a mail), that will be entered by him into the system (in CCS: a “receiving” communication action; in our approach: an event). The action component is thus as follows:


```

ask_appointment($Lecturer, $Subject, $StudNo) :
∂ proposed_appointment($Lecturer, $Subject, $DateTime) :
find_room($DateTime, $Room) :
inform($StudNo, $Subject, $DateTime, $Room) :
inform($Lecturer, $Subject, $DateTime, $Room)

```

In this example,

proposed_appointment(\$Lecturer, \$Subject, \$DateTime) is an event – for this, it is allowed to be delayed (∂). In contrast, all other items are actions that are actually executed by the process as soon as possible.

Note that entering the grade and further consequences are not covered by this action. Instead, it is appropriate to have a separate rule that reacts (again) on entering grades and, if the grade was established by such an additional assessment, take appropriate actions.

Conditions. In CCS and other process algebras, there is no explicit notion of states, the properties of a state are given by the (sequences of) actions which can be executed. When representing a stateful process, queries and values are represented e.g., as “read that $A > 0$ ”, or by explicit messages (as the account balance in Example 28). We omit the “read”, and allow queries and conditions as regular components of a process:

- “executing” a query means to evaluate the query, extend the variable bindings, and continue.
- “executing” a condition means to evaluate it, and to continue for all tuples of variable bindings where the condition evaluates to “true”. For a conditional alternative $((c : a_1) + (\neg c : a_2))$, all variable bindings that satisfy c will be continued in the first branch, and the others are continued with the second branch.

Example 30 (Processes with Conditions)

1. Consider again the scenario from Example 29, but now only one room is suitable for such assessments. Here, the process in the action part must iterate asking the lecturer for an alternative date/time until the room is available. This is done by combining CCS’s fixpoint operator with a conditional alternative:

```

fix X.(ask_appointment($Lecturer,$Subj,$StudNo) :
∂ proposed_appointment($Lecturer,$Subj,$DateTime) :
(available(room,$DateTime) +
(¬ available(room,$DateTime) : X))) :
inform($StudNo,$Subj,$DateTime) :
inform($Lecturer,$Subj,$DateTime)

```

Here, *ask_appointment* is an atomic action, *proposed_appointment* is an event, and *available* is a predicate (test).

2. The account check in Example 28 can also be expressed by a conditional alternative:

```

fix X.(first_of_month :
((Acc1 ≥ Am? : debit(Acc1, Am) : deposit(Acc2, Am)) +
(Acc1 < Am? : send_msg($owner,...))) : ∂ X)

```

Figure 5.3 shows the relationship between the process algebra language and the contributions of the domain languages and the event and test component languages.

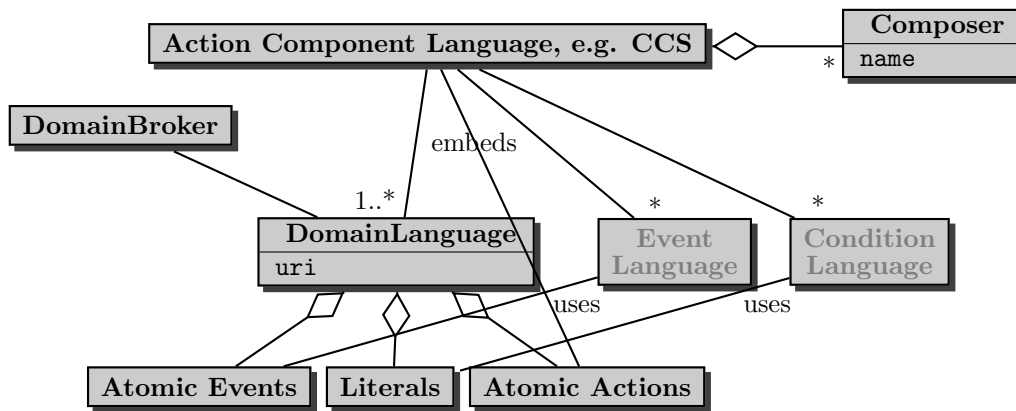


Figure 5.3: Structure of the Action Component as an Algebraic Language using CCS

XML Markup for the Action Component. According to the above considerations, processes are built over

- actions,
- events, and
- conditions

by using the CCS connectives. The language markup has the usual form of a tree structure over the CCS composers in the `ccs` namespace. The leaves are contributed by (i) atomic actions of the underlying domains, (ii) events and conditions/tests. The latter are not necessarily atomic, but are seen as black-boxes from the CCS point of view, containing markup from appropriate languages as used in the ECA event and test components (and handled by the respective services). A longer example will be given in Section 5.6.5.

5.6.3 Atomic and Leaf Items

As discussed above, the leaves on the CCS level can be atomic actions, or embedded events, queries, or test subexpressions. In accordance with ECA-ML, the latter are embedded into `ccs:Event`, `ccs:Query` and `ccs:Test` elements. The language identification is done again via the namespaces.

Atomic Actions. Atomic actions belong to some domain namespace, thus the element is in general the action in XML markup “itself” (including variables as `{ $varname }`):

```
<domain-ns:action-name attributes>
  contents
</domain-ns:action-name>
```

As an example consider an atomic action that books a given flight (flight code bound to variable `$flight`) at a given date (bound to variable `$date`):

```
<travel:book-flight code="{ $flight }" date="{ $date }" />
```

Embedded Events. Embedded events are also leaves, contained in `ccs:Event` elements (with the same semantics as `eca:Event` elements on the ECA level):

```
<ccs:Event xmlns:el="ev-uri" >
  event expression in appropriate markup
</ccs:Event>
```

Arbitrary event languages and formalisms that are supported by some service are allowed. Note that composite events integrate smoothly since they are considered to occur with the final detection of the composite event.

Embedded Conditions. Embedded tests are handled exactly in the same way:

```
<ccs:Test xmlns:c/= "cl-uri" >
  test expression in appropriate markup
</ccs:Test>
```

Embedded Opaque Items. Opaque actions (i.e., program code, mainly for queries, tests and also for actions) can be embedded as leaf elements:

```
<ccs:Opaque {url="node-url" |language="name"}>
  program code fragment
</ccs:Opaque>
```

For such fragments, either a URL where the action has to be sent to (as HTTP GET) is given, or the language is indicated (then the fragment must contain the addressing of the target node itself).

5.6.4 Example: CCS

Following a straightforward principle for term markup, the CCS operators are represented by XML elements (with parameters as attributes) according to the following nearly-DTD specification:

```
<!ENTITY % operand "(Delay | Sequence | Alternative | Concurrent | Fixpoint |
  Action | Event | Query | Test | Variable | Opaque)">
<!ELEMENT Delay EMPTY>
<!ELEMENT Sequence (%operand;, %operand;+)>
  <!ATTLIST Sequence mode "async">
<!ELEMENT Alternative (%operand;, %operand;+)>
<!ELEMENT Concurrent (%operand;, %operand;+)>
<!ELEMENT Fixpoint (%operand;+)>
  <!ATTLIST Fixpoint
    variables NMTOKENS #REQUIRED
    index NMTOKEN #REQUIRED
    localvars NMTOKENS #IMPLIED>
<!ELEMENT Variable (#PCDATA)> <!ELEMENT Event ANY> <!ELEMENT Query
  ANY> <!ELEMENT Test ANY> <!ELEMENT Action ANY>
  <!ATTLIST Event, Test, Action
    xmlns:%name; %URI; #REQUIRED >
<!ELEMENT Opaque ANY>
  <!ATTLIST Opaque
    language CDATA #IMPLIED
    url CDATA #IMPLIED>
  <!ATTLIST all-elements group-by "">
```

The semantics of the elements is described below.

- The content of all the “simple” operators consists of at least two subelements.
- `<ccs:Delay/>` indicates a delay (for waiting, in case that synchronous context is used),
- `<ccs:Sequence mode="mode">contents</ccs:Sequence>` indicates a sequence. *mode* can be *sync* or *async* corresponding to synchronous CCS (with “:” as standard combinator) or asynchronous CCS (with “.” as standard combinator); default is *mode*="async".

- `<ccs:Alternative>contents</ccs:Alternative>` stands for “ \sum ” and “+” (alternatives),
- `<ccs:Concurrent mode=“mode”>contents</ccs:Concurrent>` represents “ \times ” and “|” (parallel),
- The `<ccs:Event xmlns:lang=“uri”>`, `<ccs:Query xmlns:lang=“uri”>`, `<ccs:Test xmlns:lang=“uri”>`, and `<ccs>Action xmlns:lang=“uri”>` elements allow for embedded events, queries, tests, or actions (the latter even allow for embedding an action/process specification in another language).

Handling of “new” Variables in Fixpoint Processes. For integration with the MARS Framework that uses logical variables (that can be bound only once), variables that are bound during the evaluation of the fixpoint part must be considered to be local to the current iteration, and only the final result is then bound to the actual logical variable:

- `<ccs:Fixpoint variables=“var1 . . . varn” index=“j”
localvars=“list of variables”>
contents, usually containing the process variables
</ccs:Fixpoint>`

provides the markup for fixpoint constructs. The var_i are the process variables, j is the index of the one of the processes that is chosen, and the variables distinguished to be local can be bound in each iteration; after reaching the fixpoint they keep the value of the last iteration.

Example 31 (Variables in Fixpoint Processes) *Consider again Example 30(1). There, each iteration of the fixpoint process searching for a date where the room is available binds \$DateTime. The actual semantics is easy to understand and implement: just keep the last value.*

Grouping. For each subexpression, it can be specified if it is executed for the whole set, or separately for each tuple, or some grouping (in the same way as grouping in SQL) is applied. Clearly, subactions can only have finer granularity than the outer expressions). For specifying grouping, each action element has an optional attribute

`group-by=“variable list”`

that indicates grouping. E.g., given variables X, Y, Z, `group-by=“X Y”` means to execute the subexpression separately for all sets that have X and Y in common. Default is `group-by=“”` which means to have one group with all tuples. For convenience, `group-by=“-separately”` means to process every tuple separately, and `group-by=“-bulk”` also means to have one group with all tuples.

5.6.5 Examples in CCS

Example 32 (Basic CCS) *Consider a rule that does the following: if a flight is first delayed and then canceled (note: use of a join variable), make a reservation for each passenger at the airport hotel, and send each business class passenger an SMS.*

```
<eca:Rule xmlns:eca="http://www.semwebtech.org/languages/2006/eca-ml#" >
  <eca:Event xmlns:snoopy="http://www.semwebtech.org/languages/2006/snoopy#" >
    <snoopy:Sequence>
      <travel:delayed-flight flight="{ $flight}" date="{ $date}"/>
      <travel:canceled-flight flight="{ $flight}" date="{ $date}"/>
    </snoopy:Sequence>
  </eca:Event>
  <eca:Query>
    <eca:opaque language="sparql" domain="travel" >
      SELECT $booking, $name WHERE {
        $x rdf:type travel:flight .
```

```

    $x travel:code $flight . $x travel:date $date .
    $x travel:booking $booking .
    $booking travel:name $name }
  </eca:opaque>
</eca:Query>
<eca:Action xmlns:ccs="http://www.semwebtech.org/languages/2006/ccs#" >
  <ccs:Concurrent>
    <ccs:Action>
      <travel:reserve-room hotel="hotel uri" name="$name" />
    </ccs:Action>
    <ccs:Sequence>
      <ccs:Test >
        <eca:Opaque language="sparql" domain="travel" >
          ($booking travel:class travel:business-class)
        </eca:Opaque>
      </ccs:Test>
      <ccs:Query>
        <eca:Opaque language="sparql" domain="travel" >
          SELECT $phone WHERE
            ($booking travel:contact-phone $phone)
        </eca:Opaque>
      </ccs:Query>
      <ccs:Action>
        <comm:send-sms to= "$phone" >
          "we are very sorry ... and booked a room in ... for you"
        </comm:send-sms>
      </ccs:Action>
    </ccs:Sequence>
  </ccs:Concurrent>
</eca:Action>
</eca:Rule>

```

- after the event part, $\$flight$ is bound to the flight number,
- after the query, for each booking there is a tuple of bindings $\$flight$, $\$date$, $\$booking$, $\$name$.
- The action component does two things in parallel: reserve rooms, and for each tuple (i.e., for each booking) check if it is a business class booking (test), if yes, get the contact phone number (query) and send an SMS.

Example 33 (CCS with a fixpoint) Consider the following scenario of a university:

- if a student fails twice in a written exam (composite event), it is required that another oral assessment takes place for deciding upon final passing or failure.
- Action component of the rule: Ask the responsible lecturer for a date and time. If a room is available, the student and the lecturer are notified. If not, ask for another date/time.

The action component is expressed in CCS extended with events and queries as follows:

```

fixX.(ask_appointment($Lecturer,$Subj,$StudNo) :
  ∂ proposed_appointment($Lecturer,$Subj,$DateTime) :
  (available(room,$DateTime) +
  (¬ available(room,$DateTime) : X))) :
inform($StudNo,$Subj,$DateTime) :
inform($Lecturer,$Subj,$DateTime)

```

The XML Markup of the rule is then as follows:

```
<eca:Rule xmlns:uni="http://www.education.de" >
  <eca:Event> failed twice – binds $student ID and $course </eca:Event>
  <eca:Query> binds e-mail addresses of the student and the lecturer </eca:Query>
  <eca:Action xmlns:ccs="..." >
    <ccs:Sequence>
      <ccs:Fixpoint variables="X" index="1" localvars="$date $time $room" >
        <ccs:Sequence>
          <ccs:Action> send asking mail to lecturer </ccs:Action>
          <ccs:Event> answer binds $date and $time</ccs:Event>
          <ccs:Query> any room $room at $date $time available? </ccs:Query>
          <ccs:Alternative>
            <ccs:Test> yes </ccs:Test>
            <ccs:Sequence>
              <ccs:Test> no</ccs:Test> <ccs:Variable name="X" />
            </ccs:Sequence>
          </ccs:Alternative>
        </ccs:Sequence>
      </ccs:Fixpoint>
      <ccs:Action> send message ($date, $time, $room) to student </ccs:Action>
      <ccs:Action> send message ($date, $time, $room) to lecturer </ccs:Action>
    </ccs:Sequence>
  </eca:Action>
</eca:Rule>
```

5.6.6 Processing

Thus, the implementation of the CCS engine is only concerned with the actual CCS operators (i.e., the elements in the `ccs:` namespace. Atomic actions are executed “immediately” by submitting them to the domain nodes (if specified by an URL) or to a domain broker (see Chapter 8) that is responsible for the domain, forwarding them to appropriate domain nodes).

The handling of embedded `<ccs:Event>`, `<ccs:query>`, `<ccs:test>`, and `<ccs:action>` elements with embedded fragments of other languages is done in the same way as the evaluation of components by the ECA engine. This will be discussed in Section 9.5.

5.6.7 ECA Rules vs. CCS

In general, it is possible to decompose a CCS description completely into ECA rules with atomic actions (or with restricting the action component to BPA process specifications), or to express ECA rules as a special form of CCS processes over the above-mentioned atomic items. The advantage of supporting both formalisms in the MARS Framework lies in the appropriateness of modeling: there is behavior that is preferably and inherently formalized as ECA rules, and there is behavior that is preferably formalized in CCS. Providing both formalisms thus eases the modeling, and with this also the understandability and maintenance of behavior:

“If something happens (event) in a certain situation (condition), then proceed according to a given policy (action, as CCS process).”

Having ECA rules and processes allows to model both *reactive* and *continuous* behavior in an appropriate way.

5.7 Summary

So far, the component languages as “intermediate” level between the surrounding ECA language and the domain languages have been discussed. The Web architecture that provides the global cooperation and communication, i.e., given a rule, finding appropriate services for the components will be discussed in Section 9. The next section first deals with the domain level. The domain level is made up from autonomous domain nodes, e.g., airlines, train companies, and car rentals, and *domain brokers* that provide integrating functionality for domains.

Chapter 6

Domain Ontologies in the MARS Framework

Each Semantic Web application uses one or more domains. Often, an application (and its rules) has a core domain, and also touches several other domains.

Example 34 *Consider a travel agency. It “lives” in the travel domain, i.e., its behavior and user interface are dealing with this domain, and for serving this central purpose it communicates intensively with many other nodes in this domain. Additionally, it touches e.g. the banking domain, and sometimes also the medical domain (if travelers request information about mandatory or recommended vaccinations) etc.*

Currently, (domain) ontologies are usually described in OWL. Most people currently consider an ontology to be the description of the *static* concepts of the domain, but not of its dynamics. The situation is comparable to UML: the static part is the class diagram (which, in contrast to OWL ontologies, at least also describes the signature of the methods), while the dynamic notions are described by other diagrams.

In the same way, ontology descriptions in the MARS Framework have to be extended to cover the behavioral notions of events and actions; preferably also the semantics of actions in terms of changes and raised events must be described to allow for *reasoning* in such an ontology. In the following sections, we describe:

- generic ontology issues of domains,
- generic interfaces and functionality of domain nodes (Chapter 7); architecture samples are given in Chapter 10.
- communication and *domain brokering* issues (Chapter 8).

6.1 Domain Ontologies

As discussed in Section 3.1.3, a domain ontology is partitioned into (static) literal notions, named actions, and named events, as shown in Figure 6.1. Note that events and actions can also be structured into a class hierarchy (which is orthogonal to a “defined”-relationship between composite events/actions and simpler ones).

In the MARS Framework, domain namespaces (e.g., `travel`) are usually associated with an URL where an RDF document can be found that provides information about the domain. This consists of the RDF/RDFS/OWL definition of the domain ontology itself.

Example 35 (Definition of the Travel Ontology) *The travel domain ontology provides –among others– the following classes, properties, events and actions:*




Figure 6.1: Components of Domain Ontologies

```

<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:owl="http://www.w3.org/2002/07/owl#"
  xmlns:mars="http://www.semwebtech.org/2006/mars#"
  xml:base="http://www.semwebtech.org/domains/2006/travel">

  <mars:Domain rdf:about="#"/>
  <owl:Class rdf:ID="Airline">
    <rdfs:subClassOf rdf:resource="mars:DomainService"/>
  </owl:Class>
  <owl:Class rdf:ID="Car-Rental">
    <rdfs:subClassOf rdf:resource="mars:DomainService"/>
  </owl:Class>

  <owl:Class rdf:ID="Flight-Event">
    <rdfs:subClassOf rdf:resource="mars:Event"/>
  </owl:Class>
  <owl:Class rdf:ID="Flight-Action">
    <rdfs:subClassOf rdf:resource="mars:Action"/>
  </owl:Class>
  <owl:Class rdf:ID="Cancellation">
    <rdfs:subClassOf rdf:resource="mars:Event"/>
  </owl:Class>
  <owl:Class rdf:ID="Canceled-Flight">
    <rdfs:subClassOf rdf:resource="#Cancellation"/>
    <rdfs:subClassOf rdf:resource="#Flight-Event"/>
  </owl:Class>

  <owl:Class rdf:ID="Cancel-Flight">
    <rdfs:subClassOf rdf:resource="#Flight-Action"/>
  </owl:Class>

  <owl:Class rdf:ID="Delay-Flight">
    <rdfs:subClassOf rdf:resource="#Flight-Action"/>
  </owl:Class>

  <owl:Class rdf:ID="Fully-Booked">
    <rdfs:subClassOf rdf:resource="#Flight-Event"/>
  </owl:Class>
</rdf:RDF>

```

[Filename:travel.rdf]

With such an (event) class hierarchy, there could be a rule that informs about any cancellations of any means of transportation. Note that it is not possible to express that *has-flight* is a (mandatory) property of *flight-events*, *flights*, *flight-actions* etc.

Application services (e.g., airlines) use (one or more) domain ontologies, and they are supported by *application services* that are available at some URL. An application service can support multiple domains, e.g., an airline service will support *travel* and *business* (for running its business, paying taxes), etc. For supporting the work of domain brokers (see Chapter 8), the information about an application service specifies which notions of a domain are actually supported by the service (e.g., an airline supports *travel:airport*, but not *travel:cancel-train*).

6.2 Description of Application Services

Example 36 (Metadata on Application Services) Consider application services *Orafly* and *Exist-Cars* in the traveling area (running the respective business and being implemented on some infrastructure):

```
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:owl="http://www.w3.org/2002/07/owl#"
  xmlns:mars="http://www.semwebtech.org/2006/mars#"
  xmlns:travel="http://www.semwebtech.org/domains/2006/travel#">

  <mars:Service rdf:about="http://www.semwebtech.org/nodes/2006/orafly">
    <business:has-business rdf:resource="travel:airline"/>
    <mars:uses-domain rdf:resource="travel:"/>
    <mars:uses-domain rdf:resource="business:"/>
    <business:has-taxnumber>eur-0815</business:has-taxnumber>
    <mars:supports rdf:resource="travel:Airport"/>
    <mars:supports rdf:resource="travel:Flight-Connection"/>
    <mars:supports rdf:resource="travel:is-delayed"/>
    <mars:supports rdf:resource="business:has-taxnumber"/>
  </mars:Service>
</rdf:RDF>
```

[Filename:oraflly.rdf]

The behavior and interfaces of application services are described next.

6.3 Rules in Ontologies

There are several kinds of *defining* rules in ontologies:

1. logical derivation rules for deriving concept membership or instances of properties,
2. ECE rules: derive composite events from simpler ones (e.g., “flight 50% booked” from “if a booking occurs such that 50% are reached”),
3. ACA rules: map higher-level actions to simpler ones “book a return ticket to *X* on dates *A* and *B*” to “book a ticket on date *A*” and “book a ticket back on date *B*”. (note that this defines the meaning of “return ticket” without having an explicit concept of a “return ticket” in the ontology.)

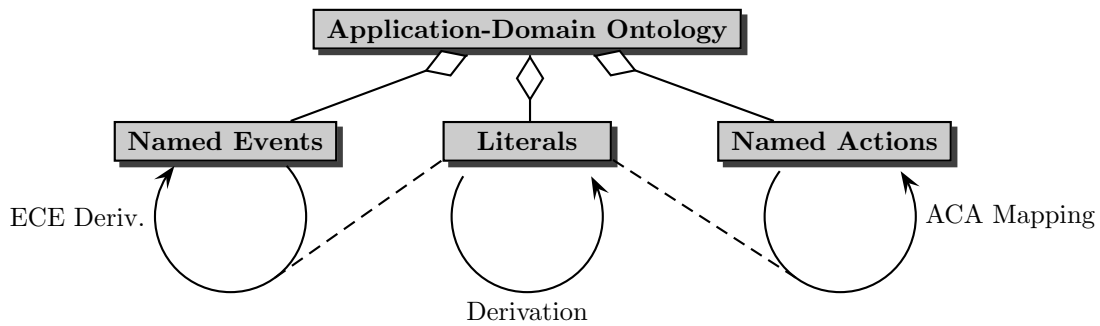


Figure 6.2: Derivation and Mapping Rules for Events, Literals and Actions

Note that some instances of such rules belong to the ontology, while others possibly only belong to certain services. Here, the rules that belong to the ontology are of interest.

ECE rules define derived, usually composite, events. ACA rules define composite actions. Both have a close similarity to ECA rules: both make use of component languages, and their implementation is preferably based on ECA rules. In both cases, the condition component often serves not only as a condition, but as a query part that extends the variable bindings (both by new variables, and also by duplicating tuples during this extension).

6.3.1 Derivation Rules

Derivation rules can be formulated as OWL axioms (subclasses, definitional axioms of classes as intersections/unions/restrictions of others), or by full (first-order) derivation rules. For the latter, there are several proposals:

- the RuleML initiative that works upon markup for derivation rules
- derivation rules in Jena.

The actual evaluation of such rules can be located in the domain broker and optionally also in the application node itself. Note that individual nodes can additionally have own local rules.

6.3.2 ECE Rules

ECE rules define derived, usually composite, events.

Example 37 (Derived Event) Consider an airline that raises the price for flights after 50% of a plane are booked. This can be done by a rule reacting on a derived event: “when 50% of the seats of flight number f on date d are booked then ...”. The derived event “50% of the seats of flight number f on date d are booked” has to be defined in the ontology, and it is raised by another ECE rule “when a flight is booked and this is just the booking that exceeds the 50% of the seats quota, then raise the event” *half-booked*.

```

<mars:definition syntax= "xml" >
  <mars:defined>
    <!-- pattern of the event to be derived -->
    <travel:half-booked flight= "$f" date= "$d" />
  </mars:defined>
  <mars:defined-as>
    <!-- E/C components how to derive it -->
    <eca:event>
      <travel:booking flight= "$f" date= "$d" />
    </eca:event>
    <eca:test>
  
```

```

    <eca:opaque language="xpath" >
      <!-- note: XML schema of the (local) database assumed to be known -->
      count($flight[@date="$d"]/booking) = $flight/id(@aircraft)/@number-of-seats div 2
    </eca:opaque>
  </eca:test>
</mars:defined-as>
</mars:definition>

```

Note that event types *event-type* that are defined as derived events also have to be “declared” in the OWL part of the ontology as `<event-type rdf:type mars:Event>`.

The implementation of derived events in course of event brokering will be discussed in Section 8.2.

6.3.3 ACA Rules

ACA rules define complex actions in terms of simpler ones. ACA rules exist on two levels: (i) abstract specification of composite named actions as processes over simpler named, still abstract actions, and (ii) local implementation of named actions by operations on a data model.

6.3.3.1 Definitory ACA Rules on the Ontology Level

Composite abstract named actions of the ontology can be specified as processes consisting of less abstract named actions. In this case, the specification is actually of the same structure as the action components of ECA rules (and the query component is also the same as for ECA rules). Such rules correspond to *definitions* of higher-level actions in the ontology. They usually hold ontology-wide.

```

<mars:Definition syntax="xml" >
  <mars:defined>
    <banking:Money-Transfer amount="$amount" from="$from" to="$to" />
  </mars:defined>
  <mars:defined-as>
    <eca:Action>
      <ccs:Concurrent xmlns:ccs="..." >
        <banking:debit account="$from" />
        <banking:deposit account="$to" />
      </ccs:Concurrent>
    </eca:Action>
  </mars:defined-as>
</mars:Definition>

```

Note that action types that are defined as composite actions also have to be “declared” in the OWL part of the ontology as `<action-type rdf:type mars:Action>`.

The implementation of composite actions in course of action brokering and ACA rules will be discussed in Section 8.4.2.

6.3.3.2 Implementational ACA Rules from the Ontology Level to the Data Model Level

Application Domain nodes receive action invocations that have to be carried out. Usually, such actions are implemented by classical Web Service calls: invoking an action at a node means invoking a method of a Web Service that starts a Java (or perl or PHP or whatever) execution. The interface and implementation of such application nodes is then very specific to a certain domain.

In the MARS Framework, we target to support a *generic* architecture which is also continued at the domain node level: Atomic actions are represented by XML (or RDF fragments) that are sent to domain nodes that implement these actions. In Section 7.3.2, we show how such a generic domain node architecture supports ACA rules in this lowest abstraction level.

6.3.4 Discussion: Comparison with RuleML Proposal

RuleML proposes a rule markup as `ruleml:imp`, `ruleml:head`, `ruleml:body` for representing rules. We did adhere to this for the following reasons;

- `imp`, `head`, `body` are clear for derivation rules and for event definition rules (since these are also a kind views), but *not* for ACA rules.
- `imp`, `head`, `body` refer to syntactic notions of the representation of rules as formal languages grammars, *not* to their semantics as derivation or reduction rules.

Making rules *full citizens* of an ontology (like `owl:class` and `owl:property`) allows for a more appropriate modeling wrt. the different kinds of rules as

- derivation rules (logical rules, ECE rules),
- reduction rules (ACA rules),
- integrity rules as assertions.

Chapter 7

Domain Application Nodes

Domain application nodes are the “leaves” of the Semantic Web architecture. They represent services in the application domain, e.g., airlines, train companies, car rentals, universities. As such, they usually have some kind of database or knowledge base. According to their task as services, they offer to execute actions (e.g., answering queries, booking a flight, or registering as a student). So far, they are closely related to common Web Services. When considering reactive behavior in the Semantic Web for the MARS Framework where *events* are also of central importance, domain nodes should also support this by *emitting* events according to the events defined by the domain ontology (e.g. that a flight is now fully booked). Nodes that do not emit events are “sinks” of behavior – which is often sufficient (e.g., for a car rental). Then, they may have local (reactive) behavior, but nothing reacts externally upon them – at least not directly; they still can be queried and by this participate in processes.

7.1 External Interface of Domain Nodes

7.1.1 Providing Static Data

Nodes serving as databases or knowledge bases can have a rather primitive functionality: providing data for lookup (that can be updated by bulk load, e.g., once per month).

Files. The most simple domain “node” is just a data source that can be read as a whole. Such standalone XML or RDF files on the Web can e.g. be accessed via a generic XQuery or SPARQL wrapper service:

- XQuery: `let $doc := document("url") ...`
- SPARQL: `select ... from url where ...`

Opaque Queries. While the above files are the simplest kind of “Web resources”, many data sources in the current Web provide a simple querying interface via HTTP GET or POST (e.g., an SQL, XPath, XQuery, RDQL, or SPARQL query interface). Wrt. the MARS Framework, such queries are handled as opaque queries whose answer are e.g. sets of XML nodes. The ECA engine deals with this via a generic wrapper for HTTP or XQuery etc. (cf. Section 10.1.2).

Framework-Aware Querying. In full exploitation of the MARS Framework, also the query component will be stated in structured way by an algebraic query language in XML or RDF markup. Only literal queries will occur at the domain nodes (in RDF/OWL: for extension of classes and relationships). The communication (especially the formatting of the results) is then done on the communication format given in Section 4.1.4.

7.1.2 Executing Action Requests

Behavior for Web nodes means usually to be able to process actions. Full-fledged Semantic Web domain nodes support actions of the corresponding domain ontologies (i.e., things $?X$ such that $(?X, \text{rdf:type}, \text{mars:Action})$ holds) that are communicated according to the downward communication format given in Section 4.1.4, consisting of the action to be executed as an XML fragment, and optionally tuples of variable bindings.

For nodes that are not framework-aware, actions can be given as opaque code in the local programming language (e.g., PL/SQL or XUpdate) or method calls (using SOAP markup; methods must be implemented locally as procedures or by ACA rules). For such nodes, wrappers can be used for adapting them to the framework level.

The incoming (atomic) action requests must be mapped to the local knowledge base as an internal behavior of the node.

7.1.3 Reporting About Behavior: Providing Atomic Events

The current Web architecture does not consider “events” communicated by *push* communication. For the MARS Framework, nodes are expected to emit atomic events (i.e., things $?X$ such that $(?X, \text{rdf:type}, \text{mars:Event})$ holds). They are communicated in XML format as described in Section 5.2.1, or later in RDF format.

Note that there are many legacy nodes that do not support events. They have to be monitored to detect their events. Communication of “events” or “news” is currently partially supported by *RSS Feeds* that provide events for *pull* communication (see Section 8.2.1).

In the MARS Framework, events are forwarded to domain brokers (as separate services) where filtering and communication to the outside takes place.

The actual emission of events upon changes in the local database must be provided by internal behavior of the node.

7.1.4 Summary

Thus, domain nodes provide a *generic*, application-domain-independent interface:

- receiving queries; answering them,
- receiving action requests, executing them,
- administrative stuff: registering an event broker.

The actual architecture that satisfies these requirements will be discussed next.

7.2 Domain Node Local Behavior

The internal behavior of nodes can be black-box behavior, implemented in any programming languages, or given by ACA, ECA and ECE rules (often by database-level triggers). The information flow between events and actions is depicted in Figure 7.1 and contains the following types of rules:

1. low-level ECA rules as triggers for local integrity maintenance in a knowledge base,
2. ECE rules: derive and raise application-level events based on internal changes knowledge base level,
3. ACA mapping: map high-level actions to lower-level (e.g. *INSTEAD OF* triggers).

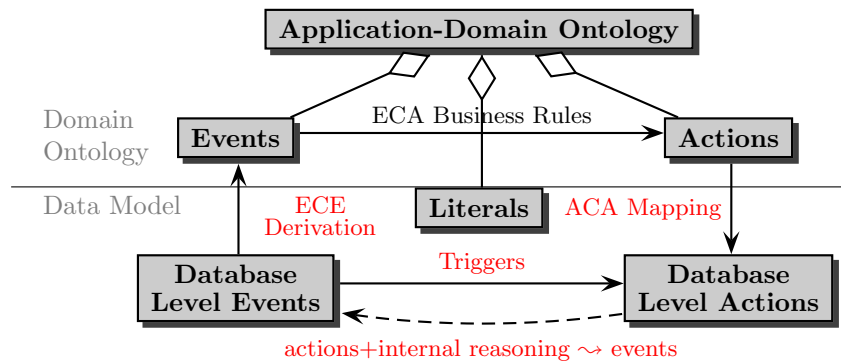


Figure 7.1: Interference of Events, Actions, and Literals

7.3 Domain Node Architecture Proposal

The Jena [38] framework provides an API for dealing with RDF and OWL data. In our architecture, Jena uses an external database (e.g. PostgreSQL) for storing base data. If Jena is used for RDFS or OWL data, an appropriate reasoner can be used. For this, Jena provides both a (restricted) internal reasoner and it can use a separate DL reasoner like or Pellet [64] via the DIG interface [26] conventions or via a Java API. The core domain node architecture is based on a service using the Jena framework with PostgreSQL and Pellet. The handling of ontology-level actions and mapping them to the RDF model level is located in a wrapper around the core as shown in Figure 7.2.

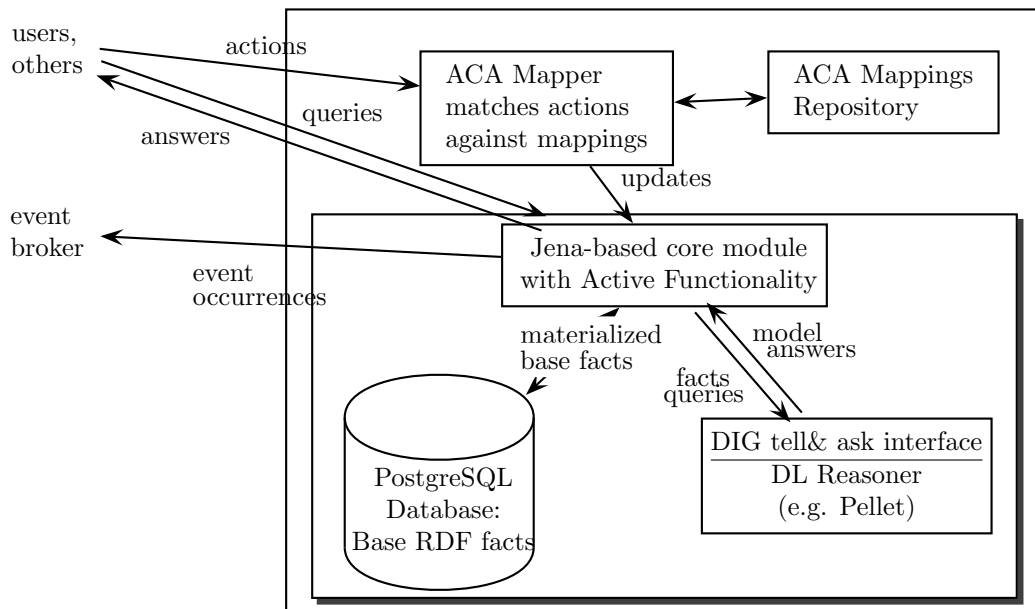


Figure 7.2: Architecture of the Domain Node

7.3.1 Functionality of the Node Core

RDF Database. The base of the node is a pure RDF database, as e.g. provided by the on Jena [38] framework with an underlying relational database. It provides storage and SPARQL querying.

OWL Knowledge Base. Using the Jena framework and activating the internal reasoner or connecting to an external OWL reasoner, like e.g. Pellet [64], OWL reasoning over the RDF data is provided.

RDF/OWL Knowledge Base with Updates and Triggers. The next step is then to provide a simple update language for RDF data and to support database triggers reacting upon *database update actions* on the physical level. With this, local ECA rules *inside* the database (that are also required to support actual updates, e.g., when deleting a property that is symmetric) as well as ECE rules raising events can be implemented. An RDF/OWL domain node with such local active behavior based on Jena [38] has been implemented and is described in [55].

7.3.2 Functionality of the Node Wrapper: Mapping of Actions

Ontology-level actions that are requested from the outside arrive in the agreed communication format. Such action requests can e.g. be

```
<travel:delay-flight flight="LH123" time="30 min" />
```

or, containing tuples of variable bindings,

```
<travel:delay-flight flight="{flight}" time="{min}" />
<logvars:variable-bindings>
  <logvars:tuple>
    <logvars:variable name="flight" >LH123</logvars:variable>
    <logvars:variable name="time" >30 min</logvars:variable>
  </logvars:tuple>
  <logvars:tuple>
    <logvars:variable name="flight" >AL400</logvars:variable>
    <logvars:variable name="reason" >45 min</logvars:variable>
  </logvars:tuple>
</logvars:variable-bindings>
```

or as a complex XML structure

```
<library:add-book>
  <library:title>Algorithms</library:title>
  <library:author>Thomas Cormen</library:author>
  <library:author>Charles Leiserson</library:author>
  <library:author>Ronald Rivest</library:author>
</library:add-book>
```

In a domain-specific Web Service, such methods are implemented by a rules that update the knowledge base accordingly, similar to INSTEAD-triggers that decide if (and optionally how) they are executed. An instance of the generic node architecture is instantiated with a set of such *mappings* which are then applied to the incoming actions.

Example 38 (Domain Node Action Implementation) *Consider the action of delaying a flight. This can be implemented by an SQL-based node as follows:*

IMPLEMENT

```
<travel:delay-flight flight="{flight}" time="{min}" >
  <travel:reason>$reason</travel:reason>
flq/travel:delay-flight>
```

```
WHERE EXISTS (SELECT *
              FROM Flights
              WHERE No = getFlightNumber($flight)) % $flight is a URI
```

```

BY
LET $No = getFlightNumber($flight)
UPDATE Arrivals
SET time:=time+$min WHERE flight=$No
INSERT INTO DELAYED
VALUES($No, $min, $reason)
// and ignore delays of all other flights

```

For an RDF-based node, the implementation could be:

```

IMPLEMENT
<travel:delay-flight flight="{flight}" time="{min}" >
  <travel:reason>$reason</travel:reason>
flq/travel:delay-flight>
WHERE ($flight-uri, travel:operated-by, our-url) % flight is operated by us
BY
UPDATE ($flight-uri, travel:leaves, $time, $time+$min)
ASSERT ($flight-uri, travel:flight-is-delayed, $time)
ANNOTATE WITH (travel:reason, $reason)
// and ignore delays of all other flights

```

More complex actions, e.g., adding multiple (*book* has-author *author*) statements in a single action for a book that has multiple authors (while another internal model would map the same facts by using a blank node) requires to express more complex actions.

In general ECA rules (with an action part e.g. using CCS over atomic database updates) could be used. This would require to have local service instances for ECA-ML, an AED, and CCS.

Prototype Implementation The prototypical node developed in [55, 73] has been extended for ACA Rules in [74].

The extension defines an XML Markup for RDFUpdate actions. With this, the mapping from actions (that are atomic from the ontology level) to a sequence of database update statements for a given XML input can be seen as a transformation that can –given actions in XML markup – e.g. be expressed in XQuery or XSLT.

The above rule for `<library:add-book>` can e.g. be expressed as above:

```

let $title := /add-book/title/text() let $book :=
library:make-book-uri($title) return
  INSERT ($book has-title $title)
  for $authorname in /add-book/author/text()
  let $author := library:make-author-uri($authorname)
  INSERT ($book has-author $author)

```

As illustrated above, often the mapping includes to generate URIs for entities. An ontology may provide a function that generates (based on a key attribute) a URI that guarantees that all nodes using this ontology identify a thing by the same URI. In the same way, this function can be local (using e.g. a local prefix), when e.g. a university generates local URIs for students.

7.3.3 Installing and Running a Domain Node

A domain node for a certain domain ontology is set up by instantiating the generic architecture with a set of such ACA mappings. Note that all nodes whose internal database schema is the RDF model of the ontology use the same ACA mapping (that should thus be given when designing the ontology). In this case, also the RDF triggers for raising basic events upon database updates can be given with the ontology.

The domain node provides a method (for local use by its maintainers) where ACA mappings can be submitted. When a mapping is submitted, its language must be indicated:

```
<applnode:aca-mapping language="xslt|xquery" >  
opaque code in the indicated language  
</applnode:aca-mapping>
```

Note that when installing a new domain node, it must usually be registered at one or more domain brokers.

Processing Incoming Actions. For every incoming action, all mappings are applied and the resulting updates are executed on the database. The internal database-level triggers support the actual execution of (intensional) updates and consistency maintenance. They also raise events that are communicated to the registered domain/event brokers.

Outlook: RDF Model. When migrating to a complete RDF environment, events may be exchanged as RDF data. Then, the above mapping can be completely implemented by SPARQL – querying the event and the database and updating the database.

Optional: Language Check Upon registration of a mapping, the node could try to determine the language by its own, trying XSLT and XQuery, and, in later versions, SPARQL.

Chapter 8

Domain Brokering

8.1 Event Brokering

Event Brokering functionality provides the mediation between event providers (i.e., domain nodes) and event consumers. The latter are the AEMs as the mediators towards the ECA engines: they take events and produce answers.

The communication of events on this level is depicted in Figure 8.1. The main communication between DNs and EBs is very simple: A DN forwards all events (as XML fragments according to the domain ontology) to its peer EB(s).

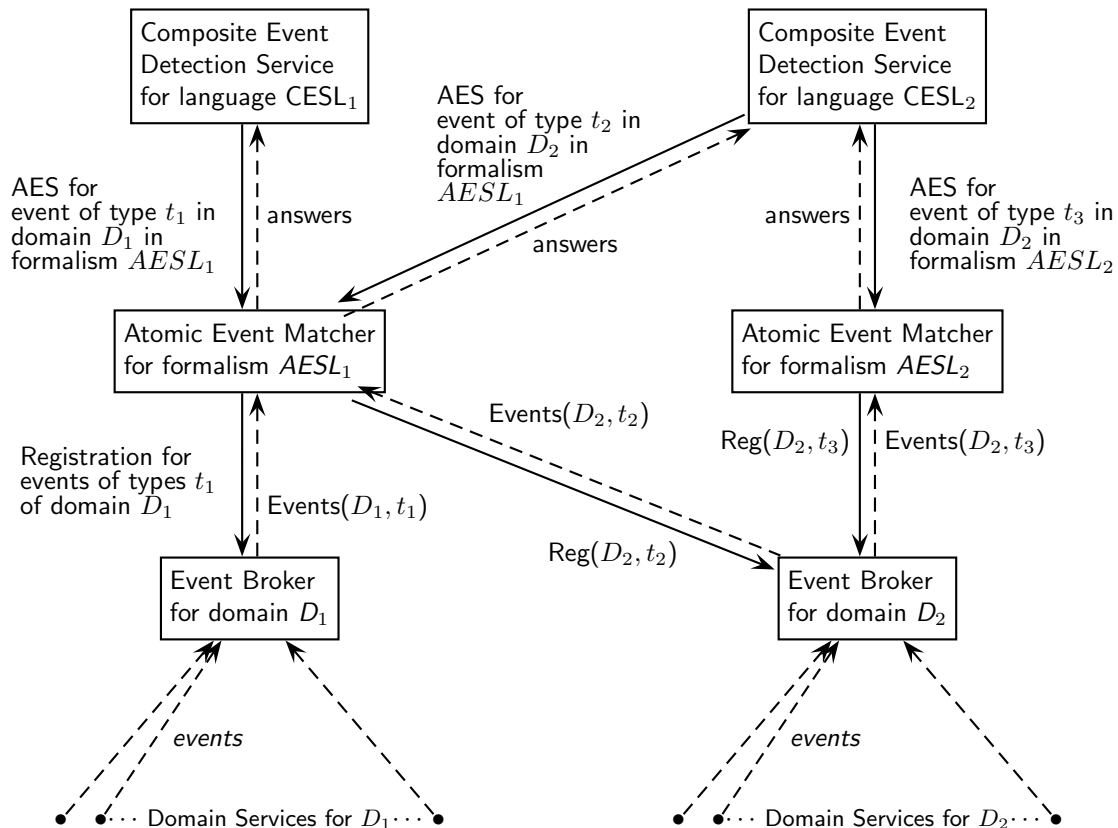


Figure 8.1: Architecture: Communication with Event Brokers

Communication of Events between AEM and EB. As described in Section 5.3, the CEDs (and in case of an atomic event component also the ECA engine) register AESs at the AEMs to be detected. The AEMs have to be aware of all relevant events. For this, they determine the domain (URI) of the event to be detected. With this URI, they can find out at the *Domain Services Registry (DSR)* (see Section 9.6) what EBs are serving this domain. The AEM then tells an EB to forward them the required events. The specification which events should be forwarded contains the domain name/URL (note that an EB may support more than one domain) and optionally also the type of the event.

The registration message follows the patterns discussed in Section 5.1, e.g.,

```
to: service-URL of the EB
<register>
  <Reply-To>URL where the events are expected at the AEM </Reply-To>
  <domain>domain-uri</domain>
  <event-type type>event type</event-type>
</register>
```

Example 39 (Registration at the Event Broker) Assume an AEM for the sample XML-QL-style event matching formalism gets the following AES to be detected:

```
<travel:canceled-flight xmlns:travel="http://www.semwebtech.org/domains/2006/travel#"
  number="{flight}" />
```

It then explores which EB serves the travel domain (for details, see Chapter 9). Using the fact that the root element of the event is *travel:canceled-flight*, it tells the EB to forward all *travel:canceled-flight* to it.

Note that when using more involved ontologies and RDF/OWL-based AES formalisms, event classes and subclasses can be defined. In such cases, simple matching of element name is not appropriate, but OWL reasoning has to be applied.

From this moment on, the EB forwards all events of this type to the AEM. In case an AES is deregistered, the AEM can also deregister (note that for this, it has to do bookkeeping if other AESs still use the respective elements).

Optional and Additional Functionality of Event Brokers. Note that the *Event Brokers (EBs)* can optionally provide AEM functionality for some formalisms. Since every service registers its functionality in the LSRs, this will be exploited correctly when the CED is searching for an AEM to match certain AESs.

There can be domain nodes that do not provide events. In this case, event brokers can also apply *continuous-query-event (CQE) rules* for detecting events by monitoring the respective sources.

8.2 Brokering of Derived Events

Derived events are specified as ECE rules whose *defined-as* subelement consists of a (composite) event specification and an optional query and/or test. From these, an ECA rule that explicitly raises the derived event can be generated in a straightforward way:

```
<eca:rule>
  <!-- contents of the body of the ECE rule definition -->
  <eca:action>
    <eca:raise-event>
      <!-- head of the ECE rule -->
    </eca:raise-event>
  </eca:action>
</eca:rule>
```

Note that there can be event types that are both directly supported by some domain nodes (in most cases, these have their own rules to derive them), and for which the ontology provides an ECE derivation rules.

8.2.1 RSS-based Event Brokering

Event Brokering for a given domain may include to raise events obtained from RSS feeds (e.g., from bioinformatics sources). This can already be done in an ECA-way, using an ontology of RSS feeds:

The broker polls the feed regularly. It is then processed by removing items that have already processed. Then, for each new item an event

```
<rss:new-feed-item url="the feed-url">one item</rss:new-feed-item>
```

is raised. From that, the contents (mainly the `description` element) is used in an *application-specific* way to raise application-dependent events. For this, it is necessary that the event part of the application ontology covers the events that are reported by the RSS feeds.

Remark *This will be applied in a Bioinformatics case study.*

8.3 Query Brokering

Clients can send queries as request either to certain domain nodes (especially opaque queries), or to the domain broker. The domain broker knows the respective ontology that consists of

- OWL statements and
- rules.

The domain broker is responsible as a mediator to answer the query by using application services. For this, a lot of algorithms for mediating and integrating information can be applied. To get the infrastructure running, a simple approach is followed first.

Prototype: Simple Approach

Example 40 *Consider a query for “connections from Göttingen to St. Malo”. The ontology specifies that “connection” is the transitive closure of train, flight, and ship connections.*

Thus, the broker can ask for all “connection” instances in any domain node, and then try to combine them.

Decomposition of the Query. Decomposition of a query means to collect all static notions (concepts and properties) that are relevant for answering the query. Note that for forwarding queries, declarations of

- `owl:inverseOf` and
- `owl:equivalentProperty`

must be considered. Moreover, if the ontology contains rules of the form $head \leftarrow body$, and the notion in the head is asked, notions occurring in the bodies have to be answered. For this, the rules have to be given in an appropriate markup (e.g., RuleML).

The relevant facts about a rule can be represented in RDF as follows (derivable from RuleML markup):

```
[Draft]
<rule, defines, notion>
<rule, uses, notion>
```

Selecting nodes to be queried. When the decomposition is computed, each of the notions is forwarded to relevant nodes, i.e., all services known to the domain broker that `mars:support` the notion.

Combining the answers and answering the original query. The domain broker collects the answers and by this obtains all instances of relevant class memberships and properties (as RDF triples). It takes the union of this (as a local RDF knowledge base) and answers the original query.

8.4 Action Brokering

Clients can request actions of domain ontologies (i.e., which are instances of subclasses of `mars:Action`) either at certain domain nodes, or at the domain broker:

- Actions that are atomic wrt. the application domain level. These actions are supported by the respective domain nodes, to which they are forwarded by the Domain Broker.
- Actions that are defined by a composition of simpler actions (using a global, ontology-defined ACA rule). The Domain Brokers are responsible for handling the ACA rules.
- note that also non-atomic actions can be supported by domain nodes if they have only local consequences.

8.4.1 Required Information for Action Brokering

As described above, the Domain Broker information (i) about which domain nodes support which (atomic) actions, and (ii) about the definitions of complex actions.

Atomic Actions. For atomic actions, in the same way as for atomic queries, the domain nodes declare by `mars:support` which actions they support. When a Domain Node registers at a Domain Broker, the support information is stored at the broker. Atomic actions are then forwarded to relevant nodes, i.e., all services known to the domain broker that `mars:support` the notion. There, they are processed as described in Section 7.3.2.

Definitions of Complex Actions As described in Section 6.3.3, the domain ontologies contain ACA Rules that define the reduction from complex actions to simpler ones (e.g., by giving a CCS expression). When initializing a Domain Broker with an ontology, these rules are stored in the broker. ACA Rules have a close similarity to ECA rules, they are actually handled by an architectural extension of ECA engines and appropriate communication; see Section 8.4.2.

Data-Dependent Forwarding Actually, most actions do not concern all nodes that `mars:support` that action type, but depending on some (key) data, only some of them.

Example 41 Consider the case that the domain node representing Frankfurt Airport decides that a given flight has to be delayed by one hour due to bad weather conditions, e.g., by a rule

```
<eca:Rule xmlns:travel="http://www.semwebtech.org/domains/2006/travel#" >
  <eca:Event>bad snow conditions detected </eca:Event>
  <eca:Query>all $flights departing in the next hour </eca:Query>
  <eca:Action>
    <travel:delay-flight code="{ $flight }" delay="1h" >
      <travel:reason>bad weather conditions</travel:reason>
    </travel:delay-flight>
  </eca:Action>
</eca:Rule>
```


Note that the travel ontology contains a triple

(travel:DelayFlight, rdf:type, mars:Action)

that indicates that this is indeed an action. Consider the action instance

```
<travel:delay-flight code="LH123" delay="1h" >
  <travel:reason>bad weather conditions</travel:reason>
</travel:delay-flight>
```

or, using an RDF URI [means, we have to go somehow from XML events to RDF events]

delay-flight(iata://flights/lh123, "1h", "bad weather conditions")

The *travel:delay-flight* action is *mars:supported* by Domain Nodes that represent airlines, but e.g. not by the train companies and hotels):

(http://lufthansa.com, mars:supports, travel:delay-flight)
(http://airfrance.com, mars:supports, travel:delay-flight) etc.

Thus, the above atomic action is forwarded to these nodes.

Actually, the cancelation of LH123 only concerns the Lufthansa node (the others will ignore it).

In a first approach, the update is sent to all nodes that support the corresponding predicate. The nodes must then decide based on the actual data, if they are actually concerned.

Note that the reader knows that the update does only concern the *Lufthansa* airline that actually operates LH123, but this requires not only to use metadata, but also data. For implementing this, the ontology must for each action (type) specify, which are the relevant nodes

```
(?A,has-relevant-node, ?Airline) :-
  (?A, rdf:type, travel:delay-flight),
  (?A, talks-about, ?Flight),
  (?Flight, travel:operated-by, ?Airline).
```

If such a specification exists in the ontology, it can be used. Otherwise the “broadcast” as above has to be done. Note that (?Flight, travel:operated-by, ?Airline) cannot be answered from the action only, thus either *all* nodes have to be asked if they operate this flight, or such “key” information must also be present in the domain broker (thus, it is quite useful that the broker also maintains a knowledge base with unchanging domain-dependent knowledge.

Actions vs. Event-Driven Architecture In an event-driven architecture, it is recommended to replace the action by raising an event “this must be done” on which the corresponding domain nodes react.

Example 42 (Actions vs. Events) Consider the following case: at some airport the weather conditions are forecasted to become bad, so that incoming flights cannot land. There is then a rule “if the forecast is ... then for all flights landing in the afternoon, cancel these flights”. All flights that are concerned can easily be selected from the flight schedule. Since these are operated by different airlines, the action “cancel-flight(\$flightno)” must be directed to several targets (airlines).

Here it is better to extend the domain ontology by an event *must-be-canceled(\$flightno)* on which the respective airlines can react.

8.4.2 Handling of Composite Actions by ACA Rules

As discussed in Section 6.3.3, ACA rules are a suitable paradigm for expressing actions on a higher abstraction level that are defined as composite actions (e.g., defining a money transfer as a debit followed by a deposit). ACA rules usually use only a single domain, but can also be extended to multiple domains.

The structure of ACA rules is closely related to ECA rules: the information flow between ACE and ECE by variable bindings is the same, and the C and E components are the same as in ECA rules (although, the C component is often empty).

Even the “ON .. DO” structure is the same: on *invocation of an action* and on *occurrence of an event*. ACA rules are triggered only upon atomic invocations (requesting some action (name) with parameters). Thus, a comparison with ECA rules with atomic events is appropriate. The specification of the invoking action is expected to use the same mechanisms as for atomic event specification. Thus, the AEMs can also be employed here.

For handling ACA rules, the ECA engine architecture can be used. When a domain broker is initialized with an ontology, it registers all ACA rules of the ontology at an ACA-aware ECA engine. Actions that are defined via ACA rules are then processed similar to atomic event patterns and atomic events:

ACA-Aware ECA-Engines. The ACA rules use a dialect of ECA-ML:

- `eca:aca-rule` for ACA rules,
 `<!ELEMENT aca-rule (%variable-decl,define-action,query*,test?,action+)>`
- `eca:define-action` for the action that is defined by the rule.

(Note that ACA rules with a markup as ECA rules will also be processed correctly since the operational semantics is the same).

Actions as Events. The domain broker registers ACA rules at an ACA-aware ECA engine. The ECA engine registers the AAS (atomic action specification, analogous to AES) at an AEM. The AEM registers at one or more domain brokers (for the domain, or for *domain:action-name*), and the domain brokers submit these actions like events to the AEM. Thus, only the domain broker must be aware if a registration by the AEM is concerned with an event or an action (which is known by the ontology). The ACA-aware ECA engine is aware whether it deals with an ECA or an ACA rule, but it does not make a difference in the processing.

Actions separated from Events. An alternative processing can be applied in a local, restricted environment: the ECA engine provides a separate method for registering ACA rules. The domain broker registers its ACA rules at such an ACA-enabled ECA engine. The ACA-enabled ECA engine registers the AAS (atomic action specification, analogous to AES) at an AAM (which is similar to an AEM, but does only matching for a given formalism and does not register the action at a domain broker). The ECA rule remembers the AEM. At runtime, the domain broker sends actions that are defined via ACA rules to the ACA/ECA-engine which forwards them to the AEMs.

Again, the AAM could be implemented as an extension of an AEM which distinguishes between events and actions by providing separate interfaces.

Discussion.

- as long as for each domain, there is only one domain broker, this is sufficient.
- if there are multiple domain brokers that support an ontology, *each* of them would register the rule. In the same way, each of them would be told to execute a composite action, and each of them would notify the ACA engine (yielding *n* invocations; even worse if they employ different ACA/ECA services).

The issue of duplicating rules, events and actions will be discussed in Section 9.7.

Chapter 9

The MARS Architecture, Ontology, Language and Service Metadata

In this section, we first describe (in fact, summarize, because its parts have already been described in the previous sections) the Web-Service-based architecture where each language is associated with one or more Web Services that are “responsible” for the language. Afterwards we will then systematically describe how the architecture is realized.

Rules are registered at some *rule evaluation service*, here, an ECA engine. The rule evaluation engine then manages the actual handling of rules based on the namespace URI references or the **language** attributes. As described above, every component (i.e., events, conditions, and actions) carries the information of the actual language it uses in its `xmlns:namespace` URI (note that this even allows for nested use of operations of *different* event algebras). Via *Languages and Services Registries (LSRs)* (cf. Section 9.5.1), the URLs and communication details of language processors can be found out. In Section 9.1, we describe a general Web-wide architecture based on autonomous Web Services. Section 9.2 describes variants that can especially be applied in “closed” environments.

Section 9.3 provides the formal ontology of concepts in MARS (that actually collects and orders the concepts discussed in the previous sections). Section 9.4 specifies the abstract interfaces of the types of services, and Section 9.5 how the concrete interfaces are described, and how the actual communication takes place.

9.1 Architecture and Processing: Cooperation between Resources

9.1.1 Event Detection

For event detection, at least two resources (or services) must cooperate: Event detection splits into matching the specification against events (*event algebra* and *atomic event matching*, and obtaining the relevant atomic events. Thus, the event component processor must be aware of the relevant atomic events.

When a rule or an event specification is submitted for registration, this has to be accompanied by information which resource(s) provide the atomic events (e.g., “@snoop: my bank is at *uri*, please supervise my account and tell me if a composite event *ev* occurs), or the detection service even has to find appropriate event sources (by the namespaces of the atomic events). The detection service then contacts them directly. The latter is e.g. appropriate for booking travels where the client is in general not aware of all relevant events (e.g., “@snoop: you know better than me who

is well-informed about events relevant for traveling, please detect the event ev_{travel} for me”), as illustrated in Figure 9.1: A client registers a rule (in the `travel` domain) at the ECA engine (Step 1.1). The ECA engine again submits the event component to the appropriate CED service (1.2), here, a SNOOP service. The SNOOP engine looks at the namespaces of the atomic events and sees that the `travel` ontology is relevant. The SNOOP service registers all atomic event patterns at the appropriate AEM (1.3). The AEM contacts a travel event broker (1.4) who keeps it informed (2.2) about atomic events (e.g., happening at Lufthansa (2.1a) and SNCF (2.1b)). The AEM matches the events against the registered patterns, and in case of a success, reports the matched event and the extracted variable bindings to the SNOOP service (3). Only after detection of the registered composite event, SNOOP submits the result to the ECA engine (4).

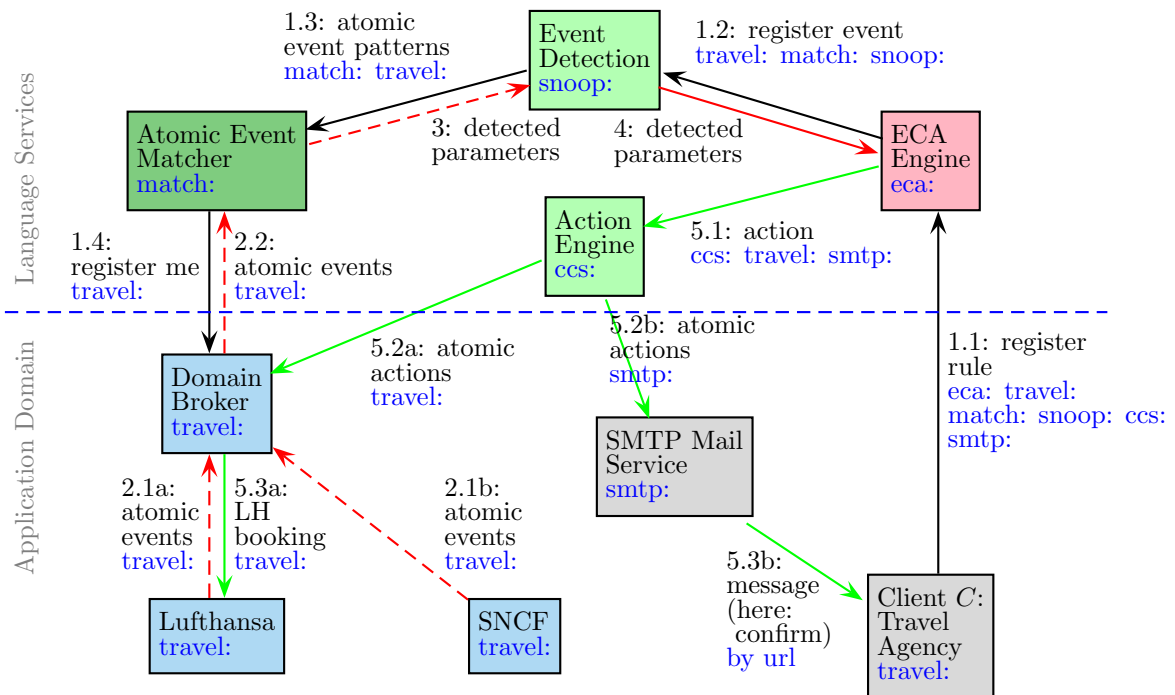


Figure 9.1: Communication: Event Processing

9.1.2 Query Processing

For queries, the opaque case using a common query language like XPath or XQuery on a given URL (XML document) is frequent. These query languages are often directly supported by the data sources. Otherwise, a “free” query service must be used. The handling of opaque queries in the prototype is described in Section 10.1.2.

For the non-opaque case, the basic schema is similar to event processing, consisting of the algebraic part and querying the actual concepts. Since every concept comes from an ontology (with namespace), the appropriate domain broker can be contacted. The graphics in Figure 9.1 does not show the evaluation of a query component.

9.1.3 Action Processing

The basic schema is symmetric to event processing, consisting of the algebraic part and executing the actual actions (which requires to determine *where* the action has to be executed): The action is submitted to an appropriate action language service (here: CCS (5.1)) that in turn submits the atomic actions to domain broker (`travel`, 5.2a) which forwards the actions to the respective

domain nodes (5.3a), and to the domain-independent services, here `smtp` (5.2b) which then sends a message (5.3b).

9.2 Architectural Variants

In addition to the above architecture, multiple intermediate variants are possible that combine functionality. In general, these are obtained by changing the communication paths and/or combining functionality in a node. The following paragraphs list (not necessarily in a systematic way) some variants.

Domain-Broker/Application-Centered. An application or portal can provide brokering service *together* with (selected) E, Q&T, and A languages (note that portals often already provide a query interface).

For event detection, a portal can offer to “process” composite event specifications. The client then submits its composite event specification to the portal. For processing it, it can either directly implement a CEL, or employ another service. The main point is that the portal is aware of all relevant events in the application domain and can feed them directly into the event detection.

Example 43 (Banking with Event Detection) *A bank can e.g. offer such functionality. Then, customers can place their composite events there and say “@bank: please trace the following composite event in language L on my account (and employ a suitable event detection service for L)”.*

The same holds for process specifications in the action part, when all actions have to be executed on the same portal or even on the same node.

Integrating AEM Services. For the actual location of the atomic event detection, there are again several alternatives: The separated architecture uses separate *Atomic Event Matchers (AEMs)* that implement an atomic event specification formalism, and that themselves are informed about the events by the application services. Instead, this “simple” matching can also be integrated either with the CED or with the domain broker:

- application services (e.g. for travel:) provide matching functionality (having an AEM for some specification formalism), or
- domain/event brokers (e.g. for travel:) can provide matching functionality (also having an AEM for some specification formalism).

ECA-Engine-Centered (without Domain Broker). In this case, there is no domain broker, but the ECA engine plays the “central” role (see Figure 9.2): clients *C* register rules to be “supervised” at the ECA service *R*. For handling the event component, *R* reads the language URI of the event component, and registers the event component at the appropriate event detection service *S* (note that a rule service that evaluates rules with events in different languages can employ several event detection mechanisms).

During runtime, the clients *C* forward all events to *R*, that in turn forwards them to all event detection engines where it has registered event specifications for *C*, amongst them, *S*. Relevant events from outside can be “imported” from an external domain broker.

S is “application-unaware” and just implements the semantics of the event combinators for the incoming, non-interpreted events. In case that a (composite) event is eventually detected by *S*, it is signalled together with its result parameters to *R*. *R* takes the variables, and evaluates the query&test (analogously, based on the respective languages), and finally executes the action (or submits the execution order to a suitable service). In the same way, defined actions (for application of ACA rules) are communicated via the ECA engine.

The dependent variant is especially preferable in closed environments where e.g. an application service (a university or an airline) wants to apply (i) only own rules to (ii) a central management of events (that can be both local events and events arriving at a given input interface).

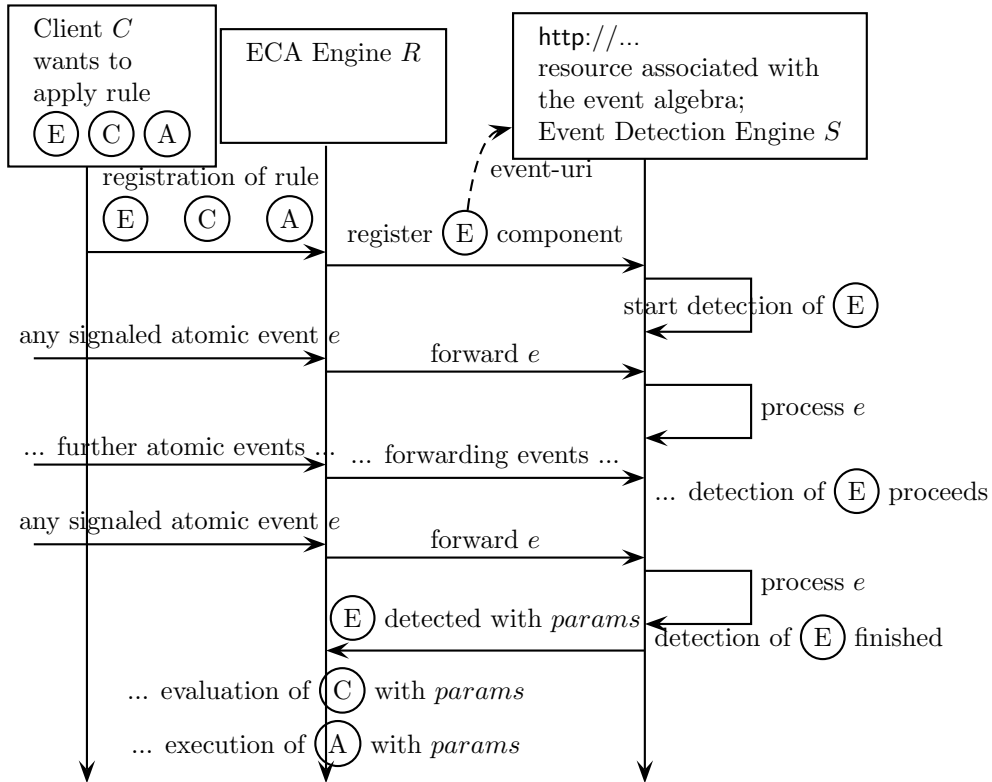


Figure 9.2: Dependent Communication

9.3 Ontology of Languages and Services

The Semantic Web as a whole, and also its dynamic aspects as implemented in the MARS Framework combine a multitude of ontologies. Domains, languages and services are themselves resources (identified by their namespace URI). The architectural relationship between these resources provides also the base for managing the actual communication.

9.3.1 The Ontology

The concepts of the MARS Framework are described in an ontology. There are four main kinds of things:

- MARS meta concepts: classes, properties, actions, events, languages, services etc.,
- *domain ontologies*, e.g. `travel:`, `uni:` that provide concepts (classes), properties/relationships, atomic events and actions. They have been described in Chapter 6.
- *languages*, that can be partitioned into *language families*: ECA languages (e.g. ECA-ML), event algebras (e.g. SNOOP), query languages (e.g. the relational algebra), and process algebras (e.g. CCS) These language families can also be seen as ontologies. Their details will be discussed again in Section 11.

Additionally, on a lower level, there are actual *programming languages* like XQuery, SQL, or Java, and formalisms e.g. for atomic event matching or first-order logic predicates that do not define a language ontology, but are just languages.

- *services* that actually implement the languages. Each kind of service offers specific *tasks* related to the kind of the language that is implemented by it.

We focus here mainly on the component and domain languages for the MARS Framework, but most considerations hold also for any language in the XML world. The ontology will be given in OWL below.

Languages in the MARS Framework. The MARS framework distinguishes several families of languages (as shown in Figure 3.9):

- the ECA language,
- languages for the event component,
- languages for the query component,
- languages for the test component,
- languages for the action component,
- domain languages.

Apart from the MARS Framework, there may be various other kinds of languages that also form subclasses of “Language”. Every language is associated as a resource with its namespace URI (as already used in the examples in the previous section; similarly the W3C languages have their namespace URIs). The current W3C proposals do not specify what is actually behind these URIs. Below, we describe what information is required for using languages in the MARS Framework.

All Languages: Information about the Language. Independent from the characteristics of a language as a “language” or as a domain language, some information about its XML representation should be given:

- for all languages that are represented by an XML markup, e.g., all above XML markup examples for component languages, or also languages like XHTML, GML (a markup language for Geography), and the Mondial language: a DTD and/or XML Schema should be provided. This can be used for validating language expressions. Note that the schema must in general care for nested expressions of other languages.

Since this information is a document (DTD or XSD), it is just a resource, without further ontology descriptions (which will be discussed below).

Generic Programming and Specification Languages. These are in the MARS Framework mainly ECA-ML and the component languages, but the same considerations also hold for other programming languages. When defining such a language, the following information related to the language has to be provided:

- specification of the language elements (e.g., the operators that an algebraic language provides).
- specification of the recommended and allowed result types (e.g., “an XML document” (XML transformation languages), or “set of tuples of variable bindings”).
- if in the future, an ontology of formal semantics will come up, the description of the semantics of a language should also be given.
- a reference processor that interprets the language. For the component languages, these are the evaluation services. For XSLT, this would be a service where one can send an XSLT document (and optionally an XML instance) and gets back the result. Note that in general, there will be multiple processors for a language. Thus, for using them throughout the Web, *registries* for processing services provide the link between languages and actual services (see Section 9.5.1).

Domain Languages. Domain languages are common ontologies, which should also contain actions and events. They are supported by domain nodes, and portals, information brokers, and event brokers.

- markup languages for application domains (e.g., Mondial, traveling, or banking): an RDF/RDFS or OWL description about the notions of that language. For languages that include events and actions, these should also be specified there.

Languages of application domains (e.g., traveling, or banking) can also be supported by services that provide access to the domain information (static and dynamic), e.g., portals, information brokers and event brokers.

Services. In the same way as there are classes of languages, there are the corresponding classes of services.

For being integrated into the MARS Framework, the component Web services must implement appropriate communication for receiving tasks (expressions of the language and variable bindings), additional information (e.g., events), and communicate results as discussed in Sections 4.1. Additionally there are several properties of services from the technical point of view that are to be indicated in a service description (see Section 9.5.1).

9.3.2 The MARS Subontology of Domain-Related Notions

Ontological information about the domains, including their *dynamic* concepts, is supported by the domain-related part of the MARS ontology:

- a meta-ontology of static and dynamic notions,
- which actual Semantic Web Nodes support a domain and its concepts, and
- which domain brokers support the domain.

```
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:owl="http://www.w3.org/2002/07/owl#"
  xmlns="http://www.semwebtech.org/2006/mars#"
  xml:base="http://www.semwebtech.org/2006/mars">

  <owl:Class rdf:ID="Domain"/>
  <owl:Class rdf:ID="DomainNotion"/>
  <owl:Class rdf:ID="Class">
    <rdfs:subClassOf rdf:resource="http://www.w3.org/2002/07/owl#Class"/>
    <rdfs:subClassOf rdf:resource="#DomainNotion"/>
  </owl:Class>
  <owl:Class rdf:ID="Property">
    <rdfs:subClassOf rdf:resource="http://www.w3.org/1999/02/22-rdf-syntax-ns#Property"/>
    <rdfs:subClassOf rdf:resource="#DomainNotion"/>
  </owl:Class>
  <owl:Class rdf:ID="Action">
    <rdfs:subClassOf rdf:resource="#DomainNotion"/>
  </owl:Class>
  <owl:Class rdf:ID="Event">
    <rdfs:subClassOf rdf:resource="#DomainNotion"/>
  </owl:Class>

  <rdf:Property rdf:ID="belongs-to-domain">
```

```

    <rdfs:domain rdf:resource="#DomainNotion"/>
    <rdfs:range rdf:resource="#Domain"/>
</rdf:Property>

<owl:Class rdf:ID="DomainService">
  <rdfs:subClassOf rdf:resource="#Service"/>
</owl:Class>

<rdf:Property rdf:ID="uses-domain">
  <rdfs:domain rdf:resource="#DomainService"/>
  <rdfs:range rdf:resource="#Domain"/>
  <inverseOf rdf:resource="#has-service"/>
</rdf:Property>

<rdf:Property rdf:ID="supports">
  <rdfs:domain rdf:resource="#DomainService"/>
  <rdfs:range rdf:resource="http://www.w3.org/2002/07/owl#Class"/>
</rdf:Property>

<owl:Class rdf:ID="DomainBroker">
  <rdfs:subClassOf rdf:resource="#Service"/>
</owl:Class>

<rdf:Property rdf:ID="has-domain-broker">
  <rdfs:domain rdf:resource="#Domain"/>
  <rdfs:range rdf:resource="#DomainBroker"/>
</rdf:Property>
</rdf:RDF>

```

[Filename:mars-domain-notions.rdf]

9.3.3 MARS Languages and Services Ontology

The MARS Languages and Services Ontology defines the types of languages and corresponding services that are known in the MARS Framework. Note that these are not exclusive to MARS, but the same classes can also be relevant to other ontologies. For each type, the type of appropriate services is given, and the allowed kinds of answers (i.e., a ResultSet (functional style), variable bindings (Datalog Style), Answers (cf. Section 4.2.3; a result set where each result is annotated by a set of tuples of variable bindings), or Boolean) are listed. From the modeling point of view, the latter could be either associated with language types or with service types, but as they belong to the semantics of the language, it is preferable to associate them with the language types. Services provides tasks (that are “typical” for their service class). The actual tasks are defined later in Section 9.4.5.

```

<?xml version="1.0"?> <!-- filename:
mars-language-and-service-types.rdf --> <rdf:RDF
xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:owl="http://www.w3.org/2002/07/owl#"
  xmlns="http://www.semwebtech.org/2006/mars#"
  xml:base="http://www.semwebtech.org/2006/mars">

  <owl:Class rdf:ID="LanguageClass"/>
  <owl:Class rdf:ID="Language">
    <owl:disjointWith rdf:resource="#LanguageClass"/>

```

```

</owl:Class>

<owl:Class rdf:ID="ResultType"/>

<ResultType rdf:ID="VariableBindings"/>
<ResultType rdf:ID="Answers"/>
<ResultType rdf:ID="ResultSet"/>
<ResultType rdf:ID="Boolean"/>

<rdf:Property rdf:ID="recommended-result-type">
  <rdfs:domain rdf:resource="#LanguageClass"/>
  <rdfs:range rdf:resource="#ResultType"/>
</rdf:Property>
<rdf:Property rdf:ID="allowed-result-type">
  <rdfs:domain rdf:resource="#LanguageClass"/>
  <rdfs:range rdf:resource="#ResultType"/>
</rdf:Property>

<LanguageClass rdf:ID="RuleLanguage">
  <rdfs:subClassOf rdf:resource="#Language"/>
</LanguageClass>
<LanguageClass rdf:ID="AlgebraicLanguage">
  <rdfs:subClassOf rdf:resource="#Language"/>
</LanguageClass>
<LanguageClass rdf:ID="ECALanguage">
  <rdfs:subClassOf rdf:resource="#RuleLanguage"/>
  <has-service-type rdf:resource="#ECAService"/>
</LanguageClass>
<LanguageClass rdf:ID="EventAlgebra">
  <rdfs:subClassOf rdf:resource="#AlgebraicLanguage"/>
  <has-service-type rdf:resource="#CompositeEventDetectionEngine"/>
  <recommended-result-type rdf:resource="#Answers"/>
  <allowed-result-type rdf:resource="#VariableBindings"/>
  <allowed-result-type rdf:resource="#ResultSet"/>
</LanguageClass>
<LanguageClass rdf:ID="AtomicEventFormalism">
  <rdfs:subClassOf rdf:resource="#Language"/>
  <has-service-type rdf:resource="#AtomicEventMatcher"/>
  <recommended-result-type rdf:resource="#Answers"/>
  <allowed-result-type rdf:resource="#VariableBindings"/>
  <allowed-result-type rdf:resource="#ResultSet"/>
</LanguageClass>
<LanguageClass rdf:ID="QueryLanguage">
  <rdfs:subClassOf rdf:resource="#Language"/>
  <has-service-type rdf:resource="#QueryService"/>
  <allowed-result-type rdf:resource="#Answers"/>
  <allowed-result-type rdf:resource="#VariableBindings"/>
  <allowed-result-type rdf:resource="#ResultSet"/>
</LanguageClass>
<LanguageClass rdf:ID="ActionLanguage">
  <rdfs:subClassOf rdf:resource="#Language"/>
  <has-service-type rdf:resource="#ActionService"/>
  <allowed-result-type rdf:resource="#VariableBindings"/>
  <allowed-result-type rdf:resource="#Boolean"/>

```

```

</LanguageClass>
<LanguageClass rdf:ID="ProcessAlgebra">
  <rdfs:subClassOf rdf:resource="#AlgebraicLanguage"/>
  <rdfs:subClassOf rdf:resource="#ActionLanguage"/>
  <has-service-type rdf:resource="#ActionService"/>
  <allowed-result-type rdf:resource="#VariableBindings"/>
  <allowed-result-type rdf:resource="#Boolean"/>
</LanguageClass>

<owl:DatatypeProperty rdf:ID="has-name">
  <rdfs:domain rdf:resource="#Language"/>
</owl:DatatypeProperty>

<rdf:Property rdf:ID="has-markup">
  <rdfs:range rdf:resource="#MarkupDescription"/>
</rdf:Property>
<!-- markup description is DTD cup XMLSchema documents -->
<rdf:Property rdf:ID="hasRDF"/>
<rdf:Property rdf:ID="hasDTD"/>
<rdf:Property rdf:ID="hasXSD"/>

<!-- note that the next also defines "Language-Class" -->

<rdf:Property rdf:ID="has-service-type">
  <rdfs:domain rdf:resource="#LanguageClass"/>
  <rdfs:range rdf:resource="#ServiceClass"/>
</rdf:Property>

<owl:Class rdf:ID="ServiceClass">
  <owl:disjointWith rdf:resource="#LanguageClass"/>
</owl:Class>
<owl:Class rdf:ID="Service">
  <owl:disjointWith rdf:resource="#ServiceClass"/>
  <owl:disjointWith rdf:resource="#Language"/>
</owl:Class>

<ServiceClass rdf:ID="ECAService">
  <rdfs:subClassOf rdf:resource="#Service"/>
</ServiceClass>
<ServiceClass rdf:ID="CompositeEventDetectionEngine">
  <rdfs:subClassOf rdf:resource="#Service"/>
</ServiceClass>
<ServiceClass rdf:ID="AtomicEventMatcher">
  <rdfs:subClassOf rdf:resource="#Service"/>
</ServiceClass>
<ServiceClass rdf:ID="QueryService">
  <rdfs:subClassOf rdf:resource="#Service"/>
</ServiceClass>
<ServiceClass rdf:ID="ProcessAlgebraService">
  <rdfs:subClassOf rdf:resource="#Service"/>
</ServiceClass>
<ServiceClass rdf:ID="DomainBroker">
  <rdfs:subClassOf rdf:resource="#Service"/>
</ServiceClass>

```

```

<ServiceClass rdf:ID="DomainService">
  <rdfs:subClassOf rdf:resource="#Service"/>
</ServiceClass>

<!-- services provide tasks (that are identified by their urls)
      that have a name and a task description -->
<owl:Class rdf:ID="Task"/>
<owl:Class rdf:ID="Task-Description"/>

<rdf:Property rdf:ID="provides-task">
  <rdfs:domain rdf:resource="#Service"/>
  <rdfs:range rdf:resource="#Task"/>
</rdf:Property>

<rdf:Property rdf:ID="meta-provides-task">
  <rdfs:domain rdf:resource="#ServiceClass"/>
  <rdfs:range rdf:resource="#Task"/>
</rdf:Property>
</rdf:RDF>

```

[Filename:mars-language-and-service-types.rdf]

[Note that the above XML/RDF fragments are valid “striped RDF” and can be validated and visualized with the RDF validator at <http://www.w3.org/RDF/Validator/>]

9.4 Service Interfaces and Functionality

The material in this section is subject to change. It will probably be validated and extended during further implementation.

For a uniform communication between the different kinds of services, each service type provides certain tasks. The Languages & Service Ontology specifies which tasks are provided by the different classes of services. The *actual* communication interfaces of *each service* are then handled by *Language & Service Registries (LSRs)* (see Section 9.5.1) which are based on the lists of tasks given next. The complete ontology is given in RDF in Section 9.4.5.

A frequent pattern is to register/submit components and subexpressions to appropriate services. For this, addresses and details of the communication format must be specified. Additionally, the answers must be received (note that the respective addresses are communicated with the Reply-To, but details of the communication format must be specified. For that, also the tasks for receiving information are listed below.

In the below list, [R] means to provide a service or to receive information while [S] means just to call a service (send information). Optionally, a service should give the description of the tasks offered by it (called when some other service wants to invoke that task); the contents of these task descriptions is discussed in Section 9.5.1 (otherwise the owner of the service has to submit this information manually to an LSR).

9.4.1 ECA Services

An ECA service must implement the ECA-ML language and adhere to the abstract semantics of ECA rules given in Sections 4.1 and 4.4.

9.4.1.1 Upper Interface:

Registration of Rules: [R] a rule can be registered as a predefined URI (e.g., if the rule is part of a rule set of a domain service, its URI will be determined there), or just as an XML

fragment or an RDF graph (e.g., by a user). In the latter case, it is associated with a URI relative to the ECA service. This URI is communicated to the registrant in the response.

Deregistration of Rules: [R] Rules can be deregistered using their URI.

Disabling and Enabling Rules: [R, Optional] Rules can be disabled and enabled using their URI. Disabling a rule means to cancel all ongoing dependent event detections. Rule instances that are under actual execution at this point will be completed. Enabling a rule means that detection processes will be started at this moment. Past atomic events do not contribute to the detection. Note that a “grey area” exists for “past” events that become known only now to a service.

9.4.1.2 Lower Interface:

Registration/Evaluation of Components: The individual components must be submitted to appropriate services. Additionally, the answers must be received:

- [S] Composite event component: register at composite event detection service (CED).
- [S] Atomic event component: in case that the event component is just an atomic event, it can either be registered via a CED, or directly at an AEM, using the same downward communication for atomic events that a CED uses (see Section 5.3).
- [R] receive answers for detected events.
- [S] Query components: submit query together with bound variables to appropriate service and receive answer.
- [S] Test component: analogous. Often the test component just uses simple (comparison) predicates between variables that can be evaluated locally.
- [S] Action component: submit action together with bound variables to appropriate service. Optionally, an answer can be received (success/failure) which can be used in transactional environments.
- [S] optionally, query and action components can be registered a priori, and actual evaluation/execution then refers to an id and just submits the current variable bindings. (In that case, the choice of the service is done once, and must be remembered.)

Validation of Components: [S, Optional] Submit component markup/code to appropriate service for validation. Note that validation has to be done hierarchically in the same way as the evaluation.

9.4.2 (Algebraic) Component Languages/Services (General)

- upper interface: receive requests:
 - CED: [R] registrations, deregistrations
 - AEM: [R] registrations, deregistrations
 - QE: [R] queries,
optional: [R] registrations, deregistrations, invocations for given variable bindings
 - CAE: [R] composite action specifications,
optional: registrations, deregistrations, invocations for given variable bindings
 - [R, optional] a method that validates a given statement,
 - [R, optional] a method that returns for a given statement the used variables, the input variables, the output variables, and the returned variables (based on the declarations of the atomic expressions in the component),
- [R] upper interface: send answers.

- lower interface: communicate with the domains
 - CED: register/deregister AEDs at AEMs [S], receive answers [R]
 - AEM: register/deregister events at EBs [S], receive events [R]
 - QE: ask domain nodes/portals/brokers [S], receive answers [R]
 - CAE: forward atomic actions to be invoked to domain nodes/portals/brokers [S]. For some formalisms: communication of embedded event patterns and queries [S/R].

9.4.3 Domain Brokers

The domain brokers act as mediators between domain nodes and component language services.

Initialization.

- [R] register an ontology (in most cases when the domain broker is initialized)
input: the ontology

Information about Domain Nodes.

- [R] receive messages from domain nodes which classes, properties, event types and action types they support
 - prototype: as a whole description (cf. Section 6.2)
 - future: as individual messages (in case a node changes its behavior)

Mediation of Requests.

- [R] receive registrations/deregistrations for events (optionally: of given types),
- [R] receive events and [S] forward them to AEMs,
- [S] send events to registered AEMs,
- [R] receive queries, and [S] evaluate them against domain, and return answers,
- [R] receive atomic actions to be executed and [S] forward them to domain nodes.

9.4.4 Domain Services

- provide a service that implements the domain and provides an appropriate communication interface (calling an atomic “thing” with some parameters).
 - [R] answering queries,
 - [S] providing events (either upon registration or to a fixed event broker),
 - [R] executing actions of the domain (submitted in XML).
 - [S, optional] send list of supported notions of the ontology to a domain broker (otherwise: support for the whole namespace is assumed).

9.4.5 The Services Ontology

The services ontology defines a “name” (= resource identifier) for each of the tasks of each kind of service. This will be used in the LSR when describing how to invoke a certain task of a certain service.

```
<?xml version="1.0"?> <!-- filename: services-ontology.rdf -->
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:owl="http://www.w3.org/2002/07/owl#"
  xmlns="http://www.semwebtech.org/2006/mars#"
  xml:base="http://www.semwebtech.org/2006/mars">

  <ServiceClass rdf:about="#ECAService">
    <meta-provides-task rdf:resource="./eca-service#register-rule"/>
    <meta-provides-task rdf:resource="./eca-service#deregister-rule"/>
    <meta-provides-task rdf:resource="./eca-service#disable-rule"/> <!-- opt -->
    <meta-provides-task rdf:resource="./eca-service#enable-rule"/> <!-- opt -->
    <meta-provides-task rdf:resource="./eca-service#receive-detected-event"/>
    <meta-provides-task rdf:resource="./eca-service#receive-query-answer"/>
    <meta-provides-task rdf:resource="./eca-service#give-service-description"/> <!-- opt -->
  </ServiceClass>
  <ServiceClass rdf:about="#CompositeEventDetectionEngine">
    <meta-provides-task rdf:resource="./ced-engine#register-event-pattern"/>
    <meta-provides-task rdf:resource="./ced-engine#deregister-event-pattern"/>
    <meta-provides-task rdf:resource="./ced-engine#validate-pattern"/> <!-- opt -->
    <meta-provides-task rdf:resource="./ced-engine#analyze-variables"/> <!-- opt -->
    <meta-provides-task rdf:resource="./ced-engine#receive-detected-event"/>
    <meta-provides-task rdf:resource="./ced-engine#give-service-description"/> <!-- opt -->
  </ServiceClass>
  <ServiceClass rdf:about="#AtomicEventMatcher">
    <meta-provides-task rdf:resource="./aem#register-event-pattern"/>
    <meta-provides-task rdf:resource="./aem#deregister-event-pattern"/>
    <meta-provides-task rdf:resource="./aem#validate-pattern"/> <!-- opt -->
    <meta-provides-task rdf:resource="./aem#analyze-variables"/> <!-- opt -->
    <meta-provides-task rdf:resource="./aem#receive-event"/>
    <meta-provides-task rdf:resource="./aem#give-service-description"/> <!-- opt -->
  </ServiceClass>
  <ServiceClass rdf:about="#QueryService">
    <meta-provides-task rdf:resource="./qs#evaluate-query"/>
    <meta-provides-task rdf:resource="./qs#register-query"/> <!-- opt -->
    <meta-provides-task rdf:resource="./qs#deregister-query"/> <!-- opt -->
    <meta-provides-task rdf:resource="./qs#invoke-query"/> <!-- opt -->
    <meta-provides-task rdf:resource="./qs#validate-pattern"/> <!-- opt -->
    <meta-provides-task rdf:resource="./qs#analyze-variables"/> <!-- opt -->
    <meta-provides-task rdf:resource="./qs#give-service-description"/> <!-- opt -->
  </ServiceClass>
  <ServiceClass rdf:about="#ActionService">
    <meta-provides-task rdf:resource="./action-service#execute-action"/>
    <meta-provides-task rdf:resource="./action-service#register-action"/> <!-- opt -->
    <meta-provides-task rdf:resource="./action-service#deregister-action"/> <!-- opt -->
    <meta-provides-task rdf:resource="./action-service#receive-query-answer"/> <!-- opt -->
    <meta-provides-task rdf:resource="./action-service#receive-detected-event"/> <!-- opt -->
    <meta-provides-task rdf:resource="./action-service#validate-pattern"/> <!-- opt -->
  </ServiceClass>
</RDF>
```

```

    <meta-provides-task rdf:resource="./action-service#analyze-variables"/> <!-- opt -->
    <meta-provides-task rdf:resource="./action-service#give-service-description"/> <!-- opt -->
</ServiceClass>
<ServiceClass rdf:about="#DomainBroker">
    <meta-provides-task rdf:resource="./domain-broker#register-ontology"/> <!-- init -->
    <meta-provides-task rdf:resource="./domain-broker#update-ontology"/> <!-- opt -->
    <meta-provides-task rdf:resource="./domain-broker#receive-bulk-support-information"/>
    <meta-provides-task rdf:resource="./domain-broker#update-support-information"/> <!-- opt -->
    <meta-provides-task rdf:resource="./domain-broker#register-for-event"/>
    <meta-provides-task rdf:resource="./domain-broker#receive-event"/>
    <meta-provides-task rdf:resource="./domain-broker#execute-query"/>
    <meta-provides-task rdf:resource="./domain-broker#receive-query-answer"/>
    <meta-provides-task rdf:resource="./domain-broker#execute-action"/>
    <meta-provides-task rdf:resource="./domain-broker#give-service-description"/> <!-- opt -->
</ServiceClass>
<ServiceClass rdf:about="#DomainService">
    <meta-provides-task rdf:resource="./domain-node#receive-query"/>
    <meta-provides-task rdf:resource="./domain-node#receive-action"/>
    <meta-provides-task rdf:resource="./domain-node#give-service-description"/> <!-- opt -->
</ServiceClass>
</rdf:RDF>

```

[Filename:services-ontology.rdf]

[Note that the above XML/RDF is valid “striped RDF” and can be validated and visualized with the RDF validator at <http://www.w3.org/RDF/Validator/>]

9.5 Locating and Contacting Language Services

Rules and component expressions consist of nested expressions in several languages/namespaces (ECA language, event language, atomic event specification language, domain vocabulary). The handover between services has to take place at the “language borders” (which are in XML markup indicated by the namespace borders). The processor that processes the surrounding language has to do the following:

- identify the embedded fragment (as XML subtree), and
- determine what processor is responsible to process it.

For performing the latter, the namespace and/or language information of the subtree contains the information about the used language:

- embedded algebraic languages: namespace of the root element of the subexpression,
- embedded atomic event specifications: namespace of the root element of the subexpression, or language attribute (cf. Section 5.2.2),
- opaque expressions: language attribute.

If a service is identified, the actual communication details (e.g., where to send the message, how to wrap it, and whether it supports multiple tuples of variable bindings) are determined. All required information is available in *Languages and Services Registries (LSRs)* that are discussed in Section 9.5.1. Note that the required algorithms are the same for every invocation (done by ECA engines, action engines (if the action embeds events or queries), and also for nested event components). This “trader” functionality is provided by a *Generic Request Handler, GRH* (Section 9.5.4) that can be implemented as a standalone service or provided as an instance of a downloadable Java class.

9.5.1 Language & Service Registries

The languages and the actual service instances are handled by *Language & Service Registries (LSRs)* [in general one LSR would be sufficient, but as there is no central thing in the Semantic Web, it cannot be expected that there is a central LSR – for running any kind of business it is sufficient to know a *good* LSR - or to know an ECA engine that knows a good LSR]. Language & Service Registries will -in the final version- be services where others can get/look up information how some task wrt. a language can be requested:

- input: a language (by its URI) and a “task”, e.g., “register an event component (in that language, for future detection)”,
- answer: the task description in XML or RDF format.

The ontology discussed in Section 9.4.5 specifies [in the future] for each kind of service all available tasks (mandatory or optional). Note that an actual service can also provide functionality of several service classes.

Identification: Mapping from Languages to Services. A first task of the LSR is to provide a mapping from the languages to actual services that can handle the tasks. This includes the *registration* of services. For each registered service, a more detailed description how to employ its tasks is then also included in the LSR.

Communication Modalities. If a service for a certain language and a certain functionality is identified, the actual communication (which URI, which format) must be determined.

9.5.2 Interface Descriptions of Individual Tasks

Framework-aware services should provide a service description for each task that deals with the following issues:

- identification of the task via its URI defined in the *Services Ontology* in Section 9.4.5.
- URLs where the respective input is expected (relative to the service url; items that are not given default to the service url),
- format of the input (wrapping),
- functionality (e.g., built-in join functionality for supporting sideways-information-passing strategies in queries).

Characteristics: Input Formats. Every request contains the information where the answer must go (Reply-To), an identifier to associate the answer (in asynchronous cases), the message contents itself (which is often a fragment in the respective language), and optional variable bindings. The input format specifies how these components are communicated with an individual given task:

- **Reply-To:** header or body (default: header as X-Reply-To)
- **Subject:** header or body (default: header as X-Subject)
- **input-format** [CED, AEM, QL, T, A]: message input formats requested by component services:
 - default: **item *** – the message body contains a sequence of elements (which elements should be clear from the context situation).

- * In most cases, this is a fragment of the respective language, and optionally variable bindings,
- * sometimes (e.g. for the event stream), it's just any elements.
- optional: the above-mentioned items are wrapped into an element with a given name: **element *name*** (with arbitrary namespace - the receiver will at most check the local-name [sometimes even not this if the element is just needed to have any root node])

Characteristics: Functionality.

- **variables** [CED, AEM, Q, T, A]: are variable bindings accepted?
 - *****: a set of tuples of variable bindings in the markup given in Section 4.2.1 is accepted.
 - **1**: one tuple of variable bindings in the markup given in Section 4.2.1 is accepted at a time (means that the calling service must iterate).
 - **no**: no variable bindings are accepted (means that the calling service must iterate and replace the variables by their values as strings in the code).
 - **ignore**: variable bindings are ignored (e.g., when deregistering something that has been registered with variable bindings).
- **join-enabled** [event, event matcher, query]: does it make sense to send variable bindings that are used for join semantics?
 - **yes**: the service implements join semantics. Variable bindings can be sent for optimized answering.
 - **no** (default): send only the input variables that will be bound by the service.
- services that accept only a single input tuple should indicate whether the returned tuple(s) are an extension of the input tuple, or if they do not “echo” all input variables:

return-input-vars [event, event matcher, query]: are the input variables returned with the answer? “echoing” these variables is in general needed for assigning the returned tuples to tuples on the ECA level; see Section 4.1.5 and requirements on Page 50.

 - **yes**: result can immediately be joined.
 - **no**: the calling service must itself care about the assignment of the returned answers with the corresponding tuples of variable bindings.
- **communication-mode** [query, test, other service queries]:
 - **synchronous**: answer is given synchronously (includes one-way communication where only a message is sent)
 - **asynchronous**: answer will be given asynchronously (in this case, the immediate answer is “OK”, and the result is latter returned with appropriate identification (Subject)).
 - Default is synchronous.
- **return-incomplete** [query, test]: does the answer contain/consider all tuples, or is it possible that several results are sent?
 - for ECA services: **yes/no** indicate if it can process incomplete answers (if this is indicated with the answer). Default: no.
 - for query/text services: **yes/no** indicate if the service probably returned incomplete answers (and indicates this with the answer). Default: no.

Example 44 (Service Description) *The following is a fragment of the service description of a CED (assumed at URL url) that allows to register CESs at url/register. A registration message must have the following format:*

- *Reply-To-address must be given in the HTTP message header,*
- *the Subject (=identifier) must be given as an element in the body,*
- *the task itself is always given in the body,*
- *and one tuple of variable bindings can be given (which is common for CEDs),*
- *and the whole body is encapsulated in a single <register> element.*

Note that line 13 refers to <http://www.semwebtech.org/2006/languages/qs/register-query> due to the setting of `xml:base`.

```

<?xml version="1.0"?> <rdf:RDF
xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:mars="http://www.semwebtech.org/2006/mars#"
  xmlns="http://www.semwebtech.org/2006/lsr#"
  xml:base="http://www.semwebtech.org/2006/lsr">
<mars:QueryService rdf:about="http://foo.nop/fancy-query-engine">
  <!-- sample: only one task description
    clients can register queries at: -->
  <has-task-description>
    <TaskDescription>
      <!-- reference to the described task according to the Service Ontology -->
      <describes-task rdf:resource="#qs/register-query"/>
      <!-- url of the service where the task can be actually invoked -->
      <provided-at rdf:resource="./register"/>
      <Reply-To>header</Reply-To>
      <Subject>body</Subject>
      <input>element register</input>
      <variables>1</variables>
      <!-- means: Subject in the message header, contents is
        <register>$event$, $var-bindings$</register> with
        max. 1 tuple of variable bindings -->
    </TaskDescription>
  </has-task-description>
</mars:QueryService>
</rdf:RDF>

```

[Filename:sample-lsr.rdf]

[Note that the above XML/RDF is valid "striped RDF" and can be validated and visualized with the RDF validator at <http://www.w3.org/RDF/Validator/>]

A message thus could look as follows, sent to `url/register`:

```

X-Reply-To: <http://bla.nop/i-want-my-answers-here>
<register>
  <Subject>this-is-my-id-for-this-task-007</Subject>
  <evt:operator xmlns:evt="...">
    contents
  </evt:operator>
  <log:variable-bindings xmlns:log="...">
    <log:tuple>...</log:tuple>
  </log:variable-bindings>
</register>

```

When a service registers itself at a LSR, it either submits its SD or a reference to it.

Non-Framework-Aware Services. In case that non-framework-aware services are used in a rule, the LSR must contain appropriate entries.

- the MARS Framework maintainers can add SDs for services that are used frequently,
- clients who submit a rule (whose namespace URI refer to such a service) can submit the specification of the service.

9.5.3 The Languages and Services Registry RDF Model

Since the metadata structure is complex, it its preferably expressed in RDF. Nevertheless, for a first prototype on the XML level, we assume that some metadata is available as an XML file. For each language, the following is given:

- language URI (mandatory),
- language name , e.g. “XQuery”, “SNOOP”, may be null,
- type of language (ECA, CEL, AESL, ...),
- List of appropriate services by URIs and their Service Descriptions.

Naming Schema. The languages developed in this MARS Framework are located in the “www.semwebtech.org” domain (which is registered by the DBIS group at Göttingen University); there also most of the reference services are running.

Namespaces refer to “http://www.semwebtech.org/*/2006” (to have room for new versions in the following years). The language namespaces are maintained as “hash namespaces” (according to the terminology in [65]) and described by RDF/RDFS files accessible at these URLs.

```
<!-- filename: lsr.rdf --> <rdf:RDF
xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:mars="http://www.semwebtech.org/2006/mars#"
  xmlns="http://www.semwebtech.org/2006/lsr#"
  xml:base="http://www.semwebtech.org/2006/lsr">

<rdf:Property rdf:ID="has-task-description">
  <rdfs:domain rdf:resource="http://www.semwebtech.org/2006/mars#Service"/>
  <rdfs:range rdf:resource="#TaskDescription"/>
</rdf:Property>

<mars:ECALanguage
rdf:about="http://www.semwebtech.org/languages/2006/eca-ml">
  <mars:shortname>eca-ml</mars:shortname>
  <mars:name>ECA Markup Language</mars:name>
  <mars:is-implemented-by>
  <mars:ECAService
    rdf:about="http://www.semwebtech.org/services/2006/eca-engine">
  <has-task-description>
  <TaskDescription>
    <describes-task rdf:resource="/eca-service/register-rule"/>
    <provided-at rdf:resource="????"/>
    <Reply-To>????</Reply-To>
    <Subject>????</Subject>
    <input>????</input>
    <variables>?</variables>
```

```

    </TaskDescription>
  </has-task-description>
<has-task-description>
  <TaskDescription>
    <describes-task rdf:resource="/eca-service/deregister-rule"/>
    <provided-at rdf:resource="????"/>
    <Reply-To>????</Reply-To>
    <Subject>????</Subject>
    <input>????</input>
    <variables>?</variables>
  </TaskDescription>
</has-task-description>
<has-task-description>
  <TaskDescription>
    <describes-task rdf:resource="eca-service/receive-detected-event"/>
    <provided-at rdf:resource="????"/>
    <Reply-To>????</Reply-To>
    <Subject>????</Subject>
    <input>????</input>
    <variables>?</variables>
  </TaskDescription>
</has-task-description>
<has-task-description>
  <TaskDescription>
    <describes-task rdf:resource="eca-service/receive-query-answer"/>
    <provided-at rdf:resource="????"/>
    <Reply-To>????</Reply-To>
    <Subject>????</Subject>
    <input>????</input>
    <variables>?</variables>
  </TaskDescription>
</has-task-description>
</mars:ECAService>
</mars:is-implemented-by>
</mars:ECALanguage>

<mars:EventAlgebra
  rdf:about="http://www.semwebtech.org/languages/2006/snoopy">
<mars:name>SNOOP (from the Sentinel Database System)</mars:name>
<mars:is-implemented-by>
  <mars:CompositeEventDetectionEngine
    rdf:about="http://www.semwebtech.org/services/2006/snoopy">
  <has-task-description>
  <TaskDescription>
    <describes-task rdf:resource="ced-engine/register-event-pattern"/>
    <provided-at
      rdf:resource="http://www.semwebtech.org/services/2006/snoopy/????"/>
    <Reply-To>????</Reply-To>
    <Subject>????</Subject>
    <input>????</input>
    <variables>ignore</variables>
  </TaskDescription>
</has-task-description>
<has-task-description>

```

```

<TaskDescription>
  <describes-task rdf:resource="ced-engine/deregister-event-pattern"/>
  <provided-at
    rdf:resource="http://www.semwebtech.org/services/2006/snoopy/????"/>
  <Reply-To>????</Reply-To>
  <Subject>????</Subject>
  <input>????</input>
  <variables>ignore</variables>
</TaskDescription>
</has-task-description>
<has-task-description>
  <TaskDescription>
    <describes-task rdf:resource="ced-engine/receive-detected-event"/>
    <provided-at
      rdf:resource="http://www.semwebtech.org/services/2006/snoopy/????"/>
    <Reply-To>????</Reply-To>
    <Subject>????</Subject>
    <input>item</input>
    <variables>*</variables>
  </TaskDescription>
</has-task-description>
</mars:CompositeEventDetectionEngine>
</mars:is-implemented-by>
</mars:EventAlgebra>

<mars:AtomicEventFormalism
  rdf:about="http://www.semwebtech.org/languages/2006/aes-xmlql">
  <mars:name> Atomic Event Detection by XML-QL</mars:name>
  <mars:is-implemented-by>
    <mars:AtomicEventMatcher
      rdf:about="http://www.semwebtech.org/services/2006/aem-xmlql">
      <!-- this aem receives each task at a separate url. Tasks are
        not wrapped into any additional element, events come as elements
        <namespace:localname> -->
    <has-task-description>
      <TaskDescription>
        <describes-task rdf:resource="aem/register-event-pattern"/>
        <provided-at
          rdf:resource="http://www.semwebtech.org/services/2006/aem-xmlql/register"/>
        <Reply-To>body</Reply-To>
        <Subject>body</Subject>
        <input>item *</input>
        <variables>1</variables>
      </TaskDescription>
    </has-task-description>
  <has-task-description>
    <TaskDescription>
      <describes-task rdf:resource="aem/deregister-event-pattern"/>
      <provided-at
        rdf:resource="http://www.semwebtech.org/services/2006/aem-xmlql/deregister"/>
      <Reply-To>body</Reply-To>
      <Subject>body</Subject>
      <input>item *</input>
      <variables>ignore</variables>
    </TaskDescription>
  </has-task-description>
</mars:AtomicEventMatcher>
</mars:is-implemented-by>
</mars:AtomicEventFormalism>

```



```

    </TaskDescription>
  </has-task-description>
<has-task-description>
  <TaskDescription>
    <describes-task rdf:resource="aem/receive-event"/>
    <provided-at
      rdf:resource="http://www.semwebtech.org/services/2006/aem-xmlql/events"/>
    <!-- other parameters are not given. Events are simply sent. -->
  </TaskDescription>
</has-task-description>
</mars:AtomicEventMatcher>
</mars:is-implemented-by>
</mars:AtomicEventFormalism>

<mars:AtomicEventFormalism
  rdf:about="http://www.semwebtech.org/languages/2006/aes-xpath">
  <mars:name> Atomic Event Detection by XML-QL</mars:name>
  <mars:is-implemented-by>
    <mars:AtomicEventMatcher
      rdf:about="http://www.semwebtech.org/services/2006/aem-xpath">
      <!-- this aem receives everything at the same url. Tasks are
        wrapped as <register> or <deregister> (without namespace), events
        come as elements <namespace:localname> -->
      <has-task-description>
        <TaskDescription>
          <describes-task rdf:resource="aem/register-event-pattern"/>
          <provided-at
            rdf:resource="http://www.semwebtech.org/services/2006/aem-xpath"/>
          <Reply-To>body</Reply-To>
          <Subject>body</Subject>
          <input>element register</input>
          <variables>1</variables>
        </TaskDescription>
      </has-task-description>
      <has-task-description>
        <TaskDescription>
          <describes-task rdf:resource="aem/deregister-event-pattern"/>
          <provided-at
            rdf:resource="http://www.semwebtech.org/services/2006/aem-xpath"/>
          <Reply-To>body</Reply-To>
          <Subject>body</Subject>
          <input>element register</input>
          <variables>ignore</variables>
        </TaskDescription>
      </has-task-description>
      <has-task-description>
        <TaskDescription>
          <describes-task rdf:resource="aem/receive-event"/>
          <provided-at
            rdf:resource="http://www.semwebtech.org/services/2006/aem-xpath"/>
          <!-- other parameters are not given. Events are simply sent. -->
        </TaskDescription>
      </has-task-description>
    </mars:AtomicEventMatcher>

```

```

</mars:is-implemented-by>
</mars:AtomicEventFormalism>

<mars:QueryLanguage rdf:about="http://www.w3.org/XQuery">
  <mars:name>XQuery (Opaque)</mars:name>
  <mars:is-implemented-by>
    <mars:QueryService
      rdf:about="http://www.semwebtech.org/services/2006/????">
      <has-task-description>
        <TaskDescription>
          <describes-task rdf:resource="query-engine/evaluate-query"/>
          <provided-at
            rdf:resource="http://www.semwebtech.org/services/2006/????/????"/>
          <Reply-To>????</Reply-To>
          <Subject>????</Subject>
          <input>????</input>
          <variables>*</variables>
        </TaskDescription>
      </has-task-description>
    </mars:QueryService>
  </mars:is-implemented-by>
</mars:QueryLanguage>

<mars:ProcessAlgebra
  rdf:about="http://www.semwebtech.org/languages/2006/ccs">
  <mars:name>CCS</mars:name>
  <mars:is-implemented-by>
    <mars:ActionService
      rdf:about="http://www.semwebtech.org/services/2006/ccs">
      <has-task-description>
        <TaskDescription>
          <describes-task rdf:resource="query-engine/evaluate-query"/>
          <provided-at
            rdf:resource="http://www.semwebtech.org/services/2006/ccs/????"/>
          <Reply-To>????</Reply-To>
          <Subject>????</Subject>
          <input>????</input>
          <variables>*</variables>
        </TaskDescription>
      </has-task-description>
    </mars:ActionService>
  </mars:is-implemented-by>
</mars:ProcessAlgebra>

<mars:Domain
  rdf:about="http://www.semwebtech.org/domains/2006/travel#">
  <mars:name>Traveling</mars:name>
</mars:Domain>

<!-- note that the relationship between domains and application
services is n:m (a service may implement multiple domains), and
the relationship between domains and domain broker services is
also n:m (a broker may support multiple domains) -->

```

```

<mars:DomainBroker
  rdf:about="http://www.semwebtech.org/services/2006/travel-broker">
  <mars:name>Travel Domain Broker</mars:name>
  <supports-domain>http://www.semwebtech.org/domains/2006/travel</supports-domain>
  <!-- to be extended -->
</mars:DomainBroker>

<mars:ApplicationNode
rdf:about="http://www.semwebtech.org/nodes/2006/orafly">
  <mars:name>Oracle Airlines</mars:name>
  <mars:uses-domain>http://www.semwebtech.org/domains/2006/travel</mars:uses-domain>
  <!-- list of all names in the domain (events, actions, concepts, properties)
        that are supported by that application node -->
  <mars:supports>flight</mars:supports>
  <mars:supports>cancel-flight</mars:supports>
  <mars:supports>flight-is-canceled</mars:supports>
  <mars:supports>canceled-flight</mars:supports>
  <!-- to be extended -->
</mars:ApplicationNode> </rdf:RDF>

```

[Filename:lsr.rdf]

The LSR can be found as an RDF/XML file at <http://www.semwebtech.org/2006/lsr/lsr.rdf>. [Note that the above XML/RDF is valid “striped RDF” and can be validated and visualized with the RDF validator at <http://www.w3.org/RDF/Validator/>]

The following SPARQL query illustrates the combination of the RDF information described in this chapter.

```

# call # jena -q -pellet -il RDF/XML -qf lsr-query.sparql

PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> PREFIX
rdfs: <http://www.w3.org/2000/01/rdf-schema#> PREFIX owl:
<http://www.w3.org/2002/07/owl#> PREFIX mars:
<http://www.semwebtech.org/2006/mars#> PREFIX lsr:
<http://www.semwebtech.org/2006/lsr#>

SELECT ?L ?S ?T ?URL FROM
<file:mars-language-and-service-types.rdf> FROM
<file:services-ontology.rdf> FROM <file:lsr.rdf> WHERE {
  ?L rdf:type mars:Language .
  ?L mars:is-implemented-by ?S .
  ?S lsr:has-task-description ?TD .
  ?TD lsr:describes-task ?T .
  ?TD lsr:provided-at ?URL
}

```

9.5.4 Service Brokering: Generic Request Handlers

Given any “task” in some language (e.g., the execution of a query component in XQuery), the LSR serves for identifying a target service, and for obtaining information about the communication. Below, we distinguish two kinds of task handling:

- component language tasks, i.e., (composite) event languages, queries, and (composite) actions,
- specialized auxiliary tasks, such as event matching.

Handling Component Language Tasks

The handling of component language fragments includes the communication of variable bindings in both directions. It occurs in the following cases:

- obviously: execution of a component of an ECA rule,
- detection of events or execution of queries in composite action specifications,
- nested subevents in another language inside a composite event,
- analogously for nested query expressions (e.g., a join between a Datalog query with the result of an XQuery expression).

Thus, it is reasonable to specify and implement this trader functionality separately (instead of intergating it within the ECA engine, which would be sufficient at a first glance). Since the functionality of the *Generic Request Handler (GRH)* is (i) simple, (ii) canonic, and (iii) frequently needed, it is designed as a class that can be instantiated for each kind of service that needs it. ECA engines or action language engines that support embedded queries and events will probably have their “own” instances. In contrast, composite event specifications that use subevents from different languages are expected to be much less frequent, so a CED may not have an own GRH but can use a “free” one (as a Web Service).

Use of the GRH: Integration within an ECA Engine. The functionality and use of the GRH is illustrated here together with an ECA engine, cf. Figure 9.3.

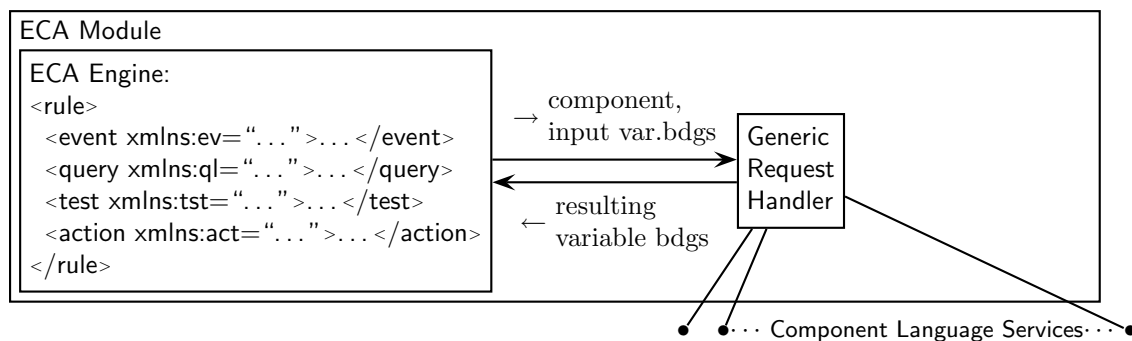


Figure 9.3: Architecture of an ECA Module with own GRH

The core *ECA engine* implements *only* the handling of variable bindings and the ECA rule semantics. It is supported by a GRH instance for communicating with autonomous component services. The GRH communicates the components to suitable services. It also receives the answers and wraps them into the agreed format of variable bindings (in case of non-framework-aware services).

GRH: Functionality

- dealing with requests coming from the ECA engine, forwarding them to component services (decoding the namespace URL and addressing an appropriate Web Service is hardcoded in the first prototype that supports only a very restricted number of component languages; this will later be separated in the LSR),
- receive answers, transform them into variable bindings (in `<variable-bindings>` format).

GRH: Interface

- in from above: requests from ECA (complete components in ECA-ML and contents),
- out to below (autonomous services): requests to framework-aware and non-framework-aware (HTTP) services,
- in from below: results in the `<logvars:answers>` format from framework-aware component services and unwrapped results from non-frame-aware services (HTTP answers).

Interface of Modules towards their GRH

If the GRH is not a separate service, but a local instance, the invocation of the GRH call can be done as a method call. Additionally, there may be “free” GRH services.

- out: requests to GRH: components in XML Markup of the individual languages,
- in: answers from GRH, in `<variable-bindings>` format.

Global Cooperation via GRHs.

As discussed above, GRH instances can be embedded into modules that have to deal with (embedded) language fragments, or can be standalone. Thus, not only the ECA engine uses a GRH, but also the actual processing of the action component with its embedded expressions makes use of a GRH. The handling of embedded `<ccs:Event>`, `<ccs:Query>`, `<ccs:Test>`, and `<ccs:Action>` elements with embedded fragments of other languages is done in the same way as described in Section 5. Atomic actions are executed “immediately” by submitting them to the domain nodes (if specified by an URL) or to a domain broker (responsible for the domain, forwarding them to appropriate domain nodes). Figure 9.4 illustrates this by an ECA engine and an action engine (CCS).

Example 45 *Consider again the sample rule given in Section 5.6.5. The rule is submitted to the ECA engine, which registers the event pattern via the GRH at the snoopy composite event detection engine. When the event pattern is detected for some flight, snoopy informs the ECA engine. The ECA engine submits the query with `flight` and `date` to the GRH which forwards it to a domain broker that supports the `travel` domain and receives a set of extended answer tuples (`flight`, `date`, `booking`, `name`). The ECA engine then submits the action component with the obtained variable bindings via the GRH to the CCS engine. The CCS engine starts two concurrent threads. The first of them just submits for each tuple of variable bindings (i.e., for each person) the atomic action `<travel:reserve-room .../>` to a domain broker that supports the `travel` domain (which will in turn invoke the reservation at the indicated hotel). The second thread first (again via the GRH and a `travel` domain broker) evaluates the test whether the given booking is a business class booking – all tuples where this is not the case are removed. Then, it extends the variable bindings by `phone`. Finally, the action `<comm:send-sms .../>` together with the (remaining) tuples is sent to the GHR which forwards it to a domain broker that supports the `comm` domain, which will then forward it to an appropriate service node that sends the message.*

Specialized Request Handling

There are also many situations that require only specialized handling of requests where a Generic Request Handler is not necessary (nevertheless, it could be applied). For instance, the handling of the communication between CEDs and AEMs can be integrated within the CEDs:

- registration/deregistration: for each atomic event, determine its namespace and formalism,
- ask the LSR for an appropriate service and where to send the event pattern, and if it has to be wrapped into a special XML element.
- send the pattern.

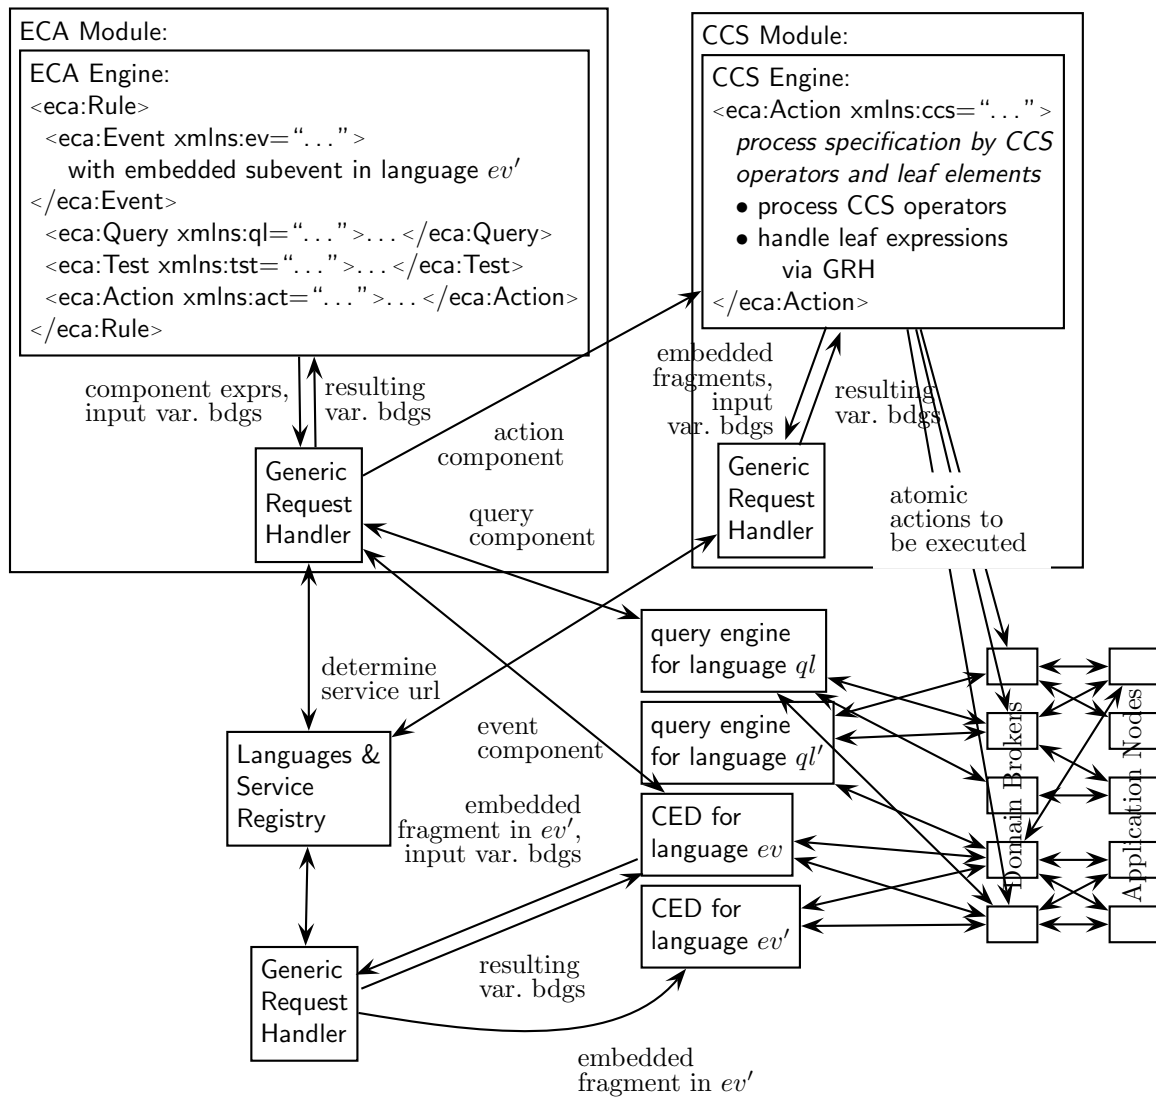


Figure 9.4: Global Cooperation using GRHs

9.5.5 Using the LSR

For “isolated” tasks, e.g. evaluating a query or executing a composite action, a suitable service can be contacted without considering past communication. On the other hand, if communication refers to a previous activity (e.g., deregistering something that has been registered before), this must obviously be done by the same service (recall that the same formalism can be implemented by several services). For that, a client (e.g., the ECA engine when registering event patterns) must do bookkeeping where something has been registered, and contact this service later again.

In many cases (also others), a client can take advantage from *caching* information about a “friend” service. Then, it will not have to ask the LSR each time.

Robustness against changing Service Descriptions. When changing the Service Description of some service, its owner must update the data about its service in the LSR (which will be a service of its own in a later version). When a client’s request that used communication bookkeeping fails, the client should ask the LSR if the service description of that service changed.

9.6 Contacting Domain Services

As described in Section 8, the communication with domain services is done via Domain Brokers. Thus, Atomic Event Matchers and Query Engines do not contact the domain nodes directly, but communicate with an appropriate Domain Broker. For that, *Domain Service Registries (DSRs)* are used.

For the current prototype, there is only a single domain and Domain Broker (whose address is hardcoded); thus the development of a DSR is postponed.

9.7 Issue: Redundancy and Duplicates in Communication

When a larger number of nodes (application nodes, domain brokers, language nodes) is considered, there is a problem of redundancy and duplicates: if an application node sends events to multiple event brokers, and each of them forwards the event to the AEMs that registered for the event type, there will be multiple detections of the event.

The duplication can be handled either by bookkeeping, or by restricting the communication. Since bookkeeping would require an agreement between multiple services that also influences the internal processing of events, a restricted communication strategy is recommended, as sketched below:

For an application node that supports domains d_1, \dots, d_n :

- for each d_i , it registers at all domain brokers for d_i that it considers relevant, trustable etc.
- for each domain d_i , it sends its events only to one domain broker, and executes actions only from that domain broker (i.e., registers its actions only at that domain broker).
- for each domain d_i , it registers for the static notions that it supports at all of the domain brokers for d_i . Queries for d_i are thus answered by all domain brokers.

For an AEM:

- for an event specification that uses an event type t in domain d , it registers for t at all domain brokers for d that it considers relevant, trustable etc.

For an action engine:

- for an action in domain d to be executed, it sends the action to all domain brokers for d that it considers relevant, trustable etc.

For query engines:

- not yet relevant since we only have opaque queries against certain sources, but neither a “global” query language nor an integrating query answering engine.

So far, the communication prevents duplication of events or actions:

- integration of events: takes place at the AEMs and CEDs,
- integration of data (queries) takes place at the Domain Brokers.

The remaining problem are the ACA rules.

For domains:

- every domain (owner) distinguishes a primary domain broker that submits the ACA rules to *one* ECA engine (which will collect relevant events and action executions via the AEMs).

Chapter 10

Implementation and Prototype

This chapter describes the current state of the prototype implementation at Göttingen University. As pointed out in the previous section about global architectures, the implementation of the MARS Framework consists of separate modules which can also exist in different instances, even in different implementations, sharing common interfaces as defined above. The interfaces can be refined during the development with the reference prototype.

The current architecture consists of modules of the following kinds:

- (language-independent) ECA engines,
- (language-independent) atomic event matching engines (AEM),
- composite event detection engines (CED),
- query language engines (QLE), optionally directly coupled with a database,
- action language engines (ALE),
- Generic Request Handlers (GRH; optionally integrated as instances into the ECA or ALE engines),
- domain information brokers/services (DB),
- language and service registries (“Yellow Pages”) (LSR),

We will usually use the term “Service” when looking from the outside (i.e., interfaces and URLs), and “Engine” when looking at the inside, i.e., algorithms and languages.

The prototype is implemented incrementally. It will contain at least one instance of every language type.

10.1 Simple Setting for an ECA Prototype

For running a simple ECA prototype, the components have to be as simple as possible. The first step is to simulate them at all by “faking” appropriate `<logvars:answers>` messages, and in the next step simple component expressions are used:

- Events: atomic events are specified in the domain ontologies. They have a name, and properties. Atomic events have been described in Section 5.2.1.
- Queries: atomic facts are different from atomic events. Considering the instance level of RDF and OWL, they are (i) class atoms and (ii) property triples. in XML, atomic facts are mixed up in the tree structure, and also the query languages are not based on atomic expressions.

The simplest case here is to use existing services, i.e., to query existing databases. This is e.g. done by XQuery. Here, *opaque expressions* are used. The use of opaque query expressions is investigated in Section 10.1.2.

- Tests: atomic tests are mainly predicates that are independent from the application domain, like “=”, “<”, “>” etc. Domain-dependent predicates have to be dealt with in the query part. Nevertheless, the simplest case of a test is – no test at all.
- Actions: atomic actions are specified in the domain ontologies. They have a name, and are invoked with arguments. Similar to the query component, also *opaque* actions can be considered which are also easily available by existing Web Services. Simple actions also include raising events (and thus introduce a feedback cycle) and sending messages (smtp).

The most simple, standalone ECA prototype works as follows:

- fake responses from event detection
- use opaque queries
- write the actions to be executed into a log.

10.1.1 ECA Module Implementation

The prototype is implemented incrementally. The first module was an ECA Engine prototype whose first version has been developed by simulating the other services [68] (first presentation in late 2005).

The ECA engine implements the functionality described in Sections 4.1-4.4: registration of rules, breaking them into parts, registering the event part at an appropriate service, receiving answers (variable bindings), invoking query services, evaluating conditions and invoking action services. It is responsible for dealing with constructs on ECA-ML, i.e., the *eca* namespace.

Architecture

The core *ECA engine* implements *only* the handling of variable bindings and the ECA rule semantics. It is supported by an instance of a *Generic Request Handler* (see Section 9.5.4) for communicating with autonomous component services. The GRH communicates the components to suitable services. It also receives the answers and wraps them into the agreed format of variable bindings (in case of non-framework-aware services). The architecture is as shown in Figure 9.3.

ECA: Functionality

- ECA semantics (handling variable bindings, joins),
- control flow in ECA rules.

ECA: Interface

- in from above: registration of rules,
- out to below: requests to GRH: components in XML Markup of the individual languages,
- in from below: answers from GRH, in <variable-bindings> format.

Lessons Learnt. During the development of the first version, separately developed in a BSc thesis [68], the GRH was “invented” and closely coupled with the ECA engine. It did not only communicate requests and answers, but also had access to the rules where it executed the methods for their logical semantics. This turned out not to be appropriate since the GRH functionality must also be used by the (later) CCS engine. A redesign took place.

10.1.2 Handling Opaque Queries

The conceptually simplest step (making a first ECA prototype run) is to investigate opaque queries. Since many repositories in the Web are non-semantic, the handling of opaque queries in XPath or XQuery is also required for the full MARS Framework.

10.1.2.1 Communicating with Primitive Services

When using opaque components, the most basic assumption is that the service is not aware of anything in the MARS Framework (i.e., neither that the language itself is marked up, nor that it receives any variable bindings according to the format described in Section 4.1.2).

In these cases, the opaque fragment is submitted as a string; variables must be included in the opaque code:

- XPath and HTTP-based queries: variable occurrences in the query string are string-replaced by the actual value (only possible for literals).
- XQuery queries: variables are bound by `let`-statements before submitting the query to the service (possible for literals and XML fragments).

For this, it must be possible to see or derive from the `eca:opaque` element which variables are used.

- general case: additionally to its text contents (which is the event/query/action statement), the `eca:opaque` element can contain elements of the form `<input-variable name="eca-var-name" use="local-var-name" />` that indicate that the value of `eca-var-name` should be replaced for the string `local-var-name` in the query (e.g., for embedding JDBC where variables are only named ?1, ?2 etc.)
- short forms: when the variable names in the opaque part coincide with those on the ECA level, `<eca:opaque input-variables="var1 ... varn">` or multiple subelements of the form `<input-variable name="var_i" />` can be used.
- XPath: if the string `"$var-name"` occurs in the query, it refers to the same variable in the ECA rule. It must be replaced.
- XQuery: if the string `"$var-name"` occurs in the query, it refers to the same variable in the ECA rule. It must be replaced, or a `let`-statement must be written in front.

In these simple cases, the query does not bind additional variables but the results are just answer sets. They are then only bound to variables on the rule level.

This mechanism can deal with most simple HTTP GET Web Services. Applications are here mainly application services, but also generic (query) language services can be used.

10.1.2.2 Framework-Aware Wrappers

Non-MARS-aware services can be wrapped into framework-aware ones. With the latter, communication can take place as described in Section 4.2. The wrapper performs the tasks as described just above. For each tuple of distinct variable bindings (after restricting to the variables actually used in the component!), the query (or action) must be evaluated and handled separately. This makes it reasonable to provide a generic wrapper functionality that fills the gap between the ECA functionality and primitive services:

- the ECA service submits opaque fragments to a wrapper using the above format for downward communication.
- the wrapper then appropriately calls the primitive service as described above,
- and collects the answers and gives them back using the above format for upward communication. Note that opaque fragments can bind variables, or they only return functional services.

Generic Opaque XQuery Wrapper. As already mentioned on page 33, framework-aware wrappers for non-framework-aware language services (e.g., XQuery) are both useful for developing a prototype and for integrating frequently used Web Services.

As a specific format we propose to support XQuery (e.g. for opaque query components against documents on the Web) by a standalone wrapper service that accepts the format given in Section 4.2.2 and employs a plain service. It iterates over the input variable bindings and evaluates queries by adding `let var := xml-fragment` statements in front of the query. It collects and joins the results and returns them in the upward communication format discussed in Section 4.2.3.

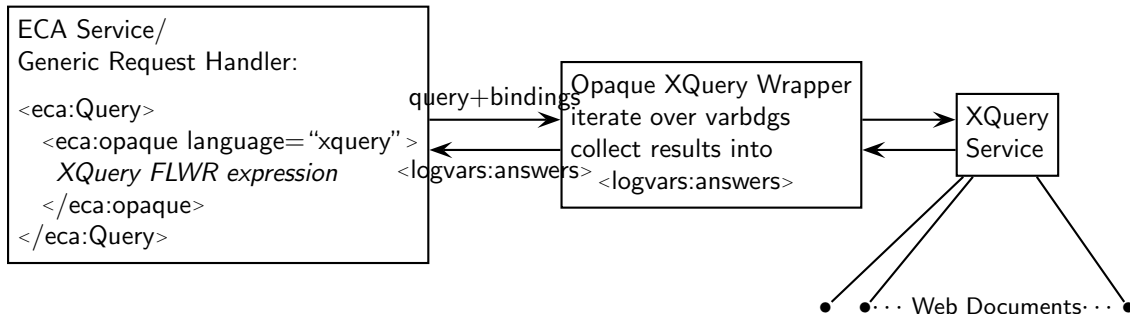


Figure 10.1: Architecture: Generic Opaque Prototype Pattern (e.g., XQuery)

Generic Action Services. In the same way as there are domain query services and generic services that e.g. evaluate XQuery queries that contain a `document(...)` expression, there can be services that execute explicit updates. Since there is not yet an update language for the Semantic Web, neither for XML, we give only a sketch of such a fragment:

```

<eca:Rule>
  :
  <eca:Action>
    <eca:atomic>
      <eca:opaque language="xquery+updates" >
        <eca:input-variables names="flight reason today" />
        update
        for $f in document("http:www.lufthansa.de/schedule.xml")//flight
        where $f/@nr=$flight and $f/@date=$today
        insert <canceled> {$reason} </canceled> into $f
      </eca:opaque>
    </eca:atomic>
  </eca:Action>
</eca:Rule>
  
```

10.1.2.3 Raising Events by Opaque Atomic Actions

As described in Section 3.3.6 for the language binding for opaque components, there the language binding is done via an explicit service URL that is contacted via an HTTP GET method. This mechanism can be “misused” especially in the prototype for raising events in the action component: the URL is the url of the service where the event should be “visible” (which is in general a domain broker).

```

<eca:Action>
  <eca:opaque method="post" uri="http://www.cs.uni-goe.de/munopag/incomingevents" >
    <uni:register subject="Databases" name="John Doe" />
  </eca:opaque>
</eca:Action>
  
```

```
</eca:opaque>
</eca:Action>
```

“sends” the event representation `<uni:register subject=“Databases” name=“John Doe”/>` to the exam administration system of the CS department.

10.1.2.4 Web Service Calls via HTTP/SOAP

In the same way as queries via HTTP GET are supported, also actions via HTTP or SOAP Web Service Calls are supported by an appropriate wrapper.

10.1.2.5 Opaque: Matching Regular Expressions

Many applications require to process non-XML contents, often text (e.g., from RSS feeds). For this, the “event” contains the RSS body, and the query part must extract variable bindings from this.

We propose to use regular expressions for extracting such data by regular expressions. The uses the same `match` predicate as defined as a built-in predicate in [46] for the Florid system in F-Logic. The proposed markup is as follows:

```
<eca:Query>
  <eca:opaque language=“regexp”>
    match(“in-variable”, “regexp”, “out-variable+”)
  </eca:opaque>
</eca:Query>
```

The *in-variable* is the input, the regular expression extracts *n* values that are bound to the *n out-variables* given in the list.

10.1.2.6 Summary

The above kinds of services have been implemented together with the ECA engine as a first, simple version of a prototype. A demo+testing version is available since early 2006 at <http://www.semwebtech.org/eca/frontend>. Since it is an experimental environment, it is sometimes down due to restructurings.

10.2 Extending the Prototype with Component Services

In the second step, component engines (CED+AEMs and ALE) have been implemented that communicate with the ECA engine: they are written wrt. the Göttingen ECA engine and are also expected to participate in cross-communication with the Lisbon one.

10.2.1 Composite Event Detection and Atomic Event Matchers

Services for composite event detection according to the general interfaces are developed separately.

A first integrated CED+AEM: Snoopy. A CED for a language closely related to the SNOOP event algebra [24] of the “Sentinel” active database system has been implemented in [70]. The CED uses the XML markup presented in Section 5.2.6.

As a separate, standalone thesis, the original version incorporated composite event detection in the SNOOP language directly with an AEM for the XML-QL Style matching formalism described in Section 5.2.2.1. The task of the thesis was specified to do this completely in XML/XSL (which is not the most efficient solution, but provides a better demonstration of the algorithm):

- an XSL stylesheet that maps a SNOOP expression into an automaton, represented in XML,
- an XSL stylesheet that, given an atomic event, transforms the XML representation of an automaton situation into the representation of the subsequent state and outputs the detected composite events.

The resulting code has been split according to the performed tasks into a CED and an AEM module as follows:

Separate XML-QL-Matching style AEM. The matching part of the stylesheet has been separated and wrapped into an XML-QL-Matching style AEM. The module is available since early 2006 in the prototype and allows to replace the “faked” answers to event detection by actually sending atomic events.

Separate Snoopy CED. The remaining part that implements the SNOOP language is installed as a separate service.

An XPath-based AEM. Another AEM using the XPath-navigation based formalism described in Section 5.2.2.2 has also been implemented as an XSL stylesheet. It is currently being wrapped.

10.2.2 Queries and Updates

Here, several alternatives already exist that only need to be wrapped according to the general interface. Pure language implementations for stating queries on independent XML instances have to be distinguished from XML databases. In most cases, *opaque* query components will be used since there are not yet query languages using XML-markup.

Saxon. Saxon [40] is an XQuery implementation that allows to state queries against XML sources on the Web. A framework-aware wrapper around saxon makes it immediately integrable. Saxon does not deal with updates.

Commercial XML-enabled and XML database systems. With Oracle, IBM DB2, MS SQL Server etc., most classical relational systems have been XML-enabled in the last years. For queries, the SQLX standard [28] is supported that embeds XPath into SQL. Updates are implemented via transformations; thus triggers on the XML level as described in Section 3.2.1 are currently not supported.

eXist. eXist [30] is an open-source XML database that runs as a Web service. For queries, XQuery is supported; also XSLT is supported for transformations. Updates are possible via XUpdate and the XQuery+Update extension in the style of [72, 45]. A framework-aware wrapper around an eXist database that supports queries and XUpdate has been implemented. An eXist database is used currently as a substitute stub for queries against domain nodes. Several domain node architectures (Oracle, Jena) and a generic Domain Broker are currently under implementation.

Jena. Jena [38] is an open-source framework that provides functionality as an RDF query language; also OWL reasoners can be integrated. A Jena-based node has been developed in [73], using Pellet [64] as attached OWL reasoner. It supports queries (SPARQL), atomic updates, and database triggers on the RDF/OWL model:

- for the event part see Section 3.2.2.
- query part: an SPARQL query that may use the variables bound in the event part (which is only one tuple in this case) and binds additional variables that occur free in the query;

- test part: included with query part (conjunctive query);
- action part:
 - update of the local RDF database (opaque) of the form


```
<eca:opaque> rdf update expression</eca:opaque>
```
 - raising an (internal) event (analogously to RAISE_EVENT for the Oracle node described below):


```
<eca:raise-event> atomic event in markup </eca:raise-event>
```

The raised event is always sent to the “hull” of the node.

Note that the bodies can contain free variables (that in the opaque case need to be declared as `input-variable` in the markup according to Section 10.1.2). The action must be carried out for *each* tuple of variable bindings (that can be several after evaluating the query).

SQL Database Node. An SQL-based node (sample scenario: car rental company) is under implementation in [33]. It is based on Oracle, especially its Rule Engine (available since version 10g).

10.2.3 Composite Actions and Processes

A CCS-based action engine implementing the behavior described in Section 5.6 is under implementation. It shares the GRH functionality with the ECA engine.

10.3 Application Domains

Domain Nodes. Application nodes consist mainly of local databases and application functionality, as much as possibly also defined by rules. Domain nodes based on SQL, XML, and RDF are under implementation.

RDF: The functionality described in Section 7 provides for a “native” Semantic Web node architecture consisting of an RDF/OWL node that additionally emits events and executes actions of the domain ontology.

XML: An eXist database is used currently as a substitute stub for queries against domain nodes. The nodes under development will provide local active functionality (e.g., by internal database triggers) and support for events, queries and actions according to Chapter 7.

SQL: The use of the Oracle Rule Engine for SQL-based node (sample scenario: car rental company) has been evaluated in [33]. A wrapper for providing domain node functionality is planned.

Domain Broker. A generic domain broker according to Chapter 8 has been implemented in [43].

10.4 Infrastructure

The first prototype used hard-coded URIs and a fixed communication format/wrapping. Its development has led to a clear understanding and isolation in a separate class of the GRH functionality. It has been recently extended with an LSR as an XML/RDF file (see Section 9.5.3) which can be found at <http://www.semwebtech.org/2006/1sr/1sr.rdf>. The current refinement and adaptation to the Lisbon Services will lead to a full specification and trading functionality of the LSR and GRH modules.

10.5 Using the Prototype Demonstrator

The demo frontend at <http://www.semwebtech.org/eca/frontend> provides a choice of predefined actions:

- register and deregister predefined rules. Some rules serve for illustrating the ECA functionality (atomic events, simple opaque queries, faked actions), others serve for testing composite events,
- send (predefined) faked event detection answers to the ECA engine,
- send (predefined) atomic events to the AEMs: they will match them and inform the CED,
- Using the form, users can also write own rules and register them,
- Using the form, users can also write own events and send them to the AEMs.
- Users can also program an own service and feed the AEMs with atomic events via HTTP following the communication specification given in the LSR.

10.6 Anticipated “Foreign” Modules

10.6.1 Wrapped Composite Event Detection Engines

XChange. XChange is currently closely connected with the Xcerpt system. Here, a wrapper that returns variable bindings instead of immediately submitting them to Xcerpt is required. A first version can be a simple Xcerpt rule that sends a message that contains simply the XML serialization of the variable bindings. For this XChange/Xcerpt must be wrapped in a Web Service.

RuleCore. The RuleCore event detection module returns variable bindings. Wrapping it as a Web Service makes it immediately integrable.

10.6.2 Wrapped Query Engines

Florid/LoPiX. Florid (F-Logic Reasoning in Databases) [32] is an implementation of the F-Logic knowledge management and database formalism [41]. Its migration to XML, LoPiX (Logic Programming in XML) [47], provides similar features based on a data model that extends XML. A wrapper into a Web Service has been implemented recently which makes Florid/LoPiX available as a service for opaque queries.

Xcerpt. Xcerpt is the XML query language developed by REVERSE WG I4. It allows to state queries against XML sources on the Web. Wrapping it as a Web Service makes it immediately integrable. A database implementation and a lifting to RDF are planned. Xcerpt supports updates of documents via XChange’s update actions. Updates are currently implemented by translating them into the definition of a view and then replacing the original document by this view. Thus, triggers on the XML level as described in Section 3.2.1 are currently not supported.

Chapter 11

The RDF Level: Language Elements and their Instances as Resources

Rules on the semantic level, i.e., RDF or OWL, lift ECA functionality wrt. two (independent) aspects: first, the events, conditions and actions refer to the ontology level as described above. On an even higher level, the above rule ontology and event, condition, and action subontologies regard rules themselves as objects of the Semantic Web. Together with the languages and their processors, this leads directly to a resource-based approach: every rule, rule component, event, subevent etc. becomes a resource, which is related to a language which in turn is related to other resources.

Describing rules in RDF will allow to reason about rules. This will support several things:

- validation and support for execution (e.g., use of variables)
- actual reasoning about the behavior of a node, including correctness issues.
- expressing rules in abstract terms instead of wrt. concrete languages. The services can then e.g. choose which concrete languages support the expressiveness required by the rule's components.

Events and Actions on the RDF Level. For the RDF/OWL level, we assume that not only the data itself is in RDF, but that also events and actions are given as XML/RDF fragments (using the same URIs for entities and properties as in the static data).

11.1 Ontologies vs. Markup Languages

MARS assumes that each language defines an own, autonomous set of element names (DTD/XML Schema, XML world) and/or ontology (OWL world), e.g. the ECA-ML, SNOOP, and CCS:

- the OWL ontology defines *all* notions of the language. This is used when representing things in RDF, e.g., for reasoning.
- an XML Markup is given by a DTD that consists of a subset of the ontology notions. The XML markup is a *stripped* version of a *specific* RDF/XML serialization of the RDF graph. It is used e.g., for data exchange between services and when a services processes a fragment.

Services may use each of the representations. Thus, there must be translations in both ways. The translation is concerned with two issues:

- translations between ontology and markup of a certain language, and
- glueing between languages/ontologies (i.e., how to compose a rule from ECA-ML, SNOOP and CCS notions).

11.1.1 Fragment Translation between both Representations

In the RDF graph, everything, including the language itself is a resource. In the XML representation, the XML tree represents only the elements of the languages as a nested structure.

Language Elements and Element Names: while in the RDF world, every class is represented by an arbitrary URI, in the XML representation, only *namespace:elementname* where *elementname* is a QName is allowed. For that, it is strongly recommended that languages do not use hierarchical namespaces, but flat *hash* namespaces.

Language Indicators: for every class, the `rdfs:isDefinedBy` property refers to the language where it belongs to. E.g., For all ECA ontology concepts, `rdfs:isDefinedBy` points to <http://www.semwebtech.org/languages/2006/eca-ml>.

In the XML representation, the language is encoded in the namespace. The namespace of an element indicates the language. Recall that the namespace declaration can be omitted in subelements that reside in their parent's namespace.

Mapping: The XML markup is based on one of the RDF/XML serializations of the RDF graph. For the mapping (in both directions), the following is given:

- a DTD that specifies the RDF/XML serialization that is used. Then, the serialization can be based on a depth-first algorithm based on the target DTD. In case an element declaration contains a possible subelement twice, any position of them can be chosen in the actual mapping.
- a specification which intermediate (mostly: properties) elements are omitted in the stripped XML markup (see below how to declare this).

Additional Properties: In addition to the language elements itself, some meta-properties of language elements have to be represented:

- `rdfs:isDefinedBy` property for an explicit link of every notion to its language
- `mars:markup-omit` property with literal value “yes” or “no” to each such ontology language) [or define `markup-omit` in the language's own namespace]

If an “imported” non-MARS-origin language specifies own properties of that kind, `rdfs:subClassOf` and `rdfs:subPropertyOf` can be used for integration.

11.1.2 Glue in XML and RDF Representations

Elements (in XML) or notions (in RDF) representation must be glued together between bordering languages or ontologies.

Glue in RDF/OWL: Glue Axioms. In the RDF/OWL representation, the glue is provided by axioms that identify notions of the ontologies to be connected (using `rdfs:subClassOf` and `rdfs:subPropertyOf`). For instance, the SNOOP ontology's central class is `snoopy:Event`, whereas the ECA-ML ontology defines `eca:Event`. For that, *glue axioms* are added that state e.g.

```
snoopy:Event rdfs:subClassOf eca:Event ,
ccs:Action rdfs:subClassOf eca:Action ,
ccs:Event rdfs:subClassOf eca:Event .
```

Note that for the instance-based *processing*, these glue axioms are not necessary, if the individual ontologies specify `rdfs:domain` and `rdfs:range` of the tree edges accordingly. By this, for the bordering nodes, class membership wrt. both bordering ontologies can be derived.

Glue in XML: Embedding. In XML, fragments of the embedded language are embedded into elements of the outside language. For instance, SNOOP expressions are embedded inside `eca:Event` elements as described in Section 5.2.4/Example 18:

```
<eca:Rule ... >
  <eca:Event xmlns:snoopy="http://www.semwebtech.org/languages/2006/snoopy#" >
    <snoopy:Cumulative>
      :
    </snoopy:Cumulative>
  </eca:Event>
  :
</eca:Rule>
```

11.1.3 Design of Ontologies, Languages and Namespaces

When looking at existing language proposals, there is not yet a homogeneous handling of namespaces and ontology descriptions by OWL files behind them. For the RDF language, the design is as follows:

- The URL `http://www.w3.org/1999/02/22-rdf-syntax-ns#` is given as RDF namespace. There, an actual RDF/RDFS/OWL document (with name `22-rdf-syntax-ns.htm`) can be found.
- The URL `http://www.w3.org/2000/01/rdf-schema#` returns a document `rdf-schema.htm`.

Note that nevertheless, the URI `http://www.w3.org/2000/01/rdf-schema#Property` as a URL does not *address* the `rdfs:Property` resource (which actually does not exist):

```
<!-- excerpt from 22-rdf-syntax-ns.htm --> <rdfs:Class
rdf:about="http://www.w3.org/1999/02/22-rdf-syntax-ns#Property">
  <rdfs:isDefinedBy rdf:resource="http://www.w3.org/1999/02/22-rdf-syntax-ns#"/>
  <rdfs:label>Property</rdfs:label>
  <rdfs:comment>The class of RDF properties.</rdfs:comment>
  <rdfs:subClassOf rdf:resource="http://www.w3.org/2000/01/rdf-schema#Resource"/>
</rdfs:Class>
```

If `rdf:ID` would have been used instead of `rdf:about`, the Class would have been the resource itself.

- Namespaces can either be defined as hash-namespace or as slash-namespace (cf. [65]). Most namespaces in MARS are hash-namespaces. Recall that XML markup allows only for qnames as identifiers within the namespaces.

The correct decision is as follows:

- The namespace of ECA-ML is `http://www.semwebtech.org/languages/2006/eca-ml#`.
- This gives two possibilities:
 - `http://www.semwebtech.org/languages/2006/eca-ml` is a directory that contains the ECA-ML OWL file and the server delivers this file when the directory is accessed (independent whether with or without hash). This is currently done by linking the file to `index.html`.
 - `http://www.semwebtech.org/languages/2006/eca-ml` is the name of the OWL file. Then, no server mapping is required.

Analogously for the other language namespaces.

11.2 Ontology of Expressions

Languages in general define *expressions* and their semantics. The expressions' ontology is divided into three parts:

General Expressions are either CompositeExpressions, AtomicExpressions (wrt. that language; this can be composite expressions according to an *embedded* language), or Variables.

Algebraic expressions are a special subclass of expression that adheres to a term structure, i.e., in addition to subexpressions they are related to an operator.

Variables: Expressions *use* variables in certain ways. Note that variables are actually *resources* that are referenced.

The below ontologies provide the required concepts. They are included e.g. from the Snoopy Ontology in Section 11.7 and from the CCS Ontology in Section 11.8.

11.3 Subontology of General Expressions

General expressions cover algebraic and non-algebraic expressions.

```
<!DOCTYPE rdf:RDF [  
  <!ENTITY xsd "http://www.w3.org/2001/XMLSchema#">  
  <!ENTITY mars "http://www.semwebtech.org/2006/mars#"> ]>  
<rdf:RDF  
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"  
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"  
  xmlns:owl="http://www.w3.org/2002/07/owl#"  
  xmlns:expr="http://www.semwebtech.org/2006/expressions#"  
  xml:base="http://www.semwebtech.org/2006/expressions">  
  
  <owl:Ontology rdf:about=""/>  
  
  <!-- general languages and expressions --> <owl:Class  
  rdf:about="Expression"/>  
  
  <owl:FunctionalProperty rdf:about="#has-language">  
    <rdfs:domain rdf:resource="#Expression"/>  
    <rdfs:range rdf:resource="#mars;Language"/>  
  </owl:FunctionalProperty>  
  
  <rdfs:Class rdf:about="#Expression">  
    <owl:unionOf rdf:parseType="Collection">  
      <rdfs:Class rdf:about="#Variable"/>  
      <rdfs:Class rdf:about="#AtomicExpression">  
        <owl:disjointWith rdf:resource="#Variable"/>  
      </rdfs:Class>  
      <rdfs:Class rdf:about="#CompositeExpression">  
        <owl:disjointWith rdf:resource="#Variable"/>  
        <owl:disjointWith rdf:resource="#AtomicExpression"/>  
      </rdfs:Class>  
    </owl:unionOf>  
  </rdfs:Class>
```

```

<rdf:Property rdf:about="#has-subexpression">
  <rdfs:domain rdf:resource="#CompositeExpression"/>
  <rdfs:range rdf:resource="#Expression"/>
</rdf:Property> </rdf:RDF>

```

[Filename:expr.rdf]

Note: expr:has-subexpression is most generic – it allows for ordered/unordered and named/unnamed relationships.

11.4 Subontology of Algebraic Expressions

```

<!DOCTYPE rdf:RDF [
  <!ENTITY mars "http://www.semwebtech.org/2006/mars#"> ]>
<rdf:RDF
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:owl="http://www.w3.org/2002/07/owl#"
  xmlns:expr="http://www.semwebtech.org/2006/expressions#"
  xml:base="http://www.semwebtech.org/2006/expressions">

  <owl:Ontology rdf:about=""/>

  <owl:Class rdf:about="#AlgebraicExpression">
    <rdfs:subClassOf rdf:resource="#Expression"/>
    <owl:equivalentClass>
      <owl:Restriction>
        <owl:onProperty rdf:resource="#has-language"/>
        <owl:allValuesFrom rdf:resource="#&mars;#AlgebraicLanguage"/>
      </owl:Restriction>
    </rdfs:equivalentClass>
  </owl:Class>

  <owl:Class rdf:about="#Operator"/> <owl:Class
  rdf:about="#Arities"/>
  <!-- should be defined as xsd:nonNegativeInteger + positive Infinity -->

  <rdf:Property rdf:about="#arity">
    <rdfs:domain rdf:resource="#Operator"/>
    <rdfs:range rdf:resource="#Arities"/>
  </rdf:Property>

  <rdf:Property rdf:ID="provides-operator">
    <rdfs:domain rdf:resource="#AlgebraicLanguage"/>
    <rdfs:range rdf:resource="#Operator"/>
  </rdf:Property>

  <rdf:Property rdf:about="#arity">
    <rdfs:domain rdf:resource="#Operator"/>
    <rdfs:range rdf:resource="#Arities"/>
  </rdf:Property>

  <owl:Class rdf:about="#AlgebraicCompositeExpression"/>

```

```

<owl:EquivalentClass>
  <owl:intersectionOf rdf:parseType="Collection">
    <owl:Class rdf:about="#AlgebraicExpression"/>
    <owl:Class rdf:about="#CompositeExpression"/>
  </owl:intersectionOf>
</owl:EquivalentClass>
<rdfs:subClassOf>
  <owl:Restriction>
    <owl:onProperty rdf:resource="#has-operator"/>
    <owl:cardinality>1</owl:cardinality>
  </owl:Restriction>
</rdfs:subClassOf>
</owl:Class>

<owl:FunctionalProperty rdf:about="#has-operator"/>
  <rdfs:domain rdf:resource="#AlgebraicCompositeExpression"/>
  <rdfs:range rdf:resource="#Operator"/>
</owl:FunctionalProperty> </rdf:RDF>

```

[Filename:algexpr.rdf]

11.5 Subontology of Variable Usage

Variables are resources, usually associated with a rule, or in case of local variables, by a fragment (see examples in later sections). Expressions *use* variables in several ways. The use and handling of variables has already been discussed in Sections 4.1.5 and 4.3. The following relationships can hold between an expression and a variable:

use-variable: a variable occurs in any way in (the context of) an expression.

free variables: a variable occurs free. Such variables belong to the “outer interface” of an expression.

bound variables: a variable occurs bound (i.e., inside in the scope of a quantifier). Such variables are not relevant to the outside of an expression.

positive variables, negative variables: used in logic-based frameworks (Datalog, F-Logic).

scoping of variables: a quantifier scopes a variable, that (usually) occurs *free* in a subexpression, and is then *bound*

input variables: these variables must be bound when evaluating an expression. In logical frameworks, *negative* variables are input variables.

output variables: these variables are bound to values when evaluating an expression (e.g., OUT variables in procedures). If such a variable is already bound before evaluating an expression, it *may* be used as a join variable. For logical frameworks, it is recommended to use the notion of *positive* variables instead.

An additional side-effect that is used in many languages is to bind the results of subexpressions to variables (e.g., in F-Logic). They are to be considered as “used”, “free” and “positive” by the expression.

The actual values of these properties for a given expression can usually be determined by an appropriate service that knows about the syntax of the language. Note that if such languages use an own ontology, an ontology mapping is required. This is usually defined by (deductive) rules.

```

<!DOCTYPE rdf:RDF [
  <!ENTITY xsd "http://www.w3.org/2001/XMLSchema#">
  <!ENTITY mars "http://www.semwebtech.org/2006/mars#"> ]>
  <!ENTITY expr "http://www.semwebtech.org/2006/expressions#"> ]>
<rdf:RDF
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:owl="http://www.w3.org/2002/07/owl#"
  xmlns:expr="http://www.semwebtech.org/2006/expressions#"
  xml:base="http://www.semwebtech.org/2006/expressions">

  <owl:Ontology rdf:about=""/>

  <!-- notions around variables --> <!-- The class "Variable" has
  been defined above -->

  <rdf:Property rdf:about="#uses-variable">
    <!-- variable occurs in some way in this expression -->
    <owl:inverseOf rdf:resource="#variable-occurs-in"/>
    <rdfs:domain rdf:resource="#Expression"/>
    <rdfs:range rdf:resource="#Variable"/>
  </rdf:Property>

  <rdf:Property rdf:about="#has-free-variable">
    <rdfs:domain rdf:resource="#Expression"/>
    <rdfs:range rdf:resource="#Variable"/>
    <rdfs:subPropertyOf rdf:resource="#uses-variable"/>
  </rdf:Property>

  <rdf:Property rdf:about="#has-bound-variable">
    <rdfs:domain rdf:resource="#Expression"/>
    <rdfs:range rdf:resource="#Variable"/>
    <rdfs:subPropertyOf rdf:resource="#uses-variable"/>
  </rdf:Property>

  <rdf:Property rdf:about="#has-positive-variable">
    <rdfs:domain rdf:resource="#Expression"/>
    <rdfs:range rdf:resource="#Variable"/>
    <rdfs:subPropertyOf rdf:resource="#uses-variable"/>
  </rdf:Property>

  <rdf:Property rdf:about="#has-negative-variable">
    <rdfs:domain rdf:resource="#Expression"/>
    <rdfs:range rdf:resource="#Variable"/>
    <rdfs:subPropertyOf rdf:resource="#has-input-variable"/>
    <rdfs:subPropertyOf rdf:resource="#uses-variable"/>
  </rdf:Property>

  <rdf:Property rdf:about="#scopes-variable">
    <rdfs:domain rdf:resource="#Expression"/>
    <rdfs:range rdf:resource="#Variable"/>
    <rdfs:subPropertyOf rdf:resource="#uses-variable"/>
  </rdf:Property>

```

```

<rdf:Property rdf:about="#has-input-variable">
  <rdfs:domain rdf:resource="#Expression"/>
  <rdfs:range rdf:resource="#Variable"/>
  <rdfs:subPropertyOf rdf:resource="#uses-variable"/>
</rdf:Property>

<rdf:Property rdf:about="#has-output-variable">
  <!-- mainly for non-logical things
  <rdfs:domain rdf:resource="#Expression"/>
  <rdfs:range rdf:resource="#Variable"/>
  <rdfs:subPropertyOf rdf:resource="#uses-variable"/>
</rdf:Property>

<rdf:Property rdf:about="#bind-to-variable">
  <rdfs:domain rdf:resource="#Expression"/>
  <rdfs:range rdf:resource="#Variable"/>
  <rdfs:subPropertyOf rdf:resource="#uses-variable"/>
</rdf:Property> </rdf:RDF>

```

[Filename:vars.rdf]

Note that the mapping of variables to an XML Markup can be done in many ways as discussed in Section 4.3 for “native XML markup”.

11.6 The ECA-ML Ontology

The ECA-ML language induces an ontology of rules. The complete OWL file can be found in Section 11.6.4. Obviously, the main classes are Rule, Event (which actually means “event specification”), Query, Test, Action (which actually means “action specification”).

11.6.1 Handling Opaque Stuff

The handling of opaque fragments differs from the above XML design since OWL allows for subclassing: `eca:OpaqueQuery` is defined as a subclass of `eca:Query`; similar for the other concepts. All are subclasses of `eca:Opaque` (which is not a subclass of any of the other concepts – being `eca:Opaque` does only tell the ECA engine that something has to be handled specifically).

11.6.2 Assertions

The RDF level of rules provides the base for reasoning about rules. As it is also preferable to allow for opaque legacy components, these can be accompanied by *assertions* that formalize the relationship between the variables.

For the XML representation, `eca:Assertion` elements can be added after each component. In contrast, on the RDF level, there is no ordering of properties of a rule. Nevertheless, as variables are bound exactly once, the semantics of `eca:has-Assertion` edges can be defined unambiguously by “*an assertion is guaranteed to hold as soon as all variables in it are bound*”. Note that this corresponds to the semantics of null values in constraints in relational databases.

Apart from using assertions for reasoning, rule execution can use them as constraints to be verified when all variables are bound – and raising an exception if the assertion is not satisfied. While a test simply discards the variable bindings that do not pass (if no bondings remain, the firing of the rule ends), for an assertion failure of any tuple must cause a serious exception.

11.6.3 Handling Atomic Stuff

Atomic subexpressions belong to some other language or formalism (e.g., an AESL) or a domain (e.g., atomic actions).

- In the RDF graph, the `rdfs:isDefinedBy` property of the class where the RDF node belongs to (which can refer to a language or an application domain) allows for identifying an appropriate service,
- in the XML Markup, the language attribute or the namespace allow to derive the “responsible” language and service.

11.6.4 The ECA Ontology, OWL File, and ECA-ML Markup

For all ECA-Ont concepts, `rdfs:isDefinedBy` points to `http://www.semwebtech.org/languages/2006/eca-ml`. Note that we define `eca:Event` as equivalent to `eca:EventSpec` (which is actually the more correct naming, as the thing is not an event, but an event specification); analogous for `eca:ActionSpec`.

```
<rdf:RDF
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:owl="http://www.w3.org/2002/07/owl#"
  xml:base="http://www.semwebtech.org/languages/2006/eca-ml">

  <owl:Ontology rdf:about=""/>

  <rdfs:Class rdf:about="#Rule">
    <rdfs:isDefinedBy rdf:resource=""/>
  </rdfs:Class>

  <rdfs:Class rdf:about="#Event">
    <rdfs:isDefinedBy rdf:resource=""/>
    <owl:disjointWith rdf:resource="#Rule"/>
    <owl:equivalentClass rdf:resource="#EventSpec"/>
  </rdfs:Class>

  <!-- note that queries, tests, and assertions are
       not required to be disjoint -->

  <rdfs:Class rdf:about="#Query">
    <rdfs:isDefinedBy rdf:resource=""/>
    <owl:disjointWith rdf:resource="#Rule"/>
    <owl:disjointWith rdf:resource="#Event"/>
  </rdfs:Class>

  <rdfs:Class rdf:about="#Test">
    <rdfs:isDefinedBy rdf:resource=""/>
    <owl:disjointWith rdf:resource="#Rule"/>
    <owl:disjointWith rdf:resource="#Event"/>
  </rdfs:Class>

  <rdfs:Class rdf:about="#Assertion">
    <rdfs:isDefinedBy rdf:resource=""/>
    <owl:disjointWith rdf:resource="#Rule"/>
    <owl:disjointWith rdf:resource="#Event"/>
```

```

</rdfs:Class>

<rdfs:Class rdf:about="#Action">
  <rdfs:isDefinedBy rdf:resource=""/>
  <owl:disjointWith rdf:resource="#Rule"/>
  <owl:disjointWith rdf:resource="#Event"/>
  <owl:disjointWith rdf:resource="#Query"/>
  <owl:disjointWith rdf:resource="#Test"/>
  <owl:disjointWith rdf:resource="#Assertion"/>
  <owl:equivalentClass rdf:resource="#ActionSpec"/>
</rdfs:Class>

<rdf:Property rdf:about="#has-event">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:domain rdf:resource="#Rule"/>
  <rdfs:range rdf:resource="#Event"/>
</rdf:Property>

<rdf:Property rdf:about="#has-query">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:domain rdf:resource="#Rule"/>
  <rdfs:range rdf:resource="#Query"/>
</rdf:Property>

<rdf:Property rdf:about="#has-test">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:domain rdf:resource="#Rule"/>
  <rdfs:range rdf:resource="#Test"/>
</rdf:Property>

<rdf:Property rdf:about="#has-assertion">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:domain rdf:resource="#Rule"/>
  <rdfs:range rdf:resource="#Assertion"/>
</rdf:Property>

<rdf:Property rdf:about="#has-action">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:domain rdf:resource="#Rule"/>
  <rdfs:range rdf:resource="#Action"/>
</rdf:Property>

<rdfs:Class rdf:about="#OpaqueRule">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subClassOf rdf:resource="#Rule"/>
  <rdfs:subClassOf rdf:resource="#Opaque"/>
</rdfs:Class>

<rdfs:Class rdf:about="#OpaqueEvent">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subClassOf rdf:resource="#Event"/>
  <rdfs:subClassOf rdf:resource="#Opaque"/>
  <owl:equivalentClass rdf:resource="#OpaqueEventSpec"/>
</rdfs:Class>

```

```

<rdfs:Class rdf:about="#OpaqueQuery">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subClassOf rdf:resource="#Query"/>
  <rdfs:subClassOf rdf:resource="#Opaque"/>
</rdfs:Class>

<rdfs:Class rdf:about="#OpaqueTest">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subClassOf rdf:resource="#Test"/>
  <rdfs:subClassOf rdf:resource="#Opaque"/>
</rdfs:Class>

<rdfs:Class rdf:about="#OpaqueAction">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subClassOf rdf:resource="#Action"/>
  <rdfs:subClassOf rdf:resource="#Opaque"/>
  <owl:equivalentClass rdf:resource="#OpaqueActionSpec"/>
</rdfs:Class>

<rdfs:Class rdf:about="#Variable"/>

<rdf:Property rdf:about="#uses-variable">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:range rdf:resource="#Variable"/>
</rdf:Property> <rdf:Property rdf:about="#bind-to-variable">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:range rdf:resource="#Variable"/>
  <rdfs:subPropertyOf rdf:resource="#uses-variable"/>
</rdf:Property> </rdf:RDF>

```

[Filename:eca-ml.rdf]

```

# call # jena -q -pellet -il RDF/XML -qf ecaml-query.sparql

PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> PREFIX
rdfs: <http://www.w3.org/2000/01/rdf-schema#>

SELECT ?X ?Y FROM <file:eca-ml.rdf> WHERE { ?X rdfs:subClassOf ?Y
}

```

[Filename:eca-query.sparql]

11.6.5 The Implicit Event Variable

Every rule has an implicit variable \$Event (its URI is *uri-of-the-rule#Event*) which can optionally also be declared explicitly in the rule as

```
<eca:Variable rdf:ID="Event" /> ,
```

and which is bound to the event instance that fired the rule. Note that the event itself is *volatile* data, thus, at the evaluation time it has to be added as a local graph, including all volatile RDF nodes, to the rule.

11.6.6 The XML Representation: ECA-ML

The XML representation of rules not containing opaque stuff if just obtained from the striped RDF by omitting the `has-event`, `has-query`, `has-test`, and `has-action` elements. Analogously, striped RDF can be obtained from the XML representation by extension with the above elements.

If a rule contains variables on the ECA level must be transformed appropriately into the XML syntax described in Section 4.3.3. Note that `eca:bind-to-variable` does most directly correspond to the handling of variables by XML attributes.

11.6.7 Some Consequences

The definitions of `rdf:range` of `has-event`, `has-query`, `has-test`, and `has-action` allow to derive that the targets are `eca:EventSpecs`, `eca:Actions` etc., even when the XML or RDF representation list them only as `snoopy:Sequence` or `ccs:Fix`. Moreover, since the SNOOP-to-MARS ontology defines `snoopy:Sequence` as an `rdfs:subClassOf`, this can also be derived.

Note that these `eca:EventSpecs` and `eca:ActionSpecs` are *not* `mars:Events` and `mars:Actions`, but in fact only event/action *patterns* or *specifications*, containing e.g. variables.

11.7 The Snoop Ontology

The Snoop classes are the operators. Note that the Snoop ontology does not refer at all to the ECA-ML ontology since it is an independent language. Note also that what SNOOP calls “Events” are actually not events, but event specifications.

Snoop defines the following classes:

- Bag-Events: And, Or, Any, MultiOccurrence. Property: `has-subevent`,
- Any-or-Multi: Composers with counting (Any and MultiOccurrence), property: `number-of-occurrences`.
- Bounded: event definitions that are expressed by an interval with something to be checked in the interval (Not, Cumulative/Non-Cumulative (A)Periodic). Subsets: Cumulative, Non-Cumulative, Not.

The interval is given by `has-start`, `has-end`. The relevant events in the interval are given by the properties `cumulate` (Cumulative), `fired-by` (Non-Cumulative) and `exit-not` (Not).

- Sequence: with its unbounded cardinality and ordered semantics, this operator is a problem in RDF (whereas in XML it is trivial since the contents is ordered).

The current proposal is that Sequence can be modeled in two ways:

- with two subevents `first` and `second` (that can be sequences again for encoding longer sequences).
- long sequences are cumbersome in that modeling. For that, subevents (of Sequences) can be serialized by a `ccs:followedBy` property that refers to the next subevent.

Classes of Expressions vs. Operators As an algebraic language, SNOOP defines operators with arities. By that, “Sequence” is also an operator, and as above a class of terms induced by that operator.

11.7.1 The Snoopy OWL File

```
<!DOCTYPE rdf:RDF [  
  <!ENTITY expr "http://www.semwebtech.org/2006/expressions#"> ]>  
<rdf:RDF
```

```

xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
xmlns:owl="http://www.w3.org/2002/07/owl#"
xmlns:expr="http://www.semwebtech.org/2006/expressions#"
xml:base="http://www.semwebtech.org/languages/2006/snoopy">

<owl:Ontology rdf:about=""/>

<owl:Class rdf:about="#Event">
  <rdfs:isDefinedBy rdf:resource=""/>
</owl:Class>

<owl:Class rdf:about="#AtomicEvent">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subClassOf rdf:resource="#Event"/>
  <owl:disjointWith rdf:resource="#CompositeEvent"/>
</owl:Class>

<owl:Class rdf:about="#CompositeEvent">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subClassOf rdf:resource="#Event"/>
  <rdfs:subClassOf rdf:resource="&expr;AlgebraicCompositeExpression"/>
  <owl:disjointWith rdf:resource="#AtomicEvent"/>
</owl:Class>

<owl:FunctionalProperty rdf:about="#has-operator">
  <rdfs:domain rdf:resource="#AlgebraicCompositeExpression"/>
  <rdfs:range rdf:resource="#Operator"/>
  <rdfs:subPropertyOf rdf:resource="&expr;has-operator"/>
</owl:FunctionalProperty>

<expr:Operator rdf:about="#AndOp">
  <expr:has-arity>2</expr:has-arity>
</expr:Operator> <expr:Operator rdf:about="#OrOp">
  <expr:has-arity>2</expr:has-arity>
</expr:Operator> <expr:Operator rdf:about="#AnyOp">
  <expr:has-arity>*</expr:has-arity>
</expr:Operator> <expr:Operator rdf:about="#MultiOccurrencesOp">
  <expr:has-arity>1</expr:has-arity>
</expr:Operator> <expr:Operator rdf:about="#SequenceOp">
  <expr:has-arity>2</expr:has-arity>
</expr:Operator> <expr:Operator rdf:about="#NotOp">
  <expr:has-arity>3</expr:has-arity>
</expr:Operator> <expr:Operator rdf:about="#PeriodicOp">
  <expr:has-arity>3</expr:has-arity>
</expr:Operator> <expr:Operator rdf:about="#CumulativePeriodicOp">
  <expr:has-arity>3</expr:has-arity>
</expr:Operator> <expr:Operator rdf:about="#AperiodicOp">
  <expr:has-arity>3</expr:has-arity>
</expr:Operator> <expr:Operator
rdf:about="#CumulativeAperiodicOp">
  <expr:has-arity>3</expr:has-arity>
</expr:Operator>

```

```

<owl:Class rdf:about="#Sequence">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subClassOf rdf:resource="#CompositeEvent"/>
  <owl:disjointWith rdf:resource="#BagEvent"/>
  <owl:disjointWith rdf:resource="#Bounded"/>
  <owl:equivalentClass>
    <owl:Restriction>
      <owl:onProperty rdf:resource="#has-operator"/>
      <owl:hasValue rdf:resource="#SequenceOp"/>
    </owl:Restriction>
  </owl:equivalentClass>
</owl:Class>

<rdf:Property rdf:about="#has-subevent">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:domain rdf:resource="#CompositeEvent"/>
  <rdfs:range rdf:resource="#Event"/>
  <owl:inverseOf rdf:resource="#subEventOf"/>
</rdf:Property>

<owl:FunctionalProperty rdf:about="#first">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subPropertyOf rdf:resource="#has-subevent"/>
  <rdfs:domain rdf:resource="#Sequence"/>
  <rdfs:range rdf:resource="#Event"/>
  <owl:inverseOf rdf:resource="#firstOf"/>
</owl:FunctionalProperty>

<owl:FunctionalProperty rdf:about="#second">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subPropertyOf rdf:resource="#has-subevent"/>
  <rdfs:domain rdf:resource="#Sequence"/>
  <rdfs:range rdf:resource="#Event"/>
  <owl:inverseOf rdf:resource="#secondOf"/>
</owl:FunctionalProperty>

<owl:FunctionalProperty rdf:about="#followed-by">
  <rdfs:domain rdf:resource="#Event"/>
  <rdfs:range rdf:resource="#Event"/>
  <owl:inverseOf rdf:resource="#follows"/>
</owl:FunctionalProperty>

<owl:Class rdf:about="#BagEvent">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subClassOf rdf:resource="#CompositeEvent"/>
  <owl:unionOf rdf:parseType="Collection">
    <owl:Class rdf:about="#And">
      <rdfs:isDefinedBy rdf:resource=""/>
      <owl:disjointWith rdf:resource="#Or"/>
      <owl:disjointWith rdf:resource="#AnyOrMulti"/>
      <owl:equivalentClass>
        <owl:Restriction>
          <owl:onProperty rdf:resource="#has-operator"/>

```

```

    <owl:hasValue rdf:resource="#AndOp"/>
  </owl:Restriction>
</owl:equivalentClass>
</owl:Class>
<owl:Class rdf:about="#Or">
  <rdfs:isDefinedBy rdf:resource=""/>
  <owl:disjointWith rdf:resource="#And"/>
  <owl:disjointWith rdf:resource="#AnyOrMulti"/>
  <owl:equivalentClass>
    <owl:Restriction>
      <owl:onProperty rdf:resource="#has-operator"/>
      <owl:hasValue rdf:resource="#OrOp"/>
    </owl:Restriction>
  </owl:equivalentClass>
</owl:Class>
<owl:Class rdf:about="#AnyOrMulti">
  <rdfs:isDefinedBy rdf:resource=""/>
  <owl:unionOf rdf:parseType="Collection">
    <owl:Class rdf:about="#Any">
      <rdfs:isDefinedBy rdf:resource=""/>
      <owl:disjointWith rdf:resource="#MultiOccurrences" />
      <owl:equivalentClass>
        <owl:Restriction>
          <owl:onProperty rdf:resource="#has-operator"/>
          <owl:hasValue rdf:resource="#AnyOp"/>
        </owl:Restriction>
      </owl:equivalentClass>
    </owl:Class>
    <owl:Class rdf:about="#MultiOccurrences">
      <rdfs:isDefinedBy rdf:resource=""/>
      <owl:disjointWith rdf:resource="#Any" />
      <owl:equivalentClass>
        <owl:Restriction>
          <owl:onProperty rdf:resource="#has-operator"/>
          <owl:hasValue rdf:resource="#MultiOccurrencesOp"/>
        </owl:Restriction>
      </owl:equivalentClass>
    </owl:Class>
  </owl:unionOf>
  <owl:disjointWith rdf:resource="#And"/>
  <owl:disjointWith rdf:resource="#Or"/>
</owl:Class>
</owl:unionOf>
<owl:disjointWith rdf:resource="#Sequence"/>
<owl:disjointWith rdf:resource="#Bounded"/>
</owl:Class>

<rdf:Property rdf:about="#number-of-occurrences">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:range rdf:resource="#AnyOrMulti"/>
  <rdfs:domain rdf:resource="http://www.w3.org/2001/XMLSchema#integer"/>
</rdf:Property>

```

```

<owl:Class rdf:about="#Bounded">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subClassOf rdf:resource="#CompositeEvent"/>
  <owl:unionOf rdf:parseType="Collection">
    <owl:Class rdf:about="#Not">
      <rdfs:isDefinedBy rdf:resource=""/>
      <owl:disjointWith rdf:resource="#Cumulative"/>
      <owl:disjointWith rdf:resource="#NonCumulative"/>
      <owl:equivalentClass>
        <owl:Restriction>
          <owl:onProperty rdf:resource="#has-operator"/>
          <owl:hasValue rdf:resource="#NotOp"/>
        </owl:Restriction>
      </owl:equivalentClass>
    </owl:Class>
  </owl:unionOf>
  <owl:Class rdf:about="#Cumulative">
    <rdfs:isDefinedBy rdf:resource=""/>
    <owl:unionOf rdf:parseType="Collection">
      <owl:Class rdf:about="#CumulativeAperiodic">
        <rdfs:isDefinedBy rdf:resource=""/>
        <owl:disjointWith rdf:resource="#CumulativePeriodic"/>
        <owl:equivalentClass>
          <owl:Restriction>
            <owl:onProperty rdf:resource="#has-operator"/>
            <owl:hasValue rdf:resource="#CumulativePeriodicOp"/>
          </owl:Restriction>
        </owl:equivalentClass>
      </owl:Class>
      <owl:Class rdf:about="#CumulativePeriodic">
        <rdfs:isDefinedBy rdf:resource=""/>
        <owl:disjointWith rdf:resource="#CumulativeAperiodic"/>
        <owl:equivalentClass>
          <owl:Restriction>
            <owl:onProperty rdf:resource="#has-operator"/>
            <owl:hasValue rdf:resource="#CumulativeAperiodicOp"/>
          </owl:Restriction>
        </owl:equivalentClass>
      </owl:Class>
    </owl:unionOf>
    <owl:disjointWith rdf:resource="#Not"/>
    <owl:disjointWith rdf:resource="#NonCumulative"/>
  </owl:Class>
  <owl:Class rdf:about="#NonCumulative">
    <rdfs:isDefinedBy rdf:resource=""/>
    <owl:unionOf rdf:parseType="Collection">
      <owl:Class rdf:about="#Aperiodic">
        <rdfs:isDefinedBy rdf:resource=""/>
        <owl:disjointWith rdf:resource="#Periodic"/>
        <owl:equivalentClass>
          <owl:Restriction>
            <owl:onProperty rdf:resource="#has-operator"/>
            <owl:hasValue rdf:resource="#PeriodicOp"/>
          </owl:Restriction>
        </owl:equivalentClass>
      </owl:Class>
    </owl:unionOf>
  </owl:Class>

```



```

</owl:Class>
<owl:Class rdf:about="#Periodic">
  <rdfs:isDefinedBy rdf:resource=""/>
  <owl:disjointWith rdf:resource="#Aperiodic"/>
  <owl:equivalentClass>
    <owl:Restriction>
      <owl:onProperty rdf:resource="#has-operator"/>
      <owl:hasValue rdf:resource="#AperiodicOp"/>
    </owl:Restriction>
  </owl:equivalentClass>
</owl:Class>
</owl:unionOf>
<owl:disjointWith rdf:resource="#Not"/>
<owl:disjointWith rdf:resource="#Cumulative"/>
</owl:Class>
</owl:unionOf>
<owl:disjointWith rdf:resource="#BagEvent"/>
<owl:disjointWith rdf:resource="#Sequence"/>
</owl:Class>

<rdf:Property rdf:about="#has-start">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:domain rdf:resource="#Bounded"/>
  <rdfs:range rdf:resource="#Event"/>
</rdf:Property>

<rdf:Property rdf:about="#has-end">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:domain rdf:resource="#Bounded"/>
  <rdfs:range rdf:resource="#Event"/>
</rdf:Property>

<rdf:Property rdf:about="#cumulate">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:domain rdf:resource="#Cumulative"/>
  <rdfs:range rdf:resource="#Event"/>
</rdf:Property>

<rdf:Property rdf:about="#fired-by">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:domain rdf:resource="#NonCumulative"/>
  <rdfs:range rdf:resource="#Event"/>
</rdf:Property>

<rdf:Property rdf:about="#exit-not">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:domain rdf:resource="#Not"/>
  <rdfs:range rdf:resource="#Event"/>
</rdf:Property>

</rdf:RDF>

```

[Filename:snoopy.rdf]

```
# call # jena -q -pellet -il RDF/XML -qf snoopy-query.sparql

PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> PREFIX
rdfs: <http://www.w3.org/2000/01/rdf-schema#> PREFIX owl:
<http://www.w3.org/2002/07/owl#> PREFIX snoop:
<http://www.semwebtech.org/languages/2006/snoopy#>

SELECT ?X ?Y FROM <file:snoopy.rdf> WHERE {
  ?X rdfs:subClassOf ?Y . ?Y rdfs:subClassOf snoop:Event
}
```

11.7.2 The XML Representation: Snoopy-ML

The XML representation follows the term structure and omits the term construction properties. The order of arguments (when relevant, such as for cumulative and not events) corresponds to that of the original SNOOP event algebra. As motivated above, the mapping is given implicitly via the target DTDs for Striped RDF/XML SNOOP and SNOOP-ML (see Appendix 13.2.4).

11.8 The CCS Process Ontology

The CCS ontology is analogous to the snoop ontology, using the CCS operators as described in the DTD given in Section 5.6.

```
<!DOCTYPE rdf:RDF [
  <!ENTITY expr "http://www.semwebtech.org/2006/expressions#"> ]>
<rdf:RDF
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:owl="http://www.w3.org/2002/07/owl#"
  xmlns:expr="http://www.semwebtech.org/2006/expressions#"
  xml:base="http://www.semwebtech.org/languages/2006/ccs">

  <owl:Ontology rdf:about=""/>

  <owl:Class rdf:about="#CCSExpr">
    <owl:unionOf rdf:parseType="Collection">
      <owl:Class rdf:about="#Process">
        <rdfs:isDefinedBy rdf:resource=""/>
      </owl:Class>
      <owl:Class rdf:about="#Event"> <!-- see also glue ontology below -->
        <owl:equivalentClass rdf:resource="#EventSpec"/>
        <owl:disjointWith rdf:resource="#Process"/>
      </owl:Class>
      <owl:Class rdf:about="#Query">
        <owl:disjointWith rdf:resource="#Process"/>
        <owl:disjointWith rdf:resource="#Event"/>
      </owl:Class>
      <owl:Class rdf:about="#Test">
        <owl:disjointWith rdf:resource="#Process"/>
        <owl:disjointWith rdf:resource="#Event"/>
        <owl:disjointWith rdf:resource="#Query"/>
      </owl:Class>
    </owl:unionOf>
```

```

</owl:Class>

<owl:Class rdf:about="#AtomicAction">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subClassOf rdf:resource="#Process"/>
  <owl:disjointWith rdf:resource="#CompositeProcess"/>
</owl:Class>

<owl:Class rdf:about="#CompositeProcess">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subClassOf rdf:resource="#Process"/>
  <rdfs:subClassOf rdf:resource="&expr;AlgebraicCompositeExpression"/>
</owl:Class>

<owl:FunctionalProperty rdf:about="#has-operator">
  <rdfs:domain rdf:resource="#AlgebraicCompositeExpression"/>
  <rdfs:range rdf:resource="#Operator"/>
  <rdfs:subPropertyOf rdf:resource="&expr;has-operator"/>
</owl:FunctionalProperty>

<expr:Operator rdf:about="#SequenceOp">
  <expr:has-arity>*</expr:has-arity>
</expr:Operator> <expr:Operator rdf:about="#AlternativeOp">
  <expr:has-arity>*</expr:has-arity>
</expr:Operator> <expr:Operator rdf:about="#ConcurrentOp">
  <expr:has-arity>*</expr:has-arity>
</expr:Operator> <expr:Operator rdf:about="#FixpointOp">
  <expr:has-arity>1</expr:has-arity>
</expr:Operator>

<owl:Class rdf:about="#Delay">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subClassOf rdf:resource="#AtomicAction"/>
</owl:Class> <owl:Class rdf:about="#Sequence">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subClassOf rdf:resource="#CompositeProcess"/>
  <owl:equivalentClass>
    <owl:Restriction>
      <owl:onProperty rdf:resource="#has-operator"/>
      <owl:hasValue rdf:resource="#SequenceOp"/>
    </owl:Restriction>
  </owl:equivalentClass>
</owl:Class> <owl:Class rdf:about="#Alternative">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subClassOf rdf:resource="#CompositeProcess"/>
  <owl:disjointWith rdf:resource="#Sequence"/>
  <owl:equivalentClass>
    <owl:Restriction>
      <owl:onProperty rdf:resource="#has-operator"/>
      <owl:hasValue rdf:resource="#AlternativeOp"/>
    </owl:Restriction>
  </owl:equivalentClass>
</owl:Class> <owl:Class rdf:about="#Concurrent">
  <rdfs:isDefinedBy rdf:resource=""/>

```

```

<rdfs:subClassOf rdf:resource="#CompositeProcess"/>
<owl:disjointWith rdf:resource="#Sequence"/>
<owl:disjointWith rdf:resource="#Alternative"/>
<owl:equivalentClass>
  <owl:Restriction>
    <owl:onProperty rdf:resource="#has-operator"/>
    <owl:hasValue rdf:resource="#ConcurrentOp"/>
  </owl:Restriction>
</owl:equivalentClass>
</owl:Class> <owl:Class rdf:about="#Fixpoint">
<rdfs:isDefinedBy rdf:resource=""/>
<rdfs:subClassOf rdf:resource="#CompositeProcess"/>
<owl:disjointWith rdf:resource="#Sequence"/>
<owl:disjointWith rdf:resource="#Alternative"/>
<owl:disjointWith rdf:resource="#Concurrent"/>
<owl:equivalentClass>
  <owl:Restriction>
    <owl:onProperty rdf:resource="#has-operator"/>
    <owl:hasValue rdf:resource="#FixpointOp"/>
  </owl:Restriction>
</owl:equivalentClass>
<owl:equivalentClass>
  <owl:Restriction>
    <owl:onProperty rdf:resource="#fix-variable"/>
    <owl:minCardinality>1</owl:minCardinality>
  </owl:Restriction>
</owl:equivalentClass>
</owl:Class>

<rdf:Property rdf:about="#fix-variable">
  <rdfs:domain rdf:resource="#Fixpoint"/>
  <rdfs:range rdf:resource="#ProcessVariable"/>
</rdf:Property> <rdf:Property rdf:about="#local-variable">
  <rdfs:domain rdf:resource="#Fixpoint"/>
  <rdfs:range rdf:resource="#Variable"/>
</rdf:Property> <owl:DatatypeProperty rdf:about="#has-index">
  <rdfs:domain rdf:resource="#Fixpoint"/>
</owl:DatatypeProperty>

<owl:Class rdf:about="#Variable"/> <owl:Class
rdf:about="#ProcessVariable">
  <owl:disjointWith rdf:resource="#Variable"/>
</owl:Class>

<rdf:Property rdf:about="#has-subprocess">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:domain rdf:resource="#CompositeProcess"/>
  <rdfs:range rdf:resource="#CCSEExpr"/>
  <owl:inverseOf rdf:resource="#subProcessOf"/>
</rdf:Property>

<owl:FunctionalProperty rdf:about="#first">
  <rdfs:isDefinedBy rdf:resource=""/>
  <rdfs:subPropertyOf rdf:resource="#has-subprocess"/>

```

```

<rdfs:domain rdf:resource="#Sequence"/>
<rdfs:range rdf:resource="#FirstInSequence"/>
<owl:inverseOf rdf:resource="#firstOf"/>
</owl:FunctionalProperty>

<owl:FunctionalProperty rdf:about="#second">
<rdfs:isDefinedBy rdf:resource=""/>
<rdfs:subPropertyOf rdf:resource="#has-subprocess"/>
<rdfs:domain rdf:resource="#Sequence"/>
<rdfs:range rdf:resource="#SecondInSequence"/>
<owl:inverseOf rdf:resource="#secondOf"/>
</owl:FunctionalProperty>

<owl:FunctionalProperty rdf:about="#followed-by">
<rdfs:domain rdf:resource="#CCSEExpr"/>
<rdfs:range rdf:resource="#CCSEExpr"/>
<owl:inverseOf rdf:resource="#follows"/>
</owl:FunctionalProperty> </rdf:RDF>

```

[Filename:ccs.rdf]

11.9 The Composite Formulas and Queries Ontology

Queries have yet mainly been discussed in terms of opaque queries (XML – XPath/XQuery, Xcerpt). In Section 5.4 we discussed that querying can be divided into two parts: *obtaining* answers from data, and *combining* answers, usually in an algebraic way. On the other hand, the combination of atomic concepts is often done in a logical way, e.g. in First-Order Logic, with logical connectives. As shown for the relational algebra it is equivalent to the (safe) relational calculus.

As an ontology for composite queries, we propose to use the Relational Algebra as it is immediately applicable to sets of tuples of variable bindings. Straightforward extensions include grouping, aggregation, and quantification (where \exists is actually a projection). Query algebras have been developed over different underlying data models, i.e., the relational one, the object-oriented one (which showed that this is not much different from the relational algebra), and F-Logic (cf. [46, Section 2.2.1]); cf. Section 5.4.

As an ontology for composite formulas, we propose to use the first-order logic notions, in consultation with the FOL-RuleML initiative [15].

We deal with atomic queries (and atomic event matching) in the next chapter.

Chapter 12

Atomic Event Matching, Queries and Atomic Actions in RDF

While in the previous section, concepts have been “ontologized” that have been developed in the earlier sections, one level is still open: the level of atomic event matching, queries, tests, assertions (that have not been used up to now, but which are useful for enabling reasoning about rules) and expressing atomic actions. These were dealt with on an opaque or merely syntactical level until now:

- The *Atomic Event Specifications* have been discussed in Section 5.2.2; actually, they are kinds of queries against incoming events.
- Atomic Queries have shortly been discussed in Section 5.4 – as either RDF triples or atomic opaque queries on the datamodel level.
- Atomic actions have been given as XML terms with variable references as e.g. used in XQuery return statements.

The above formalisms do not allow to reason about a rule, e.g., “which rules are affected from cancelling flight LH123”? – probably e.g., LH123, from Frankfurt to JFK, will trigger a rule whose firing event is “on any event that concerns a flight to New York”. For that, domain-level information together with insight into the atomic event specifications of the rules are necessary.

We first discuss in Section 12.1.1 an approach that uses RDF graph *patterns*, reminiscent of the XML-QL style matching formalism from Section 5.2.2.1 that turned out to be quite “declarative” – but it is not semantic at all. Although this looks nice, it must be discarded since it is not compatible with the inherent semantics of RDF/XML information. This is then revised into a cleaner approach which separates the patterns from *real* instance data (Section 12.1.2). Nevertheless, with this, the patterns are not accessible for reasoning about rules. In Section 12.2, we then give a totally different formalism for formulas that allows the use as queries, atomic event specifications, tests, and assertions that is based on *class specifications* and *constraints*, which then allows for DL-based reasoning about rules.

The following sample data contains

- two sample flights, LH123 from Frankfurt to JFK, and LH456 from MUC to LIS,
- a cancellation event for the first of them.

```
<!DOCTYPE rdf:RDF [  
  <!ENTITY travel "http://www.semwebtech.org/domains/2006/travel#"> ]>  
<rdf:RDF  
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
```

```

    xmlns:travel="http://www.semwebtech.org/domains/2006/travel#"
    xml:base="sample:traveldata#">
<travel:Flight
rdf:about="&travel;/airlines/lufthansa/flights/lh123">
  <travel:from rdf:resource="&travel;/airports/fra"/>
  <travel:to rdf:resource="&travel;/airports/jfk"/>
</travel:Flight> <travel:Flight
rdf:about="&travel;/airlines/lufthansa/flights/lh456">
  <travel:from rdf:resource="&travel;/airports/muc"/>
  <travel:to rdf:resource="&travel;/airports/lis"/>
</travel:Flight> <travel: canceled-flight rdf:ID="occur-cf-event">
  <travel:flight rdf:resource="&travel;/airlines/lufthansa/flights/lh123"/>
</travel: canceled-flight> </rdf:RDF>

```

[Filename: sample-event-and-data.rdf]

12.1 Discarded Pattern Approaches

12.1.1 The Embedded Pattern Approach

For some time, we favored an approach where atomic events and queries are represented by RDF graph patterns like in SPARQL, but explicit in an RDF/XML markup. This would make up for a nice rule representation: Some nodes belong to the ECA-ML URI domain/namespace, the next level belongs to the URI domain/namespace of the algebraic languages like Snoopy or CCS, and the “atomic” leaves of the algebra trees belong to the domain namespaces. Even shared logical variables easily fit into that graph approach.

But, if these rules get “freely accessible” as RDF/XML data in the Web, it is *not* known that the application domain elements stand for patterns, and not for instances. The patterns mix up with the range and domain axioms of the domain ontology would result in extremely unintended, if not even inconsistent “information”.

Example 46 (Discarded: Rules with Atomics as Graph Patterns) *We show in two examples how variables are modeled by resources: Consider the ECA rule “if a flight is cancelled that is relevant to one of our customers, then send that customer an SMS” which uses an atomic event, a simple conjunctive query/condition and an atomic action component.*

The most declarative convenient way is to “declare” variables as resources on the “global” level of a rule and to reference them in the statements; nevertheless, the declaration can also happen at arbitrary object positions:

```

<rdf:RDF
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:eca="http://www.semwebtech.org/languages/2006/eca-ml#"
  xmlns:travel="http://www.semwebtech.org/domains/2006/travel#"
  xmlns:comm="http://www.semwebtech.org/domains/2006/communication#"
  xmlns="foo:local#">
  <Variable rdf:ID="Flight"/> <Variable rdf:ID="Date"/> <Variable
rdf:ID="Mobile"/> <eca:Rule>
  <eca:has-event>
    <travel: canceled-flight rdf:ID="cf">
      <travel:flight rdf:resource="#Flight"/>
    </travel: canceled-flight>
  </eca:has-event>
  <eca:has-query>

```



```

    <travel:Customer>
      <travel:booked-for rdf:resource="#Flight"/>
      <travel:mobile-no rdf:resource="#Mobile"/>
    </travel:Customer>
  </eca:has-query>
<eca:has-action>
  <comm:send-sms>
    <comm:to rdf:resource="#Mobile"/>
    <text>Your flight {#Flight} has been cancelled</text>
  </comm:send-sms>
</eca:has-action>
</eca:Rule> </rdf:RDF>

```

[Filename: rule-example-1.rdf]

In the above rule, the node

```

    <travel:canceled-flight rdf:ID="cf">
      <travel:flight rdf:resource="#Flight">
    </travel:canceled-flight>

```

is not intended as an instance of a `travel:canceled-flight` event, but it inevitable would be parsed as such. The same holds for the RDF pattern query. Both would even result in the fact that the resource identified by `local#Flight` is not only a variable, as intended, but also a `travel:flight` instance (as implied by the range definitions of the `travel` namespace).

Below, a SPARQL query is given that illustrates the semantics and the problem of the above definitions. It uses an excerpt of the `travel` domain definitions:

```

<!DOCTYPE rdf:RDF [
  <!ENTITY mars "http://www.semwebtech.org/2006/mars#"> ]>
<rdf:RDF
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:owl="http://www.w3.org/2002/07/owl#"
  xmlns:travel="http://www.semwebtech.org/domains/2006/travel#"
  xml:base="http://www.semwebtech.org/domains/2006/travel">

  <owl:Class rdf:about="#Event">
    <rdfs:subClassOf rdf:resource="#mars;Event"/>
  </owl:Class> <owl:Class rdf:about="#FlightEvent">
    <rdfs:subClassOf rdf:resource="#Event"/>
  </owl:Class> <owl:Class rdf:about="#canceled-flight">
    <rdfs:subClassOf rdf:resource="#FlightEvent"/>
    <rdfs:subClassOf>
      <owl:Restriction>
        <owl:onProperty rdf:resource="#flight"/>
        <owl:allValuesFrom rdf:resource="#Flight"/>
      </owl:Restriction>
    </rdfs:subClassOf>
  </owl:Class> </rdf:RDF>

```

[Filename: travel-miniexcerpt.rdf]

```
# call # jena -q -pellet -il RDF/XML -qf rule-query-1.sparql
```

```
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> PREFIX
```

```
mars: <http://www.semwebtech.org/2006/mars#> PREFIX travel:
<http://www.semwebtech.org/domains/2006/travel#>
```

```
SELECT ?X ?Z ?T FROM <file:eca-ml.rdf> FROM
<file:rule-example-1.rdf> FROM <file:travel-miniexcerpt.rdf> WHERE
{{?X rdf:type mars:Event .
  ?X rdf:type ?T } UNION
  {?X travel:flight ?Z .
  ?Z rdf:type ?T }}
```

[Filename: rule-query-1.sparql]

The result of the query shows that the rule pattern contributes to instances: <file:rule-example-1.rdf#cf> is an instance of <http://www.semwebtech.org/2006/mars#Event>. The same happens for patterns in the query part.

12.1.2 Embedded Event and Query Patterns as parseType=XML

To avoid that AES patterns are “found” as RDF instance data in the Web, they can be handled as XML literals. When rule evaluation comes to the atomic level, they are “unpacked”, parsed as RDF/XML, and can then be used. The disadvantage is that they are then completely opaque, and cannot be found when e.g. searching for “all rules that react on a travel: canceled-flight event.

```
<rdf:RDF
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:eca="http://www.semwebtech.org/languages/2006/eca-ml#"
  xmlns:travel="http://www.semwebtech.org/domains/2006/travel#"
  xmlns:comm="http://www.semwebtech.org/domains/2006/communication#"
  xmlns="foo:local#">

  <Variable rdf:ID="Flight"/> <Variable rdf:ID="Date"/> <Variable
  rdf:ID="Mobile"/> <eca:Rule>
    <eca:has-event rdf:parseType="Literal">
      <travel:canceled-flight rdf:ID="cf">
        <travel:flight rdf:resource="#Flight"/>
      </travel:canceled-flight>
    </eca:has-event>
    <eca:has-query rdf:parseType="Literal">
      <travel:Customer>
        <travel:booked-for rdf:resource="#Flight"/>
        <travel:mobile-no rdf:resource="#Mobile"/>
      </travel:Customer>
    </eca:has-query>
    <eca:has-action rdf:parseType="Literal">
      <comm:send-sms>
        <comm:to rdf:resource="#Mobile"/>
        <text>Your flight {#Flight} has been cancelled</text>
      </comm:send-sms>
    </eca:has-action>
  </eca:Rule> </rdf:RDF>
```

[Filename: rule-example-2.rdf]

```
# call # jena -q -pellet -il RDF/XML -qf rule-query-2.sparql

PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> PREFIX
mars: <http://www.semwebtech.org/2006/mars#> PREFIX eca:
<http://www.semwebtech.org/languages/2006/eca-ml#> PREFIX travel:
<http://www.semwebtech.org/domains/2006/travel#>

SELECT ?X ?E ?Y ?Z FROM <file:eca-ml.rdf> FROM
<file:rule-example-2.rdf> FROM <file:travel-miniexcerpt.rdf> WHERE
{{?X eca:has-event ?E} UNION
  {?ME rdf:type mars:Event} UNION
  {?TE rdf:type travel:Event} UNION
  {?EE rdf:type eca:EventSpec}}
```

[Filename: rule-query-2.sparql]

Thus, languages that use a pattern syntax for queries, and also for updates or rule reads, as known from Prolog, Datalog, F-Logic [41, 42], XPathLog [47, 50], or Xcerpt/XChange [23, 21] are not best-suited for being used in the Semantic Web.

12.2 Formulas, Queries, Tests, Assertions, Atomic Event Matching, and Atomic Actions in RDF

The previous section showed that describing anything like queries, event specifications, or actions by “giving” the patterns explicitly is problematic wrt. the basic idea of the Semantic Web as an “open” data source since these patterns are –without understanding their context– understood as actual data. Thus, formulas must be given by using an ontology that *describes* the structure (i.e., the patterns) of the formula and its components.

Formulas can serve for all tasks that are still open: The actual use of the formula is then given by the context, i.e., by the property how it is referenced to from its environment:

- Queries are formulas (often constrained to conjunctive ones) with free variables (to be bound to answers during the evaluation). They are referenced by `eca:has-query`, or as subexpressions of a composite query according to some algebra (boolean or relational Algebra). From the definitions of the algebra ontology, it can be derived that the query is of type `eca:Query`. All nodes in the query are of types that are defined as static notions (classes or properties) by some domain ontology.
- Atomic Event Specifications are (small) queries against an RDF graph given by the event. They are referenced either as atomic events from an ECA rules (`eca:has-event`) or as subevents from SNOOP (where the property also induces that they are of type `eca:EventSpec`), *and* – which is defined below in detail – they refer to a class (of events) that is defined as an event class (subclass of `mars:Event`) by some domain ontology.
- Tests and Assertions are (sometimes closed) formulas; free variables are required to be bound before evaluation.
- Atomic Action specifications are quantifier-free formulas; free variables are required to be bound before evaluation. The formulas are kind of “facts” that are used to describe an RDF graph that represents the action to be executed. The “handle” node of the graph is an instance of a class that is defined as an action (subclass of `mars>Action`) by some domain ontology.

Note that with the above definition, the same formula (over static concepts of application domains) can act as a query (when some variables are not yet bound), as a test (all variables bound), and as an assertion (all variables bound).

We propose an RDF ontology that allows to *describe* the queries including the use of join conditions by variables (thus extending the OWL-DL queries by defining result classes that are restricted to tree-like queries as they do not allow to use variables). Although it would be sufficient to “ontologize” SPARQL by providing constructs to describe triples/statements, using resources and variables (i.e., URIs that are of `rdf:type Variable`), the language design extends this with borrowing the definition of classes from OWL.

12.2.1 The OWLQ Conjunctive Query Language

As discussed in Section 4.2.3, queries can have a functional semantics (the items addressed by a path expression), and a logical semantics (the resulting variable bindings).

- definition of “relevant” classes (by common OWL means) to find out which instances “qualify” for a complex concept description (e.g., for an Atomic Event Specification). The definition can either define a new class as e.g. an intersection of an existing class and an `owl:Restriction`, or use `owlq:baseClass` to refer to an existing one directly (which means for an OWLQ interpreter the same as `owl:equivalentClass`).
- declaration of variables ranging over instances of a class. (SPARQL equivalent: `{?CVar rdf:type :theClass}`).
- declaration or constraining of variables ranging over the values (URIs or literals) of a property wrt. a value, a class membership, or another variable, (SPARQL equivalent e.g. `{?Var1 :property ?Var2}`); for constraints associated to a class, they refer to the instances of that class (or the variable that ranges over them): `{?CVar :property ?Var2}`).

```

<!DOCTYPE rdf:RDF [
  <!ENTITY rdf "http://www.w3.org/1999/02/22-rdf-syntax-ns#">
  <!ENTITY owl "http://www.w3.org/2002/07/owl#"> ]>
<rdf:RDF
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:owl="http://www.w3.org/2002/07/owl#"
  xml:base="http://www.semwebtech.org/languages/2006/owlq">

  <!-- "entry node" is usually an owlq:Query.
   If only an owlq:Class is given, the query has a functional
   result that consists of the instances of that class.
   The latter is used as AES formalism -->

  <owl:Class rdf:about="#Query"/> <!-- defines owlq:Query -->
  <rdf:Property rdf:about="#definesClass">
    <rdfs:domain rdf:resource="#Query"/>
    <rdfs:range rdf:resource="#Class"/>
  </rdf:Property> <rdf:Property rdf:about="#hasResultClass">
    <rdfs:domain rdf:resource="#Query"/>
    <rdfs:range rdf:resource="#Class"/>
  </rdf:Property> <rdf:Property rdf:about="#hasConstraint">
    <rdfs:domain rdf:resource="#Query"/>
    <rdfs:range rdf:resource="#Constraint"/>
  </rdf:Property>

```

```

<owl:Class rdf:about="#Variable"/> <owl:Class
rdf:about="#MustBindVariable">
  <rdfs:subClassOf rdf:resource="#Variable"/>
</owl:Class> <owl:Class rdf:about="#OptionalVariable">
  <rdfs:subClassOf rdf:resource="#Variable"/>
</owl:Class> <owl:Class rdf:about="#DontCareVariable">
  <rdfs:subClassOf rdf:resource="#Variable"/>
</owl:Class>

<owl:Class rdf:about="Class"> <!-- defines owlq:Class -->
  <rdfs:subClassOf rdf:resource="#owl;Class"/>
</owl:Class>
<rdf:Property rdf:about="#baseClass">
  <rdfs:domain rdf:resource="#Class"/>
  <rdfs:range rdf:resource="#Class"/>
</rdf:Property>
<rdf:Property rdf:about="#scopesVariable">
  <rdfs:domain rdf:resource="#Class"/>
  <rdfs:range rdf:resource="#Variable"/>
  <owl:inverseOf rdf:resource="#rangesOver"/>
</rdf:Property>

<owl:Class rdf:about="Constraint"/>
  <!-- subject -->
  <rdf:Property rdf:about="#onVariable">
    <rdfs:domain rdf:resource="#Constraint"/>
    <rdfs:range rdf:resource="#Variable"/>
  </rdf:Property>
  <!-- predicate -->
  <rdf:Property rdf:about="#onProperty">
    <rdfs:domain rdf:resource="#Constraint"/>
    <rdfs:range rdf:resource="#rdf;Property"/>
  </rdf:Property>
  <!-- object -->
  <rdf:Property rdf:about="#equalsVariable">
    <rdfs:domain rdf:resource="#Constraint"/>
    <rdfs:range rdf:resource="#Variable"/>
  </rdf:Property>
  <rdf:Property rdf:about="#hasValue">
    <rdfs:domain rdf:resource="#Constraint"/>
  </rdf:Property>
  <rdf:Property rdf:about="#allValuesFrom">
    <rdfs:domain rdf:resource="#Constraint"/>
  </rdf:Property>
  <rdf:Property rdf:about="#someValueFrom">
    <rdfs:domain rdf:resource="#Constraint"/>
  </rdf:Property>

<rdf:Property rdf:about="#hasConstraint">
  <rdfs:domain rdf:resource="#Class"/>
  <rdfs:range rdf:resource="#Constraint"/>
  <!-- then, the constraint has usually no "onVariable" but
       refers to the instances of the class -->
</rdf:Property> </rdf:RDF>

```

[Filename: owlq.rdf]

Note that the above variable declarations usually allow to infer the “type” of a variable when reasoning about queries and rules.

Domain-Dependent Predicates. Like in many query languages, tests can be integrated in the constraints by appropriate predicates taken from other ontologies, operating on the variables.

12.2.2 Specification of Atomic Events by OWLQ Queries

The above is especially useful for Atomic Event Specifications, specifying relevant atomic events for an ECA rule (cf. Section 5.2.2):

- an AES is a definition of a class of events. The basic specification can thus be given by deriving an OWL class as a `owl:Restriction` from an event class of some application domain. Convention: if an AES in a rule consists only of an `owlq:Class` definition, this is considered to define the functional results.
- The event (as an RDF graph) is bound to the implicit Event variable on the rule level (cf. page 11.6.5).
- This is then extended with the use of variable bindings for join conditions etc. by using `owlq:Constraint`.

If an incoming event is a member of that class, the variables are bound according to the given specification (e.g. by evaluating an appropriate SPARQL query that yields all bindings).

Example 47 *Atomic Event Specification in RDF*

```
<!DOCTYPE rdf:RDF [  
  <!ENTITY travel "http://www.semwebtech.org/domains/2006/travel#"> ]>  
<rdf:RDF  
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"  
  xmlns:owl="http://www.w3.org/2002/07/owl#"  
  xmlns:eca="http://www.semwebtech.org/languages/2006/eca-ml#"  
  xmlns:owlq="http://www.semwebtech.org/languages/2006/owlq#"  
  xmlns:travel="http://www.semwebtech.org/domains/2006/travel#"  
  xml:base="the:rule#">  
  
  <eca:Variable rdf:ID="Flight"/> <eca:Variable rdf:ID="Date"/>  
  <eca:Variable rdf:ID="Mobile"/> <owl:Restriction  
  rdf:ID="FromFrankfurt">  
    <owl:onProperty rdf:resource="&travel;from"/>  
    <owl:hasValue rdf:resource="&travel;/airports/fra"/>  
  </owl:Restriction>  
  
  <eca:Rule rdf:ID="the-rule">  
    <eca:has-event>  
      <owlq:Class rdf:ID="RelevantEvent">  
        <owl:scopesVariable rdf:resource="#Event"/>  
        <owl:intersectionOf rdf:parseType="Collection">  
          <owl:Class rdf:about="&travel;canceled-flight"/>  
          <owl:Restriction>  
            <owl:onProperty rdf:resource="&travel;flight"/>  
            <owl:someValuesFrom rdf:resource="#FromFrankfurt"/>  
          </owl:Restriction>  
        </owl:intersectionOf>  
      </owlq:Class>  
    </eca:has-event>  
  </eca:Rule>  
</rdf:RDF>
```

```

    </owl:Restriction>
  </owl:intersectionOf>
</owlq:Class>
</eca:has-event>
<eca:has-query/>
<eca:has-action/>
</eca:Rule>

<owlq:Constraint>
  <owlq:onVariable rdf:resource="#Event"/>
  <owlq:onProperty rdf:resource="#&travel;flight"/>
  <owlq>equalsVariable rdf:resource="#Flight"/>
</owlq:Constraint> </rdf:RDF>

```

[Filename: rule-example-3.rdf]

```
# call # jena -q -pellet -il RDF/XML -qf rule-query-3.sparql
```

```

PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> PREFIX
rdfs: <http://www.w3.org/2000/01/rdf-schema#> PREFIX owl:
<http://www.w3.org/2002/07/owl#> PREFIX mars:
<http://www.semwebtech.org/2006/mars#> PREFIX eca:
<http://www.semwebtech.org/languages/2006/eca-ml#> PREFIX owlq:
<http://www.semwebtech.org/languages/2006/owlq#> PREFIX travel:
<http://www.semwebtech.org/domains/2006/travel#>

```

```

SELECT ?R ?CF ?FF ?E ?RE ?V1 ?V2 ?P FROM <file:eca-ml.rdf> FROM
<file:owlq.rdf> FROM <file:rule-example-3.rdf> FROM
<file:travel-miniexcerpt.rdf> FROM
<file:sample-event-and-data.rdf> WHERE {{{?R rdf:type eca:Rule} .
  {?R eca:has-event ?E} .
  {?E rdf:type eca:EventSpec}
  {?E rdfs:subClassOf ?RE }} UNION
{?CF rdf:type travel: canceled-flight} UNION
{?CF travel:flight ?FF} UNION
{?E rdf:type mars:Event} UNION
{?FF rdf:type <the:rule#FromFrankfurt>} UNION
{?RE rdf:type <the:rule#RelevantEvent>} UNION
{{{?E rdf:type eca:EventSpec} .
  {?E owlq:scopesVariable ?V1} .
  {?C rdf:type owlq:Constraint} .
  {?C owlq:onVariable ?V1} .
  {?C owlq:onProperty ?P} .
  {?C owlq>equalsVariable ?V2}}} UNION
{{{?E rdf:type eca:EventSpec} .
  {?RE rdf:type ?E} .
  {?E owlq:scopesVariable ?V1} .
  {?C rdf:type owlq:Constraint} .
  {?C owlq:onVariable ?V1} .
  {?C owlq:onProperty ?PP} .
  {?C owlq>equalsVariable ?V2} .
  {?RE ?PP ?V2}}}

```

[Filename: rule-query-3.sparql]

The occurred “canceled-flight” event is correctly identified as an instance of the class `<the:rule#RelevantEvent>`. Additionally, the rule is identified to have an event specification whose “main” class is `travel:canceled-flight`.

The last part of the query illustrates how the declaration of the `owlq:Constraint` is connected with the actual variables: The `sample:traveldata#occur-cf-event` is of an instance of the class specified as `eca:EventSpec`. This class is the scope of the variable `the:rule#Event` that is used in a constraint to define the variable `the:rule#Flight` as the value of its property `travel:flight`. The value for the given event is the URL of `lh123`.

12.2.3 Evaluation of OWLQ Conjunctive Query Specifications

OWLQ allows immediately for expressing conjunctive queries and binding multiple variables from them:

- define or reference relevant classes and variables ranging over them,
- additional variables are bound, and
- constraints between variables can be expressed.

From the RDF graph of the OWLQ query, a SPARQL query against the event can be derived

- For each variable V ranging over the instances of a class Cl , given by `{?CL owl:scopesVariable ?V}`, have a conjunct `?V rdf:type ?CL`.
- for each Constraint C , add conjuncts

```
{?C rdf:type owl:Constraint} . {?C owl:onVariable ?V} .
{?C owl:onProperty ?P} . {?C owl>equalsVariable ?W} .
{?V ?P ?W}
```

Atomic Event Detection in OWLQ. Atomic Event Specifications in OWLQ (i) check if the event matches the specification, and if yes, (ii) generate the variable bindings.

If an OWLQ AES does not contain an own class definition, it may be more efficient to translate the RDF event into an XML markup and to evaluate XPath/XQuery queries against it.

Embedding of OWLQ Atoms in Formulas The OWLQ are implicitly conjunctive formulas over class expressions and properties (comparable to F-Logic Molecules [41, 42], XPathLog terms [47, 50], or Xcerpt terms [23]). These can be true/false under given variable assignments and thus can be embedded into any language for expressing composite formulas (e.g., defining disjunction, classical negation, implication, and quantification), as e.g. sketched in Section 11.9.

12.2.4 Specification of Atomic Actions: RDF-CL

Atomic actions are RDF (often seen as RDF/XML) instances. For the same reason as above, RDF/XML *patterns* in the rule are not suitable, as they are interpreted as *instances*, not as *specifications of instances*. Like OWLQ, a suitable ontology for generating OWL, or sufficiently, RDF graphs is required. This language is called *RDF-CL (RDF Creation Language)*. RDF-CL provides a straightforward ontology to *describe* what is to be generated:

```
<!DOCTYPE rdf:RDF [
  <!ENTITY rdf "http://www.w3.org/1999/02/22-rdf-syntax-ns#">
  <!ENTITY owl "http://www.w3.org/2002/07/owl#"> ]>
<rdf:RDF
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  >
```



```

xmlns:owl="http://www.w3.org/2002/07/owl#"
xml:base="http://www.semwebtech.org/languages/2006/rdf-cl">

<owl:Class rdf:about="#Variable"/> <owl:Class
rdf:about="#Create"/> <rdf:Property rdf:about="#createNode">
  <rdfs:domain rdf:resource="#Create"/>
  <rdfs:range rdf:resource="#NodeSpec"/>
</rdf:Property> <rdf:Property rdf:about="#createProperty">
  <rdfs:domain rdf:resource="#Create"/>
  <rdfs:range rdf:resource="#PropertySpec"/>
</rdf:Property>

<owl:Class rdf:about="#NodeSpec"/> <rdf:Property
rdf:about="#ofClass">
  <rdfs:domain rdf:resource="#NodeSpec"/>
  <rdfs:range rdf:resource="#owl;Class"/>
</rdf:Property> <rdf:Property rdf:about="#withProperty">
  <rdfs:domain rdf:resource="#NodeSpec"/>
  <rdfs:range rdf:resource="#PropertySpec"/>
  <owl:inverseOf>
    <rdf:Property rdf:about="#has-subject">
      <rdfs:subPropertyOf rdf:resource="#subject"/>
    </rdf:Property>
  </owl:inverseOf>
</rdf:Property>

<owl:Class rdf:about="#PropertySpec"/>
  <rdf:Property rdf:about="predicate">
    <rdfs:domain rdf:resource="#PropertySpec"/>
    <rdfs:range rdf:resource="#rdf;Property"/>
  </rdf:Property>
  <rdf:Property rdf:about="#object">
    <rdfs:domain rdf:resource="#PropertySpec"/>
  </rdf:Property>
  <rdf:Property rdf:about="#value-of">
    <rdfs:domain rdf:resource="#PropertySpec"/>
    <rdfs:range rdf:resource="#Variable"/>
  </rdf:Property>
</rdf:RDF>

```

[Filename: rdf-cl.rdf]

rdf-cl:NodeSpec is a generic constructor. Its rdf-cl:ofClass property usually refers to a domain class. Each action generates *one* “handle” node that is a action of some domain ontology. The other created nodes may also be class instances of the domain that are to be generated by the action. The type of the node to be created can also be given by a variable.

rdf-cl:PropertySpec is similar to rdf:Statement, with the following deviations:

- the subject is implicit. The rdf-cl:PropertySpec has always to be linked to a rdf-cl:NodeSpec in the same action. The action must be a connected graph.
- object: variables (i.e., resources X such that X rdf:type rdf-cl:Variable) can be referenced.

Chapter 13

Rules and Rule Components as Resources

Using the above ontologies, every rule is interpreted as a network of RDF resources of the contributing ontologies (ECA, event algebras, OWLQ specifications, application domains etc.). RDF/OWL reasoning can e.g. be applied for analysis of structural correctness, safety of variables, evaluation order etc. The chapter starts with some examples and then sketches reasoning about rules (which is still future work).

13.1 Putting the Ontologies Together

The composition of rules from elements of different ontologies, such as the ECA-ML ontology, composite event ontologies, process ontologies, query language ontologies, domain ontologies and specification of patterns about domain notions follows the structure given in the previous chapters, e.g. depicted in Figures 3.12 and 5.1. The individual ontologies are independent from each other.

For composing them in a *semantic* way, they should either refer to the notions in the MARS ontology, or separate *glue* ontologies are needed to describe the relationships of their notions with the MARS notions. This is mostly done by making `rdfs:subClassOf`, `rdfs:subPropertyOf`, `owl:equivalentClass`, or `owl:disjointWith` statements.

The rule becomes then a semantic structure that allows for evaluating/executing it as already described for the XML level, and for *reasoning* about it, as will be described in Section 13.2. First, some examples of RDF-level rules are shown.

Example 48 Consider again the university example (Example 33) which combines three application-independent language ontologies. Now, we give the RDF/XML representation:

- the ECA ontology,
- the SNOOP ontology of the event algebra in the event part,
- the CCS ontology of the process algebra in the action part,

and an application-dependent ontology:

- the university application ontology: there, the semantics of the atomic events defined in this ontology must be available (diagonally crosshatched).

This point of view leads to the above service-oriented distributed architecture by associating the “responsibility” for handling the respective resources by appropriate services. Additionally, e.g. collections of (sub)events as well as complete (application-specific) rule bases can be designed, published by associating them with a URI, and reused.

```

<!DOCTYPE rdf:RDF [
  <!ENTITY rdf "http://www.w3.org/1999/02/22-rdf-syntax-ns#">
  <!ENTITY owlq "http://www.semwebtech.org/languages/2006/owlq#">
  <!ENTITY fml "http://www.semwebtech.org/languages/2006/fml#">
  <!ENTITY uni "http://www.semwebtech.org/domains/2006/uni#">
  <!ENTITY smtp "http://www.semwebtech.org/domains/2006/smtp#">
  <!ENTITY httpget "http://www.semwebtech.org/languages/2006/http-get#"> ]>
<rdf:RDF
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:eca="http://www.semwebtech.org/languages/2006/eca-ml#"
  xmlns:snoop="http://www.semwebtech.org/languages/2006/snoopy#"
  xmlns:ccs="http://www.semwebtech.org/languages/2006/ccs#"
  xmlns:owlq="http://www.semwebtech.org/languages/2006/owlq#"
  xmlns:rdf-cl="http://www.semwebtech.org/languages/2006/rdf-cl#"
  xmlns:fml="http://www.semwebtech.org/languages/2006/fml#"
  xmlns:uni="http://www.semwebtech.org/domains/2006/uni#"
  xmlns:smtp="http://www.semwebtech.org/domains/2006/smtp#"
  xml:base="foo:rule">

<eca:Rule>

  <!-- the following variable declarations are redundant
  <eca:uses-variable rdf:resource="#student"/>
  <eca:uses-variable rdf:resource="#subject"/>
  <eca:uses-variable rdf:resource="#student-email"/>
  <eca:uses-variable rdf:resource="#secondfailure"/>
  <eca:uses-variable rdf:resource="#lecturer"/>
  <eca:uses-variable rdf:resource="#lecturer-email"/> -->
<eca:has-event>
  <snoop:Sequence>
    <snoop:bind-to-variable rdf:resource="#event"/>
    <snoop:first>
      <owlq:Class>
        <owlq:baseClass rdf:resource="&uni;Failure"/>
        <!-- the constraints define join variables between the two
            relevant occurrences of the event -->
        <owlq:hasConstraint>
          <owlq:Constraint>
            <owlq:onProperty rdf:resource="&uni;student"/>
            <owlq:equalsVariable rdf:resource="#student"/>
          </owlq:Constraint>
        </owlq:hasConstraint>
        <owlq:hasConstraint>
          <owlq:Constraint>
            <owlq:onProperty rdf:resource="&uni;subject"/>
            <owlq:equalsVariable rdf:resource="#subject"/>
          </owlq:Constraint>
        </owlq:hasConstraint>
      </owlq:Class>
    </snoop:first>
    <snoop:second>
      <owlq:Class>
        <owlq:baseClass rdf:resource="&uni;Failure"/>
        <!-- the constraints define join variables between the two

```

```

    relevant occurrences of the event -->
<owlq:hasConstraint>
  <owlq:Constraint>
    <owlq:onProperty rdf:resource="&uni;student"/>
    <owlq:equalsVariable rdf:resource="#student"/>
  </owlq:Constraint>
</owlq:hasConstraint>
<owlq:hasConstraint>
  <owlq:Constraint>
    <owlq:onProperty rdf:resource="&uni;subject"/>
    <owlq:equalsVariable rdf:resource="#subject"/>
  </owlq:Constraint>
</owlq:hasConstraint>
</owlq:Class>
</snoop:second>
</snoop:Sequence>
</eca:has-event>

<!-- get e-mail addresses of student and lecturer -->
<eca:has-query>
  <owlq:Query>
    <owlq:hasConstraint>
      <owlq:Constraint>
        <owlq:onVariable rdf:resource="#student"/>
        <owlq:onProperty rdf:resource="&uni;has-email"/>
        <owlq:equalsVariable rdf:resource="#student-email"/>
      </owlq:Constraint>
      <!-- get the lecturer of the 2nd exam in the
        multi-occurrences -->
    </owlq:hasConstraint>
    <owlq:hasConstraint>
      <owlq:Constraint>
        <owlq:onVariable rdf:resource="#event"/>
        <owlq:onProperty rdf:resource="\_2"/>
        <owlq:equalsVariable rdf:resource="#secondfailure"/>
      </owlq:Constraint>
    </owlq:hasConstraint>
    <owlq:hasConstraint>
      <owlq:Constraint>
        <owlq:onVariable rdf:resource="#secondfailure"/>
        <owlq:onProperty rdf:resource="&uni;lecturer"/>
        <owlq:equalsVariable rdf:resource="#lecturer"/>
      </owlq:Constraint>
    </owlq:hasConstraint>
    <owlq:hasConstraint>
      <owlq:Constraint>
        <owlq:onVariable rdf:resource="#lecturer"/>
        <owlq:onProperty rdf:resource="&uni;has-email"/>
        <owlq:equalsVariable rdf:resource="#lecturer-email"/>
      </owlq:Constraint>
    </owlq:hasConstraint>
  </owlq:Query>
</eca:has-query>

```

```

<!-- was the second failure a written or oral exam? -->
<!-- this could alternatively be encoded already in the
      second failure event -->
<eca:has-test>
  <owlq:Constraint>
    <owlq:onVariable rdf:resource="#secondfailure"/>
    <owlq:onProperty rdf:resource="&rdf;type"/>
    <owlq:hasValue rdf:resource="&uni;oral-exam"/>
  </owlq:Constraint>
</eca:has-test>

<!-- start process -->
<eca:has-action>
  <ccs:Sequence>
    <ccs:has-subprocess>
      <ccs:Fixpoint>
        <ccs:fix-variable rdf:resource="#X"/>
        <ccs:has-index>1</ccs:has-index>
        <ccs:local-variable rdf:resource="#datetime"/>
        <ccs:local-variable rdf:resource="#time"/>
        <ccs:local-variable rdf:resource="#room"/>
        <ccs:has-body>
          <ccs:Sequence>
            <ccs:has-subprocess>
              <!-- from the smtp domain metadata it is clear that this is a mars:Action -->
              <rdf-cl:NodeSpec rdf:nodeID="seq2.1">
                <rdf-cl:ofClass rdf:resource="&smtp;Send-Mail"/>
                <rdf-cl:with-property>
                  <rdf-cl:PropertySpec>
                    <rdf-cl:property rdf:resource="&smtp;to"/>
                    <rdf-cl:value-of rdf:resource="#lecturer-email"/>
                  </rdf-cl:PropertySpec>
                </rdf-cl:with-property>
                <rdf-cl:with-property>
                  <rdf-cl:PropertySpec>
                    <rdf-cl:property rdf:resource="&smtp;text"/>
                    <rdf-cl:value>Please suggest appointment for {#student}, {#subject}</rdf-cl:value>
                  </rdf-cl:PropertySpec>
                </rdf-cl:with-property>
                <ccs:followedBy rdf:nodeID="seq2.2"/>
              </rdf-cl:NodeSpec>
            </ccs:has-subprocess>
            <ccs:has-subprocess>
              <owlq:Query rdf:nodeID="seq2.2">
                <!-- evaluation binds $datetime -->
                <owlq:hasResultClass>
                  <owlq:Class xmlns:owlq="http://www.semwebtech.org/languages/2006/owlq#"
                    rdf:about="&uni;AppointmentProposal">
                    <owlq:hasConstraint>
                      <owlq:Constraint>
                        <owlq:onProperty rdf:resource="&uni;at-time"/>
                        <owlq>equalsVariable rdf:resource="#datetime"/>
                      </owlq:Constraint>
                    </owlq:hasConstraint>

```

```

    </owlq:Class>
  </owlq:hasResultClass>
  <ccs:followedBy rdf:nodeID="seq2.3"/>
</owlq:Query>
</ccs:has-subprocess>
<ccs:has-subprocess>
  <ccs:Query rdf:nodeID="seq2.3">
    <ccs:has-language rdf:resource="&httpget;"/>
    <ccs:input-variable rdf:resource="#datetime"/>
    <ccs:bind-to-variable rdf:resource="#room"/>
    <ccs:has-specification rdf:parseType="Literal">
      http://localhost/free-rooms?date=$date&time=$time
    </ccs:has-specification>
    <ccs:followedBy rdf:nodeID="seq2.4"/>
  </ccs:Query>
</ccs:has-subprocess>
<ccs:has-subprocess>
  <ccs:Assertion rdf:nodeID="seq2.4"> <!-- assertion expressed by an owl constraint -->
    <rdf:type rdf:resource="&owlq;Constraint"/>
    <owlq:onVariable rdf:resource="#room"/>
    <owlq:onProperty rdf:resource="&uni;available-at"/>
    <owlq:hasValue rdf:resource="#datetime"/>
    <ccs:followedBy rdf:nodeID="seq2.5"/>
  </ccs:Assertion>
</ccs:has-subprocess>
<ccs:has-subprocess>
  <ccs:Alternative rdf:nodeID="seq2.5">
    <ccs:has-subprocess>
      <ccs:Test> <!-- test expressed by an owl constraint -->
        <rdf:type rdf:resource="&owlq;Constraint"/>
        <!-- not satisfied when #room is unbound -->
        <owlq:onVariable rdf:resource="#room"/>
        <owlq:onProperty rdf:resource="&rdf;type"/>
        <owlq:hasValue rdf:resource="&uni;Room"/>
      </ccs:Test>
      <!-- then go on-->
    </ccs:has-subprocess>
    <ccs:has-subprocess>
      <ccs:Sequence>
        <ccs:has-subprocess>
          <ccs:Test rdf:nodeID="seq3.1"> <!-- test expressed by a fml Negation -->
            <rdf:type rdf:resource="&fml;Negation"/>
            <fml:negates>
              <owlq:Constraint>
                <!-- not satisfied when #room is unbound -->
                <owlq:onVariable rdf:resource="#room"/>
                <owlq:onProperty rdf:resource="&rdf;type"/>
                <owlq:hasValue rdf:resource="&uni;Room"/>
              </owlq:Constraint>
            </fml:negates>
            <ccs:followedBy rdf:nodeID="seq3.2"/>
          </ccs:Test>
        </ccs:has-subprocess>
      </ccs:has-subprocess>
    </ccs:Alternative>
  </ccs:has-subprocess>
</ccs:has-subprocess>

```

```

        <ccs:ContinueFixpoint rdf:nodeID="seq3.2">
          <ccs:withVariable rdf:resource="#X"/> <!-- the CCS fix variable -->
        </ccs:ContinueFixpoint>
      </ccs:has-subprocess>
    </ccs:Sequence>
  </ccs:has-subprocess>
</ccs:Alternative>
</ccs:has-subprocess>
</ccs:Sequence>
</ccs:has-body>
<ccs:followedBy rdf:nodeID="seq1.2"/>
</ccs:Fixpoint>
</ccs:has-subprocess>
<ccs:has-subprocess>
  <!-- from the smtp domain metadata it is clear that this is a mars:Action -->
  <rdf-cl:NodeSpec rdf:nodeID="seq1.2">
    <rdf-cl:ofClass rdf:resource="&smtp;Send-Mail"/>
    <rdf-cl:with-property>
      <rdf-cl:PropertySpec>
        <rdf-cl:property rdf:resource="&smtp;to"/>
        <rdf-cl:value-of rdf:resource="#lecturer-email"/>
      </rdf-cl:PropertySpec>
    </rdf-cl:with-property>
    <rdf-cl:with-property>
      <rdf-cl:PropertySpec>
        <rdf-cl:property rdf:resource="&smtp;to"/>
        <rdf-cl:value-of rdf:resource="#student-email"/>
      </rdf-cl:PropertySpec>
    </rdf-cl:with-property>
    <rdf-cl:with-property>
      <rdf-cl:PropertySpec>
        <rdf-cl:property rdf:resource="&smtp;text"/>
        <rdf-cl:value>exam with {#student}, {#subject} at
          {#datetime} in room {#room}</rdf-cl:value>
      </rdf-cl:PropertySpec>
    </rdf-cl:with-property>
  </rdf-cl:NodeSpec>
</ccs:has-subprocess>
</ccs:Sequence>
</eca:has-action>
</eca:Rule> </rdf:RDF>

```

[Filename: example-rule-exam.rdf]

The following small files illustrates how to glue notions of different languages together:

```

<!DOCTYPE rdf:RDF [
  <!ENTITY eca "http://www.semwebtech.org/languages/2006/eca-ml#">
  <!ENTITY snoop "http://www.semwebtech.org/languages/2006/snoopy#">
  <!ENTITY ccs "http://www.semwebtech.org/languages/2006/ccs#">
  <!ENTITY owlq "http://www.semwebtech.org/languages/2006/owlq#"> ]>
<rdf:RDF
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:owl="http://www.w3.org/2002/07/owl#">

```



```

<rdf:Description rdf:about="&snoop;Variable">
  <owl:equivalentClass rdf:resource="&eca;Variable"/>
</rdf:Description>
<rdf:Description rdf:about="&ccs;Variable">
  <owl:equivalentClass rdf:resource="&eca;Variable"/>
</rdf:Description>
<rdf:Description rdf:about="&ccs;EventSpec">
  <owl:equivalentClass rdf:resource="&eca;EventSpec"/>
</rdf:Description>
<rdf:Description rdf:about="&ccs;Test">
  <owl:equivalentClass rdf:resource="&eca;Test"/>
</rdf:Description>
<rdf:Description rdf:about="&ccs;Query">
  <owl:equivalentClass rdf:resource="&eca;Query"/>
</rdf:Description>
<rdf:Description rdf:about="&owlq;Variable">
  <owl:equivalentClass rdf:resource="&eca;Variable"/>
</rdf:Description>
</rdf:RDF>

```

[Filename: example-glue.rdf]

The following query shows that all rule-wide variables are identified as `eca:Variables` from the `rdfs:range` predicates. For the analysis of positive/negative variables, see Section 13.2.2.

```
# call # jena -q -pellet -il RDF/XML -qf rule-query-4.sparql
```

```

PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> PREFIX
rdfs: <http://www.w3.org/2000/01/rdf-schema#> PREFIX owl:
<http://www.w3.org/2002/07/owl#> PREFIX mars:
<http://www.semwebtech.org/2006/mars#> PREFIX eca:
<http://www.semwebtech.org/languages/2006/eca-ml#> PREFIX snoop:
<http://www.semwebtech.org/languages/2006/snoopy#> PREFIX ccs:
<http://www.semwebtech.org/languages/2006/ccs#> PREFIX owlq:
<http://www.semwebtech.org/languages/2006/owlq#> PREFIX uni:
<http://www.semwebtech.org/domains/2006/uni#>

```

```

SELECT ?X ?Y ?Z FROM <file:eca-ml.rdf> FROM <file:snoopy.rdf> FROM
<file:ccs.rdf> FROM <file:owlq.rdf> FROM <file:rdf-cl.rdf> FROM
<file:example-rule-exam.rdf> FROM <file:example-glue.rdf> WHERE
{{?X a owl:Variable} UNION
  {?Y ccs:followedBy ?Z}}

```

[Filename: rule-query-4.sparql]

Example 49 *Figure 13.1 shows an excerpt of the rule given in Example 32, now in full RDF: “If a flight is first delayed and then canceled (note: use of a join variable), then ...”*

```

<!DOCTYPE rdf:RDF [
  <!ENTITY rdf "http://www.w3.org/1999/02/22-rdf-syntax-ns#">
  <!ENTITY owlq "http://www.semwebtech.org/languages/2006/owlq#">
  <!ENTITY rdf "http://www.w3.org/1999/02/22-rdf-syntax-ns#">
  <!ENTITY travel "http://www.semwebtech.org/domains/2006/travel#"> ]>
<rdf:RDF

```

```

xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
xmlns:eca="http://www.semwebtech.org/languages/2006/eca-ml#"
xmlns:snoop="http://www.semwebtech.org/languages/2006/snoopy#"
xmlns:owlq="http://www.semwebtech.org/languages/2006/owlq#"
xmlns:travel="http://www.semwebtech.org/domains/2006/travel#"
xml:base="foo:rule">

<eca:Rule rdf:ID="rule-df-cf">
<!-- the following variable declarations are redundant -->
<eca:uses-variable rdf:resource="#flight"/>
<eca:has-event>
  <snoop:Sequence>
    <snoop:first>
      <owlq:Class rdf:ID="ev-df">
        <owlq:baseClass rdf:reference="&travel;delayed-flight">
          <!-- the constraints define join variables between the two
               relevant occurrences of the event -->
          <owlq:hasConstraint>
            <owlq:Constraint>
              <owlq:onProperty rdf:resource="&travel;flight"/>
              <owlq:equalsVariable rdf:resource="#flight"/>
            </owlq:Constraint>
          </owlq:hasConstraint>
        </owlq:Class>
      </snoop:first>
      <snoop:second>
        <owlq:Class rdf:ID="ev-cf">
          <owlq:baseClass rdf:reference="&travel;anceled-flight">
            <!-- the constraints define join variables between the two
                 relevant occurrences of the event -->
            <owlq:hasConstraint>
              <owlq:Constraint>
                <owlq:onProperty rdf:resource="&travel;flight"/>
                <owlq:equalsVariable rdf:resource="#flight"/>
              </owlq:Constraint>
            </owlq:hasConstraint>
          </owlq:Class>
        </snoop:second>
      </snoop:Sequence>
    </eca:has-event>
    <!-- ... query and action ... -->
  </eca:Rule>

```

[Filename: example-delay-cancel.rdf]

URIs of ECA Rule Instances and Components. With the above point of view, it is obvious that each rule, and also each of its components and subexpressions is an RDF object that has a corresponding RDF URI.

Rule components, e.g., process specifications or composite event specifications can also be given in standalone RDF sources, and can then be used by other specifications or rules.

These URIs are e.g. used as identifiers when registering rules or event components, and also when communicating answers to such requests.

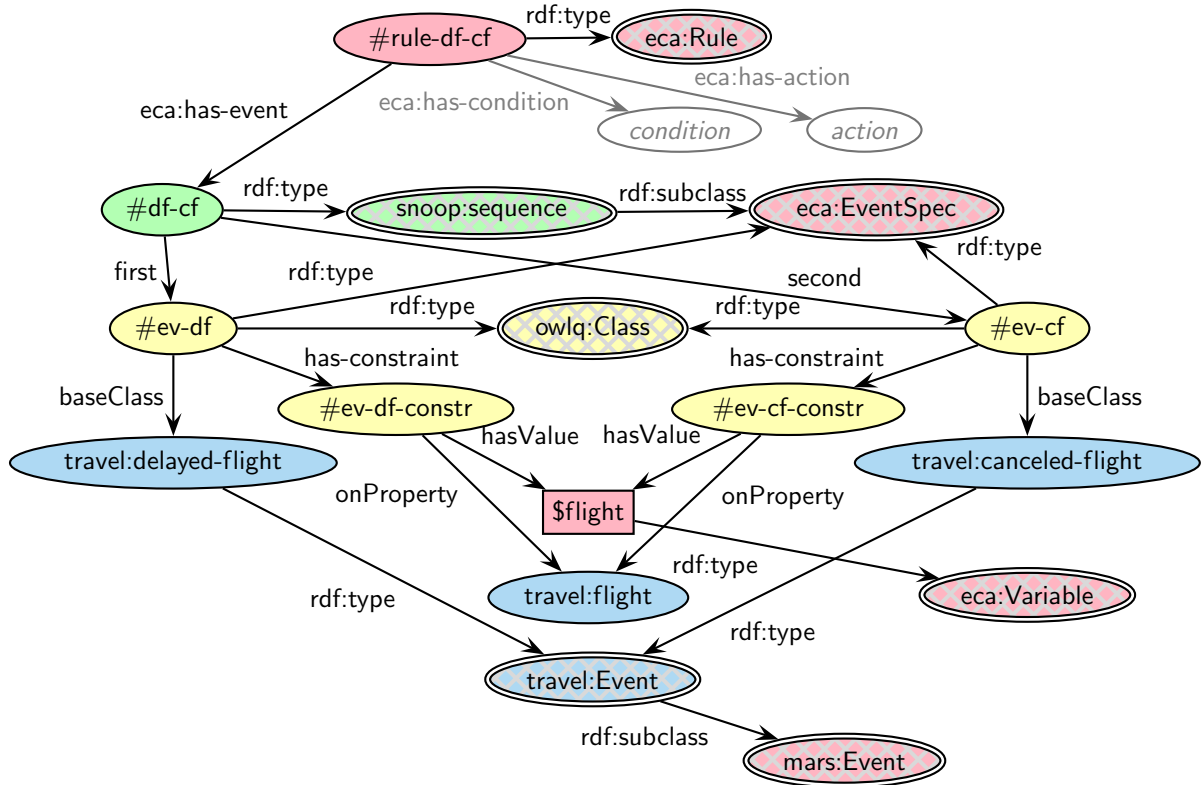


Figure 13.1: Example Rule and Event Component as Resources

A note on Variable Bindings Note that there is no RDF/OWL ontology for variable bindings. The variable bindings, answers etc stuff is purely for data exchange, which is not appropriate to be modeled and transferred in RDF format. Here, we stay with pure XML.

13.2 Reasoning about Rules

Having rules as RDF graphs allows for several kinds of reasoning about the rules. As the reference languages themselves are not yet stable, the reasoning section is still subject to change according to evaluation of the component languages.

13.2.1 Validity

The glue ontologies define how subontologies may be composed, mainly by associating classes of “higher” and “lower” ontologies (wrt. the tree representation of the rules): If a composition is not allowed, the OWL reasoner complains (e.g., when combining an action as a subexpression of an event expression). On the other hand, especially with this definition, one must be careful in the ontology design: CCS process expressions may e.g. *contain* events and queries.

Additionally, reasoning about the use of variables is a primary goal that can immediately be used for supporting the evaluation of rules.

13.2.2 Safety and Evaluation

Reasoning about the use of variable and the related notions of *safety* uses the ontology for variables that has already been introduced in Section 11.5. The usage of variables can be added as (derived) labeled edges in the RDF graph of a rule:

- usage of variables is defined by structural induction over expressions of a language,
- expressions are represented by the RDF representation of the rules and subexpressions,
- every sublanguage must provide the rules for that inductive definition for atomic expressions and for composite expressions:

*uses-relationship*₁(*expression, var*) :-
*structural-relationship*₂(*expr, subexpression*), *uses-relationship*₂(*subexpression, var*).

Note that “atomic” expressions on a higher level may be composite expressions from a lower level. As they are then members of two classes, the variable properties are propagated upwards.

- as the rule structure is a tree (or at least a directed graph, in case that subexpressions occur multiply), stratified negation is allowed in these definitions.

For the actual implementation, there are currently several alternatives:

- use Pellet with SWRL (prototype SVN branch),
- encode the expressions and the rules in F-Logic and use Florid,
- combine Jena with a Florid-based Service for F-Logic Rules.

The third alternative is anticipated, based on [39]. The resulting information allows to analyze *variable dependencies* between fragments. Then, queries can be regarded as *sets* of subqueries, and any safe evaluation strategy can be applied. Furthermore, subqueries that only need a subset of the variables can sometimes be evaluated earlier, already during the incremental event detection process.

13.2.3 Typing

As already mentioned, the handling of variables as resources and with references allows to do type reasoning about the variables, and thus also about the schema of the generated structures (e.g., when creating RDF graphs that are added to a data source).

13.2.4 Compositionality

The MARS Framework is actually a comprehensive compositional approach that allows to compose specific languages for events, processes, queries, etc. to rules. Having an ontology-level description of *semantics* would allow for inter-language, probably distributed analysis of the semantics of rules. Such a “semantic semantics” does not yet exist, but considering that “generic” notions like denotational semantics and operational semantics, and translations from many formalisms to model checking exist, show a promising research direction. First discussions with members of WG A3 (Typing and Compositionality) have taken place in Autumn 2006.

Chapter 14

MARS Appendices

14.1 Abbreviations

ECA: Event-Condition-Action

ECE: Event-Condition-Event derivation rule

ACA: Action-Condition-Action implementation rule

CQE: Continuous-Query-Event derivation rule

GRH: Generic Request Handler (ECA Engine Implementation]

CES: Composite Event Specification

CEL: Composite Event Language

CED: Composite Event Detection (Service) (for some CEL)

AES: Atomic Event Specification

AESL: Atomic Event Specification Language

AEM: Atomic Event Matcher (for some AESL)

QL: Query Language

QE: Query Engine (for some QL)

CAL: Composite Action Language

CAE: Composite Action Execution Engine (for some CAL)

EB: Event Broker

DB: Domain Broker (usually contains an Event Broker)

DN: Domain Node

DSR: Domain Service Registry

LSR: Languages and Services Registry

14.2 MARS Infrastructure URIs and Filenames

14.2.1 RDF Files

MARS Ontology Files.

- Ontology of domains: `mars-domain-notions.rdf`
Actual Location:
`http://www.semwebtech.org/2006/mars/mars-domain-notions.rdf`
- Ontology of languages and service types: `mars-language-and-service-types.rdf`
Actual Location:
`http://www.semwebtech.org/2006/mars/mars-language-and-service-types.rdf`
- Ontology of service characteristics: `services-ontology.rdf`
Actual Location:
`http://www.semwebtech.org/2006/mars/services-ontology.rdf`
- Ontology of expression structure: `expr.rdf`
Actual Location:
`http://www.semwebtech.org/2006/mars/expr.rdf`

The basic LSR data in RDF.

- The LSR RDF file `lsr.rdf`
Actual Location:
`http://www.semwebtech.org/2006/mars/lsr.rdf`

Sample and Reference Languages.

- ECA-ML ontology: `eca-ml.rdf`
Actual Location:
`http://www.semwebtech.org/languages/2006/eca-ml.rdf`
- Snoopy Ontology: `snoopy.rdf`
Actual Location:
`http://www.semwebtech.org/languages/2006/snoopy.rdf`
- CCS Ontology: `ccs.rdf`
Actual Location:
`http://www.semwebtech.org/languages/2006/ccs.rdf`
- OWLQ Ontology: `owlq.rdf`
Actual Location:
`http://www.semwebtech.org/languages/2006/owlq.rdf`
- RDF-CL Ontology: `rdf-cl.rdf`
Actual Location:
`http://www.semwebtech.org/languages/2006/rdf-cl.rdf`
- Travel Domain: `travel.rdf`
Actual Location:
`http://www.semwebtech.org/domains/2006/travel.rdf`
`orafly.rdf`
Actual Location: `http://www.semwebtech.org/nodes/2006/orafly/orafly.rdf`

14.2.2 URI Schemata

Language Types.

Used in Files `files/mars-language-and-service-types.rdf`, and `lsr.rdf`:

`http://www.semwebtech.org/2006/mars#RuleLanguage`

Service Types and Tasks.

Used in Files `services-ontology.rdf` and `lsr.rdf`:

`http://www.semwebtech.org/2006/mars#ECAService`

`http://www.semwebtech.org/2006/mars/eca-service#register-rule`

Languages.

In general, languages can have any URI (which the author assigns to them). The MARS reference languages are identified as follows (used in File `lsr.rdf`):

`http://www.semwebtech.org/languages/2006/snoopy`

Services.

In general, services can have any URI (which the author assigns to them). The services that implement the MARS reference languages are identified as follows (used in File `lsr.rdf`):

`http://www.semwebtech.org/services/2006/snoopy`

14.3 Language DTDs

14.3.1 DTD for Logical Stuff: Variable Bindings etc.

```
<!ELEMENT answers (answer*)>
<!ELEMENT answer (result?, variable-bindings?)>
<!ELEMENT result ANY>
<!ELEMENT variable-bindings (tuple+)>
<!ELEMENT tuple (variable+)>
<!ELEMENT variable ANY>
<!ATTLIST variable name CDATA #REQUIRED
                ref URI #IMPLIED> <!-- variable has either ref or content-->
```

14.3.2 DTD of ECA-ML

```
<!ELEMENT Rule (%variable-decl, Event, Query*, Test?, Action+)>
<!ELEMENT Event (%variable-decl, (atomic-event | opaque | any-element-of-a-cesl))>
<!ELEMENT Query (%variable-decl, (opaque | any-element-of-a-ql))>
<!ELEMENT Test (%variable-decl, (opaque | any-element-of-boolean-ml))>
<!ELEMENT Action (%variable-decl, (opaque | any-element-of-an-action-ml))>
  <!ATTLIST Action executed-at (URL | "owner") #IMPLIED>
```

14.3.3 Snoop DTDs

As described in Section 11.1, every language has a DTD for the actual Markup, and a DTD for the intermediate striped RDF/XML markup that connects the actual markup with the complete set of notions defined by the ontology (cf. Section 11.7.1).

14.3.3.1 The SNOOP Striped DTD

```
<!ENTITY % operand "(And | Or | Sequence | Not |
```

```

Any | MultiOccurrences |
Aperiodic | CumulativeAperiodic | Periodic | CumulativePeriodic |
AtomicEvent | ANY-from-other-namespace |
Opaque |
Variable )">
<!ELEMENT And (has-subevent, has-subevent)> <!ELEMENT Or
(has-subevent, has-subevent)> <!ELEMENT Sequence ()> <!ELEMENT Any
((has-subevent)*)> <!ATTLIST Any number-of-occurrences CDATA
#REQUIRED> <!ELEMENT MultiOccurrences (has-subevent)> <!ATTLIST
MultiOccurrences
    number-of-occurrences CDATA #REQUIRED>
<!ELEMENT Aperiodic (has-start, fired-by, has-end)> <!ELEMENT
CumulativeAperiodic (has-start, cumulate, has-end)> <!ELEMENT
Periodic (has-start, fired-by, has-end)> <!ELEMENT
CumulativePeriodic (has-start, cumulate, has-end)> <!ELEMENT Not
(has-start, exit-not, has-end)> <!ELEMENT AtomicEvent ANY>
<!ELEMENT Opaque ANY>
<!ELEMENT Variable (%operand;)>
<!ATTLIST Variable name CDATA #REQUIRED>

<!ELEMENT has-subevent (%operand;)>
<!ELEMENT has-start (%operand;)>
<!ELEMENT has-end (%operand;)>
<!ELEMENT cumulate (%operand;)>
<!ELEMENT fired-by (%operand;)>
<!ELEMENT exit-not (%operand;)>

<!ATTLIST extension for atomic-event and non-snoop subelements:
    xmlns:other-namespace CDATA #IMPLIED
    language CDATA #IMPLIED
    domain CDATA #IMPLIED>

<!ATTLIST extension for opaque:
    language CDATA #REQUIRED
    domain CDATA #REQUIRED>

<!ATTLIST extension for all class elements except variable:
    variable CDATA #IMPLIED>

```

[Filename:snoop-stripped.dtd]

14.3.3.2 The SNOOP-ML DTD

```

<!ENTITY % operand "(And | Or | Sequence | Not |
    Any | MultiOccurrences |
    Aperiodic | CumulativeAperiodic | Periodic | CumulativePeriodic |
    AtomicEvent | ANY-from-other-namespace |
    Opaque |
    Variable )">
<!ELEMENT And (%operand;, %operand;)>
<!ELEMENT Or (%operand;, %operand;)>
<!ELEMENT Sequence (%operand;, %operand;)>
<!ELEMENT Any ((%operand;)*>
<!ATTLIST Any number-of-occurrences CDATA #REQUIRED>
<!ELEMENT MultiOccurrences (%operand;)>

```



```
<!ATTLIST MultiOccurrences
  number-of-occurrences CDATA #REQUIRED>
<!ELEMENT Aperiodic (%operand;, %operand;, %operand;)>
<!ELEMENT CumulativeAperiodic (%operand;, %operand;, %operand;)>
<!ELEMENT Periodic (%operand;, %operand;, %operand;)>
<!ELEMENT CumulativePeriodic (%operand;, %operand;, %operand;)>
<!ELEMENT Not (%operand;, %operand;, %operand;)>
<!ELEMENT AtomicEvent ANY> <!ELEMENT Opaque ANY>
<!ELEMENT Variable (%operand;)>
<!ATTLIST Variable name CDATA #REQUIRED>

<!ATTLIST extension for atomic-event and non-snoop subelements:
  xmlns:other-namespace CDATA #IMPLIED
  language CDATA #IMPLIED
  domain CDATA #IMPLIED>

<!ATTLIST extension for opaque:
  language CDATA #REQUIRED
  domain CDATA #REQUIRED>

<!ATTLIST extension for all elements except variable:
  variable CDATA #IMPLIED>
```

[Filename:snoop.dtd]

Part II

r^3 : Towards a foundational ontology for reactive rules

Chapter 15

The r^3 ontology

In this part we present first results of the r^3 project towards defining an (OWL-DL) foundational ontology for reactive rules. The current proposal is at a low (structural) abstraction level; the extension of this proposal towards characterizing higher abstraction level concepts, like domain/application specific languages (vs. algebraic and general-purpose languages) is under consideration. Although focusing on reactive rules, the r^3 ontology defines a vocabulary that allows for the definition of rule (component) languages that can be extended to model other types of rules.

We start (in section 15.1) by introducing the r^3 ontology from a global point of view¹. In the following sections we detail the r^3 ontology explaining (and illustrating²) how to define different languages (in section 15.2), and how to use these languages to define heterogeneous rules (in section 15.3). The definition of the r^3 ontology is further complemented (in section 15.4) by sketching how the ontology can be used to provide additional semantic value to heterogeneous XML markups. We do this by describing how an RDF model, compliant with the r^3 ontology, may be automatically obtained from an XML markup adhering to the previously proposed general markup principles (viz. ECA-ML).

The complete r^3 ontology (<http://rewerse.net/I5/NS/2006/r3>, v0.20), as well as the implementation and demonstration of r^3 prototype (in its current state, v0.12, as described in [2]) can be found at <http://rewerse.net/I5/r3/>.

15.1 An Ontology for Reactive Rules

From a global point of view, the r^3 ontology defines a vocabulary to describe `RulePackages`³. A `RulePackage` is just a single `Rule` or a `RuleSet`, as depicted in figure 15.1. In order to cope with the language heterogeneity requirements previously stated, the r^3 ontology includes a meta-language/vocabulary to describe rule (component) languages/vocabularies. This meta-level assumes a term based structure for rule languages where atomic functors are composed using algebraic operators. To account for languages not compliant with this term structure (like purely textual languages, e.g. database triggers, external program call for action, etc) opaque/native “functors” are also allowed.

In terms of this meta-vocabulary, a `RulePackage` is a language `Construction` (exclusively a `Rule` or a `RuleSet`), which is a `RulePackageConstruct`, taking a set of `Arguments` (`Rule Components` or `RuleSet Members`). The complete definition of a `Construction` is to be provided in section 15.3,

¹The r^3 OWL-DL ontology is presented here using UML2 diagrams for the sake of readability.

²Examples illustrating RDF models: use N3 notation, omit prefix declarations, and assume r^3 as the empty(‘:’) prefix.

³For the sake of conciseness, we use generic terms, like “Rule” or “Language”, whereas more specific terms, like “Reactive Rule” or “Reactive Rule Language”, would have been more consistent with the actual motivation of the presented work. As such, it is worth stressing here that the adequacy of the r^3 ontology to describe rules in general, although not excluded, has not been properly evaluated and is, at the time of this writing, an open matter.

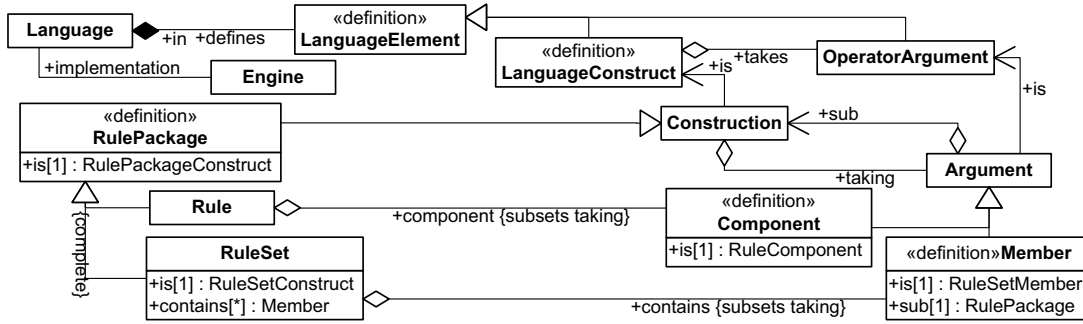


Figure 15.1: Rule Packages

but for now, cf. figure 15.1, a Construction is a LanguageConstruct taking a set of Arguments, in accordance with its LanguageConstruct that takes a set of OperatorArguments (non-empty, in case of an algebraic operator). An Argument is a specific OperatorArgument including a sub-Construction.

LanguageConstructs and OperatorArguments constitute some of the specialized forms of LanguageElements defined in a Language. The actual operational implementation of (the elements of) a Language is provided by some Engine.

A RuleSet is a RuleSetConstruct that recursively contains other RulePackages (viz. Member is a RuleSetMember including a sub-RulePackage). A Rule, as further detailed in figure 15.2, is either a NativeRule defined by a mandatory (and usually textual) opaque source; or an AbstractRule, fully defined by a non-empty set of Components (where a Component is a RuleComponent).

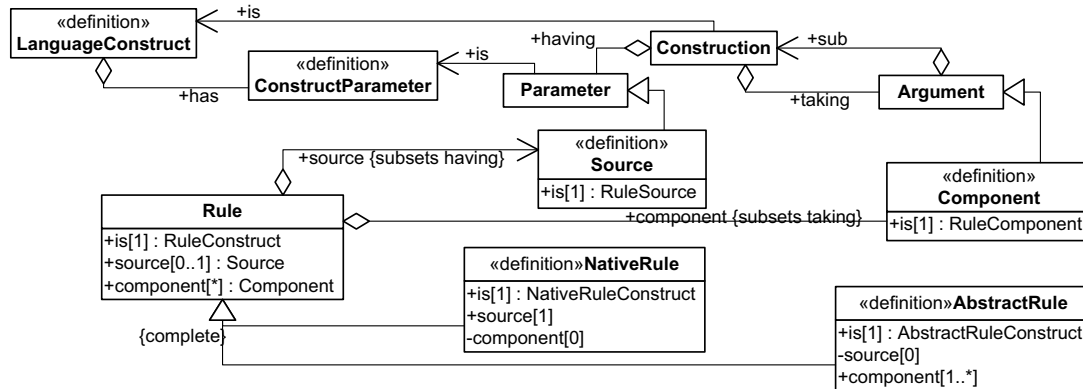


Figure 15.2: Rules

The opaque Source of a NativeRule is neither a Construction nor an Argument. Instead, it is defined as a Parameter (one that is a RuleSource). This leads to an extension of the LanguageConstruct concept as being a parametric construct. Any LanguageConstruct has an (optionally empty) set of ConstructParameters, another kind of LanguageElements. As such, the language Construction concept, besides taking a set of Arguments, is further refined, in figure 15.2, as having a set of Parameters, where a Parameter is a ConstructParameter. The full definition of an r^3 Language is presented in the next section, followed by the complete definition of an heterogenous Construction in section 15.3.

15.2 Defining Reactive Rule Languages

In order to define the different Languages involved in a RulePackage, the r^3 ontology provides a meta-vocabulary, depicted in figure 15.3 at its core level.

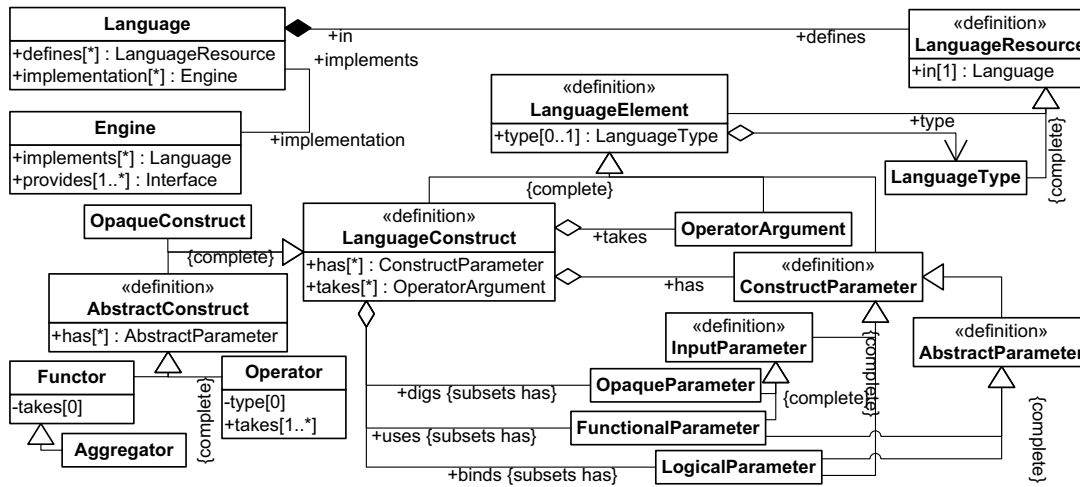


Figure 15.3: Languages

Some of the more generic terms/concepts (e.g. Language, LanguageElement) are pretty straightforward and have already been introduced before. Nevertheless, it is worth mentioning that a Language actually defines a set of LanguageResources (LanguageElements or LanguageTypes, defined in the Language itself); and that a LanguageElement has to be either a LanguageConstruct, a ConstructParameter or an OperatorArgument. A LanguageConstruct is further distinguished between an OpaqueConstruct and an AbstractConstruct (one that has only AbstractParameters), the latter being either an algebraic Operator (if it takes at least one OperatorArgument) or a Functor⁴.

The set of ConstructParameters a LanguageConstruct has, may be explicitly divided according to the specific nature of each parameter: a LanguageConstruct uses (or digs) InputParameters and binds LogicalParameters. The semantics of a LanguageConstruct is not defined unless the actual value of all its InputParameters is known. Example 50 further illustrates the concept of an r^3 Language partially presenting the definition of some Languages.

Example 50 Below we present an r^3 partial definition of several rule component languages, that will be further exploited in subsequent examples. Namely, an event algebra (snoop), two domain/application languages (travel and rental), and two general-purpose libraries (mail and text).

```

snoop:sequence a :Operator; :in snoop:event-algebra;
:takes snoop:first, snoop:other.
travel:booking-place a :Functor; :in travel:domain;
:binds travel:client, travel:flightnr, travel:seat.
travel:flight-info a :Functor; :in travel:domain;
:uses travel:flightnr;
:binds travel:date, travel:origin, travel:destination.
rental:request-quotation-for-flight a :Functor; :in rental:application
:uses rental:client, rental:flightnr.
rental:get-client a :Functor; :in rental:application
:uses rental:client;
:binds rental:client-name, rental:favorite-class, rental:max-price.
rental:get-available-cars a :Functor; :in rental:application
  
```

⁴Also distinguishable, among Functors, are Aggregators which provide the means to construct arbitrary variable Aggregations, as a specialized form of Constructions, cf. section 15.3.

```

:uses rental:office, rental:date;
:binds rental:car, rental:car-class, rental:price.
mail:send a :Functor; :in mail:library;
:uses mail:from, mail:to, mail:subject, mail:body.
text:join a :Aggregator; :in text:library;
:uses text:template, text:separator.
text:replace a :Aggregator; :in text:library;
:uses text:template.

```

An `InputParameter` may first be seen as a purely `FunctionalParameter`, whose (atomic) values, the `LanguageConstruct` transparently `uses` to define its actual semantics. However, careful observation reveals that some of these parameters may hide semantic value (e.g. variable references to be expanded) by using “internal” (sub-)languages. In these cases, the full semantics of the construction is not accessible without full knowledge of the (sub-)languages actually used inside the (consequently, non-atomic) values of those parameters. Such parameters are called `OpaqueParameters` and any `LanguageConstruct` that `has` (or, more precisely, `digs`) at least one of such `OpaqueParameters` must be an `OpaqueConstruct` (since it cannot be a semantically transparent `AbstractConstruct`).

The semantics of a `Construction` that is an `OpaqueConstruct` is known only to an `Engine` that implements the associated `Language`. Most of the times these `Engines` will only define the semantics of such an `OpaqueConstruct` in operational terms, meaning that the only form of knowing it is to submit (at runtime) an actual `Construction` to the evaluation `Interface` that the `Engine` provides. Nevertheless, static analysis may some times be possible as long as the `Engine` provides a translation `Interface` for parsing such an opaque `Construction` into one that is an `AbstractConstruct`⁵.

In general, the definition of a `LanguageElement` may include a `type`. The exceptions to this are `Operators` and `OperatorArguments`, cf. figure 15.5, given their required algebraic nature. A `LanguageType` is considered here only as some resource that implicitly defines a domain of literal values (having a lexical space compatible with `rdf:XMLLiteral`) and may refer, for instance, to a particular XML Schema (element or type, basic or not). The actual treatment of types, in the context of a general framework like [52], has no yet been studied, but it may provide, for instance, the means for a safer equality relation (e.g. `xml:space` preserving or not). As such, the concept of a `LanguageType` is already included here providing the means to define, among others⁶, the domain of a `ConstructParameter` or the range of a `Functor`.

The r^3 meta-vocabulary for defining languages (given its core level, in figure 15.3) is extended in figure 15.4 with additional `LanguageConstructs` specific to rule languages, viz. `RulePackageConstructs`, which are clearly distinguished from other constructs, viz. `ExpressionConstructs`.

`ECARuleConstructs`, which are particularly relevant to our work at this rule level, provide, together with specific `ECARuleComponents`, the means for constructing ECA rules. Also relevant are `NativeRuleConstructs` that allow modelling ECA rules at an opaque level (e.g. database triggers). Notice that a `NativeRuleConstruct`, besides its `RuleSource` (that it `digs`), additionally `has` other `InputParameters` (if needed, as is usually the case of e.g. database name and user in database triggers).

On a general level, most practical language definitions include some forms of syntactic sugar. Among the most common forms are optional/multiple parameters/arguments (e.g. in figure 15.4 a `RuleSetMember` is defined as a `CollectionArgument`). These forms of syntactic sugar very often hide major semantic decisions. In an eclectic heterogenous context like the one proposed, the “sweetness” of these syntactic forms may be misleading and hinder the semantic understanding of a set of constructions. For instance, an ECA rule may be defined as taking several condition or action components, but the actual meaning of such constructions remains unknown or opaque unless a particular composer operator is explicitly chosen and identified. Different ECA languages may follow different directions in this respect: usually a simple conjunction is assumed for conditions,

⁵The concept of an `Engine` which implements a `Language` is core to the general ECA framework and the r^3 prototype, further details about how this concept helps realize the current version of the r^3 prototype may be found in [2].

⁶`LanguageTypes` are also used to constrain the domain of logical variables and the range of expressions, as explained later in section 15.3.

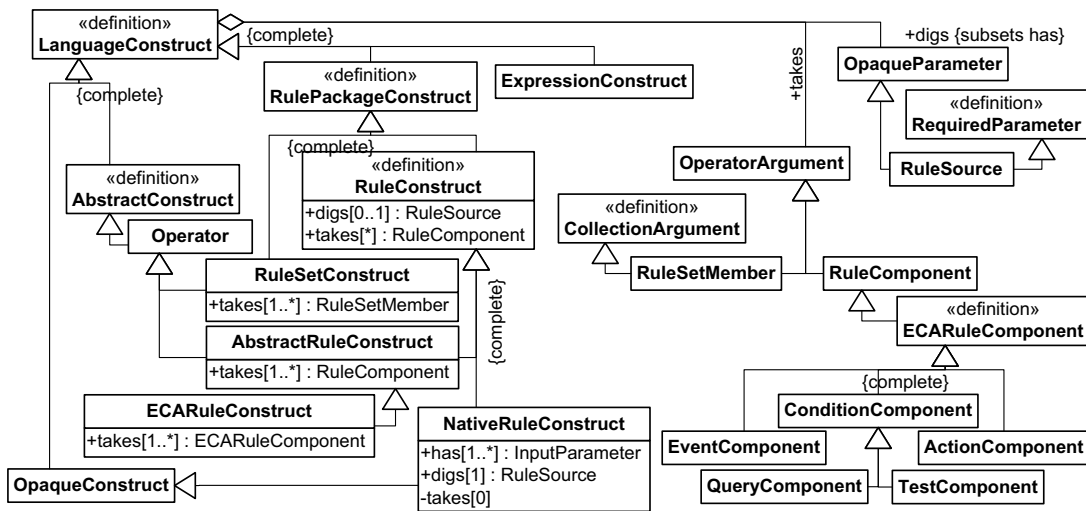


Figure 15.4: Rule Languages

but the choice becomes a bit more fuzzy for actions, particularly if transactions are considered, and things get even more involved if multiple event components are considered (where the natural choice would be to interpret them as a disjunction). Multiple/optional parameters/arguments are not actually allowed in `LanguageConstruct`. Nevertheless appropriate extension of the r^3 ontology (as presented in figure 15.5) may provide enough expressivity for the most “annoying” situations without disregarding semantic transparency.

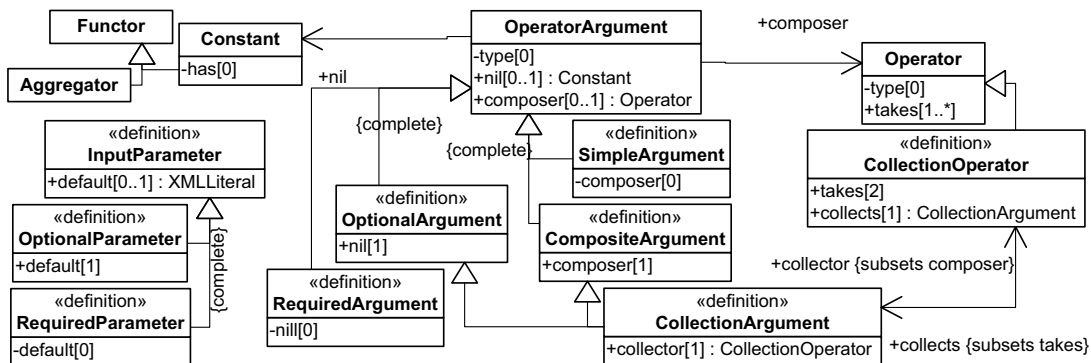


Figure 15.5: Language (not so Syntactical) Details

Optional parameters (and arguments) are modelled as usual with the introduction of `default` (and `nil`) constants that are to be used if the parameter (or argument) is omitted. For the moment, multiple parameters are not accounted for (assuming that a parameter value may as easily be a single value or a list of values using appropriate markups). Multiple arguments, as `CollectionArguments`, are considered a specialization of a more general concept: `CompositeArguments`. A `CompositeArgument`, if present, is always built using a specific `composer` `Operator` and most often it induces syntactic simplifications when defining markup languages (i.e. the `CompositeArgument` or its `composer` are omitted, as long as ambiguity issues are not introduced). A `CollectionArgument` (named after `rdf:parseType="Collection"`) has a recursive nature given by its `collector`, which, as a binary associative `CollectionOperator`, `takes` one occurrence and further `collects` additional occurrences of the `CollectionArgument` itself (terminated by a required `nil` `Constant`, which stands for the empty collection if there are no further occurrences of it). To further illustrate some of

these issues, consider example 51 where a definition of the ECA-ML Language (viz. `eca:ml`) is presented.

Example 51 *The `eca:ml` Language is defined in r^3 as:*

```

eca:rule a :ECARuleConstruct; :in eca:ml;
  :takes eca:event, eca:condition, eca:action.
eca:event a :EventComponent; :in eca:ml.
eca:condition a :ConditionComponent; :in eca:ml;
  :composer eca:collect-and-test; :nil eca:true.
eca:collect-and-test a :Operator; :in eca:ml;
  :takes eca:collect, eca:test.
eca:collect a :CollectionArgument; :in eca:ml;
  :collector eca:query-collect; :nil eca:true.
eca:query-collect a :CollectionOperator; :in eca:ml;
  :takes eca:query, eca:collect.
eca:query a :QueryComponent; :in eca:ml; :nil eca:true.
eca:test a :TestComponent; :in eca:ml; :nil eca:true.
eca:action a :ActionComponent; :in eca:ml.
eca:native a :NativeRuleConstruct; :in eca:ml;
  :uses eca:lang; :digs eca:source.
eca:opaque a :ExpressionConstruct; :in eca:ml;
  :uses eca:lang; :digs eca:literal.

```

Some of the `LanguageElements`, included in example 51, come directly from the actual ECA-ML markup and are self-explanatory, nevertheless those that don't, are worth mentioning. In the previously proposed ECA-ML markup the `eca:condition` “tree” is not explicitly included since the *collect and test composition* and the multiple *query collection* are always implicitly present and no syntactical ambiguity results from omitting them. As such, only the “leaf” arguments `eca:query` and `eca:test` were actually included in the ECA-ML markup. Furthermore, the ECA-ML markup does not include any `eca:native` element, instead it chooses to re-use a pair of other elements (`eca:rule` and `eca:opaque`). This pair of elements has no meaning in r^3 terms. In fact, the ECA-ML markup has chosen that syntactical form (of a pair) precisely because it bears no other meaning, and overloading these elements avoided the inclusion of a new term (e.g. `eca:native`). This is a good example of how a markup-based approach (even a striped one, favoring RDF/XML) differs from an ontology-based approach. The rationale behind the chosen terms for the vocabulary is simply different. Both approaches reject ambiguity, but in the former, minimizing verbosity and the number of vocabulary terms are driving forces (as long as readability and “syntactical elegance” are kept), whereas in the latter overloading of terms, without proper specialization of concepts, is to be avoided since it reduces the semantic value of the vocabulary. Both approaches are pertinent, but when defining new vocabularies for the Semantic Web, it is our stance that the former should not influence the latter, but the other way around: ontology-based vocabularies should influence and provide an appropriate abstract syntax for concrete markups, which may be improved, later, towards syntactical conciseness and elegance (although an XML markup is hardly the better serialization choice for those purposes).

As a final note about defining r^3 Languages, it is worth stressing that strict conformance to OWL-DL must be observed at all times. The `Language` meta-level of the r^3 ontology may lead to (sometimes tempting) misuses that may introduce undecidability issues. For instance, an ontology that imports the r^3 ontology and defines a particular `Language`, must do just that: define a resource that is a `Language` individual. Most frequently, such an ontology will define a set of terms (e.g. individual `LanguageElements`) in the ontology namespace. This namespace, as a URI reference, refers to *the ontology* itself which *is not* the `Language`, but instead *defines* the `Language` (which in turn, *defines*, e.g., its `LanguageElements`). Such an ontology (or its namespace, e.g. `eca:`) should never be mistaken for the `Language` itself (e.g. `eca:ml`). Another example of such (tempting) misuses, that must be avoided, is to define an OWL class as being a `LanguageType` individual, which is not a valid OWL-DL statement (one possible way around this would be to extend the `LanguageType` class with appropriate data-type properties).

15.3 Defining Reactive Rule Constructions

Given a set of Languages defined using the meta-vocabulary presented in the previous section, it is possible to define the general concept of an heterogenous CodingElement (a Construction, a Parameter or an Argument) which is a LanguageElement (respectively, a LanguageConstruct, a ConstructParameter or an OperatorArgument). The definition of a CodingElement (depicted in figure 15.6) follows a structure that closely mimics the structure of Languages, previously presented in figure 15.3).

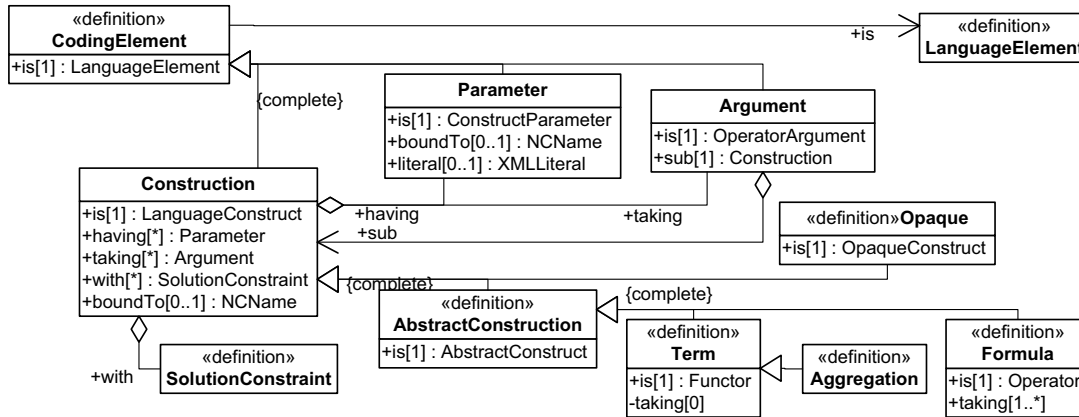


Figure 15.6: Language Constructions

A Construction may be seen as an instance of a LanguageConstruct, having a set of Parameters possibly with specific literal values (or boundTo a particular variable), and taking a concrete set of sub-Constructions as Arguments. Example 52 illustrates this concept with several concrete Construction instances.

Example 52 Find below some rule component constructions using the languages in example 50.

```

ex:event :is snoop:sequence;
:taking [:is snoop:first; :sub [
  :is travel:booking-place;
  :having [:is travel:client; :boundTo "Mail"];
  :having [:is travel:flightnr; :boundTo "Flight"]]];
:taking [:is snoop:other; :sub [
  :is rental:request-quotation-for-flight;
  :having [:is rental:client; :boundTo "Mail"];
  :having [:is rental:flightnr; :boundTo "Flight"]]].
ex:get-client :is rental:get-client;
:having [:is rental:client; :boundTo "Mail"];
:having [:is rental:client-name; :boundTo "Client"];
:having [:is rental:favorite-class; :boundTo "Class"];
:having [:is rental:max-price; :boundTo "Max-Price"].
ex:get-available-cars :is rental:get-available-cars;
:having [:is rental:office; :boundTo "To"];
:having [:is rental:date; :boundTo "Date"];
:having [:is rental:car; :boundTo "Car"];
:having [:is rental:car-class; :boundTo "Class"];
:having [:is rental:price; :boundTo "Price"].
ex:flight-info :is travel:flight-info;
:having [:is travel:flightnr; :boundTo "Flight"];
:having [:is travel:date; :boundTo "Date"];
:having [:is travel:origin; :boundTo "From"];
:having [:is travel:destination; :boundTo "To"].
  
```

Upon evaluation by an appropriate Engine, a Construction returns a set of results (given a set of input substitutions). Each returned result contains an optional literal value and a set of output

substitutions (that must be joined to one or more of the input substitutions). The returned literal values can only be used as argument values of an algebraic operator, provided the `Construction` is actually a `sub-Construction` of a `Formula` (abstract or “disguised” as an `Opaque`), unless they are `boundTo` some variable (thus extending the output substitutions).

A `Construction` may further restrict its results with a set of `SolutionConstraints` (cf. figure 15.7), each defining a *constraint domain* used to constrain the *constraint range*. The constraint range may be either the actual set of values bound to a particular variable (identified by its `name` in a `VariableDeclaration`⁷), or the set of values returned by a `TestExpression` involving several variables (where an `Expression` is understood as a `Construction` that does not contain any form of `RulePackage`, cf. figures 15.4 and 15.7).

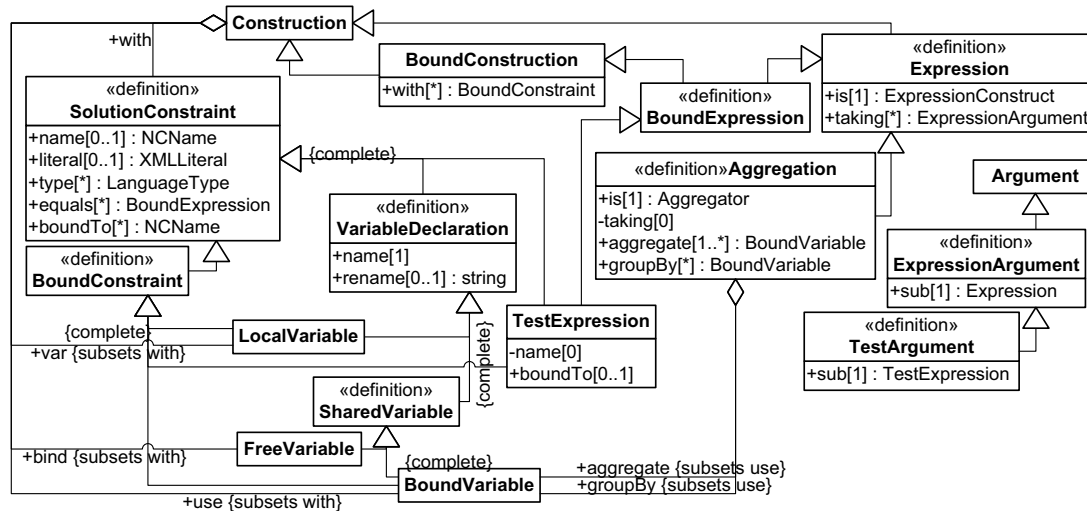


Figure 15.7: Constraints, Variables and Expressions

Any `Construction` defines a scope where `LocalVariables` may be declared. The communication between a `sub-Construction` and its ancestor/sibling `Constructions` is achieved through the `SharedVariables` it either uses or binds. A `Construction` cannot be evaluated unless all the variables it uses are actually bound to specific values.

If a `Construction` is `boundTo` a variable; that variable, if not explicitly declared, is implicitly declared as a `FreeVariable`, i.e. one that the `Construction` may bind. If a `Construction` having a `Parameter` `boundTo` a variable does not explicitly declare it, the variable is implicitly declared a `SharedVariable` according to the nature of the `Parameter`: a `FreeVariable` if it is a `LogicalParameter` that the `Construction` may bind, or a `BoundVariable` if it is an `InputParameter` that the `Construction` must use. A `Construction` implicitly inherits all the `SharedVariables` (implicit or explicit) of the `sub-Constructions` of its `Arguments`. All variables in a `BoundConstruction` must be explicitly declared with `BoundConstraints`, thus excluding the declaration of `FreeVariables`. Namely a `BoundConstruction` cannot bind `SharedVariables`.

A `SolutionConstraint` states that the intersection between the constraint domain and the constraint range (the actual set of values of a variable, for a `VariableDeclaration`, or the actual set of values returned by a `TestExpression`) must not be empty. If it does not define a constraint domain it simply states that the constraint range must not be empty. The constraint domain of a `SolutionConstraint` is defined by intersecting several sub-domains, namely a domain with a single `literal` value, some domains implicitly defined by given `types`, some domains formed by the values `boundTo` given variables, or even some domains where each one `equals` the actual set of values returned by a `BoundExpression`: a `BoundConstruction` that is an `Expression`. Operationally,

⁷A `VariableDeclaration` may additionally `rename` its variable, in order to support different naming conventions used by different `Opaque` languages.

a `SolutionConstraint` can only be checked when its constraint range is known (i.e. the variable becomes bound or is possible to evaluate the test expression). If it can be checked then, for efficiency, it must invalidate as much solutions as possible, as soon as possible (i.e. the constraint range must be checked against each of the known constraint sub-domains; this check must not be postponed until the full constraint domain is known).

Example 53 *Below we illustrate the usage of `SolutionConstraints` as `VariableDeclarations` in several `Constructions` (using the `Languages` defined in examples 50 and 51).*

```

ex:check-price :is eca:opaque;
  :boundTo "Is-True"; :var [:name "Is-True"; :literal "true"];
  # or simply, cf. :TestExpression, :literal "true";
  :use [:name "Price"];
  :use [:name "Max-Price"; :rename "MaxPrice"];
  :having [:is eca:lang; :literal "http://www.w3.org/XPath"];
  :having [:is eca:literal; :literal "$Price <= $MaxPrice"].
ex:send-quotation :is mail:send;
  :var [:name "Text"; :equals ex:quotation-message];
  :having [:is mail:from; :boundTo "Rental-Mail"];
  :having [:is mail:to; :boundTo "Mail"];
  :having [:is mail:subject; :literal "Car Rental Quotation"];
  :having [:is mail:body; :boundTo "Text"].
ex:quotation-message :is text:replace;
  :use [:name "Quotation" ];
  :var [:name "Priced-Cars";
    :equals [:is text:join;
      :aggregate [:name "Car"], [:name "Price"];
      :having [:is text:template; :literal "|Car|=|Price|"];
      :having [:is text:separator; :literal ", "];]];
  :aggregate
    [:name "Client"], [:name "Priced-Cars"],
    [:name "Flight"], [:name "Date"], [:name "To"];
  :having [:is text:template; :boundTo "Quotation"].

```

To ensure full transparency of an `AbstractConstruction`, explicit declaration of `SharedVariables` in these `Constructions` is restricted and only allowed for those variables already implicitly declared as shared: `boundTo` the `Construction` itself or to one of its `Parameters`, or declared as shared in an `Argument`. The only form of `AbstractConstruction` that does not have this restriction is a specialized form of a `Term` called an `Aggregation` (that is a purely functional `Aggregator`), which `aggregates` some variables grouped by some other variables it `uses` (`groupBy` is intended as a synonym of `use`).

Both the input and output substitutions of any `Construction` must include only the variables declared in it. Furthermore, input substitutions must be distinct and output substitutions of an `Aggregation` cannot contain the `aggregated` variables.

Given the definitions contained in the diagrams of the previous figures, an `ECARule` is finally defined as depicted in figure 15.8.

Example 54 *To further illustrate this definition a concrete `ECARule` is presented here, using the `Languages` and `Constructions` included in the previous examples.*

```

ex:rule :is eca:rule;
  :on [:is eca:event; :sub ex:event];
  :if [:is eca:condition; :sub [
    :is eca:collect-and-test;
    :var [:name "Class"];
    :taking [:is eca:collect;
      :sub [:is eca:query-collect;
        :taking [:is eca:query; :sub ex:get-client];
        :taking [:is eca:collect; :sub [:is eca:query-collect;
          :taking [:is eca:query; :sub ex:get-available-cars];
          :taking [:is eca:collect; :sub [:is eca:query-collect;
            :taking [:is eca:query; :sub ex:flight-info];
            :taking [:is eca:collect; :sub [:is eca:true]]]]]]];
        :taking [:is eca:test; :sub ex:check-price]]];
    :do [:is eca:action; :sub ex:send-quotation].

```

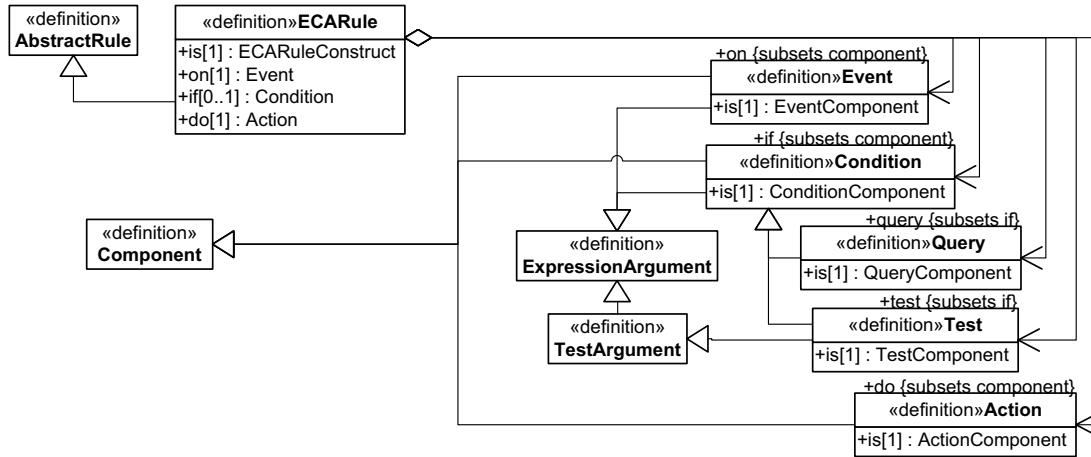


Figure 15.8: ECA Rules

15.4 From ECA-ML to r^3 DF

To complement the definition of the r^3 ontology just presented, we now informally describe how an RDF model of an ECARule, like the one included in example 54, may be automatically obtained from an XML serialization of an ECA rule. Of course, we do not simply mean RDF/XML serializations. Rather, our aim here is to translate markups that adhere to the ECA-ML proposal [52] (and to its general guidelines for component languages and logical variables). Such an ECA-ML compliant serialization of a rule is shown in example 55.

Example 55 *The ECARule of example 54 using an ECA-ML serialization:*

```

<eca:rule>
  <eca:event>
    <snoop:sequence>
      <snoop:first>
        <travel:booking-place client="$Mail" flightnr="$Flight"/>
      </snoop:first>
      <snoop:other>
        <rental:request-quotation-for-flight
          client="$Mail" flightnr="$Flight"/>
      </snoop:other>
    </snoop:sequence>
  </eca:event>
  <eca:query>
    <rental:get-client client="$Mail" client-name="$Client"
      favorite-class="$Class" max-price="$Max-Price"/>
  </eca:query><eca:query>
    <travel:flight-info flightnr="$Flight" date="$Date"
      origin="$From" destination="$To"/>
  </eca:query><eca:query>
    <rental:get-available-cars office="$To" date="$Date"
      car="$Car" car-class="$Class" price="$Price"/>
  </eca:query>
  <eca:test eca:variable="Is-True">
    <eca:variable name="Is-True">true</eca:variable>
    <eca:input-variable name="Price"/>
    <eca:input-variable name="Max-Price" use="MaxPrice"/>
    <eca:opaque lang="http://www.w3.org/XPath">
      $Price <= $MaxPrice
    </eca:opaque>
  </eca:test>
  <eca:action>
    <eca:variable name="Text"><eca:query>
      <eca:variable name="Priced-Cars"><eca:query>
        <eca:input-variable name="Car"/>
        <eca:input-variable name="Price"/>
        <text:join template="|Car|=|Price|" separator=", "/>
      </eca:query></eca:variable>
      <text:replace template="$Quotation">

```

```

    <eca:input-variable name="Client"/>
    <eca:input-variable name="Priced-Cars"/>
    <eca:input-variable name="Flight"/>
    <eca:input-variable name="Date"/>
    <eca:input-variable name="To"/>
  </text:replace>
</eca:query></eca:variable>
<mail:send from="$Rental-Mail" to="$Mail"
  subject="Car Rental Quotation" body="$Text"/>
</eca:action>
</eca:rule>

```

Given the definition of the `eca:ml Language` provided in example 51, we immediately identify the root element and its children, in the given ECA-ML example, as `LanguageElements` in the `eca:ml Language`: an `Operator`, `eca:rule`, and several `OperatorArguments`. As such, it should be possible to build a `Construction` out of it. Two of the children (`eca:event` and `eca:action`) are in fact arguments of `eca:rule`, but there are several `eca:query` and an `eca:test` that cannot be “parsed”. Furthermore, a third argument is missing: `eca:condition`. If we look closely at the definition of `eca:condition`, we see that it is a `CompositeArgument` which must be built using the particular `Operator` `eca:collect-and-test` that takes as `OperatorArguments` an `eca:collect` and an `eca:test`. We do have an `eca:test` available, so our problem is now narrowed to building an `eca:collect` out of several `eca:query`. Another look at the `eca:ml` definition tells us that `eca:collect` is also a `CompositeArgument` (a `CollectionArgument`), and that its `composer` (`collector`) takes an `eca:query`. So, we just pick one from our queries. The actual order could be relevant, but in fact that is not the case here since a partial order may be inferred, given the added semantic value provided by `Language` definitions, viz. `InputParameters`. Once we have picked a query, we are back to our problem of building an `eca:collect`, only now with one less `eca:query`. If we continue applying this process we will be reduced to the problem of building an `eca:collect`, which becomes possible (since it is an `OptionalArgument`) by simply grabbing its `nil` constant, `eca:true`. The result is an `ECARule` with the structure presented in example 54 that includes several `RuleComponents`: an `EventComponent`, several `QueryComponents`, a `TestComponent`, and an `ActionComponent`. What is still missing is the sub-`Construction` for each of those `RuleComponents`.

The same process may be easily applied to obtain the `eca:event EventComponent`, which leads us to the problem of “parsing” the `Terms` that relate to the `Functors` included in it (cf. example 50, for `travel:booking-place` and `rental:request-quotation-for-flight`). As XML elements, these `Functors`, include attributes without a namespace, which are meaningless in RDF, but nevertheless quite common in current XML markups. To avoid this problem we assume that any attribute without a namespace is to be replaced by an (equally valued) attribute with the same local name but in the namespace of its parent element. Given that, all those attributes may be identified as `LanguageElements` (viz. `ConstructParameters`). Furthermore, those `Parameters` are `boundTo` variables according to the ECA-ML guidelines, prefixing variables with ‘\$’. Given all this, not only the `eca:event EventComponent`, but also all the `eca:query QueryComponents`, may be promptly translated into RDF models containing the sub-`Constructions` of example 52.

The `eca:test TestComponent` and the `eca:action ActionComponent` both use ECA-ML nodes that are specific to logical variables (i.e. `eca:variable` and `eca:input-variable`). Although these nodes belong to the `eca:` namespace they are not defined as `LanguageElements` in `eca:ml`. Dealing with logical variables is orthogonal to language heterogeneity, in fact, logical variables provide the homogeneous “glue” that supports language heterogeneity and are not themselves part of it. Actually, this calls for a markup specific to logical variables, which is out of scope in this paper. So, for now, we stick with the ECA-ML proposal, but we do look at these XML nodes as homogeneous. As such, we map an `eca:variable` XML attribute to a `Construction` `boundTo` a variable; and the `eca:variable` (and `eca:input-variable`) XML elements to `FreeVariables` (and `BoundVariables`). Each `SharedVariable` resulting from the latter mapping (notice that ECA-ML does not include `LocalVariables`): has a `name`; is optionally `renamed` (according to its `use` attribute); and its content (if not empty) defines its value `literally` or through the (`eca:query`) sub-`Construction` that equals it.

Given this (partial) homogeneous mapping for the ECA-ML nodes related to logical variables, both the `TestComponent` and `ActionComponent` sub-`Constructions` are built as the previous `RuleComponents` (leading to the RDF model included in example 53), and our `ECARule` is finally complete and should be recognized as such by an OWL-DL reasoner (given the r^3 ontology and the appropriate `Languages`).

15.5 Related Work

To the best of our knowledge, to the present, there are only two ontology proposals for describing rules: WRL [5] and SBVR [60]. The definition of the former includes a mapping to OWL-DL, but does not include any form of reactive rules. The latter, although not excluding reactive rules (some illustrative examples are even present in the specification), explicitly chooses not to address its specificities, conditioning such specificities to the results of PRR⁸ [59]. Moreover, the latter is not formalized in OWL-DL terms, and, it is worth mentioning, that it is targeted to describe business rules in general with an emphasis on human understanding [71] (which may hinder machine computability) and it is the only of the two that actually addresses the issue of language heterogeneity (introducing the concept of business vocabularies, as a form of controlled natural language).

15.6 Final remarks

The foundational and structural nature of the r^3 ontology seems appropriate for describing rules in general (not withstanding proper extension/generalization whenever needed). For instance, it would be quite simple to provide a definition (particularly, at an algebraic level) of the Prolog language and write heterogenous Prolog rules using the r^3 vocabulary. Although our present work is not focused on rules in general, reactive rules very often use other types of rules, namely to model relationships of a more deductive nature (e.g. higher level constructions based on lower level ones, or events “derived” from actions). As such, it would be quite interesting (both from a general and reactive perspective) to collaborate with other communities towards achieving an r^2 ontology (viz. a foundational ontology for Resourceful Rules). Also, considering the extension of the r^3 ontology to higher abstraction levels (e.g. rule languages, domain/application languages, event languages, event algebras, process algebras) is to be pursued.

Regarding ECA languages, the r^3 ontology implicitly extends and generalizes the previously proposed general ECA framework of [1, 52, 54] (e.g. solution constraints and rule constructions) and calls for a revision of the ECA-ML markup [52] and of the general framework. Any revision of the ECA-ML markup should: consider an homogenous markup for logical variables, clearly identify syntactic sugar constructions (e.g. using XPath expressions to initialize variables), and be based upon the formal specification of the (revised) framework. Given the r^3 ontology (as an adequate abstract syntax for the framework), such a (fully) formal specification of the framework is now possible to achieve without resorting to general markup guidelines and principles.

Once the revised general framework is formally specified, the r^3 prototype is to be upgraded accordingly, thus supporting the final version of the ontology.

Last, but not least, it is worth mentioning that r^3 has been used to model several component languages. Namely, the languages Xcerpt [67] for queries, XChange [20] for events and actions of updates of XML data, Prova [44], XQuery/XPath and HTTP. This work actually lead to the inclusion of these `Languages` (as fully functional `Engines`) in the current version of the r^3 prototype [2]. Details about the current definition and `implementation` of these `Languages` may be found in technical reports available online [76]. This experimental work on defining `Languages` has brought relevant contributions for the r^3 ontology and must continue to be pursued and extended to other

⁸The PRR specification is not yet available. Quite recently a preliminary draft has been disclosed to RIF and is available through <http://lists.w3.org/Archives/Public/public-rif-wg/2006Jul/0023.html>. These preliminary results should be considered, asap, and related to r^3 work.

languages as, for example, the algebraic language for actions described in [9] or the event “algebra” used by [77].

Bibliography

- [1] José Júlio Alferes, Ricardo Amador, Erik Behrends, Michael Eckert, Oliver Fritzen, Wolfgang May, Paula-Lavinia Pătrânjan, and Franz Schenk. A first prototype on evolution and behaviour at the XML-Level. deliverable I5-D5, Reverse project, 2006.
- [2] José Júlio Alferes, Ricardo Amador, Erik Behrends, Michael Eckert, Oliver Fritzen, Wolfgang May, Paula-Lavinia Pătrânjan, and Franz Schenk. Completion of the prototype scenario. deliverable I5-D7, Reverse project, 2007.
- [3] José Júlio Alferes, Ricardo Amador, and Wolfgang May. A general language for evolution and reactivity in the semantic web. In *Principles and Practice of Semantic Web Reasoning (PPSWR)*, number 3703 in LNCS, pages 101–115. Springer, 2005.
- [4] José Júlio Alferes, Mikael Berndtsson, François Bry, Michael Eckert, Wolfgang May, Paula Lavinia Pătrânjan, and Michael Schröder. Use cases in evolution and reactivity. Technical Report I5-D2, REVERSE EU FP6 NoE, 2005. Available at <http://www.reverse.net>.
- [5] Jürgen Angele, Harold Boley, Jos de Bruijn, Dieter Fensel, Pascal Hitzler, Michael Kifer, Reto Krummenacher, Holger Lausen, Axel Polleres, and Rudi Studer. Web Rule Language (WRL). <http://www.w3.org/Submission/2005/SUBM-WRL-20050909/>.
- [6] James Bailey, Alexandra Poulovassilis, and Peter T. Wood. An Event-Condition-Action Language for XML. In *Int. WWW Conference*, 2002.
- [7] Erik Behrends, Oliver Fritzen, Tobias Knabke, Wolfgang May, and Franz Schenk. Rule-Based Active Domain Brokering for the Semantic Web. In *Web Reasoning and Rule Systems (RR)*, 2007.
- [8] Erik Behrends, Oliver Fritzen, Wolfgang May, and Franz Schenk. Combining ECA Rules with Process Algebras for the Semantic Web. In *Rule Markup Languages (RuleML)*, pages 29–38. IEEE Press, 2006.
- [9] Erik Behrends, Oliver Fritzen, Wolfgang May, and Franz Schenk. Combining ECA Rules with Process Algebras for the Semantic Web. In *Proceedings of Second International Conference on Rules and Rule Markup Languages for the Semantic Web, Athens, Georgia, USA (10th–11th November 2006)*, 2006.
- [10] Erik Behrends, Oliver Fritzen, Wolfgang May, Franz Schenk, and Daniel Schubert. Embedding Event Algebras and Process Algebras in a Framework for ECA Rules for the Semantic Web. *Fundamenta Informatica*, 2007.
- [11] Erik Behrends, Oliver Fritzen, Wolfgang May, and Daniel Schubert. An ECA Engine for Deploying Heterogeneous Component Languages in the Semantic Web. In *Web Reactivity (EDBT Workshop)*, number 4254 in LNCS, pages 887–898. Springer, 2006.
- [12] J. A. Bergstra and J. W. Klop. Algebra of communicating processes with abstraction. *Theoretical Computer Science*, 1(37):77–121, 1985.

- [13] Mikael Berndtsson and Marco Seiriö. Design and Implementation of an ECA Rule Markup Language. In *Rule Markup Languages (RuleML)*, number 3791 in LNCS, pages 98–112. Springer, 2005.
- [14] Eike Best, Raymond Devillers, and Maciej Koutny. The Box Algebra = Petri Nets + Process Expressions. *Information and Computation*, 178:44–100, 2002.
- [15] Harold Boley, Mike Dean, Benjamin Grosf, Michael Sintek, Bruce Spencer, Said Tabet, and Gerd Wagner. FOL RuleML: The First-Order Logic Web Language. <http://www.ruleml.org/fol/>.
- [16] Angela Bonifati, Daniele Braga, Alessandro Campi, and Stefano Ceri. Active XQuery. In *Intl. Conference on Data Engineering (ICDE)*, pages 403–418, San Jose, California, 2002.
- [17] Angela Bonifati, Stefano Ceri, and Stefano Paraboschi. Pushing Reactive Services to XML Repositories Using Active Rules. In *World Wide Web Conf. (WWW 2001)*, pages 633–641, 2001.
- [18] A. J. Bonner and M. Kifer. An overview of transaction logic. *Theoretical Computer Science*, 133(2):205–265, 1994.
- [19] Anthony J. Bonner and Michael Kifer. Transaction logic programming. In David S. Warren, editor, *Intl. Conference on Logic Programming (ICLP)*. MIT Press, 1993.
- [20] F. Bry and P.-L. Pătrânjan. Reactivity on the Web: Paradigms and Applications of the Language XChange. In *20th ACM Symp. Applied Computing*. ACM, 2005.
- [21] François Bry, Michael Eckert, and Paula-Lavinia Pătrânjan. Reactivity on the Web: Paradigms and applications of the language XChange. *J. of Web Engineering*, 5(1):3–24, 2006.
- [22] François Bry and Paula-Lavinia Pătrânjan. Reactivity on the Web: Paradigms and Applications of the Language XChange. In *20th ACM Symp. Applied Computing*. ACM, 2005.
- [23] François Bry and Sebastian Schaffert. Towards a declarative query and transformation language for XML and semistructured data: Simulation unification. In *Intl. Conf. on Logic Programming (ICLP)*, number 2401 in LNCS, pages 255–270. Springer, 2002.
- [24] S. Chakravarthy, V. Krishnaprasad, E. Anwar, and S.-K. Kim. Composite events for active databases: Semantics, contexts and detection. In *Proceedings of the 20th VLDB*, pages 606–617, 1994.
- [25] Alin Deutsch, Mary Fernandez, Daniela Florescu, Alon Levy, and Dan Suciu. XML-QL: A Query Language for XML. In *8th. WWW Conference*. W3C, 1999. World Wide Web Consortium Technical Report, NOTE-xml-ql-19980819, www.w3.org/TR/NOTE-xml-ql.
- [26] Description Logic Implementation Group (DIG). <http://dl.kr.org/dig/>.
- [27] Document object model (DOM). <http://www.w3.org/DOM/>, 1998.
- [28] Andrew Eisenberg and Jim Melton. SQL/XML and the SQLX informal group of companies. *SIGMOD Record*, 30(3):105–108, 2001. See also www.sqlx.org.
- [29] eXcelon Corp. XML application development using excelon, 2001. <http://www.exceloncorp.com/>.
- [30] eXist: an Open Source Native XML Database. <http://exist-db.org/>.
- [31] Florid homepage. <http://www.informatik.uni-freiburg.de/~dbis/florid/>, 1998.

- [32] Jürgen Frohn, Rainer Himmeröder, Paul-Th. Kandzia, Georg Lausen, and Christian Schlep-phorst. FLORID: A prototype for F-Logic. In *Intl. Conf. on Data Engineering (ICDE)*, 1997.
- [33] Carsten Gottschlich. To be extended. Master Thesis, Univ. Göttingen, 2006.
- [34] Eric N. Hanson and Samir Khosla. An introduction to the triggerman asynchronous trig-ger processor. In *Rules in Database Systems (RIDS)*, number 1312 in LNCS, pages 51–66. Springer, 1997.
- [35] D. Harel. *First-Order Dynamic Logic*. Number 68 in LNCS. Springer, 1979.
- [36] C.A.R. Hoare. *Communicating Sequential Processes*. Prentice Hall, 1985.
- [37] R. Janicki and P. E. Lauer. *Specification and Analysis of Concurrent Systems – the COSY Approach*. EATCS Monographs on Theoretical Computer Science. Springer, 1992.
- [38] Jena: A java framework for semantic web applications. <http://jena.sourceforge.net>.
- [39] Heiko Kattenstroth. Combining description logic and f-logic reasoning. Diploma Thesis, Univ. Göttingen, 2007.
- [40] Michael Kay. SAXON: an XSLT processor. <http://saxon.sourceforge.net/>.
- [41] Michael Kifer and Georg Lausen. F-Logic: A higher-order language for reasoning about ob-jects, inheritance and scheme. In *ACM Intl. Conference on Management of Data (SIGMOD)*, pages 134–146, 1989.
- [42] Michael Kifer, Georg Lausen, and James Wu. Logical foundations of object-oriented and frame-based languages. *Journal of the ACM*, 42(4):741–843, 1995.
- [43] Tobias Knabke. Development of a domain broker. Master’s Thesis, Univ. Göttingen, 2006.
- [44] Alexander Kozlenkov and Michael Schroeder. PROVA: Rule-based Java-Scripting for a Bioin-formatics Semantic Web. In *Proceedings of First International Workshop, DILS 2004*, LNCS, 2004.
- [45] Patrick Lehti. Design and Implementation of a Data Manipulation Processor for an XML Query Language (diploma thesis), August 2001. Technische Universität Darmstadt.
- [46] Bertram Ludäscher, Rainer Himmeröder, Georg Lausen, Wolfgang May, and Christian Schlep-phorst. Managing semistructured data with florid: A deductive object-oriented perspective. *Information Systems*, 23(8):589–612, 1998.
- [47] Wolfgang May. LoPiX: A system for XML data integration and manipulation. In *Intl. Conf. on Very Large Data Bases (VLDB), Demonstration Track*, pages 707–708, 2001.
- [48] Wolfgang May. The LOpix system, 2001. <http://dbis.informatik.uni-goettingen.de/lopix/>.
- [49] Wolfgang May. A rule-based querying and updating language for XML. In *Workshop on Databases and Programming Languages (DBPL 2001)*, number 2397 in LNCS, pages 165–181, 2001.
- [50] Wolfgang May. XPath-Logic and XPathLog: A logic-programming style XML data manipu-lation language. *Theory and Practice of Logic Programming*, 4(3):239–287, 2004.
- [51] Wolfgang May, José Júlio Alferes, and Ricardo Amador. Active rules in the semantic web: Dealing with language heterogeneity. In *Rule Markup Languages (RuleML)*, number 3791 in LNCS, pages 30–44. Springer, 2005.

- [52] Wolfgang May, José Júlio Alferes, and Ricardo Amador. Active rules in the Semantic Web: Dealing with language heterogeneity. In *International Conference on Rules and Rule Markup Languages for the Semantic Web (RuleML)*, volume 3791 of *LNCS*, pages 30–44. Springer, 2005.
- [53] Wolfgang May, José Júlio Alferes, and Ricardo Amador. An ontology- and resources-based approach to evolution and reactivity in the semantic web. In *Ontologies, Databases and Semantics (ODBASE)*, number 3761 in *LNCS*, pages 1553–1570. Springer, 2005.
- [54] Wolfgang May, José Júlio Alferes, and Ricardo Amador. An ontology- and resources-based approach to evolution and reactivity in the Semantic Web. In *Ontologies, DataBases, and Applications of Semantics (ODBASE)*, volume 3761 of *LNCS*, pages 1553–1570. Springer, 2005.
- [55] Wolfgang May, Franz Schenk, and Elke von Lienen. Extending an owl web node with reactive behavior. In *Principles and Practice of Semantic Web Reasoning (PPSWR)*, number 4187 in *LNCS*, pages 134–148. Springer, 2006.
- [56] R. Milner. Calculi for synchrony and asynchrony. *Theoretical Computer Science*, pages 267–310, 1983.
- [57] R. Milner. *Communication and Concurrency*. Prentice Hall, 1989.
- [58] R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes. *Information and Computation*, 1(100):1–77, 1992.
- [59] Object Management Group. *Production Rule Representation (PRR)*. OMG, 2003. <http://www.omg.org/cgi-bin/doc?br/2003-9-3>.
- [60] Object Management Group. *Semantics of Business Vocabulary and Business Rules (SBVR)*. OMG, 2006. <http://www.omg.org/cgi-bin/doc?dtc/2006-03-02>.
- [61] Y. Papakonstantinou, H. Garcia-Molina, and J. Widom. Object exchange across heterogeneous information sources. In *Intl. Conference on Data Engineering (ICDE)*, pages 251–260, 1995.
- [62] George Papamarkos, Alexandra Poulouvassilis, and Peter T. Wood. Event-Condition-Action Rule Languages for the Semantic Web. In *Workshop on Semantic Web and Databases (SWDB'03)*, 2003.
- [63] George Papamarkos, Alexandra Poulouvassilis, and Peter T. Wood. RDFTL: An Event-Condition-Action Rule Language for RDF. In *Hellenic Data Management Symposium (HDMS'04)*, 2004.
- [64] Pellet: An OWL DL reasoner. Maryland Information and Network Dynamics Lab, <http://www.mindswap.org/2003/pellet>.
- [65] Best Practice Recipes for Publishing RDF Vocabularies. <http://www.w3.org/TR/2006/WD-swbp-vocab-pub-20060314/>, 2006.
- [66] Rule markup language (ruleml). <http://www.ruleml.org/>.
- [67] S. Schaffert and F. Bry. A practical introduction to Xcerpt. In *Int. Conf. Extreme Markup Languages*, 2004.
- [68] Daniel Schubert. Development of a prototypical event-condition-action engine for the semantic web. Bachelor Thesis, Univ. Göttingen, 2005.
- [69] Software AG. Tamino – an internet database system, 2001. <http://www.tamino.com/>.

- [70] Sebastian Spautz. Automatenbasierte Detektion von Composite Events gemäss SNOOP in XML-Umgebungen. Diplomarbeit, TU Clausthal (in german), 2006.
- [71] Silvie Spreeuwenberg and Rick Gerrits. Business Rules in the Semantic Web, Are There Any or Are They Different? In *Proceedings of Summer School Reasoning Web 2006, Lisbon, Portugal (4th–8th September 2006)*, volume 4126 of *LNCS*, pages 152–163. REWERSE, 2006.
- [72] Igor Tatarinov, Zachary G. Ives, Alon Halevy, and Daniel Weld. Updating XML. In *ACM Intl. Conference on Management of Data (SIGMOD)*, pages 133–154, 2001.
- [73] Elke von Lienen. Entwicklung eines RDF-Web-Services mit Trigger-Funktionalität. Diplomarbeit, TU Clausthal (in german), 2006.
- [74] Thomas Westphal. Translation of semantic level actions in an rdf web service. Bachelor Thesis, Univ. Göttingen, 2007.
- [75] XML Syntax for XQuery 1.0 (XQueryX). <http://www.w3.org/TR/xqueryx>, 2001.
- [76] Resourceful Reactive Rules (r^3). <http://rewerse.net/I5/r3/>. Centro de Inteligência Artificial - CENTRIA, Universidade Nova de Lisboa.
- [77] ruleCore: The Active Rule Engine. <http://www.rulecore.com/>.

Acknowledgements

This research has been co-funded by the European Commission and by the Swiss Federal Office for Education and Science within the 6th Framework Programme project REVERSE number 506779 (cf. <http://reverse.net>).