A General Language for Evolution and Reactivity in the Semantic Web

José Júlio Alferes[1], Ricardo Amador[1], and Wolfgang May[2]

[1] Centro de Inteligência Artificial - CENTRIA, Universidade Nova de Lisboa
[2] Institut für Informatik, Universität Göttingen

Abstract. In this paper we define the basic concepts for a general language for evolution and reactivity in the Semantic Web. We do this by exposing an UML model that specifies an ontology for the language. The proposed language is based on Event-Condition-Action rules, where different languages for events (including languages for composite events), for conditions (queries) and actions (including complex actions) may be composed, this way catering for language heterogeneity (besides heterogeneity on the data-model) that we think is essential for dealing with evolution and reactivity in the Semantic Web.

1 Introduction

The Web and the Semantic Web, as we see it, can be understood as a "living organism" combining autonomously evolving data sources, each of them possibly reacting to events it perceives. The dynamic character of such a Web requires declarative languages and mechanisms for specifying the evolution of the data. This vision of the Web, as well as a state of the art overview of related areas, is described in our previous work [17].

Rather than a Web of data sources, we envisage a Web of Information Systems, where each such system, besides being capable of gathering information (querying, both on persistent data, as well as on volatile data such as occurring events), is capable of updating persistent data, communicating the changes, requesting changes of persistent data in other systems, and being able to react to requests from other systems. As a practical example, consider a set of data (re)sources in the Web of travel agencies, airline companies, train companies, etc. It should be possible to query the resources about timetables, availability of tickets, etc. But in such an evolving Web, it should also be possible for a train company to report on late trains, and travel agencies (and also individual clients) be able to detect such an event and react upon it, by rescheduling travel plans, notifying clients that in turn could have to cancel hotel reservations and book other hotels, or try alternatives to the late trains, etc.

The importance of being able to update the Web has long been acknowledged, and several languages exist (e.g. XUpdate [24], XML-RL [15], XPathLog [16]) for just that. More recently some reactive languages have been proposed, that not only allow for updating Web data as the above ones, but are also capable

F. Fages and S. Soliman (Eds.): PPSWR 2005, LNCS 3703, pp. 101–115, 2005.
© Springer-Verlag Berlin Heidelberg 2005

of dealing-with/reacting-to some forms of events, evaluate conditions, and upon that act by updating data. These are the cases of the XML active rules of [6], of Active XQuery [5], of the Event-Condition-Action (ECA) language for XML defined in [2], and RDFTL [18], which is an ECA reactive language on RDF data. The common aspect of all of these languages is the use of ECA (declarative) rules for specifying reactivity and evolution. Such kind of rules (also known as triggers, active rules, or reactive rules), that have been widely used in other fields (e.g. active databases [19,23]) have the general form **on** *event* **if** *condition* **do** *action*. They are intuitively easy to understand, and provide a well-understood formal semantics: when an event (atomic or composite) occurs, evaluate a condition, and if the condition is satisfied then execute an action (or a sequence of actions, a program, a transaction, or even start a process).

In fact, we agree with the arguments exposed for the definition of the above languages in what regards adopting ECA rules for dealing with evolution and reactivity in the Web (declarativity, modularity, maintainability, etc). But in our opinion, these languages fall short in various aspects, when the goal is aimed at the general view of an evolving Web as described above. Namely, they do not provide for more complex events and actions and, most important, they do not deal with heterogeneity at the level of the language. Autonomous web nodes will use different formalisms for ECA rules, and also different formalism for events, conditions and actions, depending on the requirements of their applications.

In general, actions are more than just simple updates to Web data (be it XML or RDF data). As said above, besides that, actions can be notifications to other resources, update requests of other resources, can be composition of simpler actions (like: do this, and then do that), or even transactions whose ACID properties ensure that either all actions in a transaction are performed, or nothing is done. In our view, a general language should cater for such richer actions. Moreover, events may in general be more than simple atomic events in Web data, as in the above languages. First, there are atomic events other than physical changes in Web data: events may be received messages, or even "happenings" in the global Web, which may require complex event detection mechanisms (e.g (once) any train to St. Wendel is delayed ...). Moreover, as in active databases [10,25], there may be more complex (composite) events. For example, we may want a rule to be triggered when there is a flight cancellation and then the notification of a new reservation whose price is much higher than the previous (e.g. to complain to the airline company). In this respect, there is some preliminary work on composite events in the Web [3], but that only considers composition of events of modification of XML-data in a single document.

The quite recent work on the language XChange [8] already aims at having more complex actions and events for evolution and reactivity on the Web and, in our opinion, is an important contribution in this direction. However, having in mind the requirements we set up for the general evolving Semantic Web, there are still some important aspects, that are not yet dealt with by XChange, namely that of language heterogeneity.

The problem of language heterogeneity will definitely appear when dealing with evolution and reactivity on the Web. This calls for more general languages. In such an open and heterogeneous environment as the Web, it is difficult to assume that there will be *a single* event language, or *a single* way to deal with actions. Our view is that a general language for evolution and reactivity in the Web should allow for the usage of different event languages, different condition languages, and different action languages, considering ontological descriptions and mappings for these languages. Each of these different (sub)languages should adhere to some minimal requirements (e.g. dealing with variables), but it should be as free as possible. The task of the general ECA language is then to combine these various (sub)languages for reacting and performing evolution in the (Semantic) Web. This requirement is far from the goals of XChange, which is based on a concrete language for all the components of ECA.

Moreover, the ECA rules do not only operate on the Semantic Web, but are themselves also part of it. In general, especially if one wants to reason about evolution, ECA rules (and their components) must be communicated between different nodes, and may themselves be subject to being updated. For that, the ECA rules themselves must be represented as data in the (Semantic) Web. This need calls for a (XML) Markup Language of ECA Rules. A markup proposal for active rules can be found already in RuleML [4], but it does not tackle the complexity of events, actions, and the generality of rules, as described here. Moreover, to deal with the requirements of heterogeneity and of reasoning about rules, an ontology of ECA rules and (sub)ontologies for events, conditions and actions, with rules possibly specified in RDF/OWL, is required.

In this paper we define the basic concepts of a general language for evolution and reactivity in the Semantic Web that responds to the requirements just exposed. Rather than presenting an RDF/OWL ontology, in this paper we present a UML 2.0 [14] model. By doing this, we not only consubstantiate the language concepts, but also provide an abstract syntax for it, which is already a step for having a markup (XML) language for general ECA rules. For defining such a markup, it is worth noting that the UML model we present is mappable into XMI [13], this directly providing an XML representation. The modelling of the language starts in Section 2, where the global aspects are spelled out, and the composition between the components is discussed. The common structure of the (sub)languages for the rule components discussed in Section 3. Then, in Section 4, the specific aspects of each of the E, C, and A components is discussed and an illustrative concrete (instantiation) of each of these (sub)languages is given. We end the paper by mentioning ongoing and future work.

2 Global Aspects for a General ECA Language

In order to cope with the Semantic Web heterogeneity, the target of development and definition of languages for (ECA) rules, for events, for conditions and for actions should be a semantic approach, i.e. an approach based on an (extensible) ontology for rules, events, conditions and actions that also allows for reasoning about these concepts.

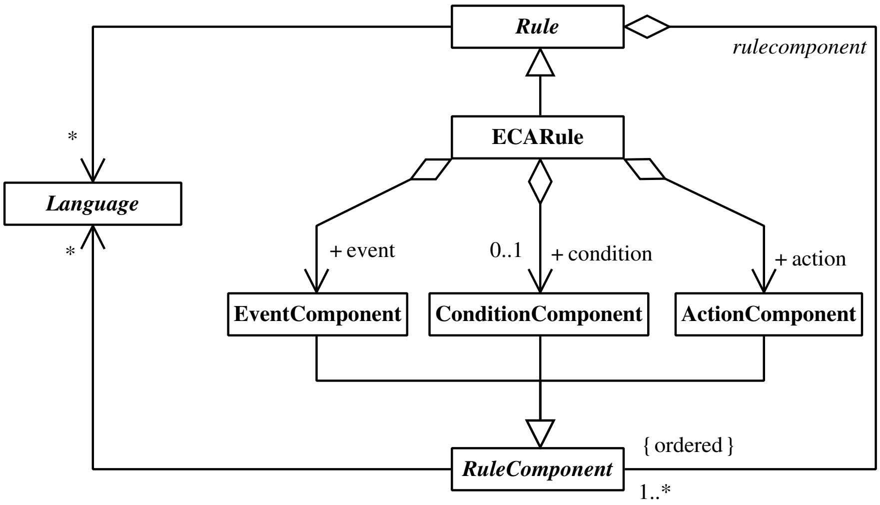

Fig. 1. General UML model for ECA rules

At a quite abstract level a rule is an aggregation of rule components, and an ECA rule can be described by the UML diagram given in Figure 1. As expected, an ECA rule has 3 different components: event, condition and action. The condition is optional, in the sense that it can be omitted, or that languages may allow to integrate the evaluation of the condition with the event component, or with the action component. This model can be readily extended by adding a fourth component (also optional) – the post-condition (another *Condition*) – 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 part. But it may have particular relevance when considered together with cascading reactions and transactional rules, in which case the post-condition allows the declarative specification of restrictions that must apply after the whole transaction, given by the action, is successfully executed. This will be further detailed in Section 4.3 below, when discussing languages for the action part.

Current databases already support active concepts by triggers (e.g., SQL) where the distinction between events, conditions and actions is not necessarily explicit. Such rules can be handled as *opaque* rules of a given language that are understood as a whole by an underlying system. Note that there exist well-defined mappings into the above ECA model.

When defining (ECA) rules, language heterogeneity has to be considered not only at the global rule level, but also at the rule component level. As stated before, several reactive rule languages have already been proposed (e.g. XChange [8], RDFTL [18]), introducing heterogeneity at the rule level. A generic approach for rules in the Semantic Web must be able to cover *all* such explicit language proposals. In most of these proposals there exists a pattern of language reuse, usually a query (sub)language that already exists (e.g. XQuery) is

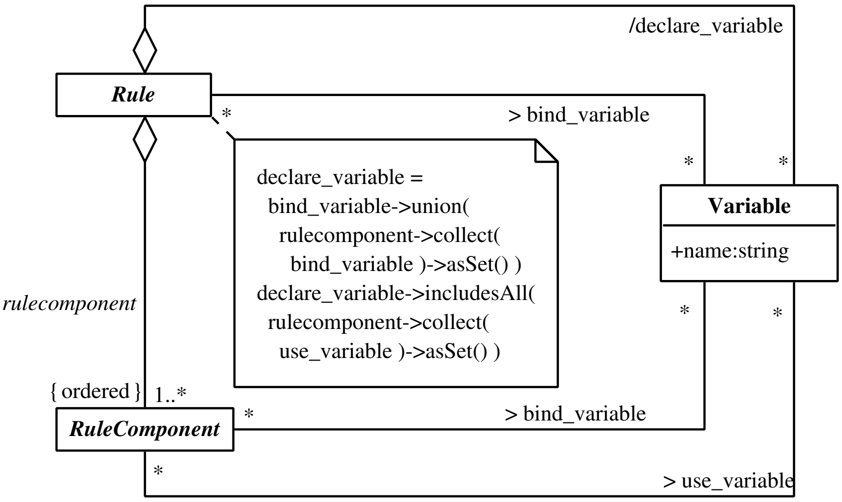

Fig. 2. Rules and Variables

chosen for the condition and either an existing update (sub)language is chosen (e.g. XQuery+Updates [21]) or an extension is built over the query language (e.g. Xcerpt [20]) in order to obtain a new (sub)language (e.g. XChange [8]) for the action component. Finally, an event (sub)language is defined (often based on an existing one from the field of Active Databases, e.g. the SQL3 standard).

A general approach requires a clean distinction between the three components E, C, and A on the ontology level of the rules (as shown in Figure 1). Given this, additional heterogeneity is provided by using and combining different event, condition, and action (sub)languages according to a global ECA schema. Each language is identified by its namespace which then contains markup elements of the specific languages. To achieve language heterogeneity at the rule component level, there must be a precise convention between all languages how the different parts of a rule can exchange information and interact with each other. This is achieved by a set of "bindable" names (logical variables), cf. Figure 2.

A variable must be bound only once to a value; in case that an already bound variable is "bound" again, the values must coincide, i.e., yielding an analogous semantics as in logic programming (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 others is bound by the latter). The OCL constraints in Figure 2 guarantee that the variables of the rule are exactly those bound either in the *Rule* or in a *RuleComponent*. The binding mechanism can be extended with a type system.

The actual handling (and its markup) of variables will be discussed below for a simple case of rules, and later in Sections 3 and 4. We currently recommend to be explicit with declaring variables that are used or bound in a rule component, although in most cases it will be possible to derive this from the markup and the languages' ontologies.

Rules with Opaque Components. At the most basic level, each rule component has a textual specification (again called *opaque*) that is to be understood by some language engine. In this case it is marked up as text content of an `eca:opaque` element that references the language via its `lang` attribute (see Example 1). Additionally, since variable bindings must be communicated and binding/using cannot be derived from the opaque text, this information must be given explicitly. Thus, the `eca:opaque` element must list all variables that are used or bound by it (i.e., whose bindings must be exchanged with the engine), also optionally giving their names (e.g., for embedding JDBC where variables are only named `?1, ?2` etc.). Variables can be bound by opaque parts by (i) matching them – logic programming style, or (ii) assigning the result set of a query to them.

Example 1. Consider an ECA rule expressing the idea that whenever a flight is cancelled, every customer who has a reservation for this flight must be notified, preferably by SMS. Using XML rule markup with opaque rule components (using different languages), this rule could take the following form:

```
<eca:rule xmlns:eca="http://www.eca.org/eca-ml"
    xmlns:datalog="http://www.lp.org/datalog"
    xmlns:xpath="http://www.w3.org/XPath"
    xmlns:pseudocode="http://www.pseudocode-actions.nop">
  <eca:bind-variable name="Reservations">
    http://www.reservations.nop/actual.xml
  </eca:bind-variable>
  <eca:variable name="Flight"/>
  <eca:variable name="Customers"/>
  <eca:event>
    <eca:opaque lang='datalog'>
      <eca:bind-variable name="Flight" use="F" mode="match"/>
      flight_cancelation(F)   <!-- matches literal against event in datalog -->
    </eca:opaque>
  </eca:event>
  <eca:condition>
    <eca:opaque lang='xpath'>
      <eca:use-variable name="Flight" use="$Flight"/>
      <eca:use-variable name="Reservations" use="$Reservations" />
      <eca:bind-variable name="Customers" mode="result-set"/>
      document($Reservations)//flight[@id=$Flight]/reservation/customer
    </eca:opaque>
    <!-- evaluates XPath expression, binds the result to the variable 'Customers'
      and checks if it is not empty -->
  </eca:condition>
  <eca:action>
    <eca:opaque lang='pseudocode'>
      <eca:use-variable name="Customers"/>
      <eca:use-variable name="Flight"/>
      for each C in Customers do
        notify_cancelation(Flight, sms:C)
```

```
          otherwise notify_cancelation(Flight, mail:C)
            otherwise signal_failure(notify_cancelation(Flight, C))
      done
    </eca:opaque>
  </eca:action>
</eca:rule>
```

Thus, the specification of a rule component in its simplest (opaque) form is just some *opaque* text associated with a set of variables that can either be bound or simply used in that component. When this text is given to the respective language engine together with a set of variables, some of them already bound, the engine interprets this text, optionally producing new bindings for some of the yet unbound variables.

3 Common Structure of E, C and A Sublanguages

The level of reasoning that can be performed with the model defined so far is yet restricted. In order to do deeper reasoning, one must go inside of the rule components. For this, instead of a simple text like the opaque specifications above, these specifications may also be given as structured ones. The generic structure of these (sub)languages, independently of whether they are event, condition or action languages, is modelled in Figure 3. Each such language consists of a set of *composers*; actual rules then combine it with a separate language of *atomic* elements (events, literals, actions) that are part of *domain languages*, and in most cases come from application-dependent ontologies. Expressions of the language are then (i) atomic elements, or (ii) composite expressions recursively obtained

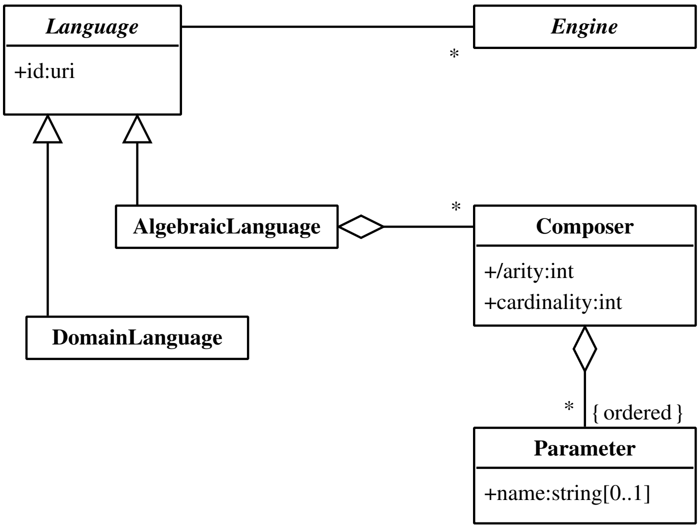

Fig. 3. Language Structure

by applying composers to expressions. Due to their structure, these languages are called *algebraic languages*, e.g. used in *event 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.*

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. A language for *atomic* events could define an event "received_message(M)" for receiving a message. Together with an action language that provides an action for sending a message, one could easily define a negotiation dialog between two systems by means of a set of reactive rules.

As mentioned above, actual rules combine these languages with appropriate languages for atomic elements. Usually, these are provided by *domain languages* (e.g. languages for the domain of travels, or banking, or ...) that are induced by an ontology, and define atomic events, predicates or literals (for conditions), and actions of that specific domain (e.g. events of train schedule changes, actions of reserving tickets, ...). There exist also domain-independent primitive constructs (with arguments), e.g. for general communication, such as the above received_message(M) (where M in turn contains domain-specific content). Note that the markup must also provide the handling of variables; here we propose to borrow from XSLT: use variables by {$*var-name*}, and bind them by <variable name="..." select="..."/> elements; but the actual decision is up to the language designers.

Each of these languages has an associated engine that captures the semantics of the (composers of the) language. The engines provide the (expected) interfaces

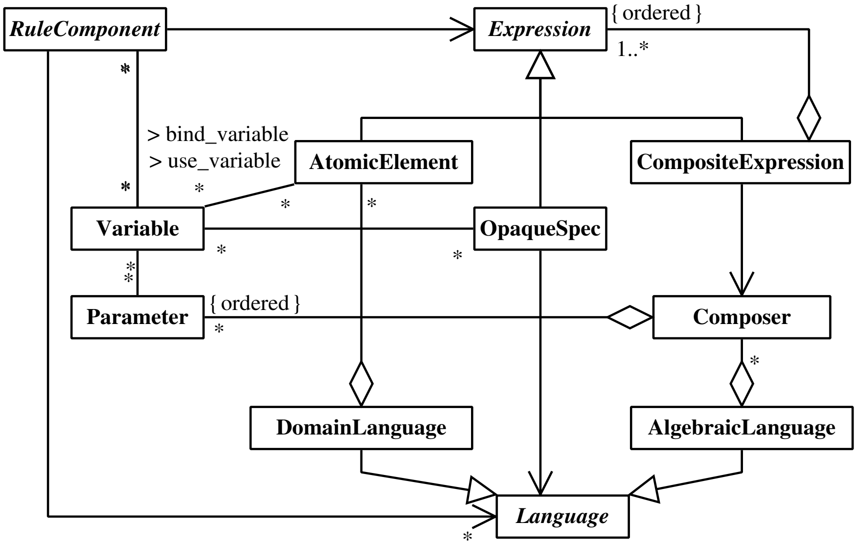

Fig. 4. Model of the Rule Components

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). The issue of transactions is of particular importance (see Section 4.3).

Note that since each subtree of such a specification is a specification of the same kind (e.g., subevents composed into more complex events), it is also possible to nest composers and expressions from different languages. Thus, languages are associated on the *expression* level. In the XML markup, this is done by the namespaces (from which the tree's markup is taken).

Given the additional level of knowledge about the structure of a sublanguage, the modelling of the rule component specifications can be more detailed, as shown in Figure 4. This raises the level of reasoning that may be performed about ECA rules, regardless of the degree of language heterogeneity that may be present.

4 Concrete Languages for Events, Conditions, and Actions

In the previous sections we have defined a general model for ECA rules and their components. In this section we discuss specific issues of the languages of events, of conditions and of actions.

4.1 A Language for Events in the Web

In the context of the Semantic Web, an (atomic) event is in general any detectable occurrence. Events in the Web can be local events, e.g. updates of local data (that can be used for deriving/raising global events), but also incoming messages, and changes in other nodes.

Atomic Events on Web data. On the most basic (physical) level, there are constructs to deal with the detection of changes on local data, be it on XML or RDF data, similar to those found in database triggers. Work on triggers for XQuery has e.g. been described in [5] with *Active XQuery* and in [2], 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 former uses the same syntax and switches as SQL. For modifications of an XML tree, we proposed in [1] the following basic constructs for atomic events of modifications of XML data:

- `ON {DELETE|INSERT|UPDATE} OF` *xsl-pattern*: if a node matching the *xsl-pattern* is deleted/inserted/updated,
- `ON MODIFICATION OF` *xsl-pattern*: if anything in the subtree is modified,
- `ON INSERT INTO` *xsl-pattern*: if a node is inserted (directly) into a node matching the *xsl-pattern*,
- `ON INSERT [IMMEDIATELY] BEFORE|AFTER` *xsl-pattern*: if a node is inserted (immediately) before or after a node matching the *xsl-pattern*.

In all these constructs, *xsl-pattern* is a (typically input) argument. Moreover, these events should make relevant values accessible, e.g., `OLD AS ...` and `NEW AS ...` (like in SQL), both referencing the complete node to which the event happened, additionally `INSERTED AS`, `DELETED AS` referencing the inserted or deleted node. These relevant values are additional arguments of the above constructs, typically (output) to be bound with variables. The implementation of these events in XML repositories is probably to be based on the *DOM Level 2/3 Events* [11].

Regarding RDF data, RDF triples, describing properties/values of a resource, are much more similar to SQL. In contrast to XML, there is no assignment of data with subtrees, which makes it impossible to express "deep" modifications in a simple event. Proposals can be found in [18], and in [1]; in the latter, we considered the following basic constructs:

- `ON {DELETE|INSERT|UPDATE} OF` *property* `[OF` *class*`]`: if a property is removed from/added to/updated of a resource of a given class, then such an event is raised;
- `ON CREATE OF` *class*: it is raised if a new resource of a given class is created;
- `ON NEW CLASS`: is raised if a new class is introduced,
- `ON NEW PROPERTY [OF CLASS` *class*`]`: is raised, if a new property (optionally: to a specified class) is introduced.

Besides the `OLD` and `NEW` values mentioned for XML, these events should consider as arguments (to bind variables) `RESOURCE AS ...` and `PROPERTY AS ...`, referring to the modified resource and the property (as URIs), respective.

Communication events. Besides the above events that react on updates on a given data model level, communication events are raised by messages, independent from the abstraction level of the rule. We propose the following basic construct:

- `ON MESSAGE [OF` *sender*`] [AT` *time*`] [MATCHING` *pattern* `]`

In this construct, the metadata about *sender* and *time* are to be bound to variables upon receipt of the message, as well as the actual *content*. However, one might want to trigger such an event only when a message with a specific sender, time, or content is received. In this case a methodology for testing the content must be specified. This can be done by (regular expression) matching, or by querying the (XML or RDF) content. For the above (opaque) syntax as triggers, we restrict it to matching; a markup version may include more detailed conditions (as illustrated below), where also more elaborate constructs for incoming messages are possible, e.g. with parameters for specifying an ontology describing the language of the message, or along the lines of the FIPA language for communication among agents [12].

Composite events. For dealing with composite events in the context of the ECA rules proposed here, the event languages must define several composers. We propose at least the following composers of events: "E_1 `OR` E_2", "E_1 `AND` E_2"

(in arbitrary order), and “E_1 `AND THEN` E_2 `[AFTER PERIOD` $\{< \mid >\}$ *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. The actual semantics of composers must be defined similarly to that of operators in event algebras in the context of active databases [25]. In it, detection of a composite event means that its “final” atomic subevent is detected. Event algebras contain not only the aforementioned straightforward basic conjunctive, disjunctive and sequential connectives, but also additional operators. Several event algebras have been defined that provide also e.g. “negative events” in the style that “when E_1 happened, and then E_3 but not E_2 in between, then do something”, “aperiodic” and “cumulative” events, e.g., the SNOOP event algebra [9] of the “Sentinel” active database system. A quite rich set of composers for events in the Web is being also considered in the language XChange [8], where exclusions, repetitions, and cardinality are also explored.

Example 2. 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 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 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:event xmlns:xmlsnoop="http://xmlsnoop.nop"
  xmlns:msg="http://www.messages.msg/messages"
  xmlns:mytravel="http://www.trains.tr">
  <eca:bind-variable name="newArrival"/>
  <!-- The 2 variables below are bound on the rule level -->
  <eca:use-variable name="myTravelAgent" use="$myAgent"/>
  <eca:use-variable name="myTrain" use="$myTrain"/>
  <xmlsnoop:or>
    <xmlsnoop:atomic detect="xml-pattern">
      <msg:receive-message sender="$myAgent">
        <content> <delayed train={$myTrain}/> </content>
      </msg:receive-message>
      <xmlsnoop:variable name="newArrival" <!-- borrowed from xsl:variable -->
         select="$event/content/delayed/@arrivalTime" />
    </xmlsnoop:atomic>
    <xmlsnoop:atomic detect="xpath">
      <xmlsnoop:cond test="$event/name()='mytravel:changeTime'"/>
      <xmlsnoop:cond test="$event/@trainId=$myTrain"/>
      <xmlsnoop:variable name="newArrival" select="$event/@newTime"/>
    </xmlsnoop:atomic>
  </xmlsnoop:or>
</eca:event>
```

The composite event is an “or” of two atomic events: the first one is receiving a message (marked-up in XML) with an attribute `sender` which is equal to the value of the variable myAgent, and with a content with a delayed element with an attribute train conciding with that of myTrain. The mechanism used here for testing this matching with the event that occurred is an xml-pattern. If so, the variable newArrival is bound to the value of the attribute arrivalTime of that delayed element. The second one is a domain specific event travel:changeTime (that occurs "somewhere in the Web" and has to be detected by Semantic Web mechanisms. It is implicitly bound to $event). The details are then checked by XPath expressions against $event: If its attribute trainId equals the value of the variable myTrain, then newArrival is bound to the value of the newTime attribute of the event.

4.2 Conditions in ECA Rules for the Web

Conditions in ECA rules basically amount to queries in the (Semantic) Web, that possibly bind rule variables to be then used in the action component. For this purpose, and in case reasoning is not required inside the condition, one can envisage the condition language specification simply as opaque. This way, e.g. XPath, XQuery, RDQL, or Xcerpt can be used in the condition.

In case reasoning about the condition component is desired, an ontology for the query language(s) is needed, that models the basic constructs and composers of the language in the terms described above. XQueryX is an example for an XML markup of a query language. Deeper work in the direction of modelling query languages for the Web already exists, e.g. in [22] where a UML modelling of the language Xcerpt [20] is shown.

4.3 Actions and Transactions

As for events, also (atomic) actions in the Web can be considered at various levels: there can be local actions of updating web data; event raising; external update requests to other nodes; general (local or remote) method calls.

Local update actions can be specified in any appropriate language for changing web data, such as XUpdate [24], XML-RL [15], or XPathLog [16]. Their integration in the ECA framework can be done as just described for conditions, i.e. either as opaque specification, or by providing a proper ontology, based on constructs and composers, specifying those update languages.

Activities of remote nodes can be invoked by sending a message with an update (request) statement. Here a basic construct for sending a message is required, the simplest one being: `SEND MESSAGE` *message* `TO` *recipient*. This message sending can also be used for event-raising actions, in this case making sure that the event raised is then collected by a corresponding `ON MESSAGE` construct. As for events, more elaborate action constructs can be defined. General action constructs that can be defined may be those for (remote) procedure/method calls to Web Services, where the SOAP protocol can be used.

The execution of an action may in general succeed or fail. Considering failure of actions is important e.g. in the case of remote update requests: once the request is issued, it is important to be able to receive feedback on whether the update was actually done, or not. For example, upon request of a flight reservation, it is important to know whether the reservation was accepted or not. The operational semantics of a general language for actions, and a corresponding processor, should thus allow for failure of atomic actions. Moreover, when used with non-deterministic condition languages, the failure of an action should somehow "backtrack" into the condition to check for alternative bindings of variables that may result in successful actions.

Complex actions can be defined by composing atomic actions. This is done by enriching the action language with appropriate composers. The most basic composers for actions are those of (parallel) conjunction of actions (A_1 `AND` A_2) sequential execution of actions (A_1 `;` A_2). Other more elaborate composers can be defined in action languages, such as if-then-else composers (`IF` *test* `THEN` A_1 `ELSE` A_2), while-iterations (`WHILE` *test* `DO` A), and forall-iterations (`FORALL` *variable* `DO` A). Note that some of these complex actions already require the use of a condition language in the action language for evaluating conditions. This idea can be further exploited by introducing an action construct – `TEST CONDITION` *condition* which tests the condition and either fails if the condition is false, or does nothing in case it is true but possibly binding some extra variables). With such a rich action language, similar to Transaction Logic [7], combining condition testing with (trans)actions, the condition of rules can be omitted.

In general, each of the complex actions should be allowed to be specified as a transaction with ACID properties, in particular where either all of the actions are executed, or the whole composite action fails, and no action is performed. This can be done by having a composer `TRANSACTION` *id* A, where *id* is a parameter for storing a unique identifier of the transaction, and A is the (complex) action. Note here how some form of post-condition, in the line of those mentioned in Section 2 may be specified by combining these transactions with the above condition testing. While the transaction composer is easy to understand in case all atomic actions in A consist of local updates, this is not the case when A involves actions like e.g. sending messages, or remote method calls. In fact, in these cases, what should be the meaning of rolling back over such an action? When a message is sent, what does it mean to rollback on sending it? It is our stance that in these cases, compensating actions must be specified, to be executed when rolling back is not possible. This, and a deeper study of transactions in this context (including considering transactions that are not limited to a single rule), is not detailed further here, and is subject of ongoing work.

5 Conclusions

In this paper we describe the basic concepts and a UML modelling of a general ECA-rule framework for the Semantic Web. Moreover, we discuss concrete languages for events, conditions, and actions to be composed in this general language. This framework sets the ground for a general framework for evolution and reactivity on the Web, where heterogeneity of languages is taken into account, and reasoning about rules is possible. The integration of other ECA-based languages in this framework, such as the ones mentioned in the introduction, is a subject of ongoing and future work. In this respect, special attention will be paid to the language XChange, as it is the one which already consider richer events and actions.

Lack of space prevents us from elaborating here on further ongoing work that is being developed by us in the context of the general language. Namely, further detailing the concepts involved in the definition of domain languages, and also the definition of a general architecture for executing the ECA rules are left out. This general architecture also raises the issue of communication strategies regarding events and actions (are events that are raised by actions "pushed" into respective nodes? or do nodes periodically "pull" for events that may have occurred?). Another important issue that is also related with the execution, and that was only briefly addressed here is that of transactions. It is our belief that the issue of transactions on the Web is an important and difficult subject, that will gain increasing importance and interest in a near future. It is in our agenda to continue working in this subject, along the lines exposed above.

Acknowledgements

This research has been funded by the European Commission within the 6th Framework Programme project REWERSE, number 506779.

References

1. J. J. Alferes, M. Berndtsson, F. Bry, M. Eckert, N. Henze, W. May, P. L. Pătrânjan, and M. Schroeder. Use-cases on evolution. Technical Report IST506779/Lisbon/I5-D2/D/PU/a1, REWERSE, 2005.
2. James Bailey, Alexandra Poulovassilis, and Peter T. Wood. An Event-Condition-Action Language for XML. In *Int. WWW Conference*, 2002.
3. M. Bernauer, G. Kappel, and G.Kramler. Composite Events for XML. In *13th Int. Conf. on World Wide Web (WWW 2004)*. ACM, 2004.
4. Harold Boley, Benjamin Grosof, Michael Sintek, Said Tabet, and Gerd Wagner. *RuleML Design.* RuleML Initiative, http://www.ruleml.org/, 2002.
5. Angela Bonifati, Daniele Braga, Alessandro Campi, and Stefano Ceri. Active XQuery. In *"Intl. Conference on Data Engineering (ICDE)"*, pages 403–418, 2002.
6. Angela Bonifati, Stefano Ceri, and Stefano Paraboschi. Pushing Reactive Services to XML Repositories Using Active Rules. In *WWW'01*, pages 633–641, 2001.
7. A. J. Bonner and M. Kifer. An overview of transaction logic. *Theoretical Computer Science*, 133(2):205–265, 1994.
8. 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.
9. S. Chakravarthy, V. Krishnaprasad, E. Anwar, and S.-K. Kim. Composite events for active databases: Semantics, contexts and detection. In *20th VLDB*, 1994.
10. Sharma Chakravarthy and D. Mishra. Snoop: An expressive event specification language for active databases. *Data & Knowledge Engineering*, 14:1–26, 1994.
11. Document object model (DOM). `http://www.w3.org/DOM/`, 1998.
12. Foundation for Intelligent Physical Agents. FIPA ACL Message Structure Specification. Technical Report SC00061G, http://www.fipa.org, Dec. 2002.

13. Object Management Group. *XML Metadata Interchange (XMI) 2.0 Specification*. OMG, 2003. http://www.omg.org/cgi-bin/doc?formal/2003-05-02.
14. Object Management Group. *OMG Unified Modelling Language (UML) 2.0 Superstructure*. OMG, 2004. http://www.omg.org/cgi-bin/doc?ptc/2004-10-02.
15. Mengchi Liu, Li Lu, and Guoren Wang. A Declarative XML-RL Update Language. In *Proc. Int. Conf. on Conceptual Modeling*, pages 506–519. Springer, 2003.
16. Wolfgang May. XPath-Logic and XPathLog: A logic-programming style XML data manipulation language. *Theory and Practice of Logic Programming*, 4(3), 2004.
17. Wolfgang May, José Júlio Alferes, and François Bry. Towards generic query, update, and event languages for the Semantic Web. In *Principles and Practice of Semantic Web Reasoning (PPSWR)*, number 3208 in LNCS, pages 19–33. Springer, 2004.
18. George Papamarkos, Alexandra Poulovassilis, and Peter T. Wood. RDFTL: An Event-Condition-Action Rule Languages for RDF. In *HDMS'04*, 2004.
19. N. W. Paton, editor. *Active Rules in Database Systems*. Monographs in Computer Science. Springer, 1999.
20. S. Schaffert and F. Bry. A practical introduction to Xcerpt. In *Int. Conf. Extreme Markup Languages*, 2004.
21. 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.
22. G. Wagner, C. V. Damásio, and S. Lukichev. First-version rule markup languages. Technical Report IST506779/Eindhoven/I1-D3/D/PU/ab1, REWERSE, 2005.
23. Jennifer Widom and Stefano Ceri, editors. *Active Database Systems: Triggers and Rules for Advanced Database Processing*. Morgan Kaufmann, 1996.
24. XML:DB Initiative, http://xmldb-org.sourceforge.net/. *XUpdate - XML Update Language*, September 2000.
25. D. Zimmer and R. Unland. On the Semantics of Complex Events in Active Database Management Systems. In *15th International Conference on Data Engineering*, pages 392–399. IEEE Computer Society Press, 1999.