Rule-Based Composite Event Queries:
The Language XChange$^{\mathrm{EQ}}$ and its Semantics

François Bry and Michael Eckert

University of Munich, Institute for Informatics
Oettingenstr. 67, 80538 München, Germany
{bry,eckert}@pms.ifi.lmu.de
http://www.pms.ifi.lmu.de

Abstract. Reactive Web systems, Web services, and Web-based publish/subscribe systems communicate events as XML messages, and in many cases require composite event detection: it is not sufficient to react to single event messages, but events have to be considered in relation to other events that are received over time.

Emphasizing language design and formal semantics, we describe the rule-based query language XChange$^{\mathrm{EQ}}$ for detecting composite events. XChange$^{\mathrm{EQ}}$ is designed to completely cover and integrate the four complementary querying dimensions: event data, event composition, temporal relationships, and event accumulation. Semantics are provided as model and fixpoint theories; while this is an established approach for rule languages, it has not been applied for event queries before.

1 Introduction

Emerging Web technologies such as reactive Web systems [9, 4, 7, 23], Web-based publish/subscribe systems [25, 15], and Web services communicate by exchanging messages. These messages usually come in an XML format such as SOAP [20] or Common Base Event (CBE) [14] and signify some application-level event, e.g., an update on a Web document, publication of new information, a request for some service, or a response to a request.

For many applications it is not sufficient to query and react to only single, atomic events, i.e., events signified by a single message. Instead, events have to be considered with their relationship to other events in a stream of events. Such events (or situations) that do not consist of one single atomic event but have to be inferred from some pattern of several events are called *composite events*.

Examples for such composite events are omnipresent. An application for student administration might require notification when "a student has both handed in her thesis and given the defense talk." A library application might send a monition when "a book has been borrowed and not returned or extended within one month." A stock market application might require notification if "the average of the reported stock prices over the last hour raises by 5%."

This article describes work on the rule-based high-level event query language XChange$^{\mathrm{EQ}}$ for the Web, focusing on language design and formal semantics.

XChange^{EQ} has been introduced in [3]; we extend on this work by providing formal semantics in the form of model and fixpoint theories for stratified programs. XChange^{EQ} is developed as a part (sub-language) of the reactive, rule-based Web language XChange [9].[1] It is however designed so that it can also be deployed as a stand-alone event mediation component in an event-driven architecture [16] or in the General Semantic Web ECA Framework described in [23].

The contributions of this article are as follows. (1) We discuss language design issues of event query languages for the Web (Section 2). We identify four complementary dimensions that need to be considered for querying events. While they might have been implicit in some works on composite event queries, we are not aware of any works stating them explicitly before.

(2) We shorty introduce XChange^{EQ} (Section 3). XChange^{EQ} is significantly more high-level and expressive than previous (composite) event query languages. To the best of our knowledge, XChange^{EQ} is the first language to deal with complex structured data in event messages, support rules as an abstraction and reasoning mechanism for events, and build on a separation of concerns that gives it ease-of-use and a certain degree of expressive completeness.

(3) We provide formal semantics for XChange^{EQ} in the form of model and fixpoint theories (Section 4). While this approach is well-explored in the world of rule-based and logic programming, its application to an event query language is novel and should be quite beneficial for research on composite event queries: semantics of earlier event query languages often have been somewhat ad hoc, generally with an algebraic and less declarative flavor, and did not accommodate rules. In our discussion, we highlight where we deviate from traditional model theories to accommodate the temporal notions required by event queries.

2   Design Considerations

Our work on XChange^{EQ} is motivated by previous work on XChange [9], a language employing Event-Condition-Action rules to program distributed, reactive Web applications. Similar to composite event detection facilities found in active databases [19, 18, 13, 12, 1], XChange provides composition operators such as event conjunction, sequence, repetition, or negation. Our experiences with programming in XChange [10, 7] has taught us that there is a considerable gap between the requirements posed by applications and the expressivity of composition operators. Further, event querying based on composition operators is prone to misinterpretations as discussions in the literature show [29, 17, 1]. This experience has lead us to reconsider and analyze the requirements for event query languages, which we present here, and to the development of XChange^{EQ}.

A sufficiently expressive event query language should cover (at least) the following four complementary dimensions. How well an event query language covers each of these dimensions gives a practical measure for its expressiveness.

---

[1] Accordingly, the superscript EQ stands for **E**vent **Q**ueries. XChange^{EQ} replaces the original composite event query constructs [8] of XChange. It has a different design and is an improvement both in expressivity and ease-of-use.

Data extraction: Events contain data that is relevant for applications to decide whether and how to react to them. For events that are received as XML messages, the structure of the data can be quite complex (semi-structured). The data of events must be extracted and provided (typically as bindings for variables) to test conditions (e.g., arithmetic expressions) inside the query, construct new events, or trigger reactions (e.g., database updates).

Event composition: To support composite events, i.e., events that consist out of several events, event queries must support composition constructs such as the conjunction and disjunction of events (more precisely, of event queries). Composition must be sensitive to event data, which is often used to correlate and filter events (e.g., consider only stock transactions from the *same* customer for composition). Since reactions to events are usually sensitive to timing and order, an important question for composite events is *when* they are detected. In a well-designed language, it should be possible to recognize when reactions to a given event query are triggered without difficulty.

Temporal (and causal) relationships: Time plays an important role in event-driven applications. Event queries must be able to express temporal conditions such as "events A and B happen within 1 hour, and A happens before B." For some applications, it is also interesting to look at causal relationships, e.g., to express queries such as "events A and B happen, and A has caused B." In this article we concentrate only on temporal relationships since causal relationships can be queried in essentially the same manner.[2]

Event accumulation: Event queries must be able to accumulate events to support non-monotonic features such as negation of events (understood as their absence) or aggregation of data from multiple events over time. The reason for this is that the event stream is (in contrast to extensional data in a database) unbounded (or "infinite"); one therefore has to define a scope (e.g., a time interval) over which events are accumulated when aggregating data or querying the absence of events. Application examples where event accumulation is required are manifold. A business activity monitoring application might watch out for situations where "a customer's order has *not* been fulfilled within 2 days" (negation). A stock market application might require notification if "the *average* of the reported stock prices over the last hour raises by 5%" (aggregation).

3   The Language XChange$^{\text{EQ}}$

XChange$^{\text{EQ}}$ is designed on the following foundations.

(1) Its syntax enforces a separation of the four querying dimensions described above, yielding a clear language design, making queries easy to read and understand, and giving programmers the benefit of a separation of concerns. Even more importantly, this separation allows to argue that the language reaches a

---

[2] While temporality and causality can be treated similarly in queries, causality raises interesting questions about how causal relationships can be *defined* and *maintained*. Investigation of these issues is planned for the future.

certain degree of expressive completeness. Our experience, stemming from attempts to express queries with existing event query languages, shows us that without such a separation not all dimensions are fully covered.

(2) It embeds the Web and Semantic Web query language Xcerpt [28] to specify classes of relevant events, extract data (in the form of variable bindings) from them, and construct new events.

(3) It supports rules as an abstraction and reasoning mechanism for events, with the same motivation and benefits of views in traditional database systems.

These foundations lead to improvements on previous work on composite event query languages in the following ways: XChange$^{\text{EQ}}$ is a high-level language with a clear design that is easy to use and provides the appropriate abstractions for querying events. It emphasizes the necessity to query data in events, which has been neglected or over-simplified earlier. Being targeted for semi-structured XML messages as required for CBE, SOAP, and Web Services, it is particularly suitable for use in business applications domains. We make an attempt towards expressive completeness by fully covering all four query dimensions explained earlier using a separation of concerns in XChange$^{\text{EQ}}$. Arguably, in previous languages that do not use such a separation, some (usually simple) queries might be expressed more compactly. This compactness then however leads easily to misinterpretations (as discussed in [29, 17, 1]) and comes in previous work at the price of a serious lack in expressiveness (incomplete coverage of the four dimensions), where less simple queries cannot be expressed.

Using the example of a stock market application, we now introduce the syntax of our event query language XChange$^{\text{EQ}}$.

3.1 Querying Atomic Events

Application-level events are nowadays often represented as XML, especially in the formats Common Base Event [14] and SOAP [20]. Skipping details of such formats for the sake of brevity, we will be using four atomic events in our stock market example: stock *buy*s, stock *sell*s, and *order*s to buy or sell stocks. Involved applications may also generate further events without affecting our examples.

The left side of Figure 1 depicts a *buy order* event in XML. For conciseness and human readability, we use a "term syntax" for data, queries, and construction of data instead of the normal tag-based XML syntax. The right side of Figure 1 depicts the XML event as a (data) term. The term syntax is slightly more general than XML, indicating whether the order of children is relevant (square brackets []), or not (curly braces {}).

Querying such single event messages is a two-fold task: one has to (1) specify a class of relevant events (e.g., all *buy* events) and (2) extract data from the events (e.g., the price). XChange$^{\text{EQ}}$ embeds the XML query language Xcerpt [28] for both. Figure 2 shows an exemplary *buy* event (left) and an event query that recognizes such *buy* events with a price total of \$10 000 or more (right).

Xcerpt queries describe a pattern that is matched against the data. Query terms can be partial (indicated by double brackets or braces), meaning that a matching data term can contain subterms not specified in the query, or total

```
<order>                                        order [
   <orderId >4711</orderId >                     orderId  {  4711  },
   <customer>John</customer>                     customer  {  "John"  },
   <buy>    <stock>IBM</stock>                   buy  [    stock  {  "IBM"  },
            <limit >3.14</limit >                          limit  {  3.14  },
            <volume>4000</volume>                          volume  {  4000  }      ]
</buy> </order>                                ]
```

Fig. 1. XML and term representation of an event

```
buy  [    orderId  {  4711  },              buy  {{   tradeId  {  var  I  },
          tradeId  {  4242  },                        customer  {  var  C  },
          customer  {  "John"  },                     stock  {  var  S  },
          stock  {  "IBM"  },                         price  {  var  P  },
          price  {  2.71  },                          volume  {  var  V  }
          volume  {  4000  }    ]           }}  where  {  var  P  ∗  var  V  >=  10000  }
```

Fig. 2. Atomic event query

(indicated by single brackets or braces). Queries can contain variables (keyword
`var`), which will be bound to the matching data, and a `where`-clause can be
attached to specify non-structural (e.g., arithmetic) conditions. In this article,
we will stick to simple queries as above. Note however that Xcerpt supports
more advanced constructs for (subterm) negation, incompleteness in breadth
and depth, and queries to graph-shaped data such as RDF. An introduction to
Xcerpt is given in [28].

The result of evaluating an Xcerpt query on an event message is the set Σ
of all possible substitutions for the free variables in the query (non-matching is
signified by $\Sigma = \emptyset$). Our example query does not match the *order* event from
Figure 1, but matches the *buy* event on the left of Fighre 2 with $\Sigma = \{\sigma_1\}$,
$\sigma_1 = \{I \mapsto 4242, C \mapsto \text{John}, S \mapsto \text{IBM}, P \mapsto 2.71, V \mapsto 4000\}$.

In addition to event messages, XChange$^{\text{EQ}}$ event queries can query for timer
events. Absolute timer events are time points or intervals (possibly periodic)
defined without reference to the occurrence time of some other event. They are
specified in a similar way as queries to event messages and we refer to [3] for
details. Relative timer events, i.e., time points or intervals defined in relation to
some other event, will be looked at in Section 3.3 on event composition.

3.2   Reactive and Deductive Rules for Events

XChange$^{\text{EQ}}$ uses two kinds of rules: deductive rules and reactive rules. Deductive
rules allow to define new, "virtual" events from the events that are received. They
have no side effects and are analogous to the definition of views for database data.
Figure 3 (left) shows a deductive rule deriving a new *bigbuy* events from *buy*
events satisfying the earlier event query of Figure 2. Deductive rules follow the
syntax DETECT *event construction* ON *event query* END. The event construction
in the rule head is simply a data term augmented with variables which are
replaced during construction by their values obtained from evaluating the event

```
DETECT bigbuy {
        tradeId { var I },                          RAISE
        customer { var C },                            to(recipient=
        stock { var S }  }                                "http://auditor.com",
ON buy {{                                                transport=
     tradeId { var I },                                 "http://.../HTTP/")
     customer { var C },                              {
     stock { var S },                                   var B
     price { var P },                                 }
     volume { var V }                             ON var B -> bigbuy {{ }}
  }} where { var P * var V >= 10000 }          END
END
```

Fig. 3. Deductive rule (left) and reactive rule (right)

query in the rule body. (Several variables bindings will lead to the construction of several events if no grouping or aggregation constructs are used.) The event construction is also called a construct term; more involved construction will be seen in Section 3.5 when we look at aggregation of data. Recursion of rules is restricted to stratifiable programs, see Section 4.2 for a deeper discussion.

Reactive rules are used for specifying a reaction to the occurrence of an event. The usual (re)action is constructing a new event message (as with deductive rules) and use it to call some Web Service. Note that this new event leaves the system and that it is up to the receiver to decide on the occurrence time (typically such events are considered to happen only at the time *point* when the corresponding message is received). For tasks involving accessing and updating persistent data, our event queries can be used in the Event-Condition-Action rules of the reactive language XChange.

An example for a reactive rule is in Figure 3 (right); it forwards every *bigbuy* event (as derived by the deductive rule on the left) to a Web Service `http://auditor.com` using SOAP's HTTP transport binding. The syntax for reactive rules is similar to deductive rules, only they start with the keyword `RAISE`; in the rule head `to()` is used to indicate recipient and transport.

The distinction between deductive and reactive rules is important. While it is possible to "abuse" reactive rules to simulate or implement deductive rules (by sending oneself the result), this is undesirable: it is difficult with events that have a duration, misleading for programmers, less efficient for evaluation, and could allow arbitrary recursion (leading, e.g., to non-terminating programs or non-stratified use of negation).

3.3 Composition of Events

So far, we have only been looking at queries to single events. Since temporal conditions are dealt with separately, only two operators, `or` and `and`, are necessary to compose event queries into *composite event queries*. (Negation falls under event accumulation, see Section 3.5.) Both composition operators are multi-ary, allowing to compose any (positive) number of event queries (without need for nesting), and written in prefix notation. Disjunctions are a convenience in prac-

```
DETECT buyorderfulfilled {  orderId { var O },
                            tradeId { var I },
                            stock { var S } }
ON and {
      order  {  orderId { var O },
                buy {{  stock { var S }  }}  },
      buy    {{  orderId { var O },
                 tradeId { var I }  }}  }
END
```

Fig. 4. Conjunction of event queries

```
and {  event o: order {{ orderId { var O } }},
       event t: extend[o, 1 min]  }
```

Fig. 5. Composition with relative timer event

tical programming but not strictly necessary: a rule with a (binary) disjunction can be written as two rules. We therefore concentrate on conjunctions here.

When two event queries are composed with `and`, an answer to the composite event query is generated for every pair of answers to the constituent queries. If the constituent queries share free variables, only pairs with "compatible" variable bindings are considered. (This generalizes to composition of three and more event queries in the obvious manner.) Figure 4 illustrates the use of the `and` operator. The *buy order fulfilled* event is detected for every corresponding pair of *buy order* and *buy* event. The events have to agree on variable O (the `orderId`). The occurrence time of the detected *order fulfilled* event is the time interval enclosing the respective constituent events.

Composition of events gives rise to defining relative timer events, i.e., time points or intervals defined in relation to the occurrence time of some other event. Figure 5 shows a composite event query asking for an *order* event and a timer covering the whole time interval from the *order* event until one minute after. This timer event will be used later in Section 3.5 when querying for the absence of a corresponding *buy* event in this time interval.

An event identifier (o) is given to the left of the event query after the keyword `event`. It is then used in the definition of the relative timer `extend[o, 1 min]` which specifies a time interval one minute longer than the occurrence interval of o. (The time point at which o occurs is understood for this purpose as a degenerated time interval of zero length.) The event identifier t is not necessary here, but can be specified anyway. Event identifiers will also be used in temporal conditions and event accumulation (Sections 3.4 and 3.5).

Further constructors for relative timers are: `shorten[e,d]` (subtracting d from the end of e), `extend-begin[e,d]`, `shorten-begin[e,d]` (adding or subtracting d at the begin of e), `shift-forward[e,d]`, `shift-backward[e,d]` (moving e forward or backward by d).

```
DETECT  earlyResellWithLoss  {   customer  {  var  C  },
                                 stock  {  var  S  }   }
ON and  {
    event  b:  buy  {{   customer  {  var  C  },
                        stock  {  var  S  },
                        price  {  var  P1  }   }},
    event  s:  sell  {{   customer  {  var  C  },
                        stock  {  var  S  },
                        price  {  var  P2  }   }}
    } where  {  b  before  s,  timeDiff(b,s)<1hour,  var  P1>var  P2  }
END
```

Fig. 6. Event query with temporal conditions

3.4   Temporal Conditions

Temporal conditions on events and causal relationships between events play an important role in querying events. We concentrate in this paper on temporal conditions, though the approach generalizes to causal relationships. Just like conditions on event data, temporal conditions are specified in the `where`-clause of an event query and make use of the event identifiers introduced above.

The event query in Figure 6 involves temporal conditions. It detects situations where a customer first buys stocks and then sells them again within a short time (less than 1 hour) at a lower price. The query illustrates that typical applications require both qualitative conditions (`b before s`) and quantitative (or metric) conditions (`timeDiff(b,s) < 1 hour`). In addition, the query also includes a data condition for the price (`var P1 > var P2`).

In principle, various external calendar and time reasoning systems could be used to specify and evaluate temporal conditions. However, many optimizations for the evaluation of event queries require knowledge about temporal conditions. See [6] for an initial discussion of temporal optimizations.

XChange^{EQ} deals with non-periodic time intervals (time points are treated as degenerated intervals of zero length), periodic time intervals (i.e., sequences of non-periodic intervals), and durations (lengths of time). An overview of the built-in constructs for temporal conditions can be found in [3].

Note that there is an important difference between timer events used in queries and references to time as part of `where`-conditions. Timer events have to happen for the event query to yield an answer (i.e., they are waited for), while time references in conditions can lie in the future and only restrict the possible answers to an event query.

3.5   Event Accumulation

Event querying displays its differences to traditional querying most perspicuously in non-monotonic query features such as negation or aggregation. For traditional database queries, the data to be considered for negation or aggregation is read-

```
DETECT  buyOrderOverdue {              RAISE  to ( . . . )  {
          orderId  {  var  I  }    }      reportOfDailyAverages {
ON and  {                                   all  entry  {
     event o:  order {{                        stock  {  var  S  },
       orderId  {  var  I  }                   avgPrice  {  avg( all  var  P )  }
       buy  {{  }}   }},                   }  group−by  var  S    }    }
     event t:  extend [ o ,  1  min ] ,  ON and  {
     while  t:  not  buy  {                  event  t:  tradingDay {{  }},
       orderId  {  var  I  }    }            while  t:  collect  sell  {
   }                                            stock  {  var  S  },
END                                             price  {  var  P  }    }
                                       END
```

Fig. 7. Event accumulation for negation (left) and aggregation (right)

ily available in the database and this database is *finite*.[3] In contrast, events are received over time in an event stream which is unbounded, i.e., potentially infinite. Applying negation or aggregation on such a (temporally) infinite event stream would imply that one has to wait "forever" for an answer because events received at a later time might always change the current answer. We therefore need a way to restrict the event stream to a finite temporal extent (i.e., a finite time interval) and apply negation and aggregation only to the events collected in this accumulation window.[4]

It should be possible to determine the accumulation window dynamically depending on the event stream received so far. Typical cases of such accumulation windows are: "from event a until event b," "one minute until event b," "from event a for one minute," and (since events can occur over time intervals, not just time points) "while event c." Here we only look at the last case because it subsumes the first three (they can be defined as composite events).

Negation is supported by applying the `not` operator to an event query. The window is specified with the keyword `while` and the event identifier of the event defining the window. The meaning is as one might expect: the negated event query `while` t: `not` q is successful if no event satisfying q occurs during the time interval given by t. An example can be seen in Figure 7 (left): it detects buy orders that are overdue, i.e., where no matching buy transaction has taken place within one minute after placing the order. The accumulation window is specified by the event query t, which is a timer relative to the *order* event. Observe that the negated query can contain variables that are also used outside the negation; the example reveals the strong need to support this.

Following the design of the embedded query language Xcerpt, aggregation constructs are used in the *head* of a rule, since they are related to the construction of new data. The task of the *body* is only *collecting* the necessary data or events. Collecting events in the body of a rule is similar to negation and indicated by

---

[3] Recursive rules or views may allow to define infinite databases intensionally. However, the extensional data (the "base facts") is still finite.

[4] Keep in mind that accumulation here refers to the way we specify queries, not the way evaluation is actually performed. Keeping all events in the accumulation windows in memory is generally neither desirable nor necessary for query evaluation.

the keyword `collect`. The rule in Figure 7 (right) has an event query collecting *sell* events over a full *trading day*. The actual aggregation takes place in the head of the rule, where all sales prices (P) for the same stock (S) are averaged and a report containing one entry for each stock is generated. The report is sent at the end of each trading day; this is reflected in the syntax by the fact that `tradingDay{{ }}` must be written as an event, i.e., must actually occur.

Aggregation follows the syntax and semantics of Xcerpt (see [27] for a full account), again showing that it is beneficial to base an event query language on a data query language. The keyword `all` indicates a structural aggregation, generating an `entry` element for each distinct value of the variable S (indicated with `group-by`). Inside the `entry`-element an aggregation function `avg` is used to compute the average price for each individual stock.

Aggregation has rarely been considered in work on composite events, though it is clearly needed in many applications, including our stock market example. A notable exception is [24], which however applies only to relational data (not semi-structured or XML) and does not have the benefits of a separation of the query dimensions as XChange^{EQ}.

4 Formal Semantics

Having introduced XChange^{EQ} informally above, we now supply formal, declarative semantics for stratified programs in the form of model and fixpoint theories. While this is a well-established approach for rule-based languages [22, 2], including traditional database query languages supporting views or deductive rules, it has not been applied to event query languages before. Related work on semantics for event queries usually has an "algebraic flavor" (as the languages themselves do), where the semantics for operators are given as functions between sequences (or histories or traces) of events, e.g., [30, 21]. Further, these approaches often neglect *data* in events (especially semi-structured data) and it is not clear how they could be extended to support deductive *rules* (or views) over events.

In addition to accommodating both rules and data, the model theoretic approach presented here can be argued to be more declarative than previous algebraic approaches, expressing *how* an event is to be to detected rather than *what* event is to be detected, making programs easier to understand and optimize.

The following specifics of querying events as opposed to pure (database) data have to be arranged for in our semantics and make it novel compared their counterpart in the logic programming literature [22, 2]: (1) in addition to normal variables, event identifiers are accommodated, (2) answers to composite event queries have an occurrence time, (3) temporal relations have a fixed interpretation. Finally, the model theory must be (4) sensible for potentially *infinite* streams of events (this also entails that negation and aggregation of events must be "scoped" over a time window as we have seen earlier in Section 3.5).

4.1 Model Theory

Our model theory is Tarskian-style [11], i.e., it uses a valuation function for free variables and defines an entailment relation between an interpretation and sentences (rules and queries) from the language *recursively over the structure* of the sentences.[5] Tarskian model theories have the advantage of being highly declarative, theoretically well-understood, and relatively easy to understand.

An **event** happens over a given time interval and has a representation as message (as data term). Formally it is a tuple of a (closed and convex) time interval t and a data term e, written e^t. The set of all events is denoted *Events*.

Time is assumed to be a linearly[6] ordered set of time points $(\mathbb{T}, <)$. The time intervals over which events happen are closed and convex, i.e., have the form $t = [b, e] = \{p \mid b \leq p \leq e\}$ (where $b \in \mathbb{T}$ and $e \in \mathbb{T}$). For convenience we define: $begin([b, e]) = b$, $end([b, e]) = e$, $[b_1, e_1] \sqcup [b_2, e_2] = [min\{b_1, b_2\}, max\{e_1, e_2\}]$, and $[b_1, e_1] \sqsubseteq [b_2, e_2]$ iff $b_2 \leq b_1$ and $e_1 \leq e_2$.

Matching of **Atomic Event Queries** against single incoming events is based on a non-standard unification that is especially designed for the variations and incompleteness in semi-structured data. Atomic Event Queries are single query terms q that match only for the data term part e of events e^t; this does not involve time or multiple events. Note that the query terms usually contain free variables. The matching of query terms and data terms is based on **Simulation**, which is a relation between ground terms, denoted $\preceq$. Intuitively, $q \preceq d$ means that the nodes and structure of q can be found in d. Simulation naturally extends to a non-ground query term q' by asking whether there is a (grounding) substitution σ for the free variables in q' such that the ground query term $q = \sigma(q')$ obtained by applying the substitution σ to q' simulates with the given data term d. Further details can be found in [27]; they are not important for understanding the presented model theory and thus not discussed here.

Substitution sets Σ rather than single substitutions σ are used in our model theory to accommodate grouping and aggregation in the construction in rule heads. Application $\Sigma(c)$ of Σ to a construct term c results in a set of data terms. For convenience we also define the application to query terms q with $\Sigma(q) = \{\sigma(q) \mid \sigma \in \Sigma\}$.

An **interpretation** for a given XChange^{EQ} query, rule, or program is a 3-tuple $M = (I, \Sigma, \tau)$, where (1) $I \subseteq Events$ is the set of events e^t that "happen," i.e., are either in the stream of incoming events or derived by some deductive rule. (2) $\Sigma \neq \emptyset$ is a grounding substitution set containing substitutions for the "normal" variables (i.e., data variables, but not event identifiers). (3) τ is a substitution for the event identifiers, i.e., a mapping from event variables to *Events*. The substitution τ for event identifiers (cf. Section 3.3) is the first unusual features of our model theory. Since τ signifies the events that contributed to the answer of some query, we also call it an "event trace."

---

[5] This recursive definition over the structure allows to consider sub-formulas of a formula in isolation, which is beneficial for both understanding and evaluation.

[6] Linear time is chosen because we are interested in event that actually happened, not in potential futures (where a branching time would be more apt).

$$I, \Sigma, \tau \models (\texttt{event } i : q)^t \qquad \text{iff exists } e^{t'} \in I \text{ with } \tau(i) = e^{t'},\ t' = t,$$
$$\text{and for all } e' \in \Sigma(q) \text{ we have } e' \preceq e$$

$$I, \Sigma, \tau \models (\texttt{event } i : \texttt{extends}[j, d])^t \quad \text{iff exists } e^{t'} \text{ with } \tau(j) = e^{t'},\ \tau(i) = e^t,\ t = t' + d$$
$$\ldots \qquad (\text{Definitions for other temporal events are similar and skipped.})$$

$$M \models (q_1 \wedge q_2)^t \qquad \text{iff } M \models q_1^{t1} \text{ and } M \models q_2^{t2} \text{ and } t = t_1 \sqcup t_2$$
$$M \models (q_1 \vee q_2)^t \qquad \text{iff } M \models q_1^t \text{ or } M \models q_2^t$$
$$I, \Sigma, \tau \models (Q \texttt{ where } C)^t \qquad \text{iff } I, \Sigma, \tau \models Q^t \text{ and } W_{\Sigma,\tau}(C) = true$$
$$I, \Sigma, \tau \models (\texttt{while } j : \texttt{ not } q)^t \qquad \text{iff exists } e^{t'} \text{ with } \tau(j) = e^{t'},\ t' = t,$$
$$\text{and for all } t'' \sqsubseteq t \text{ we have } I, \Sigma, \tau \not\models q^{t''}$$

$$I, \Sigma, \tau \models (\texttt{while } j : \texttt{ collect } q)^t \quad \text{iff exists } e^{t'} \text{ with } \tau(j) = e^{t'},\ t' = t, \text{ and exist } n \geq 0,$$
$$\Sigma_1, \ldots \Sigma_n,\ t_1 \sqsubseteq t, \ldots t_n \sqsubseteq t \text{ with } \Sigma = \bigcup_{i=1..n} \Sigma_i,$$
$$\text{and for all } i = 1..n \text{ we have } I, \Sigma_i, \tau \models q^{t_i}$$

$$I, \Sigma, \tau \models (c \leftarrow Q)^t \qquad \text{iff (1) } \Sigma'(c)^t \subseteq I \text{ for } \Sigma' \text{ maximal (w.r.t. } FreeVars(Q)) \text{ and } \tau'$$
$$\text{such that } I, \Sigma', \tau' \models Q^t, \text{ or (2) } I, \Sigma', \tau' \not\models Q^t \text{ for all } \Sigma', \tau'$$

$$W_{\Sigma,\tau}(i \texttt{ before } j) = \textbf{true} \quad \text{iff } end(\tau(i)) < begin(\tau(j))$$
$$W_{\Sigma,\tau}(i \texttt{ during } j) = \textbf{true} \quad \text{iff } begin(\tau(j)) < begin(\tau(i)) \text{ and } end(\tau(i)) < end(\tau(j))$$
$$W_{\Sigma,\tau}(i \texttt{ overlaps } j) = \textbf{true} \text{ iff } begin(\tau(j)) < begin(\tau(i)) < end(\tau(j)) < end(\tau(i))$$

Fig. 8. Model Theory for XChange$^{\text{EQ}}$

The **satisfaction** $M \models F^t$ of an XChange$^{\text{EQ}}$ expression F over an occurrence time t in an interpretation M is defined recursively in Figure 8. The time stamping of expressions is the second unusual feature of our model theory.

Given an XChange$^{\text{EQ}}$ program P and a stream of incoming events E, we call an interpretation $M = (I, \Sigma, \tau)$ a **model** of P under E if (1) M satisfies all rules $(c \leftarrow Q) \in P$ for all time intervals t and (2) contains the stream of incoming events, i.e., $E \subseteq I$. Note that here the event stream simply corresponds to the notion of base facts or extensional data found of traditional model theories.

The satisfaction relation uses a fixed interpretation W for all conditions that can occur in the where-clause of a query. This includes the temporal relations like before and is the third unusual feature of our model theory. W is a function that maps a substitution set Σ, an event trace τ, and an atomic condition C to true or false; we usually write Σ and τ in the index. $W_{\Sigma,\tau}$ extends straightforwardly to boolean formulas of conditions. The definition of W is left outside the "core model theory" to make it more modular and allow to easily integrate different temporal reasoners. In Figure 8, we have given only the definitions for before, during, and overlaps for space reasons.

Our fourth requirement on the model theory was that it is sensible on (potentially) infinite streams of events. The basic idea for this is that to evaluate a program P over a time interval t, we only have to consider events happening during t. We will state this formally after giving the fixpoint theory.

4.2 Fixpoint Theory

A model theory, such as the one presented above, has the issue of allowing many models for a given program. A common and convenient way to obtain a unique model is to define it as the solution of a fixpoint equation (which is based on the model theory). A fixpoint theory also describes an abstract, simple, forward-chaining evaluation method, which can easily be extended to work incrementally as is required for event queries [4].

Our fixpoint theory requires XChange$^{\text{EQ}}$ programs to be stratifiable [2]. **Stratification** restricts the use of recursion in rules by ordering the rules of a program P into so-called strata (sets P_i of rules with $P = P_1 \uplus \cdots \uplus P_n$) such that a rule in a given stratum can only depend on (i.e., access results from) rules in lower strata (or the same stratum, in some cases). The restriction to stratifiable programs could be partially lifted at the cost of a more involved semantics (and evaluation). This is however outside the scope of this paper.

Three types of stratification are required: (1) Negation stratification, i.e., events that are negated in the query of a rule may only be constructed by rules in lower strata, events that occur positively may only be constructed by rules in lower strata or the same stratum. (2) Grouping stratification, i.e., rules using grouping constructs like `all` in the construction may only query for events constructed in lower strata. (3) Temporal stratification, i.e., if a rule queries a relative temporal event like `extends[i, 1min]` then the anchoring event (here: `i`) may only be constructed in lower strata. While negation and grouping stratification are fairly standard, temporal stratification is a requirement specific to complex event query programs like those expressible in XChange$^{\text{EQ}}$. We are not aware of former consideration of the notion of temporal stratification. For a formal definition of our stratification, we refer to [5].

The **fixpoint operator** T_P for an XChange$^{\text{EQ}}$-Program P is defined as:

$$T_P(I) = I \cup \{e^t \mid \text{there exist a rule } c \leftarrow Q \in P, \text{ a maximal substitution set } \Sigma,$$
$$\text{and a substitution } \tau \text{ such that } I, \Sigma, \tau \models Q^t \text{ and } e \in \Sigma(c)\}$$

The repeated application of T_P until a fixpoint is reached is denoted T_P^ω.

The **fixpoint interpretation**[7] $M_{P,E}$ of a program P with stratification $P = P_1 \uplus \cdots \uplus P_n$ under and event stream E is defined by computing fixpoints stratum by stratum: $M_0 = E = T_\emptyset^\omega(E)$, $M_1 = T_{\overline{P_1}}^\omega(M_0) \ldots$, $M_{P,E} = M_n = T_{\overline{P_n}}^\omega(M_{n-1})$. Here, $\overline{P_i} = \bigcup_{j \leq j} P_j$ denotes the set of all rules in strata P_i and lower.

Theorem 1 justifies our definition as usual for fixpoint semantics: For a stratifiable program P and an event stream E, $M_{P,E}$ is a minimal model of P under E. Further, $M_{P,E}$ is independent of the stratification of P.

More interestingly, we can show that the model theory and fixpoint semantics are sensible on infinite event streams. The next theorem justifies a streaming evaluation, where answers to composite event queries are generated "online" and we never have to wait for the stream to end (which it will not if infinite). This is the last feature of our semantics that is peculiar for event queries.

Theorem 2: Let $E \mid t$ denote the restriction of an event stream E to a time interval t, i.e., $E \mid t = \{e^{t'} \in E \mid t' \sqsubseteq t\}$. Similarly, let $M \mid t$ denote the restriction of an interpretation M to t. Then the result of applying the fixpoint procedure to $E \mid t$ is the same as applying it to E for the time interval t, i.e., $M_{P,E\mid t} \mid t = M_{P,E} \mid t$. In other words to evaluate a program over a time interval t, we do not have to consider any events happening outside of t.

Proofs for both theorems are presented in an extended version of this paper [5]. The proof for theorem 1 is an adoption of a proof in [22].

---

[7] Since we consider whole programs P now, only the set I of events that happen is relevant for the fixpoint interpretation of P; Σ and τ are thus skipped from now on.

5 Conclusions and Future Work

This article has introduced the high-level event query language XChange^{EQ}, emphasizing language design and formal semantics. XChange^{EQ} deviates from previous event query languages in a separation of the query dimensions data extraction, event composition, temporal relationships, and event accumulation. This separation allows a complete coverage of each of the dimensions, yielding a language that can be argued to have reached a degree of expressive completeness.

The ability to query events represented in XML and other Web formats, makes XChange^{EQ} suited for use in service-oriented and event-driven architectures based on Web Services. Important for practical use, rules are supported as an abstraction and reasoning mechanism for events. Rule-based reasoning about events is also expected to become relevant in efforts to bring rules, including reactive rules, to the (Semantic) Web [26, 4].

Efficient evaluation methods that utilize temporal conditions [6] and query optimization for large numbers of event queries are the current focus of our research. Implementation of our language in the scope of XChange is ongoing work.

Acknowledgments

References

1. R. Adaikkalavan and S. Chakravarthy. SnoopIB: Interval-based event specification and detection for active databases. *Data and Knowledge Eng.*, 2005. In press.
2. K. R. Apt, H. A. Blair, and A. Walker. Towards a theory of declarative knowledge. In *Foundations of Deductive Databases and Logic Programming*. Morgan Kaufmann, 1988.
3. F. Bry and M. Eckert. A high-level query language for events. In *Proc. Int. Workshop on Event-driven Architecture, Processing and Systems*, 2006.
4. F. Bry and M. Eckert. Twelve theses on reactive rules for the Web. In *Proc. Int. Workshop Reactivity on the Web*, 2006.
5. F. Bry and M. Eckert. Rule-based composite event queries: The language XChange^{EQ} and its semantics [extended version]. Technical report, Inst. f. Informatics, U. of Munich, 2007. Available at www.pms.ifi.lmu.de/publikationen/.
6. F. Bry and M. Eckert. Temporal order optimizations of incremental joins for composite event detection. Technical report, Inst. f. Informatics, U. of Munich, 2007. Available at www.pms.ifi.lmu.de/publikationen/.
7. F. Bry, M. Eckert, H. Grallert, and P.-L. Pătrânjan. Evolution of distributed Web data: An application of the reactive language XChange. In *Proc. Int. Conf. on Data Engineering (Demonstrations)*, 2006.
8. F. Bry, M. Eckert, and P.-L. Pătrânjan. Querying composite events for reactivity on the Web. In *Proc. Int. Workshop on XML Research and Applications*, 2006.

9. F. Bry, M. Eckert, and P.-L. Pătrânjan. Reactivity on the Web: Paradigms and applications of the language XChange. *J. of Web Engineering*, 5(1), 2006.

10. F. Bry, M. Eckert, P.-L. Pătrânjan, and I. Romanenko. Realizing business processes with ECA rules: Benefits, challenges, limits. In *Proc. Int. Workshop on Principles and Practice of Semantic Web*, 2006.

11. F. Bry and M. Marchiori. Ten theses on logic languages for the Semantic Web. In *Int. Workshop on Principles and Practice of Semantic Web Reasoning*, 2005.

12. A. P. Buchmann, J. Zimmermann, J. A. Blakeley, and D. L. Wells. Building an integrated active OODBMS: Requirements, architecture, and design decisions. In *Proc. Int. Conf. on Data Engineering*, 1995.

13. S. Chakravarthy, V. Krishnaprasad, E. Anwar, and S.-K. Kim. Composite events for active databases: Semantics, contexts and detection. In *Proc. Int. Conf. on Very Large Data Bases*, 1994.

14. Common Base Event. www.ibm.com/developerworks/webservices/library/ws-cbe.

15. Y. Diao, S. Rizvi, and M. J. Franklin. Towards an internet-scale XML dissemination service. In *Proc. Int. Conf. on Very Large Data Bases*, 2004.

16. O. Etzion. Towards an event-driven architecture: An infrastructure for event processing (position paper). In *Proc. Int. Conf. on Rules and Rule Markup Languages for the Semantic Web*, 2005.

17. A. Galton and J. C. Augusto. Two approaches to event definition. In *Proc. Int. Conf. on Database and Expert Systems Applications*, 2002.

18. S. Gatziu and K. R. Dittrich. Events in an active object-oriented database system. In *Proc. Int. Workshop on Rules in Database Systems*, 1993.

19. N. H. Gehani, H. V. Jagadish, and O. Shmueli. Composite event specification in active databases: Model & implementation. In *Proc. Int. Conf. on Very Large Data Bases*, 1992.

20. M. Gudgin et al. SOAP 1.2. W3C recommendation, 2003.

21. A. Hinze and A. Voisard. A parameterized algebra for event notification services. In *Proc. Int. Symp. on Temporal Representation and Reasoning*, 2002.

22. J. W. Lloyd. *Foundations of Logic Programming*. Springer-Verlag, 1993.

23. W. May, J. J. Alferes, and R. Amador. Active rules in the Semantic Web: Dealing with language heterogeneity. In *Proc. Int. Conf. on Rules and Rule Markup Languages for the Semantic Web*, 2005.

24. I. Motakis and C. Zaniolo. Temporal aggregation in active database rules. In *Proc. Int. Conf. on Management of Data (SIGMOD)*, 1997.

25. J. Pereira, F. Fabret, H.-A. Jacobsen, F. Llirbat, and D. Shasha. WebFilter: A high-throughput XML-based publish and subscribe system. In *Proc. Int. Conf. on Very Large Databases*, 2001.

26. Rule Interchange Format WG Charter. www.w3.org/2005/rules/wg/charter.

27. S. Schaffert. *Xcerpt: A Rule-Based Query and Transformation Language for the Web*. PhD thesis, Inst. f. Informatics, U. of Munich, 2004.

28. S. Schaffert and F. Bry. Querying the Web reconsidered: A practical introduction to Xcerpt. In *Proc. Extreme Markup Languages*, 2004.

29. D. Zhu and A. S. Sethi. SEL, a new event pattern specification language for event correlation. In *Proc. Int. Conf. on Computer Communications and Networks*, 2001.

30. D. Zimmer and R. Unland. On the semantics of complex events in active database management systems. In *Proc. Int. Conf. on Data Engineering*, 1999.