

Extending Complex Event Processing to Graph-structured Information

Gala Barquero
Universidad de Málaga
Málaga, Spain
gala@lcc.uma.es

Javier Troya
Universidad de Sevilla
Sevilla, Spain
jtroya@us.es

Loli Burgueño
Universidad de Málaga
Málaga, Spain
loli@lcc.uma.es

Antonio Vallecillo
Universidad de Málaga
Málaga, Spain
av@lcc.uma.es

ABSTRACT

Complex Event Processing (CEP) is a powerful technology in real-time distributed environments for analyzing fast and distributed streams of data, and deriving conclusions from them. CEP permits defining complex events based on the events produced by the incoming sources in order to identify complex meaningful circumstances and to respond to them as quickly as possible. However, in many situations the information that needs to be analyzed is not structured as a mere sequence of events, but as graphs of interconnected data that evolve over time. This paper proposes an extension of CEP systems that permits dealing with graph-structured information. Two case studies are used to validate the proposal and to compare its performance with traditional CEP systems. We discuss the benefits and limitations of the CEP extensions presented.

ACM Reference Format:

Gala Barquero, Loli Burgueño, Javier Troya, and Antonio Vallecillo. 2018. Extending Complex Event Processing to Graph-structured Information. In *Proceedings of ACM/IEEE 21th International Conference on Model Driven Engineering Languages and Systems (MODELS'18)*. ACM, New York, NY, USA, 10 pages.

1 INTRODUCTION

Stream processing systems are becoming widespread due to the steadily growing number of information sources that continuously produce and offer data. Among them, Complex Event Processing (CEP) is currently a mature technology for analyzing and correlating streams of information about real-time events that happen in a system, and deriving conclusions from them [4, 12, 20, 21]. A distinguishing feature of CEP, not present in many stream processing systems, is that it permits defining complex events or patterns on top of the primitive events in order to identify elaborate meaningful circumstances and to respond to them as quickly as possible. Such event types and event patterns are defined using Event Processing Languages (EPLs).

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

MODELS'18, October 2018, Copenhagen, Denmark

© 2018 Copyright held by the owner/author(s).

In many applications, however, the information that needs to be analyzed is not structured as a mere sequence of partially ordered timed events, but as graphs of highly interconnected datasets that evolve over time—e.g., contagious disease spreading data or social media networks. In these systems, not only the type of event and the moment in time at which it occurs are relevant, but also its connections with other surrounding objects, the state of these objects, and the current network topology.

This paper proposes an extension of CEP systems that permits dealing with graph-structured information, in addition to their timed-events stream processing capabilities, and its implementation using graph technologies based on cluster computing platforms. Two case studies are used to illustrate the proposal, and to discuss the benefits and limitations of this kind of extension of CEP systems.

Our solution makes use of the fact that the structure of a CEP event stream can be generalized from a sequence of time-ordered elements to a model (i.e., a graph of interrelated elements). Besides, the behavior of a CEP system can be naturally considered as a special case of in-place model transformation (MT) [9], and therefore most of the MT machinery can be reused. To tackle the stringent requirements of CEP systems regarding their scalability and performance, we propose the use of the recent graph-parallel computation technologies (such as the GraphX component of Spark [15]), which provide the required supporting storage and access infrastructure.

One of our major contributions is the generalization of the concept of CEP *windows* to the more general concept of graph *vicinity*. CEP makes use of windows in order to restrict the matching space to the substream of events with a certain *time* or *length* window, i.e., with timestamps inside a temporal window, or limited to a consecutive number of events, respectively. Our approach also permits restricting the pattern matching to a subset of the whole space. In addition to the temporal or length windows from the original CEP, we introduce the concept of *spatial windows*, which are defined in terms of *vicinity graphs* [10], composed of elements related to those in the rules by means of the graph arcs, and which constitute their local context. It is important to note that, in our proposal, time becomes just one of the possible dimensions through which we can navigate the graph.

The structure of the paper is as follows. First, Section 2 briefly introduces CEP systems and their main features, as well as Apache

Spark and graph-parallel computation technologies, the three underlying technologies over which our approach is built. Then, Section 3 describes our extension to CEP systems for dealing with graph structures, beyond sequences of timed events. Section 4 discusses how we have implemented our proposal using cluster computing platforms, and Section 5 describes the validation experiments we have performed to assess our approach. Finally, Section 6 relates our work to other similar approaches and Section 7 concludes and outlines some future lines of work.

2 BACKGROUND

2.1 CEP in a nutshell

CEP [12, 20] is a form of Information Processing [4] whose goal is the definition and detection of situations of interest, from the analysis of low-level event notifications [5]. According to the Event Processing Technical Society [13], we use the *simple* events term to refer to the low-level primitive event occurrences, and *complex* events to those that summarize, represent, or denote a set of other events. *Derived* events are a particular kind of complex events, which are generated as a consequence of applying a process or method to one or more other events. In summary, CEP systems analyze streams of *simple* events to detect occurrences of *complex* events, that represent the high-level situations of interest to the CEP modeler, using declarative *rules* that define the *derived* events in terms of *patterns* of (simple or derived) events, their contents, and temporal relations.

Although several CEP systems and languages exist, they all share the same basic concepts, mechanisms and structure. These are briefly introduced below.

- **Events.** In CEP, every event (simple or derived) has a *type* and a set of *attributes*. Events are atomic, happen instantaneously, and they all have an attribute with information about the moment in time at which they occur.
- **Patterns.** A CEP rule defines a derived event, by means of a *pattern* that combines other events. Whenever the pattern is detected in the stream (i.e., it is *satisfied* by the stream events), the derived event is created. The simplest kind of pattern (called *selection pattern*) permits creating derived events every time a given simple event is detected in the stream.
- **Elements of patterns.** There are several representative elements of CEP patterns that are commonly used in rules.
 - **Windows:** We can assign *windows* to patterns, restricting their scope. Windows could refer to specific time intervals (time windows) or to number of occurrences of particular events (length windows). Moreover, we can distinguish two types of windows: batch and sliding. The first ones have fixed starting and ending points. The latter ones move the interval so its ending point coincides with the current position of the pointer traversing the stream (e.g., the current time).
 - **Temporal sequencing of events:** A key CEP operator is *followedBy* (“->”), which introduces a temporal ordering between two events. Events related by this operator do not need to be consecutive: “A -> B” only implies that

event A occurs some time before B, i.e., the timestamp of A precedes that of B.

- **Pattern combination:** Patterns can be combined in different ways by using logical operators (or, and, etc.) and temporal connectors (until, while, etc.), among others. Negation is also possible, representing the fact that an event has not happened. In addition, windows can be combined, restricting their scope.
- **Higher-order events.** Higher-order derived events can be defined when the pattern that specifies the event makes use of other previously defined derived events.

Such elements are indeed very close to those of in-place model transformation rules, with two main additions. First, given that the stream can be considered infinite, or at least too large to be handled in full, CEP patterns use the concept of *window* to restrict their scope. Therefore all matches are *local*, within the scope defined by the window. Second, all events have a timestamp and therefore they can be (totally) ordered; in other words, the stream can be considered as a linear sequence of ordered events. Relationships between the events are not explicitly modeled, they are normally specified either by time-order information, or by name/id matching.

2.2 Apache Spark

Apache Spark is a general-purpose cluster computing platform that extends the popular MapReduce model to efficiently support more types of computations, including interactive queries and stream processing [17]. One of the main features Spark offers for speed is the ability to run computations in memory, while being more efficient than, e.g., MapReduce for complex applications running on disk.

Spark combines different processing types and it is designed to be fast, highly accessible, and providing simple APIs in different programming languages such as Python, Java, Scala, and SQL, and a set of rich built-in libraries. It also integrates smoothly with other Big Data tools.

Spark is built on the concept of distributed datasets, which contain arbitrary Java, Python or Scala objects, upon which parallel operations can be efficiently applied. The Spark API is built around the concept of a *resilient distributed dataset* (RDD), which is a fault-tolerant collection of elements that can be operated on in parallel. The RDD API supports two types of operations: transformations, which define a new dataset based on previous ones, and actions, which execute operations on a cluster. RDDs can be created in two ways: parallelizing an existing collection, or referencing a dataset in an external storage system (e.g. as a shared filesystem, HDFS, HBase, or any other data source offering a Hadoop InputFormat).

Spark also provides different tools that are built as libraries on top of Spark, such as *Spark Streaming*, which enables processing of live streams of data; *Spark SQL*, for working with structured data; *Mlib*, containing common machine learning functionality, and *Graphx* for manipulating graphs.

2.3 Graph processing systems

Graph structures are commonly used for representing the data of models, and for performing patterns on them. Graphs are composed of nodes, edges and properties for this representation. Nodes

S4: A user, with an *h-index* higher than 50¹, posts three tweets in a row in the last hour containing a hashtag that describes a photo. In this case, we generate a `NiceTwitterPhoto` object.

S5: A Flickr user favorites a photo identified as `NiceTwitterPhoto` but none of her followers do. If so, we generate a `Misunderstood` object identifying the user.

S6: Taking into account the 10,000 most recent tweets, a user with *h-index* higher than 70 and more than 50K followers who writes a tweet raises an `InfluencerTweeted` event, pointing to the user.

To record such new objects (that correspond to CEP derived events), we have included in the metamodel 6 extra classes (shaded in gray). Therefore, a model of this application will contain both the information from the sources (Twitter and Flickr) and also the information we generate during its lifetime, as the system evolves.

Note that in order to respond to these patterns, we need to take into consideration not only the more static information about the social network users, but also the dynamic and rapidly changing information appearing such as tweets or likes. Some situations (S4 and S5) also imply traversing the model through its relationships, including the closure of a relationship (e.g., `Follow`). This means that the search space can be extremely high. Other situations (S5) require negations, which also imply searches in a very large model.

3.2 Extending CEP elements and operators

In this subsection we present our proposal for extending CEP elements and operators to deal with graph-structured information, such as models.

3.2.1 Extending the structure. In order to extend the structure of CEP we make use of conceptual modeling concepts for representing information items and their relationships, i.e., our information system will be represented by a model, and not just by a sequence of events. Elements of such a model that represent CEP events will have an attribute, *timestamp*, with information about the moment in time at which the event occurs.

Apart from the relationships defined among the nodes, a derived directed relation `FollowedBy` can be defined between the model elements (i.e., graph nodes) with timestamps if the timestamp of the first element is less than that of the second. In this way, a CEP event stream is just a graph with one single dimension.

3.2.2 Vicinity: extending CEP Windows. In CEP, windows could refer to specific time intervals or to number of occurrences of particular events. There are two types of CEP windows: batch and sliding. The first ones have fixed starting and ending points. The latter ones move the interval so its ending point coincides with the current position of the pointer traversing the stream (e.g. the current time). In other words, CEP windows permit defining *segments* of the linear stream of events that determine the *local* context of a match.

When the linear nature of the event stream is generalized to a graph, instead of stream segments (time or length windows) we can use *vicinity graphs* [10], which are subgraphs composed of elements related to one or more objects. They represent the *neighborhood*

of these objects, and they can be expressed by means of graph patterns.

In this way, we can talk about *spatial* vicinity, in addition to the *temporal* vicinity used in CEP. Then, given one object, its vicinity will be composed of other objects related to it by links (directly, or through a sequence of consecutive links, that we shall call *hops*), i.e., its nearby objects. To determine the window that establishes the local context of a rule, we can simply define the vicinity of the objects in the rule.

There are different strategies to define the vicinity graph of an object, depending on how we navigate through the graph structure, and the goal we pursue. Representative examples of algorithms for creating relevant vicinity graphs of nearby objects are used for finding related pages in the WWW [10, 18]. These algorithms use different strategies, e.g., going through the parents and children of a page, and then visiting the children and parents of those—using a backward-forward and forward-backward strategy. We could also traverse the graph moving only forward or backward (as in the case of CEP sliding windows), or using any other traversal strategy: in-breadth, in-depth, topological, hybrid, etc. Traversal could be done through any kind of link, or we could navigate the graph through some selected kinds of relations. The information on the nodes could also be considered when building the vicinity graph, assigning weights to nodes and to links, instead of just bean-counting them.

Note that if we collapse a graph into the one-dimensional sequence that represents a CEP event stream, by focusing only on the timestamp attribute of the objects that represent events, and the `followedBy` derived relation, the concept of vicinity faithfully corresponds to the CEP concept of window, time becoming just one of the possible dimensions by which we can navigate the graph.

3.2.3 Extending the behavior. Given that now we have a model (i.e., a graph), recursive in-place model transformation rules can be naturally used to represent how derived information items (e.g. complex events) are created and added to the system [9], hence being able to extend the traditional CEP behavior. Each CEP rule can be mapped to an in-place model transformation rule that specifies the matching pattern and creates the derived event. Non-timed derived events can also be generated by other MT rules, permitting further situations of interest to be detected as the graph evolves—i.e., new nodes or links are created or destroyed.

4 IMPLEMENTATION

4.1 Architecture

This section describes the prototype implementation that we have developed to validate our proposal, and to analyze its strengths and weaknesses.

In the first place, we have used Apache Spark to process the models. In order to create the graph dataset, it is necessary to use one Spark RDD for the vertices and another one for the edges, thus storing the simple events and the relations among them in two RDDs. The graph dataset is populated using the corresponding APIs from the different sources; in our example, the Twitter and Flickr APIs. A thread is in charge of obtaining the batches of raw

¹In Twitter, a user has an *h-index* of *h* if she has at least *h* followers who have more than *h* followers each.

```

def hotTopic {
  // Filters photos, tweet, hashtags and incoming links to hashtags
  val hashtagTweetPhotoFilter = graph.subgraph(
    vpred = (id, attr) =>
      ( attr.isInstanceOf[Hashtag]
        || (attr.isInstanceOf[Tweet] && (System.currentTimeMillis() - attr.asInstanceOf[Tweet].date) < 3600000)
        || (attr.isInstanceOf[Photo] && (System.currentTimeMillis() - attr.asInstanceOf[Photo].dateLastUpdate) < 3600000)
      ),
    epred = t =>
      t.toString().contains("contains") || t.toString().contains("tags")
  )
  // Computes incoming links to each node
  val hashtagTweetPhotoFilterInDegrees =
    hashtagTweetPhotoFilter.outerJoinVertices(hashtagTweetPhotoFilter.inDegrees) {
      (id, _, counter) => counter match {
        case Some(counter) => counter
        case None           => 0 // No inDegree means zero inDegree
      }
    }
  // Gets Hashtags with more than 100 incoming links. Result: HotTopic graph with Hashtags
  val graphHotTopicAux = hashtagTweetPhotoFilterInDegrees.subgraph(
    vpred = (id, attr) => attr >= 100).mapVertices((id, attr) =>
    (System.currentTimeMillis()))
  //Result
  verticesHotTopic = (graphHotTopicAux.vertices.map(rdd =>
    (rdd._1, new HotTopic(rdd._1, rdd._2))) ++ verticesHotTopic).reduceByKey((a, b) =>
    if (a.currentTimestamp > b.currentTimestamp) {a} else {b})
    .filter { v =>
      (System.currentTimeMillis() - v._2.asInstanceOf[VertexProperty].currentTimeStamp) < 3600000
    }
  }
}

```

Figure 2: Scala code for the HotTopic event.

data, give them the proper format of RDD, and store them in the dataset.

CEP patterns are implemented in terms of Spark and Graphx functions written in Scala—which perform the matching process of the rule— followed by the generation of the new objects—, which correspond to the derived elements. We use one dedicated thread per rule that is continuously running, and each rule has a corresponding RDD to store the derived events.

In order to represent the objects in the graph dataset we need to distinguish between *transient* and *permanent* information. The former refers to the time-sensitive information coming from sources with high-volume of data that do not need to persist in memory for a long time. Basically, this kind of information corresponds to traditional CEP data streams; once they are analyzed by the CEP rules, they can be discarded. In our example, these can be the tweet feeds. Permanent (or static) information corresponds to data which is stable in time, and that we want to persist in the database for longer periods. For instance, the information about the Twitter and Flickr users, or the Flickr photos, in our example.

Both kinds of data are stored in the dataset, although the manner in which their lifecycles are managed is different. Objects representing permanent information are stored as single nodes in the graph dataset, and relationships among them as relationships in the dataset, too. Transient data (i.e., events) need to be processed as soon as possible by the CEP rules, and each rule should try to be as efficient as possible. Then, we simulate a time window and every object that represents a time-sensitive event has a timestamp and it will be removed when its timestamp expires. The time window value will be set by the rule with the biggest time window, 60 minutes in our example.

When appropriate, our implementation can consider complex events like transient data too, so they will have a timestamp attribute and will be removed when they expire. However, if the same complex event is detected in a time window several times, for efficiency reasons its timestamp will be updated instead of duplicating the event, or discarding the last detection.

4.2 Writing the patterns for the example

As mentioned above, CEP patterns can be implemented by means of in-place model transformation rules, which in turn we express in this version as Spark patterns, and then creation of the resulting object(s).

To illustrate how these patterns are expressed, in this section we show some of the patterns that correspond to the situations described in the running example. The rest of them, together with the rest of the project's models, artifacts and programs, can be downloaded from [1].

In the first place, pattern **S1** generates a new HotTopic object in the dataset every time a hashtag is used by both Twitter and Flickr users at least 100 times in the last 60 minutes. This can be expressed in terms of a Spark pattern as the Listing in Fig. 2 shows.

In the first part, the rule selects groups of three objects (Hashtag, Tweet, and Photo) that are related by the incoming Hashtag associations (contains and tags). Then, our code computes the incoming links for each node, and selects those results that have more than 100 incoming links. Finally, it creates a HotTopic object for each matching, and stores it in a Spark RDD with a timestamp and the id of the hashtag attribute.

Note how in the last instruction, the pattern checks if the new detected event is in the RDD and in that case updates the dataset

with the new event. In turn, the last filter discards all `HotTopic` events created more than one hour ago (3600000 milliseconds), as mentioned in the description of the event (Sect. 4.1).

Pattern **S4** (whose code is not listed here) requires a user, with an *h-index* higher than 50, to post three tweets in a row containing a hashtag that corresponds to a photo in the last hour. If such a situation is detected, a `NiceTwitterPhoto` object identifying the photo is created. Observe how this pattern specifies that the user publishes a tweet that *contains* a hashtag. In this case we use a 2-hops *spatial* window, meaning that we will traverse all objects of class `TwitterUser` that have a direct relationship to class `Tweet`, and the `Tweet` class is directly connected with class `Hashtag`, but there is not a direct relation between `TwitterUser` and `Hashtag` classes. Additionally, the pattern checks that the hashtag *describes* a photo, therefore it uses 3 hops since it will traverse relation *describes* or *tags* in the metamodel.

In pattern **S5**, a Flickr user favorites a photo which has been identified as `NiceTwitterPhoto`, but no other of her followers does. In this case, we want to generate a `Misunderstood` object identifying user. This pattern shows an example of a NAC (Negative Application Condition) and complex event from another complex event.

Finally, pattern **S6** creates `InfluencerTweeted` events considering the 10000 most recent tweets. For this we need a length window of tweets and an aggregation function for calculating the users' *h-index*.

As a final remark, note that in this paper we are just concerned with the expression of such patterns, not on how to determine the optimal size of the spatial or temporal windows. These are decisions that depend on the problem domain and on the particular patterns, and therefore they fall outside the scope of this paper.

5 DISCUSSION

Once we are able to write CEP patterns in terms of Spark ones, this section discusses some aspects of our proposal, as well as some of its main benefits and limitations.

5.1 Performance

5.1.1 Comparison with CEP solutions. Performance is a key issue in data processing systems, specially when dealing with high volumes of information. CEP systems are optimized for this, and permit processing large streams of events very efficiently. Apart from the internal optimizations to handle the events in parallel, they normally use in-memory structures for storing the events being processed. This is possible because every CEP pattern only deals with a substream of events (those inside the pattern window) whose size is bounded.

However, in our case we need to deal with much larger models, because the graph should contain not only the portion of the *transient* data, but also the application's *static* information—note that we cannot deal with the events independently, because one of the assumptions in our case is that they are highly connected to other objects. Hence, the performance of the kinds of applications that we have in mind can not be expected to be as optimal as conventional CEP systems. However, we are interested in knowing the performance penalty that we obtain if we had to deal with similar

data, i.e., our overhead. In other words, we are interested in the following research question:

R1. Is the performance of our approach similar to the performance of existing off-the-shelf CEP engines, when dealing with just streams of information (i.e., with no graph structure)?

To respond to this question we have conducted one experiment, consisting on replicating one typical CEP application using our approach, and measuring the overhead that we obtain in the response times with regards to those obtained by the CEP system.

We conducted two versions of this experiment. The first one consisted in sending a sequence of events to our system and to the CEP system, every time they occurred. In this was we were simulating a realistic behavior of the applications. Secondly, we sent the same sequence, but all at once, i.e., without any delay between the events. All tests were run on a machine whose operating system is Ubuntu 64 bits, with 64Gb of RAM memory, and an Intel Xeon CPU E5-2680 processor with 16 cores of 2.7 GHz. We used Spark version 2.3.0 in our implementation, and Esper version 4.7.0 [11] as the CEP engine.

In CEP, performance is mainly dependent on the window size defined for the patterns. Therefore we checked with different window sizes (note that, as mentioned in Section 4.1, only events with timestamps inside the windows are maintained in the dataset).

The example used to conduct the tests² assumes a fleet of motorbikes equipped with sensors that produce real-time information about their state, including the timestamp of the event, the motorbike id, the name of the current location, the speed in *Km/h*, the pressure of the two tyres measured in BARS, and whether the driver is on the seat or not. The following tuples are examples of these simple Motorbike events. Timestamps are expressed using the POSIX time convention, which is roughly the number of seconds that have elapsed since January 1, 1970 [16].

```
Motorbike(1520674009,1,"Seville",100,3.1,3.1,true)
Motorbike(1520674010,1,"Seville",90,3.1,3.1,true)
Motorbike(1520674008,2,"Malaga",62,3.01,3.01,true)
Motorbike(1520674011,1,"Seville",0,3.1,3.1,true)
Motorbike(1520674009,2,"Malaga",70,3.01,3.01,true)
Motorbike(1520674012,1,"Seville",0,3.07,3.07,false)
```

We are interested in monitoring these periodic events, generating derived events when some situations of interest occur. In particular, we are interested in the following CEP patterns:

- **BlowOutTyre:** The pressure of one of the tyres of a moving motorbike goes down from more than 2.0 BAR to less than 1.2 BAR in less than 5 seconds.
- **Crash:** The speed of a motorbike goes from more than 50 km/h to 0 km/h in less than 3 seconds.
- **DriverLeftSeat:** The seat sensor detects that the driver has left the seat.
- **OccupantThrownAccident:** A motorbike suffers a blow out of one of its tyres, then a `Crash` event is detected, and the driver is thrown out, everything within less than 3 seconds.

To show how these patterns are mapped into Spark patterns, Fig. 3 shows the Scala code corresponding to pattern `DriverLeftSeat`, whose expression in Esper is shown in Fig. 4.

²The complete description of the example is available from http://atenea.lcc.uma.es/index.php/Main_Page/Resources/CEP

```
def driverLeftSeat(){
  // Group events by motorbikeId
  val groupMotorbike = graph.vertices.groupBy(attr => attr._2.asInstanceOf[MotorbikeEvent].motorbikeId)
  // Sort events by timestamp
  val groupMotorbikeOrder = groupMotorbike.map(f =>
    (f._1, f._2.toList.sortBy(f => f._2.asInstanceOf[MotorbikeEvent].currentTimestamp))).collect
  // Add new events to dataset DriverLeftSeat
  verticesDriverLeftSeat = verticesDriverLeftSeat ++ groupMotorbikeOrder.flatMap(f =>
    isDriverLeftSeat(f._2))
  val driverLeftSeatVertices = sc.parallelize(verticesDriverLeftSeat).zipWithIndex().map(r =>
    (r._2.asInstanceOf[VertexId], r._1))
  //Filter events that are in the defined window
  verticesDriverLeftSeat = verticesDriverLeftSeat.filter { attr =>
    (System.currentTimeMillis() - attr.currentTimeStamp) < 3000 }
  graphDriverLeftSeat = Graph(driverLeftSeatVertices, sc.parallelize(ListBuffer[Edge[(Long)]]() ))
}
//Function for checking complex event
def isDriverLeftSeat(list: List[(VertexId, VertexProperty)]) : ListBuffer[(VertexProperty)] = {
  val logger: Logger = LoggerFactory.getLogger(this.getClass)
  var result = ListBuffer[(VertexProperty)]()
  for (a <- list; b <- list) {
    if (a._2.asInstanceOf[MotorbikeEvent].seat == true
      && b._2.asInstanceOf[MotorbikeEvent].seat == false
      && b._2.asInstanceOf[MotorbikeEvent].currentTimestamp >
        a._2.asInstanceOf[MotorbikeEvent].currentTimestamp
      && sc.parallelize(verticesDriverLeftSeat).filter(attr =>
        attr.asInstanceOf[DriverLeftSeatEvent].currentTimestamp1 ==
          a._2.asInstanceOf[MotorbikeEvent].currentTimestamp
      && attr.asInstanceOf[DriverLeftSeatEvent].currentTimestamp2 ==
        b._2.asInstanceOf[MotorbikeEvent].currentTimestamp).count == 0) {
      val driverLeftSeatEvent = new DriverLeftSeatEvent(
        System.currentTimeMillis(),
        a._2.asInstanceOf[MotorbikeEvent].motorbikeId,
        b._2.asInstanceOf[MotorbikeEvent].location,
        a._2.asInstanceOf[MotorbikeEvent].seat,
        b._2.asInstanceOf[MotorbikeEvent].seat,
        a._2.asInstanceOf[MotorbikeEvent].currentTimestamp,
        b._2.asInstanceOf[MotorbikeEvent].currentTimestamp)
      result += driverLeftSeatEvent
    }
  }
  return result
}
```

Figure 3: GraphX/Scala code for the DriverLeftSeat pattern.

```
@Name('DriverLeftSeat')
insert into DriverLeftSeat
select current_timestamp() as timestamp,
       a2.motorbikeId as motorbikeId,
       a2.location as location
from pattern [every
  (a1 = Motorbike(a1.seat = true)
  -> a2 = Motorbike((a2.seat = false)
  and (a1.motorbikeId = a2.motorbikeId))]
```

Figure 4: Esper code for the DriverLeftSeat pattern.

The performance figures obtained for experiments are shown in Tables 1 and 2.

Table 1 compares the results of the Esper system and our implementation for a different number of input elements, using a realistic simulation between motorbike events whereby each motorbike generates one event every second. Both systems are able to process the events with almost no delay.

In turn, Table 2 compares results of Esper system and our implementation for a different number of input elements, when we feed the complete stream of events to the system without any delay between them. We can see that our system takes approximately 6

Table 1: Performance Figures for the Motorbike Example with stream simulation (in seconds).

#Events	5K	10K	20K	30K
Esper	5,001	10,001	20,001	30,002
Spark	5,001	10,002	20,001	30,001

Table 2: Performance Figures for the Motorbike Example without stream simulation (in milliseconds).

#Events	5K	10K	20K	30K
Esper	1,214	2,208	3,610	5,386
Spark	31,039	60,472	119,038	176,424

milliseconds to process each event, whereas Esper only takes some nanoseconds per event.

In any case, as mentioned above we do not try to compete with Esper nor other dedicated CEP systems, because in order to be able to deal with graph-structured information and high volumes of data, we have to pay the price of storing a large amount of data. Even so, for some kinds of applications where the event occurrence

rate is not very high (e.g., more than one event per millisecond), the response time of our solution is acceptable.

5.1.2 Absolute performance. Apart from estimating the potential overhead of our solution when compared to pure CEP systems, we were also concerned about the response times of our prototype implementation.

R2. Is the performance of our approach acceptable for dealing with large datasets?

Table 3 shows the response times we have obtained when running the Flickr and Twitter example with different patterns and datasets of different sizes. Columns headers indicate the number of nodes/associations of the models used in each test. For example, the first one consisted on 127,000 objects and 65 millions of associations (links) among them, which accounts for a total of 6.627 millions of elements in the graph. Columns 2 and 3 correspond to models with 14.26 and 28.554 million elements, respectively. The high number of links between the model objects reflects the high number of possible relations defined in the metamodel (Fig. 1). Executions were carried out using the same hardware as we did for the motorbike experiments. To conduct this experiment we designed a program that populates our dataset with a specific number of events (coming from synthetic data previously created) and then run the patterns in parallel. The complete code for these patterns, and the rest of the artefacts, can be found at [1].

In general, most results are in the range of milliseconds, something that can be acceptable for applications with a non very high event occurrence rate, as explained before.

Furthermore, in a previous approach we used neo4j and Cypher, respectively, to represent the graph dataset and to express the queries. Although the usability of that approach was much better than the one we obtain with GraphX (see next section), the performance was not acceptable: on the one hand, the response times were of two orders of magnitude with respect to the ones we get with GraphX (minutes instead of milliseconds); and on the other hand, the number of nodes we were able to handle in reasonable time were below 20K (whilst now we are able to deal with models of 28 millions elements in a reasonable time). This is why we decided to abandon the Neo4j path and to use Spark.

5.2 Expressiveness

In addition to getting an initial assessment of the performance of our proposal, we also wanted to check its expressiveness for describing the kinds of queries that we initially devised.

R3. How expressive is our proposal with regard to CEP, i.e., can we write all kind of CEP patterns with GraphX?

We still need to perform a detailed analysis, to see if all CEP patterns can be expressed with our proposal, but initially GraphX is based on a general purpose and very expressive language like Scala, so we do not foresee major obstacles here. In particular, after having written the patterns for the case studies used to validate our proposal, and learned how they can be expressed using GraphX.

R4. How easy is it to use our proposal, i.e., how easy is it to write our patterns with GraphX?

```

match
  (m1{seat:'true'}), (m2{seat:'false'})
where
  m1.id=m2.id and (m2.timestamp-m1.timestamp)=1
create
  (m1)-[:EVENT]-(:DriverLeftSeat{
    timestamp:m2.timestamp,
    id:m2.id,
    location:m2.location})-[:EVENT]->(m2)

```

Figure 5: Cypher code for the DriverLeftSeat pattern.

This is what we consider one of the weakest aspects of our proposal at this moment. Writing GraphX patterns in Scala is not a trivial task. For instance, compare the code shown in figures 3 and 4, corresponding to the DriverLeftSeat pattern expressed in GraphX/Scala and in Esper, respectively. It is easy to see the benefits of using a dedicated language such as Esper, specifically tailored to express these kinds of patterns.

Something that we also realized when writing the patterns is that current model transformation languages may have difficulties specifying spatial windows in a natural manner. This is not the case for temporal windows, since we could use the object's timestamps to restrict the selection and matching processes. However, graph database languages such as Cypher or Gremlin provide very interesting mechanisms for establishing spatial windows, and greatly simplify the expression the patterns we are interested in. For example, Figure 5 shows the DriverLeftSeat pattern expressed in Cypher.

We are working (as part of our future work) on an automatic mapping from Cypher to GraphX, which will allow users to write their patterns in Cypher, and then get the corresponding GraphX queries. Note that Cypher has a very special feature for representing spatial windows in a very easy way, since it is possible to specify both a specific number of hops (over one type, or over *any* type of relation in the graph), or a bound on it, in the query [25]. Extending model transformation languages with these mechanisms could be an interesting line of research, too.

5.3 Threats to Validity

In this subsection, we elaborate on several factors that may hinder the validity of our results.

Internal validity — are there factors which might affect the results in the context the case study?

Concerning the performance figures, we would need to experiment with more case studies to confirm that the processing times of our proposal are acceptable when dealing with very large datasets of different kinds. Of course, we can never compete with Esper or any other CEP engine because they only deal with *transient* information. In the kinds of applications we are interested in, the *permanent* data needs to be managed, too, and therefore we need to be able to store it and navigate through the model's relationships in the patterns. This imposes some performance penalties, but we have seen how our proposal has much better performance results than other solutions based on graph databases, and the results we obtain are still acceptable.

The performance figures obtained also depend on the concrete expression of the patterns. Therefore, there might be alternative

Table 3: Performance Figures for the Twitter and Flickr example, for each pattern (in milliseconds).

	127K/6500K	260K/14000K	554K/28000K
HofTopic	62.00	45.00	35.66
PopularTwitterPhoto	69.00	77.00	68.00
PopularFlickrPhoto	76.67	51.33	55.67
NiceTwitterPhoto	108.00	101.66	107.67
Misunderstood	46.33	57.33	53.66
InfluencerTweeted	226.00	334.33	424.33

ways of writing the same patterns which would yield better response times. Anyway, our current figures are reasonable enough to show the feasibility of our proposal.

With respect to the expressiveness of our proposal, we would need to conduct more studies, although our initial experiments have shown that the use of a general-purpose language such as Scala and GraphX has the expressiveness required to represent all kinds of patterns. There is the issue of how easy it is to write queries with our proposal. This is another line of research that we need to address, because the more complex the patterns are, the more difficult they become to debug and prove correct.

External validity — to what extent is it possible to generalize our findings in general?

So far, we cannot claim any performance results outside the context of the presented case studies. Nevertheless, the evaluation method used in the case study can indeed be applied on other examples as well. In addition, the case study may be repeated on other hardware platforms to analyze, e.g., the impact of the number of cores or the memory size on the performance.

6 RELATED WORK

Our work is mainly related to CEP systems, and probably the closest work to ours is [9], where the authors use VIATRA to implement a very efficient CEP system using model transformations techniques. We depart from that work in three main aspects. First, their focus is on traditional CEP systems, whereby the information to process is not graph-structured but a sequence (stream) of events. Second, the size of the models used in [9] was rather limited, probably due to the way in which models were stored and handled. Our current solution permits dealing with much larger datasets, because of the use of the Spark technologies. Finally, they do not consider the idea of spatial windows because there is not need for it in these kinds of stream processing applications.

The use of spatial windows in models have been already suggested for dealing with infinite [2] or streaming models [3], where only a portion of elements (e.g., the ones inside a *sliding window*) is available at any given moment in time, or in the case of *approximate* model transformations [24]. However, these approaches did not satisfactorily resolve one important issue: they had to consider that some connections between model elements could be broken because not all the elements within a window are available at the same time, and therefore the models could not be correct. Here we have used the concept of vicinity graph, which permits tackling these issues, although also with some associated costs: as in CEP systems, we trade efficiency for accuracy, given that decisions

are made based on the local context of the query, and not on the complete model.

Regarding the use of graph databases for storing large models, NeoEMF [7] provides a multi-database model persistence framework for very large models, and [8] defines a language to perform OCL queries on graph databases that outperforms (in terms of memory footprints and time) other existing solutions. In addition, [6] defines a mapping from ATL to Gremlin, which enables model transformations on large models stored in graph databases. As mentioned above, we tried with neo4j and also with Gremlin, but the results were not good enough when compared to CEP systems.

7 CONCLUSIONS AND FUTURE WORK

This work has explored the extension of CEP concepts and mechanisms to the real of models that contain graph-structured information, beyond sequential streams of events. The concept of *vicinity* has been introduced, and a prototypical solution has been developed to validate the proposal and to analyze its advantages and limitations.

The proposal has just set the basis and initial ideas for this line of work, and opens the path for different research alternatives. In particular, we are interested in investigating what kinds of changes in the graph structure, patterns and in the architecture of our prototypical implementation would serve to optimize the response time of our current solution. Furthermore, a possible next step for improving the response time could be implementing Spark Streaming in our system, and combining it with Graphx and Spark SQL to optimize the patterns.

We are also planning to develop a mapping from Cypher to GraphX, so that patterns could be expressed in Cypher, which is much more compact, easy to use, and specific to this type of queries, and then translated into the more efficient GraphX queries.

Regarding the concept of vicinity, we are planning to extend MT transformation languages, and in particular ATL, with the concept of local context of a query (i.e., vicinity), and then generate the queries, in a similar way to what [6] do. Moreover, we would like to study how different strategies (in the style of [10, 18]) could be defined to improve the accuracy of the results, and how to measure the error we are incurring by selecting a given spatial window when specifying rules that contain NACs.

Finally, many more experiments and case studies are need to better assess the expressiveness and performance of our proposal, and to improve its current results. Although our initial results are currently encouraging, there is room for improvement in several aspects, as we have discussed in this paper.

REFERENCES

- [1] Gala Barquero, Loli Burgueño, Javier Troya, and Antonio Vallecillo. Last accessed: May 2018. *The Graph-CEP Project*. <http://atenea.lcc.uma.es/projects/GraphCEP.html>.
- [2] Benoit Combemale, Xavier Thirioux, and Benoit Baudry. 2012. Formally Defining and Iterating Infinite Models. In *Proc. of MODELS 2012 (LNCS)*, Vol. 7590. Springer, 119–133.
- [3] Jesús Sánchez Cuadrado and Juan de Lara. 2013. Streaming Model Transformations: Scenarios, Challenges and Initial Solutions. In *Proc. of International Conference on Model Transformation (ICMT 2013) (LNCS)*, Vol. 7909. Springer, 1–16.
- [4] Gianpaolo Cugola and Alessandro Margara. 2012. Processing flows of information: From data stream to complex event processing. *ACM Comput. Surv.* 44, 3 (2012), 15:1–15:62.
- [5] Gianpaolo Cugola, Alessandro Margara, Mauro Pezzè, and Matteo Pradella. 2015. Efficient analysis of event processing applications. In *Proc. of DEBS'15*. ACM, 10–21.
- [6] Gwendal Daniel, Frédéric Jouault, Gerson Sunyé, and Jordi Cabot. 2017. Gremlin-ATL: a scalable model transformation framework. In *Proc. of ASE'17*. IEEE Computer Society, 462–472.
- [7] Gwendal Daniel, Gerson Sunyé, Amine Benellallam, Massimo Tisi, Yoann Vernageau, Abel Gómez, and Jordi Cabot. 2017. NeoEMF: A multi-database model persistence framework for very large models. *Sci. Comput. Program.* 149 (2017), 9–14. <https://doi.org/10.1016/j.scico.2017.08.002>
- [8] Gwendal Daniel, Gerson Sunyé, and Jordi Cabot. 2016. Mogwai: A framework to handle complex queries on large models. In *Proc. of RCIS'16*. IEEE, 1–12.
- [9] István Dávid, István Ráth, and Dániel Varró. 2018. Foundations for Streaming Model Transformations by Complex Event Processing. *Software and System Modeling* 17, 1 (2018), 135–162. <https://doi.org/10.1007/s10270-016-0533-1>
- [10] Jeffrey Dean and Monika R. Henzinger. 1999. Finding Related Pages in the World Wide Web. *Comput. Netw.* 31, 11-16 (May 1999), 1467–1479.
- [11] EsperTech. [n. d.]. Esper - Complex Event Processing. <http://www.espertech.com/esper/>. (accessed 18 November 2017).
- [12] Opher Etzion and Peter Niblett. 2010. *Event Processing in Action*. Manning Publications.
- [13] Event Processing Technical Society. 2011. *Event Processing Glossary, Version 2.0*. http://www.complexevents.com/wp-content/uploads/2011/08/EPTS_Event_Processing_Glossary_v2.pdf.
- [14] Joseph E. Gonzalez, Yucheng Low, Haijie Gu, Danny Bickson, and Carlos Guestrin. 2012. PowerGraph: Distributed Graph-parallel Computation on Natural Graphs. In *Proc. of OSDI'12*. USENIX Association, Berkeley, CA, USA, 17–30.
- [15] Joseph E. Gonzalez, Reynold S. Xin, Ankur Dave, Daniel Crankshaw, Michael J. Franklin, and Ion Stoica. 2014. GraphX: Graph Processing in a Distributed Dataflow Framework. In *Proc. of the 11th USENIX Conference on Operating Systems Design and Implementation, OSDI'14*. USENIX Association, Berkeley, CA, USA, 599–613.
- [16] IEEE Std 1003.1-2008. 2016. *The Open Group Base Specifications. Issue 7, Sect. 4.16, Seconds Since the Epoch*.
- [17] Holden Karau, Andy Konwinski, Patrick Wendell, and Matei Zaharia. 2015. *Learning Spark*. O'Reilly.
- [18] Jon M. Kleinberg. 1999. Authoritative Sources in a Hyperlinked Environment. *J. ACM* 46, 5 (Sept. 1999), 604–632.
- [19] Jure Leskovec, Kevin J. Lang, Anirban Dasgupta, and Michael W. Mahoney. 2009. Community Structure in Large Networks: Natural Cluster Sizes and the Absence of Large Well-Defined Clusters. *Internet Mathematics* 6, 1 (2009), 29–123.
- [20] David C. Luckham. 2002. *The Power of Events: An Introduction to Complex Event Processing in Distributed Enterprise Systems*. Addison-Wesley.
- [21] David C. Luckham. 2012. *Event Processing for Business: Organizing the Real-Time Enterprise*. Wiley.
- [22] Grzegorz Malewicz, Matthew H. Austern, Aart J.C Bik, James C. Dehnert, Ilan Horn, Naty Leiser, and Grzegorz Czajkowski. 2010. Pregel: A System for Large-scale Graph Processing. In *Proc. of the 2010 ACM SIGMOD International Conference on Management of Data, SIGMOD'10*. ACM, New York, NY, USA, 135–146. <https://doi.org/10.1145/1807167.1807184>
- [23] L. Page, S. Brin, R. Motwani, and T. Winograd. 1998. *The PageRank Citation Ranking: Bringing Order to the Web*. Technical Report. Stanford Digital Library Technologies Project. citeseer.ist.psu.edu/page98pagerank.html
- [24] Javier Troya, Manuel Wimmer, Loli Burgueño, and Antonio Vallecillo. 2014. Towards Approximate Model Transformations. In *Proc. of the Workshop on Analysis of Model Transformations (AMT'14) (CEUR Workshop Proceedings)*, Vol. 1277. 44–53. <http://ceur-ws.org/Vol-1277/5.pdf>
- [25] Jim Webber, Ian Robinson, and Emil Eifrem. 2013. *Graph Databases*. O'Reilly Media.