

# Compliance with Environmental Regulations through Complex Geo-Event Processing

Federico Herrera, Laura González, Daniel Calegari, Bruno Rienzi  
 Universidad de la República, Facultad de Ingeniería,  
 Montevideo, Uruguay, 11300,  
 {fherrera,lauragon,dcalegar,brienzi}@fing.edu.uy

## Abstract

In a context of e-government, there are usually regulatory compliance requirements that support systems must monitor, control and enforce. These requirements may come from environmental laws and regulations that aim to protect the natural environment and mitigate the effects of pollution on human health and ecosystems. Monitoring compliance with these requirements involves processing a large volume of data from different sources, which is a major challenge. This volume is also increased with data coming from autonomous sensors (e.g. reporting carbon emission in protected areas) and from citizens providing information (e.g. illegal dumping) in a voluntary way. Complex Event Processing (CEP) technologies allow processing large amount of event data and detecting patterns from them. However, they do not provide native support for the geographic dimension of events which is essential for monitoring requirements which apply to specific geographic areas. This paper proposes a geospatial extension for CEP that allows monitoring environmental requirements considering the geographic location of the processed data. We extend an existing platform-independent, model-driven approach for CEP adding the geographic location to events and specifying patterns using geographic operators. The use and technical feasibility of the proposal is shown through the development of a case study and the implementation of a prototype.

**Keywords:** e-government, regulatory compliance, environmental regulations, complex event processing, geographic information systems.

## 1 Introduction

A good management of natural resources and the environment requires the establishment of legal frameworks through laws and regulations. Those laws are continuously evolving in a complex way. Governments are increasingly interested in ensuring compliance with these laws, so the monitoring of variables involved in those laws (e.g. level of air pollution) is necessary.

The expansion of social networks, the Internet of Things (IoT) and the increasing amount of sensors available throughout the territory, has led to large amounts of data related to natural resources and the environment. For example, many organizations use sensors to measure air quality and water pollution, among others. It is also common to find applications that allow ordinary citizens to share, report and provide information, such as reporting the existence of illegal dumps. This information is generally referred to as Volunteered Geographic Information (VGI) [1].

Since the environmental regulatory framework may have rules involving specific geographic areas (e.g. a river, a country, a municipality, etc), it is essential to consider the geographic dimension of data to monitor compliance with these laws.

Complex Event Processing (CEP) technologies allow processing, analyzing and correlating large amounts of heterogeneous data from many sources. This data comes in the form of events for detecting, in real time, relevant and critical situations [2]. While CEP is ideal for monitoring this kind of laws, current CEP technologies rarely provide support for geo-events, i.e. events which includes the geographic dimension (e.g. the point or coordinate in which the event occurred). Similarly, current CEP technologies do not provide native operators to correlate geo-events (e.g. measuring the distance between two coordinates, computing the intersection of two areas).

In this paper we propose a geographic extension to an existing platform-independent, model-driven approach for CEP [3] in order to facilitate monitoring and enforcement of environmental laws. Specifically, the extension allows: (a) defining events with a property of type Geometry (e.g. line, point, polygon, etc.), through which it is possible to specify the geographic location of the event; and (b) specifying complex events using geographic operators (e.g. intersection), which allows correlating events given their geographic locations. The ideas are validated through the development of a case study based on Uruguayan environmental regulations. In addition, we present a prototype using a CEP engine (Esper), which validates the technical feasibility of the solution.

This paper is a substantially extended and thoroughly revised version of [4]. Additional material includes:

- further background information
- a conceptual analysis of environmental regulations and how they can be monitored with CEP
- an extension of the case study and more details of the technological prototype

The rest of this document is organized as follows. In Section 2 we introduce the main concepts, technologies and standards related to this work. In Section 3 we describe the characteristics of environmental regulations and how they can be monitored using a geographic extension for CEP. In Section 4 we show how this extension can be used in the context of a real-world case study and in Section 5 we describe the implementation of a prototype for running such case study. In Section 6 we present some related work. Finally, in Section 7 we present conclusions and an outline of future work.

## 2 Background

In this section we present background information on the three main topics of this work: regulatory compliance, geographic information standards and complex-event processing.

### 2.1 Regulatory Compliance

Nowadays, organizations are faced with an increasing number of requirements originating from different sources, as legislative bodies (e.g. personal data protection laws), standards, regulatory bodies and specific contracts (e.g. service-level agreements), among others [5]. Regulatory compliance mainly refers to being in accordance with a specific set of these requirements [6]. To this end, organizations have to establish internal controls to ensure that these requirements are met, since otherwise they face the risk of litigation and sanctions, and even criminal complaints [7]. This section describes the main activities involved in regulatory compliance management and specific issues concerning the compliance with environmental regulations.

#### 2.1.1 Regulatory Compliance Management

Regulatory compliance management involves different activities, including compliance modeling, compliance monitoring, compliance enforcement, compliance reporting and compliance enactment [8]. Compliance modeling deals with extracting compliance requirements from regulations and representing them in a formal and machine-readable way. Compliance monitoring continuously checks the compliance of systems, services and processes while they are running [9]. This allows the timely detection of non-compliance as well as the possibility to apply reactive and proactive countermeasures [10]. Compliance enforcement comprises performing tasks in order to avoid the violation of a given regulation [8]. Compliance reporting focuses on generating documentation, traces and analytics based on the data generated by the different compliance activities (e.g. compliance monitoring) [8]. Finally, compliance enactment deals with performing activities in order to recover from a non-compliance (e.g. executing a compensation action) or to resolve it (e.g. remove the cause of a non-compliance) [8].

Automating compliance validation presents an opportunity for organizations given that current studies show that, in general, the frequency of compliance audits, monitoring and reporting leads to a more successful compliance management [11]. Automation of compliance requires methods and mechanisms at different abstraction levels ranging from laws and regulations to the supporting software systems [11].

In [11] two approaches for achieving compliance are described: compliance by design and compliance by detection. The "by design" approach is preventative as it focuses on enforcing the correct behavior of systems and on preventing the occurrence of damaging events. This approach comprises refining regulations to the deepest system layer so that, if the refinement is complete and correct, non-compliance becomes technically impossible. In turn, the "by detection" approach is based on monitoring the operation of systems as well as data usage in order to compare the activities being currently executed with the applicable compliance

requirements. A comprehensive solution for automating compliance may benefit from both approaches given that applying only one of the approaches may not be technically or economically feasible [11].

### 2.1.2 Environmental Regulatory Compliance

Environmental regulations comprise guidelines, specifications or laws designed to manage environmental resources and amenities in order to protect the natural environment and mitigate the effects of pollution on human health and ecosystems [12, 13].

The development of environmental regulations has turned into a major activity for most governments. For example, the Environmental Protection Agency (EPA)<sup>1</sup> of the United States of America has the mission to protect human health and the environment. To this end, the agency develops and enforces environmental regulations<sup>2</sup> spanning many environmental topics, from acid rain reduction to wetlands restoration. In addition, the Directorate-General for Environment is the European Commission department in charge of the EU policy on the environment. This department "aims to protect, preserve and improve the environment for present and future generations, proposing and implementing policies that ensure a high level of environmental protection and preserve the quality of life of EU citizens"<sup>3</sup>.

Environmental compliance can be defined as the state of being in accordance with environmental regulations [12, 13]. As the number and complexity of these regulations raise, achieving environmental compliance becomes an increasingly arduous task for governments. In this context, environmental monitoring is critical in order to ensure public safety and to provide information for decision support systems [14]. Environmental monitoring involves activities to continuously observe and regularly measure environmental parameters of specific areas in order to identify environmental changes and aid decision making for the site [15]. Indeed, this activity is crucial for controlling environmental compliance given that it provides the values of different parameters (e.g. pH) that may be included in regulations (e.g. water regulations).

## 2.2 Geographic Information Standards

The Open Geospatial Consortium<sup>4</sup> (OGC) is an international organization integrated by private companies, government agencies, research centres and universities for the promotion of worldwide available open-geospatial information. To this end, OGC develops interface specifications which are publicly available. These specifications, known as OGC standards, provide a framework for developers to create software that enables users to access, visualize, analyze and process geographic information from a variety of sources across a generic computing interface, using Web services.

The Web services architecture defined by OGC (Figure 1) is based on the following principles:

- **Service components are organized into multiple tiers.** The tiers are the Client tier, the Application Services tier, the Processing Services tier and the Information Management Services tier. The tier hierarchy is loose in that a certain tier can bypass the tier that is directly underneath. A service in the Processing tier can also use another service from its same tier.
- **Services are self-describing and chainable:** All services can be discovered and consumed by means of a publish-find-bind interaction typical of Service Oriented Architectures (SOA). Services can be chained with other services.
- **Services use open Internet Standards:** Only widely-used Web standards such as HTTP GET, HTTP POST, URLs, MIME types and XML are required to implement these services. SOAP and WSDL are optional.
- **Service interfaces offer a few static operations.** Web service interfaces provide only a few operations per service specification, either mandatory or optional. Since operations are defined in a standard interface specification, services can be implemented by commercial off-the shelf (COTS) software, such as Internet Map Servers.
- **Services are stateless.** Service operations are stateless (unless specified otherwise), not retaining session data between invocations.

Each tier of services includes multiple specific types of services. The majority of those types of services have an associated OGC Web service standard (e.g. Web Map Service) as is described below. It is worth noting that many OGC standards have also become ISO standards.

<sup>1</sup><https://www.epa.gov/>

<sup>2</sup><https://www.epa.gov/laws-regulations>

<sup>3</sup><http://ec.europa.eu/dgs/environment>

<sup>4</sup><http://www.opengeospatial.org/>

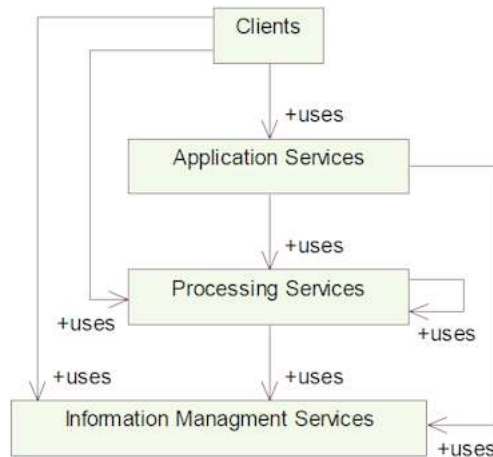


Figure 1: OGC Web Services Architecture [16]

- Application Services.** The Application Services are meant to support basic Clients, especially thin clients such as Web browsers, lifting up some heavy tasks from the client side to the server side. These services allow the users to interact with services from the underlying tiers. Some examples are Web portal services, geographic data discovery services, access control services. etc. Services in this tier do not necessarily follow a standard OGC specification, so they could be implemented as an application-specific SOAP or REST Web service. Nevertheless, they could also be implemented as an OGC Web Processing Service (WPS) since it is the most generic of all the OGC Web services.
- Processing Services.** The Processing Services tier contains services designed to process data. The Web Processing Service<sup>5</sup> (WPS) and the Coordinate Transformation Service<sup>6</sup> (CT) fall under this category.
- Information Management Services.** The Information Management Services are designed to store and provide access to data or metadata, often from multiple datasets. The Web Map Service<sup>7</sup> (WMS), Web Feature Service<sup>8</sup> (WFS), Sensor Observation Service<sup>9</sup> (SOS) and Catalog Service<sup>10</sup> (CSW) are examples of these services.

OGC also defines a data model for spatial objects known as Simple Features Standard (SFS) [17]. SFS specifies an object model in which the Geometry class (Figure 2) is at the base of the hierarchy and can represent planar geometries (up to 2 dimensions). A Geometry is linked to a Spatial Reference System (SRS) that gives meaning to the geometry's coordinates in reference to the surface of the Earth, by specifying the coordinate system, ellipsoid, projection, datum and other relevant geodetic parameters.

The Geometry class definition is shown in Figure 3. Some of its methods allow

- accessing basic information about the object (e.g. dimension, geometry type, and reference system (SRID))
- testing spatial relationships between geometric objects (e.g. operators like crosses, contains and intersects)
- performing some analysis (e.g. operators like union, intersection and distance)

Some widely used concrete subclasses of Geometry are Point, LineString and Polygon. The following examples show how each of them can be represented and stored using the Well-Known Text (WKT) [17] representation in a two-dimensional space.

<sup>5</sup><http://www.opengeospatial.org/standards/wps>

<sup>6</sup><http://www.opengeospatial.org/standards/ct>

<sup>7</sup><http://www.opengeospatial.org/standards/wms>

<sup>8</sup><http://www.opengeospatial.org/standards/wfs>

<sup>9</sup><http://www.opengeospatial.org/standards/sos>

<sup>10</sup><http://www.opengeospatial.org/standards/cat>

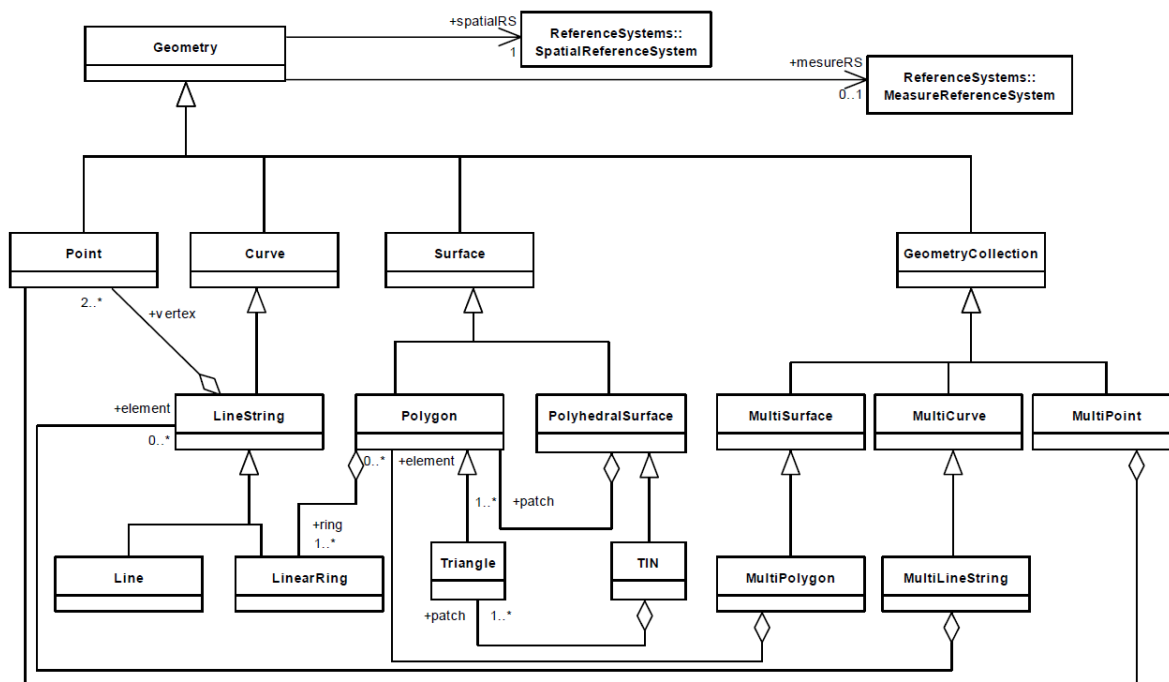


Figure 2: SFS Class Hierarchy [17]

- 'POINT(x y)' is any Point
- 'LINESTRING(x1 y1, x2 y2, x3 y3)' is a three-point LineString
- 'POLYGON ((x1 y1, x2 y2, x3 y3, x4 y4, x1 y1))' is a Polygon with a four-point exterior ring

Since WKT does not provide information of the SRS, the SRS's identifier (SRID) must be provided elsewhere. For instance, if geographic coordinates (latitude and longitude) are used with the WGS84<sup>11</sup> SRS, the SRID will be 4326.

The SFS standard is implemented in many Application Programming Interfaces (API's), such as the Java Topology Suite<sup>12</sup> (JTS) for the Java programming language, and in many Database Management Systems (DBMS) such as PostgreSQL<sup>13</sup> with PostGIS<sup>14</sup> extension.

### 2.3 Complex Event Processing

This section describes the Complex Event Processing (CEP) technology. Section 2.3.1 explains general aspects of CEP and Section 2.3.2 presents a CEP metamodel proposed in [3, 18]

#### 2.3.1 CEP General Aspects

CEP refers to methods, techniques and tools to process simple events as they occur and to permanently and timely derive complex events from a combination of the simple ones [2]. To this end, the flow of incoming events is continuously monitored by event queries which specify situations as a combination of events occurring or not occurring, over time. CEP platforms provide support for various types of event patterns (e.g. logic-based, temporal) that facilitate the specification of event queries [19, 20].

There are multiple approaches and classifications of CEP solutions. Two main approaches exist on how to model the domain: those using ontologies for providing semantics to raw information [21, 22, 23], and those using Model-Driven Development (MDD) techniques providing a model for each domain [24]. Both approaches have pros and cons. In particular, in [25] the author discusses their differences and argues that the MDD option is best suited to design a generic and scalable solution in cases where a model with a large and extensible number of operands and operators is required.

<sup>11</sup><http://gisgeography.com/wgs84-world-geodetic-system/>

<sup>12</sup><http://tsusiatsoftware.net/jts/main.html>

<sup>13</sup><https://www.postgresql.org/>

<sup>14</sup><http://www.postgis.net/>

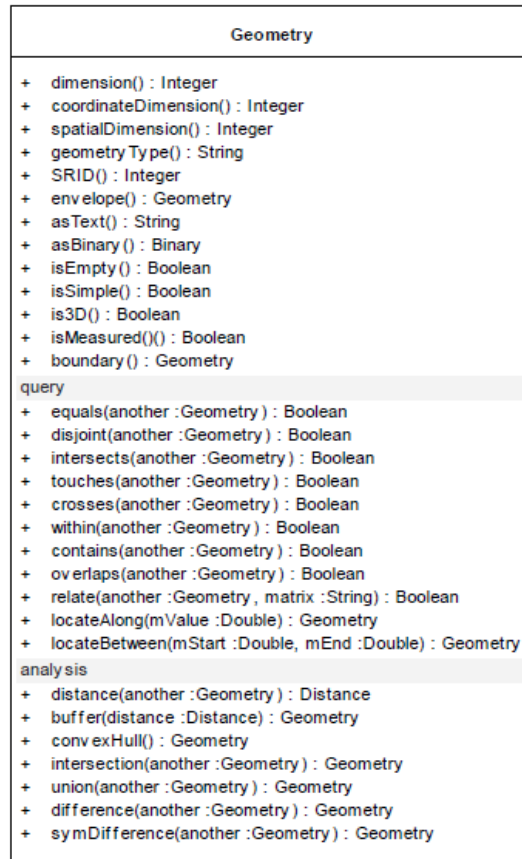


Figure 3: Geometry Class Properties and Methods [17]

Currently, there are various available products that support CEP, e.g. Esper<sup>15</sup>, an open-source CEP engine available in Java and .NET. Esper allows running custom actions when certain conditions on event streams are met. These queries (or patterns) handled by Esper are written using the Event Processing Language (EPL), an SQL-based language allowing the creation of complex event patterns to monitor. For example, assuming that the event called `WaterQualityEvent` represents that new data about water quality arrives, the EPL expression in Figure 4 specifies that we are interested in the moment in which the measurement took place (i.e. `log_time`) and the level of pH of the water (i.e. `ph_water`) when the pattern `ph_water>10` occurs.

```
select x.log_time, x.ph_water
from pattern [every x = WaterQualityEvent (ph_water>10)]
```

Figure 4: Examplng EPL expression

### 2.3.2 CEP Metamodel

In [3, 18] the authors propose a platform-independent, model-driven approach for CEP. The proposal is based on a generic metamodel that can represent any type of event and operators on them. The user must define two different models: one for the domain (the metamodel is depicted in Figure 5) and the other for the CEP patterns (the metamodel is depicted in Figure 6). The proposal is extensible but does not involve a geographic component for the events and patterns.

The CEP domain metamodel (depicted in Figure 5) is composed by the following classes:

- `CEPDomain`. It is the main metaclass. It represents a specific CEP domain which may be composed of at least one type of event (`Event`).
- `Event`. It represents an event within a specific CEP domain. It has a concrete type (`typeName`) and it is composed of properties (`EventProperty`).

<sup>15</sup><http://www.espertech.com/products/esper.php>

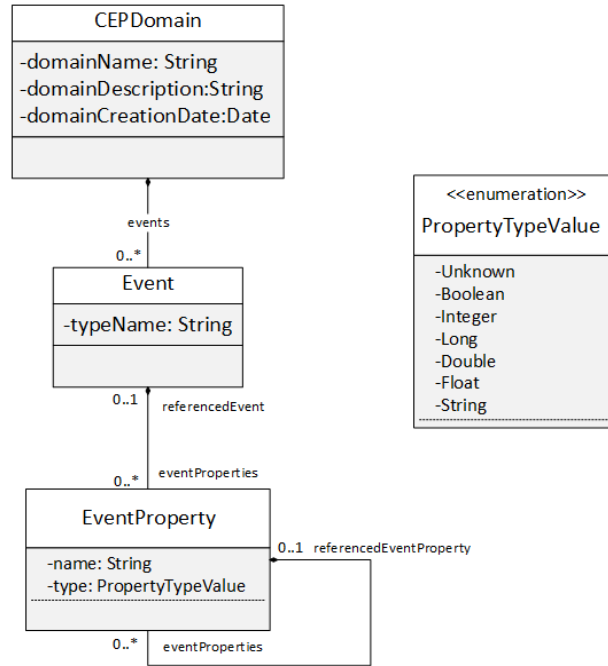


Figure 5: CEP Domain Metamodel [3]

- **EventProperty.** It represents a property or feature of an event. Each property has a name and one of the following types: Unknown, Boolean, Integer, Long, Double, Float or String.

Once the specific CEP domain is defined, it is necessary to define the rules or patterns which must be detected. The patterns can be defined based on the metamodel in Figure 6 (simplified version). The main concepts here are:

- **ComplexEvent:** represents the type of complex event to be created whenever the pattern is detected. Each complex event has a particular type and will be composed of one or more properties. The type name is the name of the event pattern (patternName in CEPEventPattern), thus by determining the name of the event it is possible to infer the pattern that has been detected.
- **CEPEventPattern:** represents the main metaclass. An event pattern is an instance of it. The pattern can be formed by: links (Link) with the rest of the components, elements (EventPatternElement) to define the detection conditions of a pattern, a complex event (ComplexEvent) representing the type of event generated when the pattern is detected, and actions (Action) which are triggered when the pattern is activated.
- **Link:** represents the relation and arity between operands and operators.
- **Operand:** represents a piece of data that is used for performing an operation.
- **Operator:** represents a concrete operation that can be performed for expressing a pattern. There are three different kinds of operators: conditions (ConditionOperator), patterns (PatternOperator) and aggregations (AggregationOperator).

The *Pattern Operand* in Figure 7 allows the representation of different kind of event patterns, for example related to a Timer event. In particular, *WithinTimer* is permanently evaluated to false if the contained pattern expression does not turn to true during the specified time period (years, months, weeks, days, hours, minutes, seconds and milliseconds)

The proposal also provides a graphical tool allowing the modeling of complex events without advanced programming skills. At the implementation level, the tool translates the modeled events into EPL patterns, and evaluate them using Esper.

### 3 Monitoring Environmental Regulations with CEP

Environmental regulations (e.g. laws, decrees), as any other regulation, are usually organized in large documents which embed constraints that apply to specific environmental elements (e.g. water, air, waste)

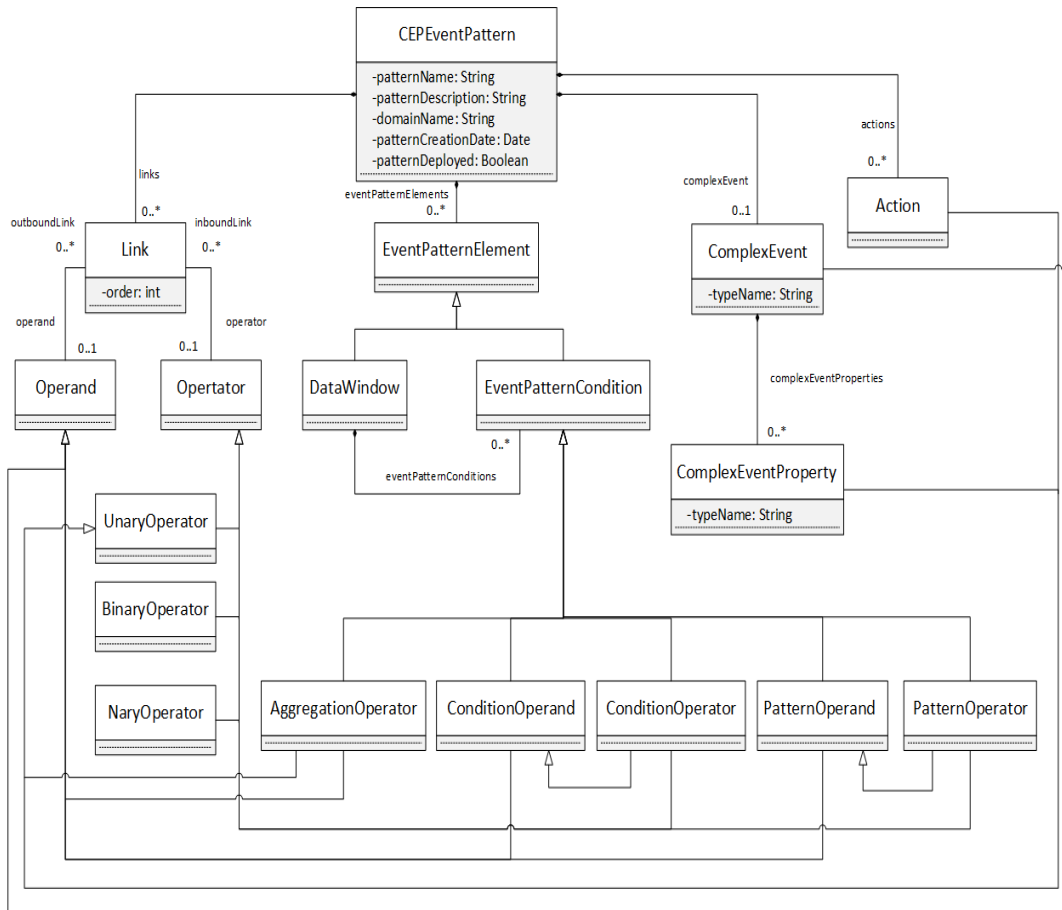


Figure 6: CEP Patterns Metamodel [3]

and may establish, for instance, the allowed level of various substances (e.g. chrome, mercury) that these elements may contain. For example, EPA establishes, in the Secondary Drinking Water Standards, that the maximum contaminant level for sulfate is 250 mg/L.

In turn, CEP technologies have extensively been used to check at runtime the compliance of systems and processes with various type of regulations [26, 27, 28, 29]. We claim that the constraints established by environmental regulations can also be monitored with CEP technologies.

In what follows we take a model-driven approach by defining a high-level conceptual model for environmental regulations as well as its relationship with the CEP metamodel presented in Section 2.3.2 extended with the geographic dimension. The model-driven approach was taken based on the analysis performed in [25] and described in Section 2.3.1. In addition, we base our proposal in the metamodel presented in Section 2.3.2 given that, to the best of our knowledge, there are not other metamodels in the area as complete and generic as this one.

Later on, in Section 4, we develop a case study which shows this approach in practice.

### 3.1 Geographic Extension for CEP

The geographic dimension of events may be essential for monitoring some of these environmental regulations given that they may involve specific geographic areas (e.g. a river, a country). In what follows we present a geographic extension for CEP which expands the metamodel for CEP [3] described in Section 2.3.2 with geographic elements.

In the new metamodel, a geometry defined by the OGC SFS reference model is a native event property (typed as Geometry). This is shown in Figure 8 which extends the metamodel presented in Figure 5.

Additionally, we extend the existing operators with geographic operators (e.g. Union, Intersection and Distance) in order to be able to support geographic operations. The original model (Figure 6) defines logical operators (And, Or, Not), arithmetic operators (Addition, Subtraction) and comparison operators (Equal, LessThan, GreaterThan). The geographic extension depicted in Figure 9 includes other operators with geographic characteristics that are classified as follows:



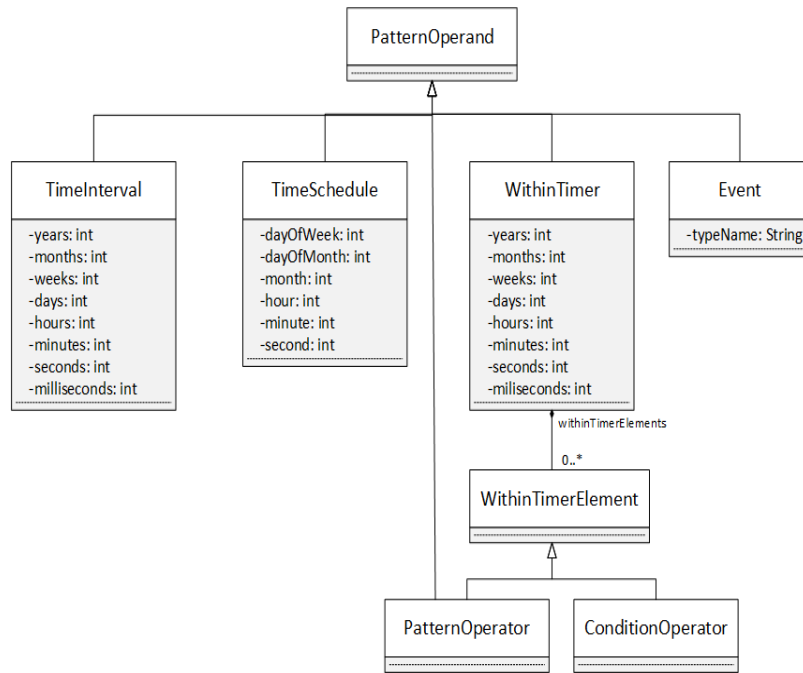


Figure 7: Pattern Operand [3]

- **GeoArithmeticOperators**: These operators receive n-operands of the type Geometry and return another geometry. This category includes operators such as Union, Intersection and Difference.
- **GeoBooleanOperators**: These operators receive n-operands and return true or false. This category includes operators such as Equals, Distance, Intersects and Contains.

As an example, the geographical Distance operator indicates that an event occurred at a given distance. This is based in the ideas proposed in [17] where the OGC defines distance as a native operator of the Geometry class as follows: *distance(another: Geometry): Distance*. For example, if the Geometry is a Point, a specific Distance operator for Point is defined by the following function: *distancePoint(Point, Distance): Boolean*. This could be useful in a situation which requires knowing if the events occurred in a specific area.

In summary, our proposal supports the representation of geographical events and operators of any kind, independently of a particular CEP technology or product. A complete example is shown in Section 4. For these events to be evaluated by a specific CEP engine, a translation process is required in order to express this representation in the particular language that the engine supports. Section 5 shows how some parts of this representation can be expressed in the language used by the Esper CEP Engine.

### 3.2 Modeling Environmental Regulations

Environmental regulations are usually organized in sections (e.g. articles in laws). Each section may include constraints that apply to specific environmental elements (a.k.a. factors). Assuming that data regarding the state of environmental elements are received in the form of events, constraints may be monitored by specifying event patterns using these events. For example, a pattern can be specified to detect if the level of chrome in water has been continuously increasing during the last hour. This may allow taking preventive actions before the level of chrome reaches a not allowed value.

Figure 10 presents a metamodel for environmental regulations. This metamodels relates the problem space provided by the environmental regulations and the solutions space provided by the extended CEP metamodel presented before.

The metamodel comprises the following concepts for representing environmental regulations. We exemplify this with the model of our case study depicted in Figure 12, which is deeply developed in Section

- **MainDocument**: it represents the main document to be considered, it could be a *Law*, a *Decree*, or another kind of regulation. In our case study it corresponds to the Decree 256/79 which also refers to the Law 14850. A document is related to a concrete CEP domain which will define events and patterns for the monitoring of regulations.

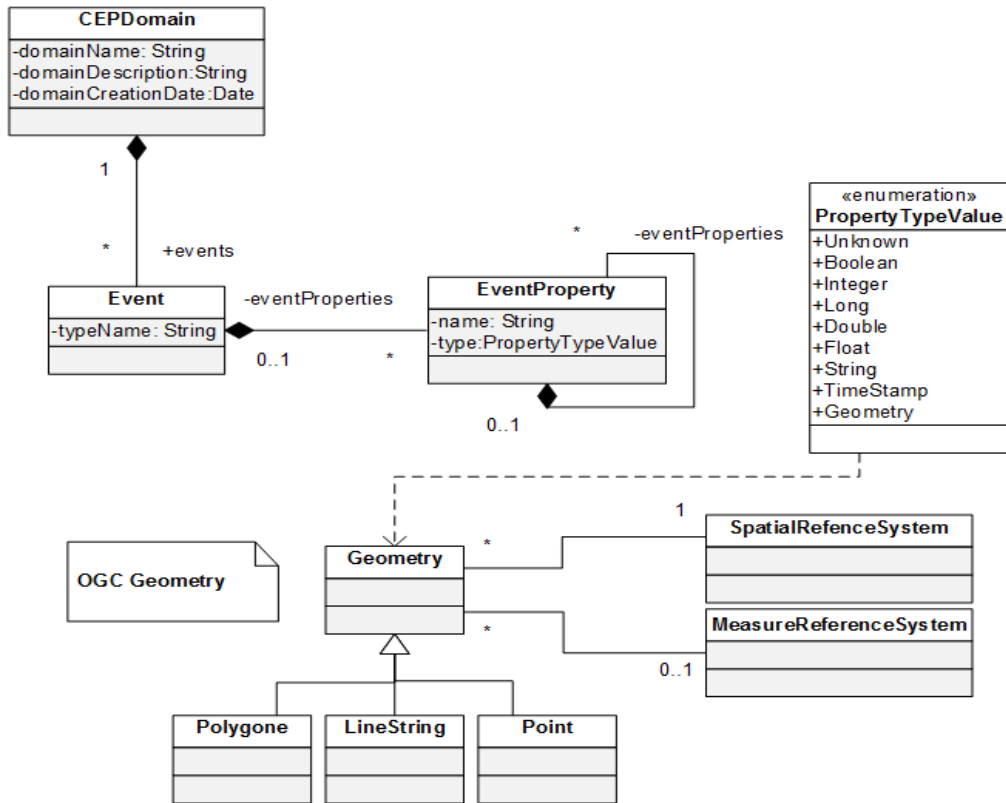


Figure 8: Model for GeoEvents

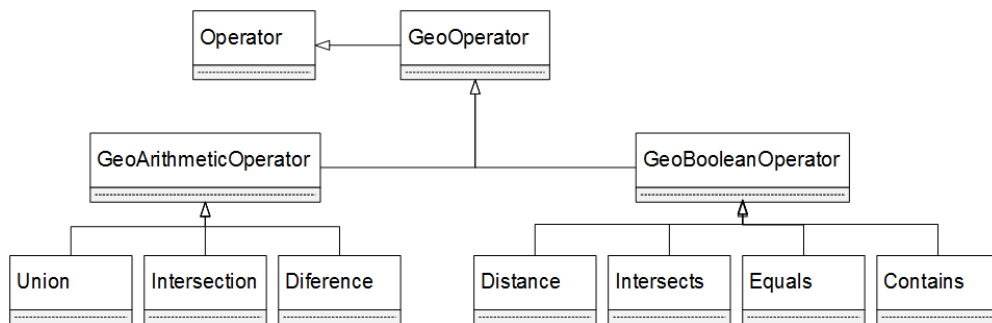


Figure 9: GeoEvents Operators

- **Section:** it is a part of the main document, for example, an article can be a section if the main document is a Law. In our case study there are two different sections: article 3 and 5.
- **Factor:** it is an environmental element of the reality which the law aims to control. Some examples of these elements are: air, water and waste. In our case study, water is the factor.
- **Constraint:** represent the requirements established in a Section of a MainDocument for a Factor. For example, constraints allow fixing ranges or limits values for the parameters (e.g. the substances that may be present in the Factors) considered in the law. In our case study there are two constraints with respect to water pollution. A constraints is related to a CEP domain patterns which is used for its monitoring. This pattern uses events defined within the CEP domain which is related to the document in which the constraint is defined.
- **DocType:** it is used to give the name to the different types of legal documents. In our case study we have a decree and a law.
- **FactorType:** it is used to identify the different types of environmental factors. In our case study we have water. aa

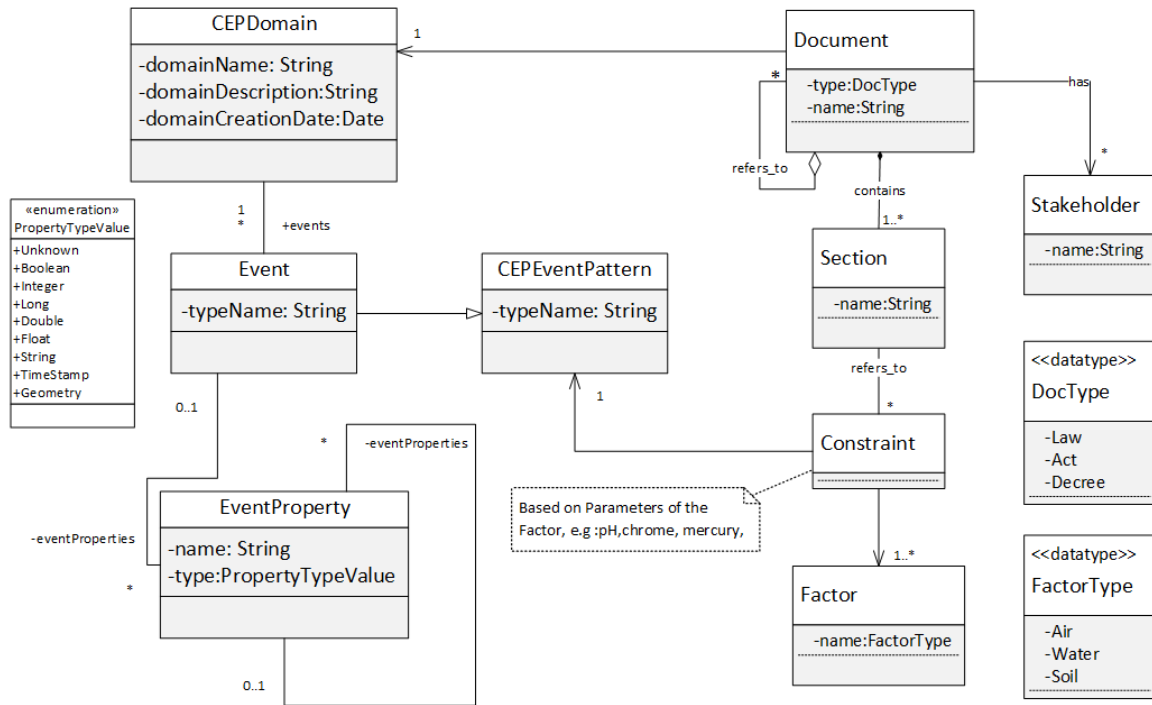


Figure 10: Conceptual Model for Monitoring Environmental Regulations with CEP

This metamodel provides a well-structured and simpler version of environmental regulations. For example, it can be used for organizational issues and the generation of end-user monitoring software: as an abstraction of the regulations focused on the constraints they contain, as well as the existing connection with the CEP patterns that are used for their monitoring.

## 4 Case Study

As described in Section 2.1.2, environmental compliance is the state of being in accordance with environmental regulations [12, 13]. These regulations may, for example, specify expected values for different environmental parameters. In particular, water quality criteria have been established for a number of traditional water quality variables such as pH and dissolved oxygen, among others. Such criteria guide decision makers in the establishment of control strategies in order to monitor and be able to guarantee the quality of water for different purposes [30]. This section presents a case study in the context of the Uruguayan legal system for water resources.

### 4.1 General Context: Legal System for Water Resources in Uruguay

In Uruguay, the Water Code (DL No. 14,859) is the foundation of the legal system for water resources. It establishes the competencies of public and private organizations as well as the mechanisms for their proper management. Besides, Law No. 14.440 regulated by Decree 216/76, establishes rules on the discharge of waste water. On the other hand, the national Water Policy, defined in Law No. 18.610 and in its Decree 78/10, includes the management of water resources as the services and applications related to the water. The quality of water resources is regulated by the Decree 253/79. This Decree sets the parameters to evaluate the water in Uruguay (e.g. pH, chrome) as well as their allowed values which depend on the use for which the water is intended. In particular, the different uses of water are classified in the following classes:

- Class 1: Water that can be used for human drinking.
- Class 2: a) Water intended for the irrigation of plants within the wetting system and intended for human consumption in general. b) Water used in recreational activities by direct contact with the human body.
- Class 3: Water destined to the preservation of hydrological flora and fauna, water intended for the irrigation of cultivation whose product is not consumed in its natural form or water for irrigation systems in which the water does not take contact with the product.

- Class 4: Water flows in general that should keep the harmony with the environment and water to irrigate crops whose products are not consumed by human beings.

Table 1 presents the allowed values<sup>16</sup> for each class.

Table 1: Allowed Values for the Different Uses of Water

Parameter	Class 1	Class 2a	Class 2b	Class 3	Class 4
PH (min-max)	6,5 - 8,5	6,5 - 9,0	6,5 - 8,5	6,5 - 9,0	6,0 - 9,0
Chrome Total (max)(mg/l)	0,05	0,005	0,05	0,05	0,5
Cyanide (max)(mg/l)	0,005	0,005	0,005	0,005	0,05
Plumb(max)(mg/l)	0,03	0,03	0,03	0,03	0,05
Cadmium (max)(mg/l)	0,001	0,001	0,005	0,001	0,01
Mercury (max)(mg/l)	0,0002	0,0002	0,0002	0,0002	0,0002

## 4.2 Application Scenario

The Santa Lucía river is the main course of water at the South of Uruguay. It is 248 kilometers long and its basin covers 12.300 square kilometers. The water purification plant of OSE<sup>17</sup>, located in the city of Aguas Corrientes, supplies potable water to approximately 60% of the population of Uruguay.

For the purpose of this case study, it is assumed that the Municipality of Sauce is willing to monitor the quality of the water in the Santa Lucía River along the stretch that lies within Sauce's territorial jurisdiction. The main goal of this monitoring is to evaluate whether the water is safe for human drinking, for what it must comply with the requirements established by the Decree 253/79 presented in the previous section. The water for this purpose falls into the category Class 1 in Table1.

In order to evaluate the parameters established by the decree (e.g. pH, chrome), the municipality plans to use open data collected through sensors located along the river. These sensors could be provided by a government agency, such as DINAMA<sup>18</sup>. These data would include all the information related to the measurements of the different parameters as well as the location of the sensor that captured that information. Data provided by individuals (e.g. shared through social networks) could also be used. These data should, however, go under a certain validation process carried out by a government agency, for example, by MTOP<sup>19</sup>.

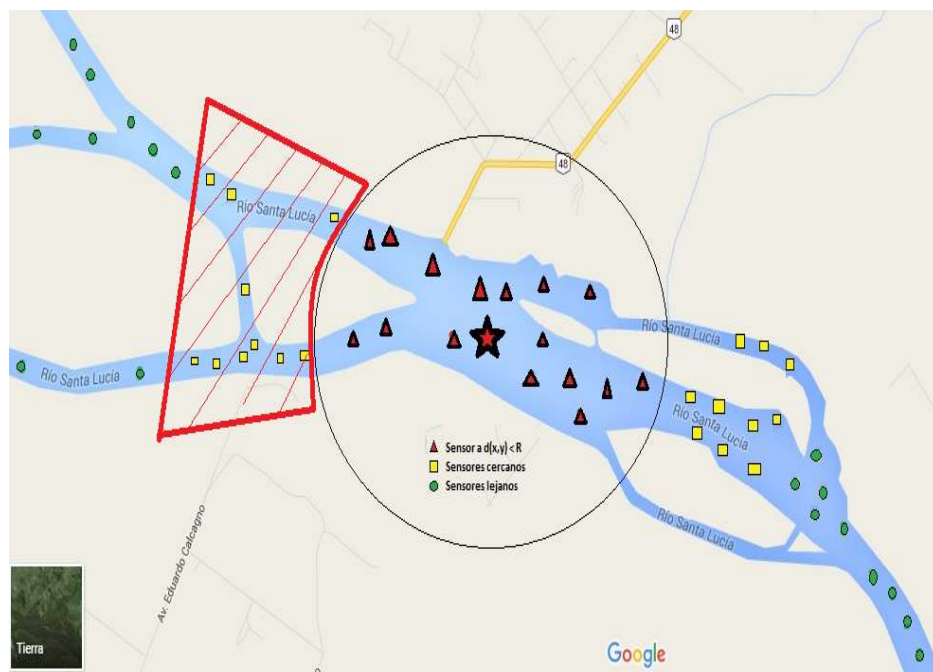


Figure 11: Case Study: Application Scenario

<sup>16</sup>A subset of the values established in Decree 253/79

<sup>17</sup>OSE: Obras Sanitarias del Estado is an uruguayan public company responsible for the supply of drinking water and sanitation throughout the country

<sup>18</sup>DINAMA: National Environment Directorate of Uruguay

<sup>19</sup>MTOP: Ministry of Transport and Public Works of Uruguay

Figure 11 shows how the sensors are located along the river. The sensors depicted as triangles are the ones in the area of interest (i.e. Sauce's territorial jurisdiction). In addition, the sensors depicted as squares are sensors which are close to the area of interest. Finally, the rest of the sensors are far from this area.

In this context, the municipality of Sauce wants to monitor the different values received from the sensors in order to be notified of possible anomalies and act accordingly. In particular, the municipality takes as reference the values specified for Class 1 and it wants to be notified in case of the following two situations:

- **Pollution in the Area of Interest:** The sensors in the area of interest report out-of-range values for pH and chrome.
- **Potential Pollution in the Area of Interest:** Some sensors close to the area of interest (on the left in Figure 11) report out-of-range values for pH or chrome. In addition, these sensors report the occurrence of rains.

### 4.3 Modeling the Application Scenario

For the purpose of the case study, only Class 1 is considered, given that this class sets the parameters related to water for human drinking. Those parameters include smell, color, pH, nitrate, cyanide, coliforms, lead, chrome, among others. In the rest of the case study only two parameters are going to be considered: pH and Total Chrome. However, others parameters may be included.

In Figure 12 the conceptual model presented in Figure 10 is leveraged for representing the aspects of the environmental regulations described in the previous sections (i.e. the figure depicts an instance of that model). In particular, the following concepts were identified:

- **MainDocument:** The Law 14859 and Decree 256/79.
- **Section:** Article 5 and Article 3, where the different classes and values to be controlled are defined.
- **Factor:** Water, which is the factor addressed by the law.
- **Constraint:** The rules in Article 5 as shown in Table 1.

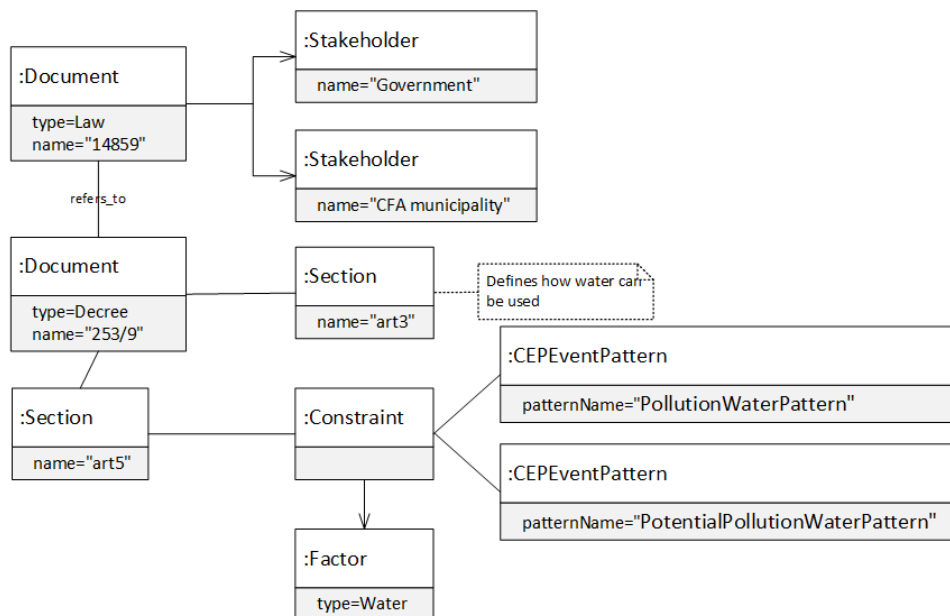


Figure 12: Relevant aspects of the regulations considered in the case study

In addition, in order to detect the relevant situations for Sauce's Municipality, as described in Section 4.2, two CEP patterns are identified:

- **PollutionWaterPattern:** This pattern detects pollution in water in the area of interest using the pH and chrome values reported by sensors in this area.
- **PotentialPollutionWaterPattern:** This pattern aims to predict the pollution of water in the area of interest by using the values reported by sensors which are close to this area.

#### 4.4 Detecting Pollution in Water

In order to detect when the levels of pH and chrome in water get out of permitted ranges in the area of interest, the event pattern named PollutionWaterPattern was designed. This event pattern is based on two simple events and a geographic operator. The simple events that were identified are:

- **WaterPH:** Data with the measured values of pH in water are received.
- **WaterTotalChrome:** Data with the measured values of total chrome in water are received.

These events must have a geographic component so they were modeled as GeoEvents, inheriting their properties as depicted in Figure 13.

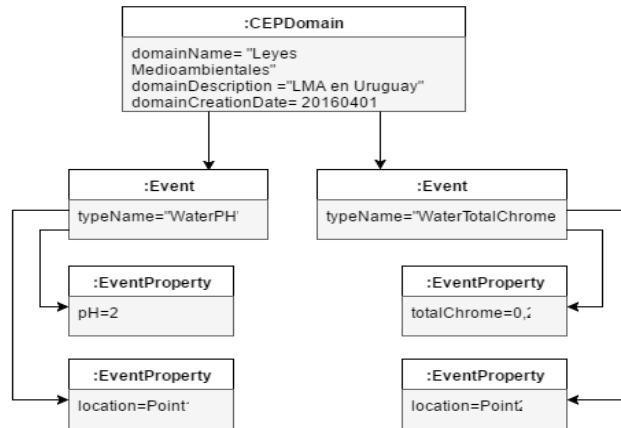


Figure 13: Domain model for Detecting Pollution in Water

In addition, the operator *Distance* was also required. This operator provides the capability of detecting if data is relevant according to the distance to a given location. In the proposed model, this operator can receive any GeoEvent Operand. For the case study, the Event is associated to a coordinate and implemented as a Geometry:Point, as presented in Figure 14.

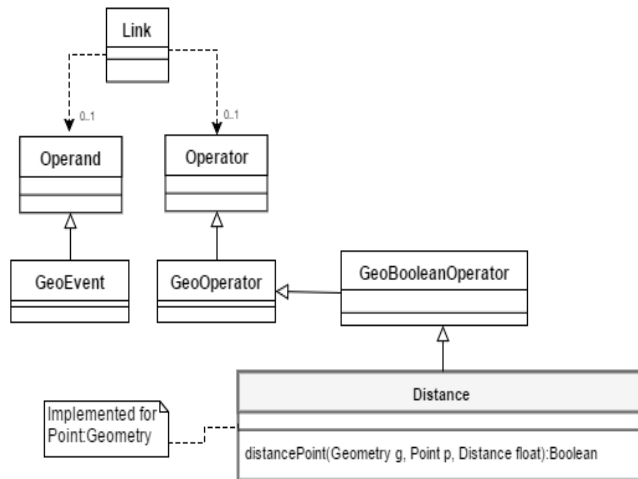


Figure 14: Geo Distance Operator

Table 2 presents a summary of the Detecting Pollution in Water issue and how a possible solution was modeled.

#### 4.5 Predicting Pollution in Water

It is known that a river has multiple torrents that carry the water from its different sources until it flows into another major channel. In the case of the Santa Lucía River, the convergence of these flows is into the Rio de la Plata. For this reason, the Municipality of Sauce could use sensor data collected near the area of

Table 2: Summary of the Detecting Pollution in Water Issue

	<b>Water Pollution CLASS 1</b>
Law/Decree/Policy	Law 14859, Decree 253/79
Stakeholders	Ministerio de Transporte y Obras Públicas (MTO) y Ministerio de Vivienda, Ordenamiento Territorial y Medio Ambiente (MVOTMA )
Monitored Factor	Water
Monitored Parameter	PH, Total Chrome
Measuring tools	The parameters can be monitored with specific sensors for every purpose.
Alert criteria	An alert is activated when one of the two parameters is out of its control range and the measurement detected is within a distance of less than 100 meters from a certain position which is specified by its coordinates (x, y) and its spatial reference identifier (srid)
Simple Event A	PH out of range
Simple Event B	Total chrome exceeds 0.05 mg/L
Complex Event to detect	If Point(x,y,srid) is the reference point: A.geometry.distancePoint(Point(x,y,srid), 100) and B.geometry.distancePoint(Point(x,y,srid),100)
Action performed	Alert

interest to monitor tendencies in measured values. In addition, other complementary data can be used (e.g. the temperature of the water, the amount of daily rainfall).

Real-time analysis of water measurements is vital for generating alerts to indicate when water is no longer potable according to environmental legislation. Indeed, for this type of vital resource it is important to prevent problems so that predictive analysis becomes an important tool to help detecting unwanted patterns in the monitored values. Specifically, tendencies can be detected by using measurements over a given period of time from the upstream and downstream sensors by interpolating values, with the goal of trying to keep values according to environmental laws.

Concretely, the PotentialPollutionWaterPattern could be defined in order to trigger alerts when sensors located in the proximity of the main monitored radio detect irregular values for a period of time. This solution allows the timely detection of potential anomalies in the area of interest according to the expected quality of water established in Class 1 in Decree 253/79.

As shown in Figure 15, the domain of events is modeled as was detailed in Figure 7 considering the interval of time and a new event associated with the amount of daily precipitations. This simple event is described as follows.

**WaterRainfall:** Data of rainfalls are received. This event has a geographical component that is usually a polygon corresponding to the area of the precipitation.

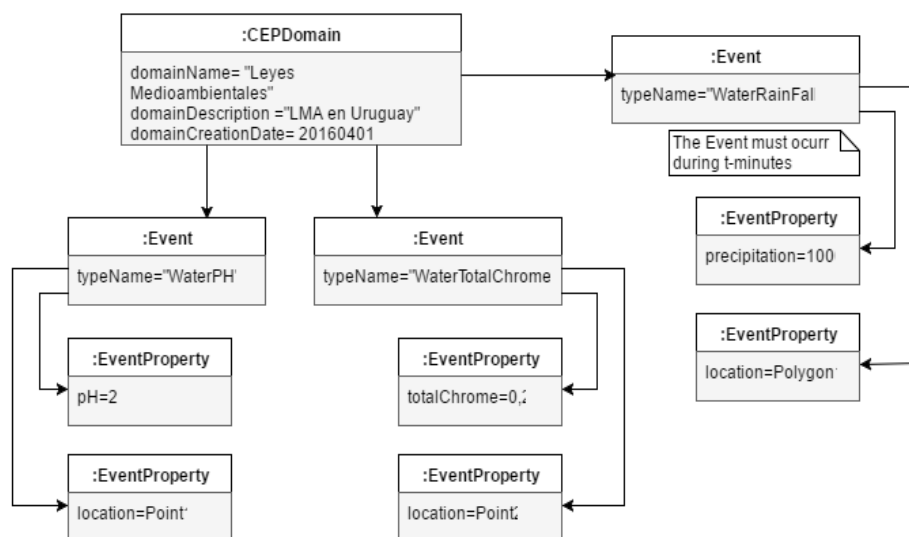


Figure 15: Domain Model for Predicting Pollution in Water

This way, the event pattern `PotentialPollutionWaterPattern` is defined as in Figure 16, where: `Z` is the area close to the area of interest, `A` is `WaterPH`, `B` is `WaterCromoTotal`, and `C` is `WaterRainfall`.

```
WithinTimer[(Z.contains(A.geometry) or Z.contains(B.geometry)) and Z.contains(B.geometry)]
```

Figure 16: `PotentialPollutionWaterPattern` event pattern

As previously presented, this new pattern requires the geo-operator `Contains`, which given a geometry and a geo-event indicates if the event occurred in that geometry. At a conceptual level, the `PotentialPollutionWaterPattern` uses the generic operator `Contains`. This operator presented, in Figure 17, is defined as follows: *contains(geometry, Geo-Event): Boolean*. Consequently, it was also required to implement an operator that supports the Geometry of type Polygon given the characteristics of the monitored zone (`Z`).

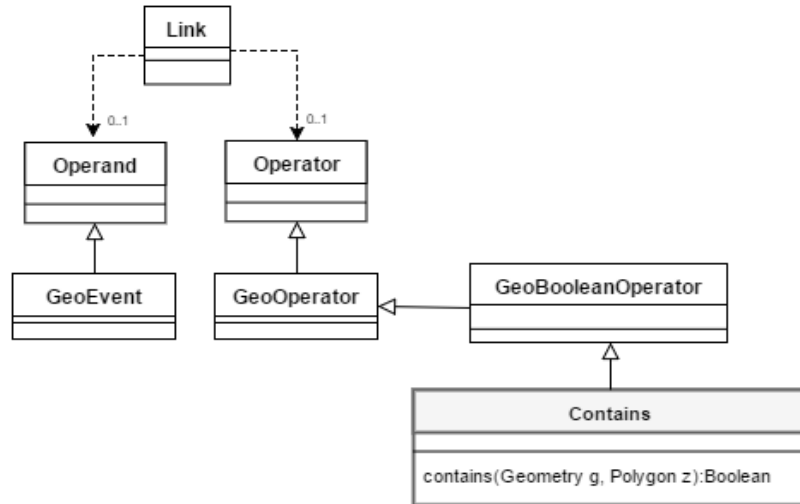


Figure 17: Geo Contains Operator

In summary, we showed how our model-driven approach allows to represent environmental laws as well as how their monitoring and enforcement can be performed by means of a CEP-based solution, strongly dependent on geographical information. In the next section we present a concrete implementation of these ideas as a stronger validation of the proposal.

## 5 Prototype Implementation

As presented in Section 3.1, the proposed metamodel allows representing geospatial events and operators independently of a particular CEP implementation. For the purpose of a technological validation of the proposal, a prototype was implemented in order to develop the case study presented in Section 4. In particular, some of the event patterns described in the case study were translated to EPL so that they can be interpreted by the Esper engine.

The prototype shown in Figure 18 allows managing the registration of sensors, data sources, complex events and alarms which are generated when complex events are detected.

For the prototype's implementation, complex event patterns were specified using the EPL language for Esper. In particular, the EPL code for the `PollutionWaterPattern` is shown in Figure 19. This code uses the `distancePoint` operator, which was introduced in Section 3.1, in a procedural style. The `pointFromText` function is used to create a fixed Point whose coordinates are coded following the WKT representation explained in Section 2.2, in a certain spatial reference system identified by the `srId` parameter.

The implementation of the geospatial operators, such as `distancePoint` was achieved through the use of the JTS API mentioned in Section 2.2. Other geographic technologies were used in this prototype, such as 52 North<sup>20</sup>, which serves as an implementation of the SOS standard for sensor observations, Geoserver<sup>21</sup> to provide the geographic layers that are seen in Figure 18, using WMS and WFS standards, and PostgreSQL DBMS with PostGIS extension to store geographic information.

<sup>20</sup><http://52north.org/>

<sup>21</sup><http://geoserver.org/>



Figure 18: Software to Manage Sensors and Alarms

```

select * from SensorEvent
where (distancePoint(geom,pointFromText(POINT(x y),srid),100)
      and (1=idParameterType and 1=sourceElement and (value > 8.5 or value < 6.5))
      or (10=idParameterType and 1=sourceElement and value > 0.05))

```

Figure 19: EPL of the PollutionWaterPattern for Esper

Figure 20 shows the behavior of the data through the prototype for the detection of patterns of complex events and the generation of the corresponding alarms. Once the user defines patterns or rules to detect, the execution is as follows:

- *Data reception:* events come from different data sources and sensors in a standard format and are caught by the prototype. In the example, the events of measurements of ph and total chrome from water are received with the location where these measurements were taken.
- *Analysis:* the CEP engine with extended capabilities evaluates the rules that are defined and if some of the events satisfy the defined patterns, a complex event is then detected. Two of the rules that were discussed previously are shown in the EPL code above and in Figure 20.
- *Action performed:* when a complex event is detected, an action is performed. In this case an alert is generated by sending an email indicating that there is a measurement out of the range. Events that do not apply to any rule are discarded, such as the event e3 in Figure 20.

The implementation of the prototype is based on an Enterprise Service Bus (MuleESB<sup>22</sup>). In Figure 21 the mediation flow for a new alert is shown. This flow receives events from all the available SOS sources to do the processing. In the first flow, input data is converted to an Esper event. In the second flow, the Esper event is analyzed to verify if it matches the alert clause. In the case that it does match, a notification is sent to all the users who have subscribed to receive that alert.

<sup>22</sup><https://www.mulesoft.com/platform/soa/mule-esb-open-source-esb>

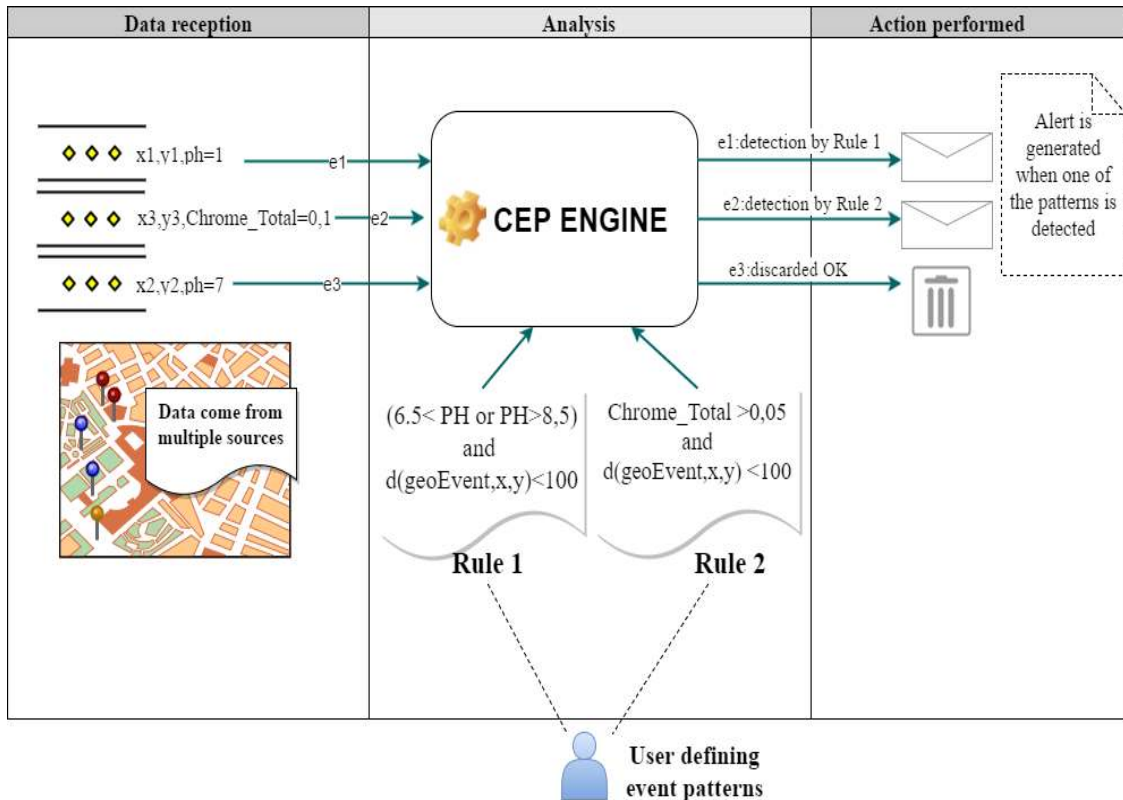


Figure 20: Data Flow and Generation of Alerts

This prototype shows that our model-driven approach can be implemented and, thus, it is possible to perform the CEP-based monitoring and enforcement of environmental laws. However, our approach is still incomplete, since we manually developed express CEP patterns within Esper, and there is no automatic translation between our regulations and CEP domain models, and Esper. This is part of our future work.

## 6 Related Work

The Infrastructure for Spatial Information in Europe project (INSPIRE, [31]) is an active project addressing issues of Spatial Data Infrastructures. Some of its sub-projects study the monitoring of environmental variables and define Domain Specific Languages (DSLs) for this task. In particular, these DSLs allow representing the various factors which are relevant for monitoring environmental laws (e.g. water, air and soil). In addition, these factors can also be related with the laws that deal with them. Compared to our proposal, this project focuses on the definition of DSLs but does not propose mechanisms to monitor environmental variables at runtime. However, the DSLs proposed by this initiative could be used in our work. Indeed, some ideas from them were taken for the model proposed in Section 3.

There are also academic projects and proposals that deal with monitoring and controlling compliance requirements arising from different types of regulations. For example, the Compliance-driven Models, Languages, and Architectures for Services project (COMPAS, [32]) focused on the compliance of business processes within an organization. Also, the Enabling Change and Compliance For Collaborative Processes project (C3PRO, [33]) focused on inter-organizational business processes. In addition, in [34] the authors propose ExpPDT: a policy-based approach for automating compliance management. ExpPDT is a formal policy language which was initially developed to deal with privacy issues but it was then extended in order to address other compliance domains (e.g. business process compliance). Also, in [35] a policy-based and model-driven approach for compliance management is proposed. Compared to our work, these proposals address compliance requirements which mainly apply to business processes. However, they do not consider regulations that specify requirements which can be monitored based on geographic information coming from sensors or VGI. In addition, these projects mainly focus on the enforcement of compliance requirements at design time (e.g. when designing a business process) and not at runtime.

The problem of managing large and diverse formats of geospatial data is related to the traditional challenges of Big Data [36]. Big Data is usually described using the “3V” Model [37] which involves the

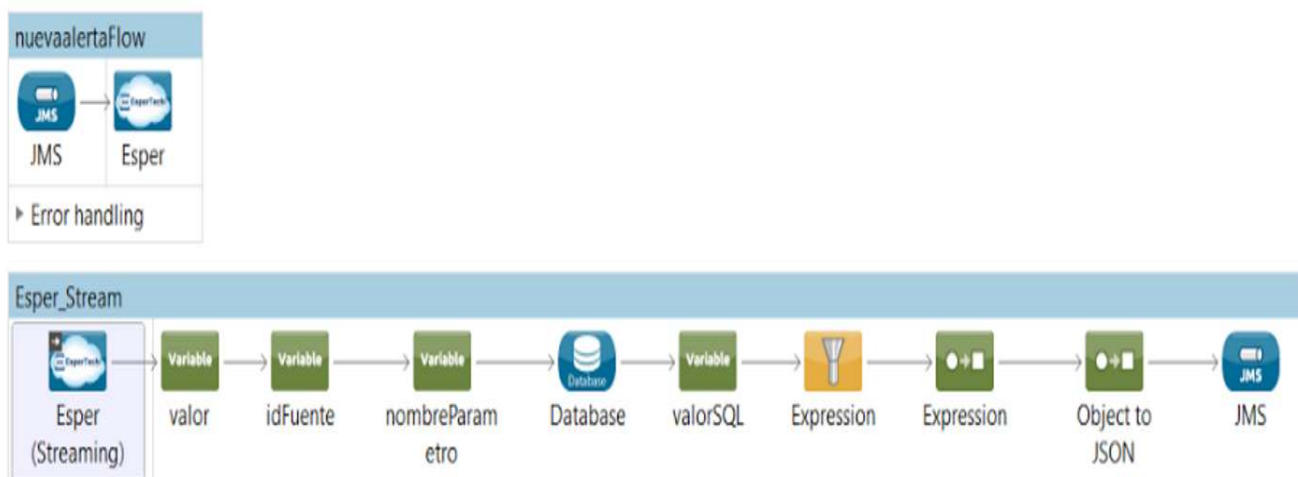


Figure 21: New Alert Mediation Flow in MuleESB

following three dimensions: "Volume", "Velocity" and "Variety". However, this approach focuses in data which were obtained a priori and stored in databases [38]. In order to process real time data coming from heterogeneous data sources, traditional Big Data solutions can be complemented with Fast Data [39, 40]. Fast data allows analyzing real time data extracting the relevant information for the business and dealing, in this way, with the "Value" dimension which considers that not all the data have a business value. Fast Data involves the use of dynamic technologies such as CEP which allows processing a large amount of events as well as relate them to detect and respond in real time to critical business situations. This way, our CEP-based proposal can be considered as a complementary approach to big data solutions. Indeed, CEP technologies are applicable to big data scenarios and we consider that CEP and big data technologies can benefit from each other [41, 42].

There are certain proposals that address specific technical aspects of collecting and processing data from multiple sensors. For instance, in [43], a Service Oriented Architecture (SOA) in which each sensor is "service enabled" is proposed, meaning that any service can publish its data through Web services (e.g. OGC SOS), as opposed to a more traditional approach in which ETL (Extraction, Transform, Load) tools are required to integrate data. This proposal acknowledges the need of some notification-triggering mechanism, following OGC's Sensor Event Service<sup>23</sup> discussion paper, but it focuses on simple events, i.e. when a threshold value is reached on a single sensor, rather than complex events. Some sensor-based SOAs have also been successfully implemented and proven in real scenarios. In [44], several monitoring systems developed by the OSIRIS Project<sup>24</sup> are described. These systems tackle air pollution, water pollution, forest fire, and buildings fire, and they all make extensive use of OGC standards and proposals under the Sensor Web Enablement<sup>25</sup> umbrella, especially the SOS standard. Even though some of these systems have specific modules for generating alerts when some conditions are met, no further details are given as to how those conditions are evaluated.

With respect to tools that extend CEP engines with geographical capabilities, although there are some proposals, they lack a generic model for specifying events, i.e. they are restricted to a specific tool. In Table 3 there is a brief comparison of some existent products, some of which are discontinued.

## 7 Conclusions and Future Work

In this paper we present how to perform the CEP-based monitoring and enforcement of environmental laws following a model-driven approach. For this purpose, we propose a geographical extension of a complex event metamodel which allows defining events considering the geographic dimension and specifying complex events using geographic operators. These extended capabilities allow correlating events based on their location. The extension is particularly useful to tackle with requirements for the monitoring and enforcement of environmental laws. In fact, we also define a metamodel for environmental regulations which connect the constraints stated within them and the CEP events and patterns that are defined for their monitoring. The

<sup>23</sup>[http://portal.opengeospatial.org/files/?artifact\\_id=29576](http://portal.opengeospatial.org/files/?artifact_id=29576)

<sup>24</sup>[http://cordis.europa.eu/project/rcn/80164\\_en.html](http://cordis.europa.eu/project/rcn/80164_en.html)

<sup>25</sup><http://www.opengeospatial.org/ogc/markets-technologies/swe>

Table 3: CEP Engines with Geographic Support

Engine	Feature
RULE CORE	Developed in Python. The tool triggers actions by detecting patterns of complex events. Currently outdated.
SPATIAL RULES	Rules according to areas. Compatible with geospatial OGC standards Currently outdated.
GCEP	Esper extension. Last update: 2010.
WSO2	Full platform with geographical extensions, also an ESB. Active project with full documentation.

metamodels are independent of any specific CEP platform, separating the specification of the problem from a particular implementation.

The paper also presents a case study concerning water legislation in Uruguay, which was modeled with the geographic extension in order to validate the proposal. One aspect to consider as future work is the development of a DSL for environmental laws, allowing to express relevant sections of the text of a law, from which the rules (or patterns) to be fulfilled can be extracted. This is in some way related to the INSPIRE project [31]. The conceptual model presented in Section 3 is a first step towards the specification of this DSL. A complete DSL will allow raising the level of abstraction in specifying compliance requirements without using the lower-level complex events model.

For validating the technical feasibility of these ideas, a prototype was implemented involving the definition of geo-events and the use of the Esper tool as a CEP platform. This was a first step towards providing geospatial support in a CEP environment. We are currently working on extending the prototype in order to support other geographic operators. This will allow the development of case studies of greater complexity. In addition, we are interested in exploring the possibility of integrating the current CEP platform with geographical information systems to ease the management of geospatial data. Finally, we are interested in exploring the definition of an automatic transformation from the proposed metamodel to EPL, so as to achieve a full and automatic deployment of the solution in Esper.

## Acknowledgment

This work was partially funded by the Comisión Sectorial de Investigación Científica (CSIC), Universidad de la República, Uruguay. We would also like to thank students Gonzalo Antúnez and Marcelo Piñeiro who participated in the development of the prototype.

## References

- [1] M. F. Goodchild, "Citizens as sensors: the world of volunteered geography," *GeoJournal*, vol. 69, no. 4, pp. 211–221, Aug 2007. [Online]. Available: <http://dx.doi.org/10.1007/s10708-007-9111-y>
- [2] M. Eckert, F. Bry, S. Brodt, O. Poppe, and S. Hausmann, "A CEP Babelfish: Languages for Complex Event Processing and Querying Surveyed," in *Reasoning in Event-Based Distributed Systems*, ser. Studies in Computational Intelligence. Springer, 2011, no. 347, pp. 47–70. [Online]. Available: [http://dx.doi.org/10.1007/978-3-642-19724-6\\_3](http://dx.doi.org/10.1007/978-3-642-19724-6_3)
- [3] J. Boubeta-Puig, G. Ortiz, and I. Medina-Bulo, "Model4cep: Graphical domain-specific modeling languages for CEP domains and event patterns," *Expert Syst. Appl.*, vol. 42, no. 21, pp. 8095–8110, 2015. [Online]. Available: <http://dx.doi.org/10.1016/j.eswa.2015.06.045>
- [4] F. Herrera, L. González, and D. Calegari, "Complex event processing with geospatial support for monitoring and controlling compliance with environmental regulations," in *XLII Latin American Computing Conference (CLEI)*. IEEE, 2016, pp. 1–10. [Online]. Available: <https://doi.org/10.1109/CLEI.2016.7833414>
- [5] M. P. Papazoglou, "Making business processes compliant to standards and regulations," in *Proceedings of the 15th IEEE International Enterprise Distributed Object Computing Conference (EDOC)*. IEEE Computer Society, 2011, pp. 3–13. [Online]. Available: <http://dx.doi.org/10.1109/EDOC.2011.37>

- [6] S. W. Sadiq, G. Governatori, and K. Namiri, *Modeling Control Objectives for Business Process Compliance*, ser. Lecture Notes in Computer Science. Springer, 2007, vol. 4714, pp. 149–164. [Online]. Available: [http://dx.doi.org/10.1007/978-3-540-75183-0\\_12](http://dx.doi.org/10.1007/978-3-540-75183-0_12)
- [7] O. Türetken, A. Elgammal, W. van den Heuvel, and M. P. Papazoglou, “Enforcing compliance on business processes through the use of patterns,” in *19th European Conference on Information Systems (ECIS)*, 2011, p. 5.
- [8] M. El Kharbili, “Business process regulatory compliance management solution frameworks: A comparative evaluation,” in *Proceedings of the Eighth Asia-Pacific Conference on Conceptual Modelling*, ser. APCCM ’12, vol. 130. Australian Computer Society, Inc., 2012, pp. 23–32.
- [9] M. Reichert and B. Weber, “Business Process Compliance,” in *Enabling Flexibility in Process-Aware Information Systems*. Springer, 2012, pp. 297–320.
- [10] L. T. Ly, F. M. Maggi, M. Montali, S. Rinderle-Ma, and W. M. P. van der Aalst, “A framework for the systematic comparison and evaluation of compliance monitoring approaches,” in *17th IEEE International Enterprise Distributed Object Computing Conference (EDOC)*. IEEE Computer Society, 2013, pp. 7–16. [Online]. Available: <http://dx.doi.org/10.1109/EDOC.2013.11>
- [11] S. Sackmann, M. Kähler, M. Gilliot, and L. Lowis, “A classification model for automating compliance,” in *10th IEEE International Conference on E-Commerce Technology (CEC) / 5th IEEE International Conference on Enterprise Computing, E-Commerce and E-Services (EEE)*. IEEE Computer Society, 2008, pp. 79–86. [Online]. Available: <http://dx.doi.org/10.1109/CECandEEE.2008.99>
- [12] “Implementing remote sensing strategies to support environmental compliance assessment: A neural network application,” *Environmental Science and Policy*, vol. 12, no. 7, pp. 948 – 958, 2009. [Online]. Available: <http://dx.doi.org/10.1016/j.envsci.2009.08.001>
- [13] L. Friesen, “Targeting enforcement to improve compliance with environmental regulations,” *Journal of Environmental Economics and Management*, vol. 46, no. 1, pp. 72 – 85, 2003. [Online]. Available: [http://dx.doi.org/10.1016/S0095-0696\(02\)00033-5](http://dx.doi.org/10.1016/S0095-0696(02)00033-5)
- [14] B. Resch, M. Mittlboeck, F. Girardin, R. Britter, and C. Ratti, “Real-time geo-awareness sensor data integration for environmental monitoring in the city,” in *2009 International Conference on Advanced Geographic Information Systems Web Services*, 2009, pp. 92–97. [Online]. Available: <http://dx.doi.org/10.1109/GEOWS.2009.31>
- [15] E. Kruijff, E. Mendez, E. Veas, and T. Gruenewald, “On-site monitoring of environmental processes using mobile augmented reality (hydrosys),” in *Proceedings of the Environmental Information Systems and Services—Infrastructures and Platforms Workshop (ENVIP 2010)*. Citeseer, 2010, pp. 6–8.
- [16] OGC, “OpenGIS Web Services Architecture Description.” [Online]. Available: [http://portal.opengeospatial.org/files/?artifact\\_{\\_}id=13140](http://portal.opengeospatial.org/files/?artifact_{_}id=13140)
- [17] OGC, “Simple Feature Access - Part 2: SQL Option | OGC.” [Online]. Available: <http://www.opengeospatial.org/standards/sfs>
- [18] J. Boubeta-Puig, G. Ortiz, and I. Medina-Bulo, “Medit4cep: A model-driven solution for real-time decision making in SOA 2.0,” *Knowl.-Based Syst.*, vol. 89, pp. 97–112, 2015. [Online]. Available: <http://dx.doi.org/10.1016/j.knosys.2015.06.021>
- [19] O. Etzion and P. Niblett, *Event Processing in Action*. Manning Publications Company, 2010.
- [20] D. C. Luckham, *Event Processing for Business: Organizing the Real-Time Enterprise*. Wiley, 2011.
- [21] B. Pernici, Ed., *Advanced Information Systems Engineering, 22nd International Conference, CAiSE 2010, Hammamet, Tunisia, June 7-9, 2010. Proceedings*, ser. Lecture Notes in Computer Science, vol. 6051. Springer, 2010. [Online]. Available: <http://dx.doi.org/10.1007/978-3-642-13094-6>
- [22] J. Bacon, P. R. Pietzuch, J. Sventek, and U. Çetintemel, Eds., *Proceedings of the Fourth ACM International Conference on Distributed Event-Based Systems, DEBS 2010, Cambridge, United Kingdom, July 12-15, 2010*. ACM, 2010.

- [23] G. Antoniou, M. Grobelnik, E. P. B. Simperl, B. Parsia, D. Plexousakis, P. D. Leenheer, and J. Z. Pan, Eds., *The Semantic Web: Research and Applications - 8th Extended Semantic Web Conference, ESWC 2011, Heraklion, Crete, Greece, May 29 - June 2, 2011, Proceedings, Part II*, ser. Lecture Notes in Computer Science, vol. 6644. Springer, 2011. [Online]. Available: <http://dx.doi.org/10.1007/978-3-642-21064-8>
- [24] C. Zang, Y. Fan, and R. Liu, “Architecture, implementation and application of complex event processing in enterprise information systems based on RFID,” *Information Systems Frontiers*, vol. 10, no. 5, pp. 543–553, 2008. [Online]. Available: <http://dx.doi.org/10.1007/s10796-008-9109-0>
- [25] J. B. Puig, “Model-driven development of domain-specific interfaces for complex event processing in service-oriented architectures,” Ph.D. dissertation, Universidad de Cádiz, 2014.
- [26] E. Mulo, U. Zdun, and S. Dustdar, “Domain-specific language for event-based compliance monitoring in process-driven soas,” *Service Oriented Computing and Applications*, vol. 7, no. 1, pp. 59–73, 2013. [Online]. Available: <http://dx.doi.org/10.1007/s11761-012-0121-3>
- [27] A. Bertolino, A. Calabrò, and G. D. Angelis, “Adaptive SLA monitoring of service choreographies enacted on the cloud,” in *7th IEEE International Symposium on the Maintenance and Evolution of Service-Oriented and Cloud-Based Systems (MESOCA)*. IEEE Computer Society, 2013, pp. 92–101. [Online]. Available: <http://dx.doi.org/10.1109/MESOCA.2013.6632741>
- [28] R. Thullner, S. Rozsnyai, J. Schiefer, H. Obwegger, and M. Suntinger, “Proactive business process compliance monitoring with event-based systems,” in *Workshops Proceedings of the 15th IEEE International Enterprise Distributed Object Computing Conference, (EDOCW)*. IEEE Computer Society, 2011, pp. 429–437. [Online]. Available: <http://dx.doi.org/10.1109/EDOCW.2011.22>
- [29] A. Baouab, O. Perrin, and C. Godart, “An optimized derivation of event queries to monitor choreography violations,” in *Proc. of Service-Oriented Computing - 10th International Conference (ICSOC)*, ser. Lecture Notes in Computer Science, vol. 7636. Springer, 2012, pp. 222–236. [Online]. Available: [http://dx.doi.org/10.1007/978-3-642-34321-6\\_15](http://dx.doi.org/10.1007/978-3-642-34321-6_15)
- [30] R. Helmer and I. Hespanhol, *Water Pollution Control*. CRC Press, 1997.
- [31] European Union, “INSPIRE (Infrastructure for spatial information in Europe) project.” [Online]. Available: <http://inspire.ec.europa.eu/>
- [32] F. Daniel, F. Casati, V. D’Andrea, E. Mulo, U. Zdun, S. Dustdar, S. Strauch, D. Schumm, F. Leymann, S. Sebah, F. D. Marchi, and M. Hacid, “Business compliance governance in service-oriented architectures,” in *The IEEE 23rd International Conference on Advanced Information Networking and Applications (AINA)*. IEEE Computer Society, 2009, pp. 113–120. [Online]. Available: <http://dx.doi.org/10.1109/AINA.2009.112>
- [33] D. Knuplesch, M. Reichert, R. Pryss, W. Fdhila, and S. Rinderle-Ma, “Ensuring compliance of distributed and collaborative workflows,” in *9th IEEE International Conference on Collaborative Computing: Networking, Applications and Worksharing*. ICST / IEEE, 2013, pp. 133–142.
- [34] S. Sackmann and M. Kähler, “ExPDT: A Policy-based Approach for Automating Compliance,” *WIRTSCHAFTSINFORMATIK*, vol. 50, no. 5, pp. 366–374, 2008. [Online]. Available: <http://dx.doi.org/10.1007/s11576-008-0078-1>
- [35] M. E. Kharbili, Q. Ma, P. Kelsen, and E. Pulvermueller, “Corel: Policy-based and model-driven regulatory compliance management,” in *2011 IEEE 15th International Enterprise Distributed Object Computing Conference*, Aug 2011, pp. 247–256. [Online]. Available: <http://dx.doi.org/10.1109/EDOC.2011.23>
- [36] S. Li, S. Dragicevic, F. Anton, M. Sester, S. Winter, A. Çöltekin, C. Pettit, B. Jiang, J. Haworth, A. Stein, and T. Cheng, “Geospatial big data handling theory and methods: A review and research challenges,” *CoRR*, vol. abs/1511.03010, 2015. [Online]. Available: <http://arxiv.org/abs/1511.03010>
- [37] M. R. Evans, D. Oliver, X. Zhou, and S. Shekhar, “Spatial big data: Case studies on volume, velocity, and variety,” 2013.

- [38] F. Herrera, R. Sosa, and T. Delgado, "Geobi and big VGI for crime analysis and report," in *3rd International Conference on Future Internet of Things and Cloud, FiCloud 2015, Rome, Italy, August 24-26, 2015*, I. Awan, M. Younas, and M. Mecella, Eds. IEEE Computer Society, 2015, pp. 481–488. [Online]. Available: <http://dx.doi.org/10.1109/FiCloud.2015.112>
- [39] D. Hansen, "Oracle fast data: Real-time strategies for big data and business analytics." Tech. Rep., Mar. 2013. [Online]. Available: <http://www.oracle.com/us/solutions/fastdata/fast-data-gets-real-time-wp-1927038.pdf>
- [40] J. Boubeta-Puig, G. Ortiz, and I. Medina-Bulo, "A model-driven approach for facilitating user-friendly design of complex event patterns," *Expert Syst. Appl.*, vol. 41, no. 2, pp. 445–456, 2014. [Online]. Available: <http://dx.doi.org/10.1016/j.eswa.2013.07.070>
- [41] S. Bothe, V. Manikaki, A. Deligiannakis, and M. Mock, "Towards flexible event processing in distributed data streams," in *Proceedings of the Workshops of the EDBT/ICDT 2015 Joint Conference (EDBT/ICDT), Brussels, Belgium, March 27th, 2015.*, 2015, pp. 111–117. [Online]. Available: <http://ceur-ws.org/Vol-1330/paper-20.pdf>
- [42] R. Bhargavi, *Complex Event Processing Framework for Big Data Applications*. Cham: Springer International Publishing, 2016, pp. 41–56. [Online]. Available: [http://dx.doi.org/10.1007/978-3-319-31861-5\\_2](http://dx.doi.org/10.1007/978-3-319-31861-5_2)
- [43] A. Kotsev, F. Pantisano, S. Schade, and S. Jirka, "Architecture of a Service-Enabled Sensing Platform for the Environment," *Sensors*, vol. 15, no. 2, pp. 4470–4495, Feb. 2015. [Online]. Available: <http://dx.doi.org/10.3390/s150204470>
- [44] S. Jirka, A. Bröring, and C. Stasch, "Applying ogc sensor web enablement to risk monitoring and disaster management," in *GSDI 11 world conference, Rotterdam, Netherlands*, 2009.