



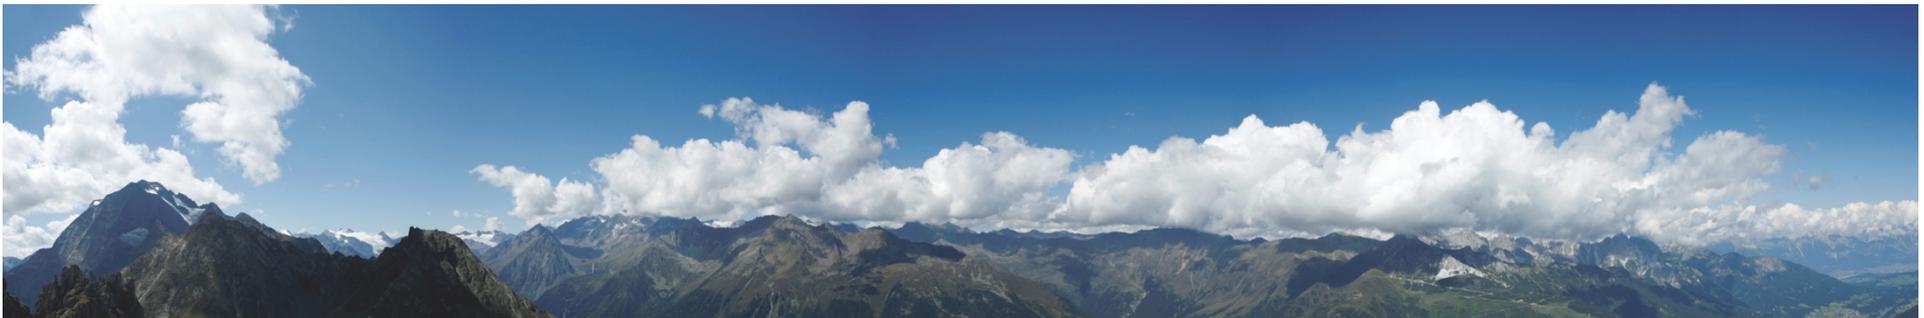
STI · INNSBRUCK

---

## Enhancing Event Processing Networks with Semantics to Enable Self-Managed SEE Federations

Srdjan Komazec

`srdjan.komazec@sti2.at`



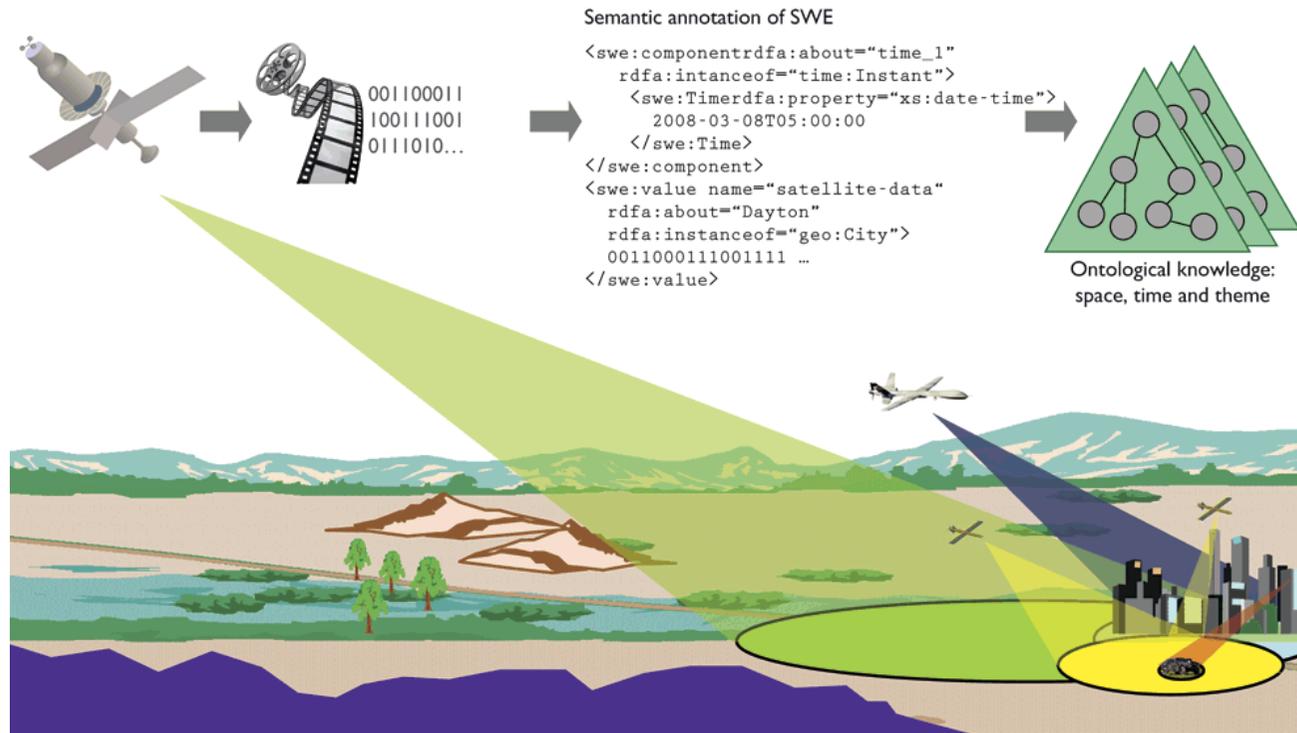
- 
- Motivation Scenarios
  - Event Processing
  - Identified Problems
  - Proposed Solution
  - Evaluation
  - Future work

# ***Motivation Scenarios***

# Motivation Scenarios

## Do we have events on the (Semantic) Web?

- Example 1: Semantic Sensor Web<sup>5</sup>
  - Collecting and processing avalanches of data about the world around us while relying on the semantic technology to increase interoperability and provide contextual information for situational knowledge.



<sup>5</sup> Amit Sheth, Cory Henson, Satya S. Sahoo, "Semantic Sensor Web," IEEE Internet Computing, pp. 78-83, July/August, 2008

# Motivation Scenarios

## Do we have events on the (Semantic) Web?

- Example 1: Semantic Sensor Web
  - Linked Sensor Data<sup>6</sup>
    - Hurricane and blizzard observations in the United States
    - Department of Meterology at the University of Utah
    - Measurements of phenomena such as temperature, visibility, precipitation, pressure, wind speed, humidity, etc

```
...
sens-obs:Observation_WindSpeed_3CLO3_2005_10_16_9_35_00
  a
    om-owl:observedProperty      weather:WindObservation ;
    om-owl:procedure             sens-obs:System_3CLO3 ;
    om-owl:result                sens-obs:MeasureData_WindSpeed_3CLO3_2005_10_16_9_35_00 ;
    om-owl:samplingTime          sens-obs:Instant_2005_10_16_9_35_00 .

sens-obs:MeasureData_WindSpeed_3CLO3_2005_10_16_9_35_00
  a
    om-owl:floatValue            "17.0"^^xsd:float ;
    om-owl:uom                   weather:milesPerHour .

sens-obs:Instant_2005_10_16_9_35_00
  a
    owl-time:Instant          ;
    owl-time:inXSDDateTime    "16-10-2005T09:35:00"^^http://www.w3.org/2001/XMLSchema#dateTime" .
...
```

A sample data from hurricane Wilma readings (October 2005)

<sup>6</sup> Harshal Patni, Cory Henson, and Amit Sheth. Linked Sensor Data. In Proceedings of 2010 International Symposium on Collaborative Technologies and Systems (CTS 2010), Chicago, IL, May 17-21, 2010.

# Motivation Scenarios

## Do we have events on the (Semantic) Web?

- Example 2: Social Networking (Twitter)
  - Adding structured annotations to a tweet<sup>7</sup>
  - A tweet can have one or more annotations.
  - Annotation set is constrained by recommended types
    - webpage, place, review, song, movie, tvshow, book, product, stock, offer, topic, event.

```
"Just saw Avatar and it was amazing"  
"annotations": [  
  {  
    'movie':  
    {  
      'title': 'Avatar',  
      'url': 'http://www.rottentomatoes.com/m/avatar/',  
      'image': '...',  
      'text': 'Avatar'  
    }  
  }  
]
```

An example of a tweet talking about Avatar movie

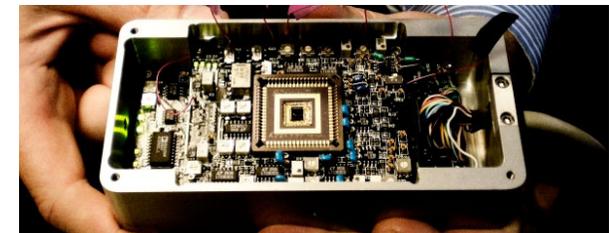
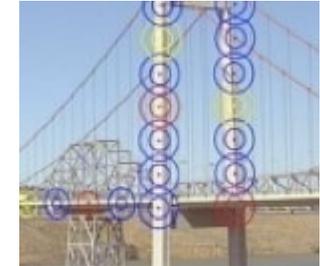
<sup>7</sup> Twitter Annotations Overview, [http://dev.twitter.com/pages/annotations\\_overview](http://dev.twitter.com/pages/annotations_overview)

# Motivation Scenarios

## Do we have events on the (Semantic) Web?

- Example 3: Internet of Things

- **Central Nervous System for the Earth (CeNSE) Project**<sup>8</sup> (HP Labs)
- A research and development program to build a **planetwide sensing network**, using billions of "tiny, cheap, tough and exquisitely sensitive detectors."
- Sensors detect vibrations, motion, light, temperature, barometric pressure, airflow and humidity.
- Possible use-cases<sup>3</sup>:
  - Warning about structural strains or weather conditions,
  - Monitor traffic, weather and road conditions,
  - Tracking hospital equipment,
  - Sniffing out pesticides and pathogens in food, etc.
- HP is hoping that at that scale, sensor nodes will cost "next to nothing, yet measure everything."<sup>9</sup> ☺



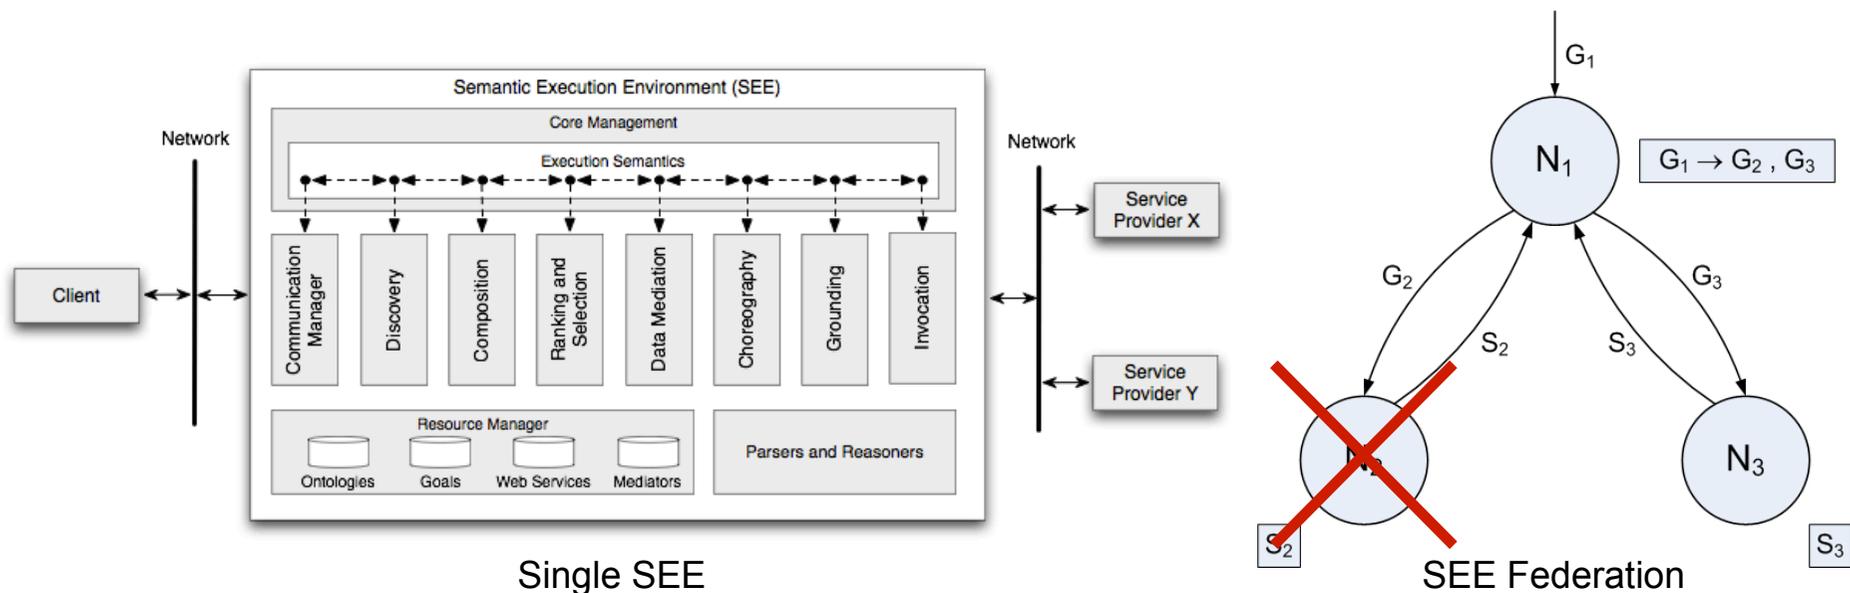
<sup>8</sup> [http://www.hpl.hp.com/research/quantum\\_systems](http://www.hpl.hp.com/research/quantum_systems)

<sup>9</sup> [http://www.readwriteweb.com/archives/cense\\_hp\\_labs.php](http://www.readwriteweb.com/archives/cense_hp_labs.php)

# Motivation Scenarios

## Do we have events on the (Semantic) Web?

- Example 4: Semantic Web Service Execution Environment
  - Closing the knowledge based control-loop in order to make the environment more robust and adaptive according to the real-time circumstances.

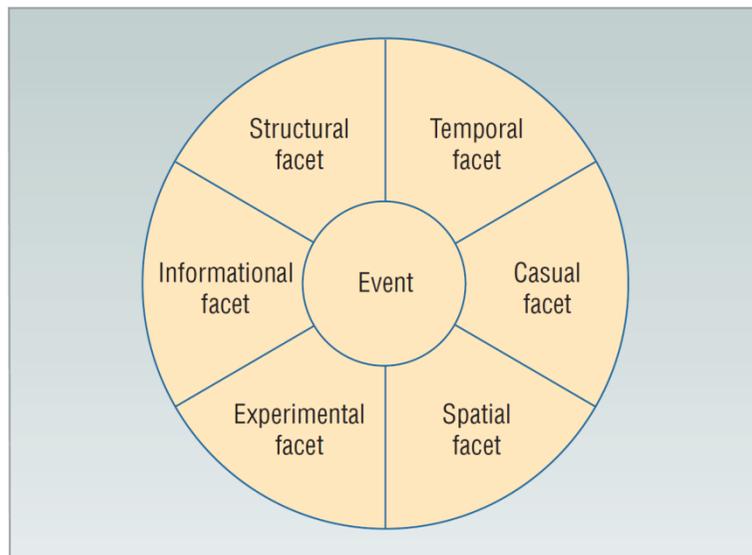


# *Event Processing*

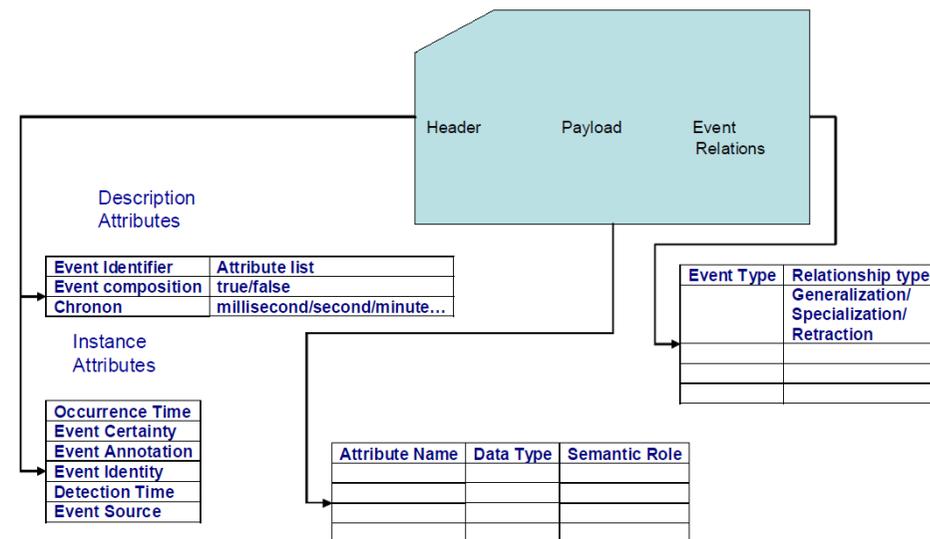
# Event Processing

## How does an event look like?

- More than one approach to describe an event



Different Facets of an Event<sup>3</sup>



The Event Type Definition Element<sup>4</sup>

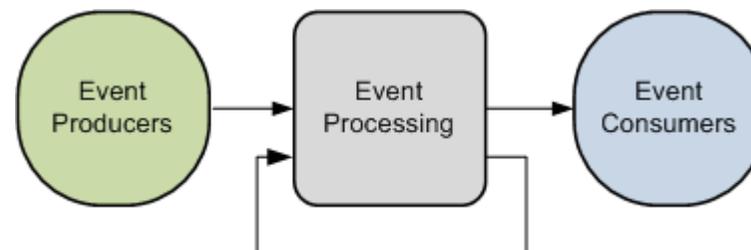
<sup>3</sup> Ramesh Jain, "EventWeb: Developing a Human-Centered Computing System," Computer, pp. 42-50, February, 2008

<sup>4</sup> Opher Etzion and Peter Niblett. Event Processing in Action. Manning Publications Co., 2010

# Event Processing

## What is event processing?

- What is *Event Processing*?
  - Event processing consists of **processing many events** happening across all the layers of an organization, **identifying the most meaningful events** within the event cloud, **analyzing their impact**, and **taking subsequent action in real time**.<sup>10</sup>
  - Operations that you can perform on events, in particular operations that take a set of one or more events as input and generate further events from them as output.<sup>11</sup>
- *Event Processing...*
  - ... usually focused on a real-world (i.e. physical) situations and problems,
  - ... provides a higher level of decoupling than the traditional processing approaches (like batch processing), and
  - ... exhibits certain level of abstraction.



<sup>10</sup> [http://en.wikipedia.org/wiki/Complex\\_event\\_processing](http://en.wikipedia.org/wiki/Complex_event_processing)

<sup>11</sup> Opher Etzion and Peter Niblett. Event Processing in Action. Manning Publications Co., 2010

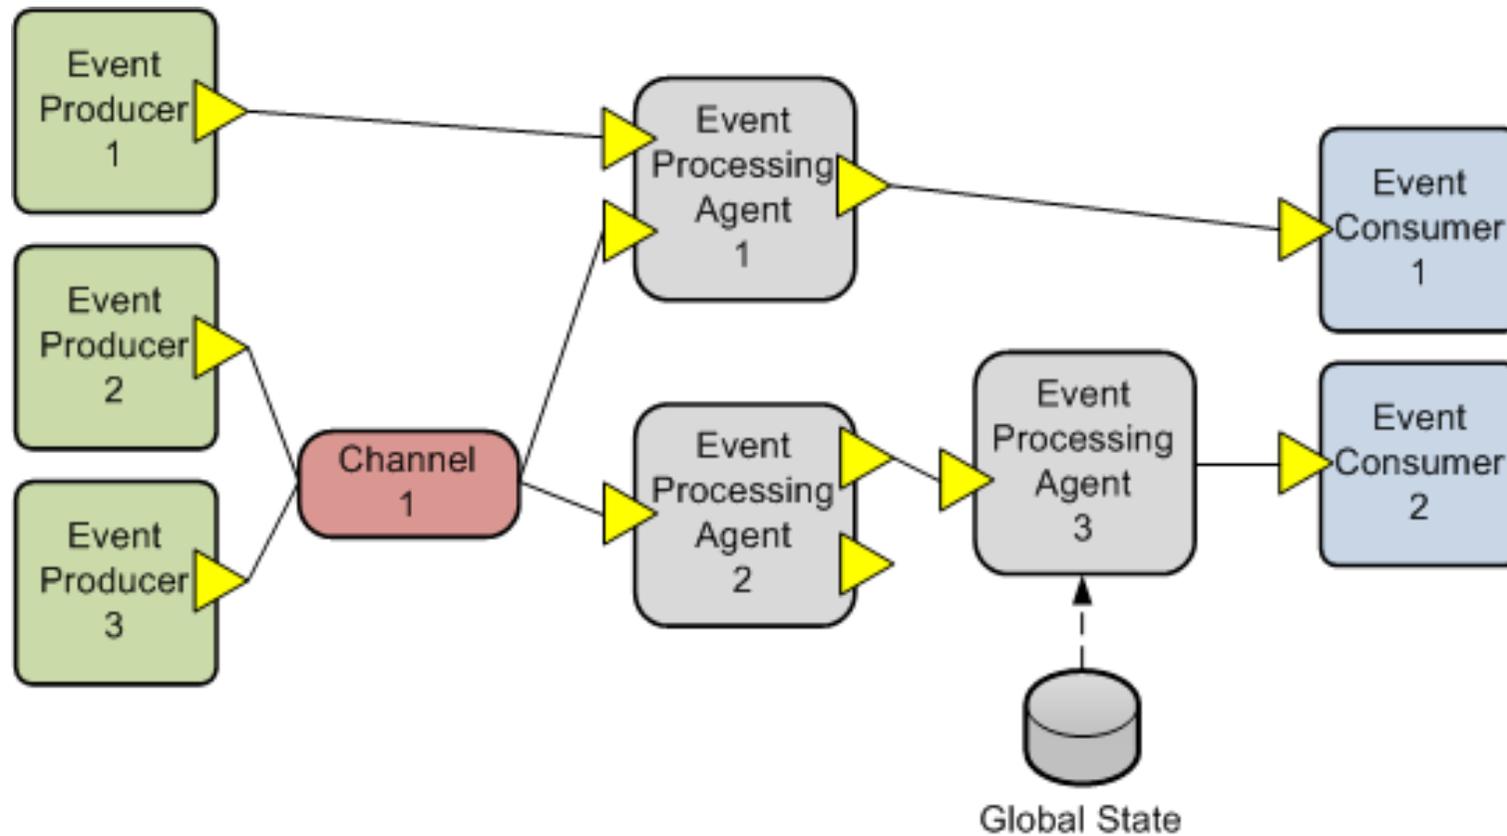
- Events filtering
  - Reducing the event stream to those fulfilling the constraints
  - E.g., filter out those events which source is geo-located outside of 10 km radius.
- Events aggregation
  - Computing values across the range of events
  - E.g., calculate average value of temperature readings in the last 10 minutes.
- Events derivation
  - Generating new events.
  - E.g., after selling the last item from stock initiate an order of new merchandise.
- Events causality
  - Calculating/exploiting causality vectors.
  - E.g., Drill down the cause of a derived event which brings system into a erroneous state.
- ...

- The approach of (Complex) Event Processing as a possible solution
  - Established by David Luckham (2002) and further refined by Etzion (2010).
  - Addressing issues of real-time event processing.
  - Focusing in particular on temporal, geo-spatial, and causal event dimensions.
  - Event pattern detection as the core task.
  - Built around the notion of Event Processing Networks.
- What is *Event Processing Network*?
  - The concept was popularized by the Luckham's book in 2002.
  - An *Event Processing Network* is a collection of event producers, consumers, event processing agents and global states which are connected by a collection of channels<sup>13</sup>.
  - The network can be built recursively out of the constituent elements.
  - The definition is abstract and platform independent (at the level of Platform Independent Model - PIM).

<sup>13</sup> Opher Etzion and Peter Niblett. Event Processing in Action. Manning Publications Co., 2010

# Problem Statement

## How to process events?



An example of Event Processing Network

# *Identified Problems*

# Problem Statement

## What happens when you apply EPN on the Web?



- Applying Event Processing Networks on the Semantic Web means:
  - Proper treatment of the notion of time and time dependent event relations
    - Can we expect synchronous time across the globe?
  - Events interoperability
    - Solutions may be segregated (current case with IoT applications).
    - Heterogeneity between the solutions may hinder the possibility to integrate different events.
  - Context-based event interpretation
    - An event interpretation depends on the context in which the event has been produced.
    - Context must be conveyed together with the event.
  - Event processing lifecycle matching the needs of the Web
    - Discovering and exploiting EPN elements on the Web.
    - Efficient pattern detection, events transformation and filtering over the Semantic Web events.

- Expected contribution

Extension of the Event Processing Network (EPN) framework with the Semantic Web technologies:

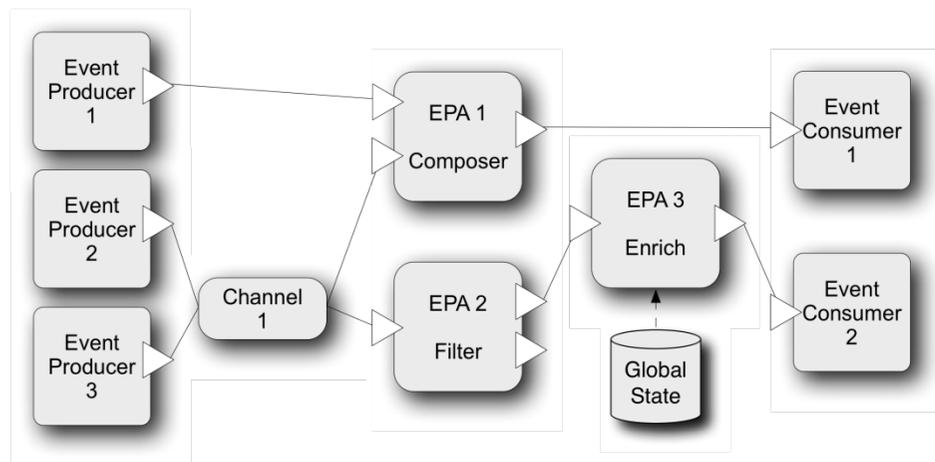
- Efficient event pattern detection mechanisms operating on top of events described in Semantic Web languages (i.e., RDF(S), OWL) taking into account the effect of inferred knowledge together with parameter, timing, geo-spatial context.
- Models built on top of the Semantic Web (Service) technologies for describing notions of event, event producers and consumers, event processing agents and their capabilities.

# ***Proposed Solution***

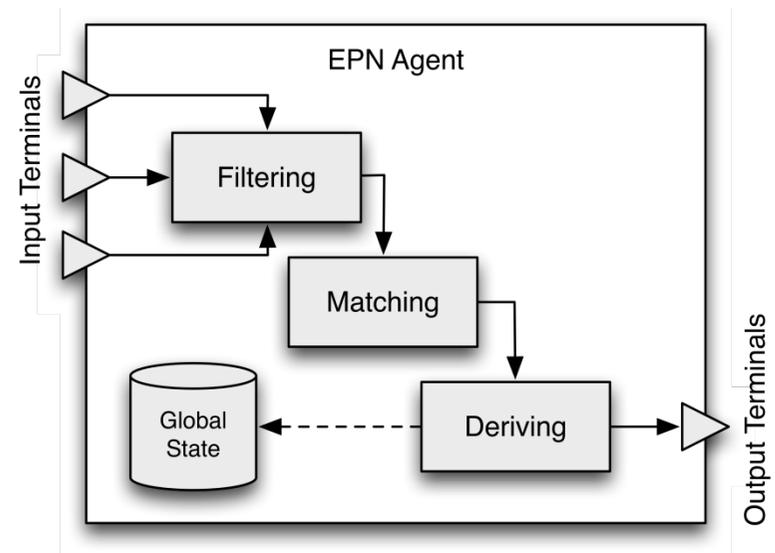
# Proposed Solution

## Event Processing Networks (EPN)

Conceptual Model of Event Processing Networks as defined by Luckham<sup>2</sup> and Entzion et al.<sup>3</sup>



An example of Event Processing Network



Architecture of EPN agent

<sup>2</sup> David Luckham. *The Power of Events: An Introduction to Complex Event Processing in Distributed Systems*. Addison-Wesley, 2001

<sup>3</sup> Opher Etzion and Peter Niblett. *Event Processing in Action*. Manning Publications Co., 2010

# Proposed Solution

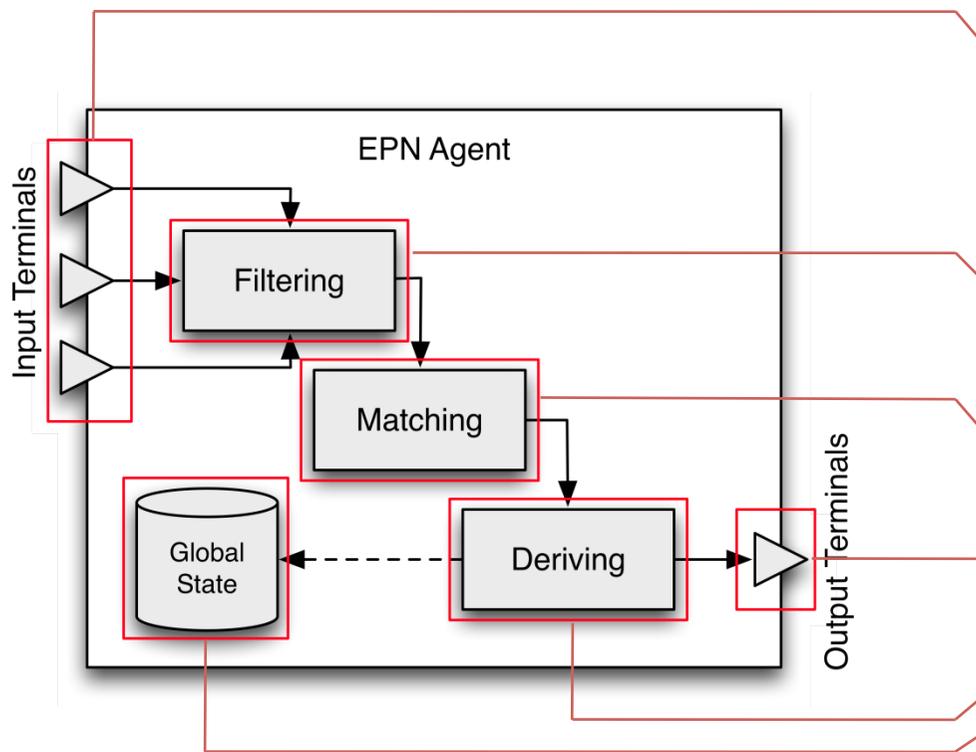
## Semantically-enhanced Event Processing Networks (sEPN)



- sEPN represents a grounding of conceptual EPN model into the Semantic Web technologies:
  - Ontological and comprehensive representation of `Event` concept.
  - Using ontology repositories to implement `Global States` and shared repositories.
  - Formal specification of `Event Producer`, `Event consumer`, `Event Channel`, `Event Processing Agent` properties, capabilities, and behavior.
- Current set of Semantic Web (Service) technologies provides a fruitful foundation to answer to some the identified gaps:
  - `SPARQL` can be used to specify behavior of EPA building blocks but it is missing general support for contexts (e.g., temporal, geo-spatial).
  - Query answering in SW data repositories is concentrating on `one-time` but not `continuous`.
  - Semantic Web Services frameworks are concentrating on `request-response` but not on `publish/subscribe` (i.e., notification) paradigm.

# Proposed Solution

## sEPN Agent Functionality Declaration



```

/* SEPN Agent functionality declaration */
SEPADecl ::= PrefixDecl*
          InTerDecl
          OutTerDecl
          FilterDecl?
          MatcherDecl?
          DeriverDecl?

/* Input and output terminals declaration */
InTerDecl ::= 'INPUT TERMINAL'
           (Term_Id Event_Type_Id)+
OutTerDecl ::= 'OUTPUT TERMINAL'
           (Term_Id Event_Type_Id)+

/* Filter function declaration */
FilterDecl ::= 'FILTER' Input_Term* Output_Term?
           Filter_Expr+
Filter_Expr ::= WhereClause

/* Matcher function declaration */
MatcherDecl ::= 'MATCHER' Input_Term* Output_Term?
           Matcher_Expr
Matcher_Expr ::= WhereClause

/* Deriver function declaration */
DeriverDecl ::= 'DERIVER' Input_Term* Output_Term
              Gl_State_Id? Deriver_Expr
Gl_State_Id ::= 'GLOBAL STATE' IRI_REF
Deriver_Expr ::= 'DERIVE' ConstructTriples
              'FROM' WhereClause

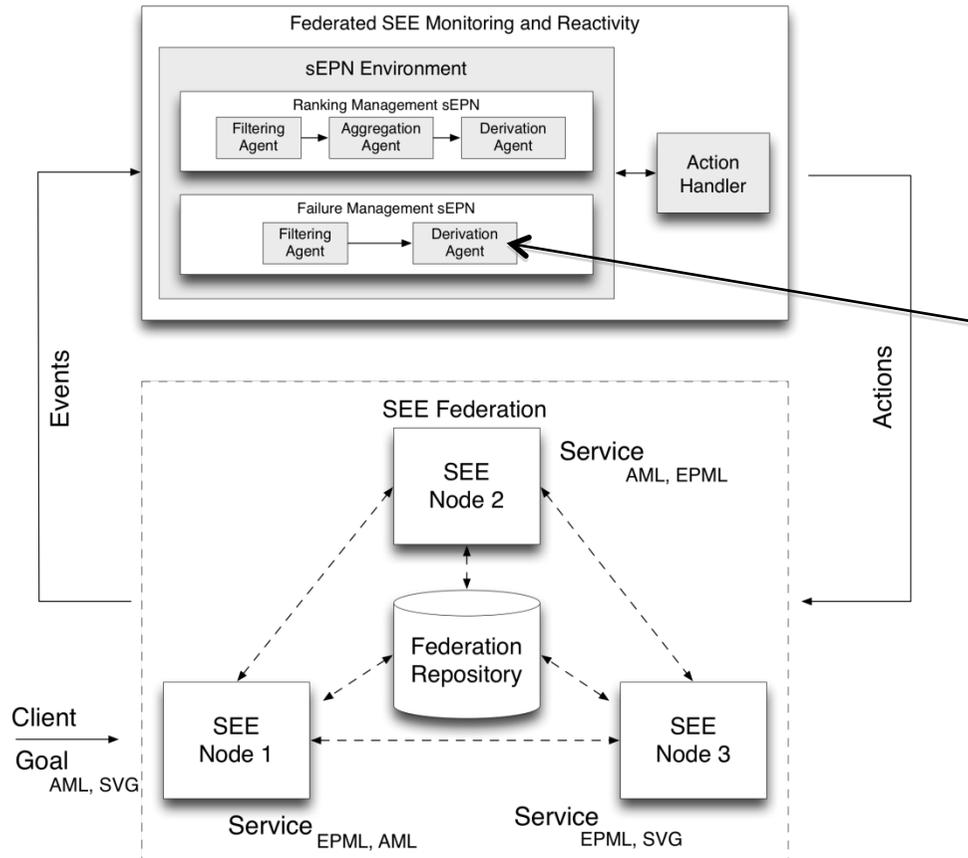
Input_Term ::= 'INPUT' Term_Id
Output_Term ::= 'OUTPUT' Term_Id
Term_Id ::= IRI_REF
Event_Type_Id ::= IRI_REF
IRI_REF ::= '<' ([^<>"{}|^\\]-[#x00-#x20])* '>'

```

Srdjan Komazec and Davide Cerri. *Enhancing Event Processing Networks with Semantics to Enable Self-Managed SEE Federations*. 3rd International Workshop on Monitoring, Adaptation and Beyond (MONA+) collocated with collocated with ECOWS, Ayia Napa, Cyprus, 2010. Under submission

# Proposed Solution

## sEPN Agent Functionality Declaration – An example



```

PREFIX
  sepn: <http://www.sepn.example/>
  seemon: <http://www.seemon.example/>
  rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
  xsd: <http://www.w3.org/2001/XMLSchema#>
  fed: <http://www.seefederation.example/>

INPUT TERMINAL
  <http://www.sepn.example/InputTerminal>
  seemon:AbortedEvent

OUTPUT TERMINAL
  <http://www.sepn.example/OutputTerminal>
  seemon:UpdateRepo

DERIVER
  INPUT <http://www.sepn.example/InputTerminal>
  OUTPUT <http://www.sepn.example/OutputTerminal>

  DERIVE {
    sepn:UpdateRepoX rdf:type seemon:UpdateRepo.
    sepn:UpdateRepoX seemon:repoId sepn:federatedRepo.
    sepn:UpdateRepoX seemon:update ?seeNode. ?seeNode
    fed:availabilityNFP "false"^^xsd:boolean.}
  FROM {
    ?abortedEvent rdf:type seemon:AbortedEvent.
    ?process seemon:event ?abortedEvent.
    ?process rdf:type seemon:FederatedWSInvocation.
    ?process seemon:agent ?seeNode.
    ?seeNode rdf:type seemon:SEE.
  }
  
```

COIN Enterprise Interoperability Use Case In Federated SEE Environment

# Proposed Solution

## Efficient Detection of Event Patterns on the Semantic Web



- $sEPA$  building block functions can be reduced to graph pattern detection problems.
- Current approaches like C-SPARQL<sup>4</sup> are separating the truth maintenance step (i.e., materialization of inferred knowledge over dynamically changing knowledge base) from the query answering step.
- Possible approach is to make knowledge derivation the integral part of graph pattern detection.
  - Taking RETE<sup>5</sup> as the starting point.
  - During Beta memory construction enrich the network with additional nodes injected as outcome of the analysis of inference rules which influence graph pattern structure.

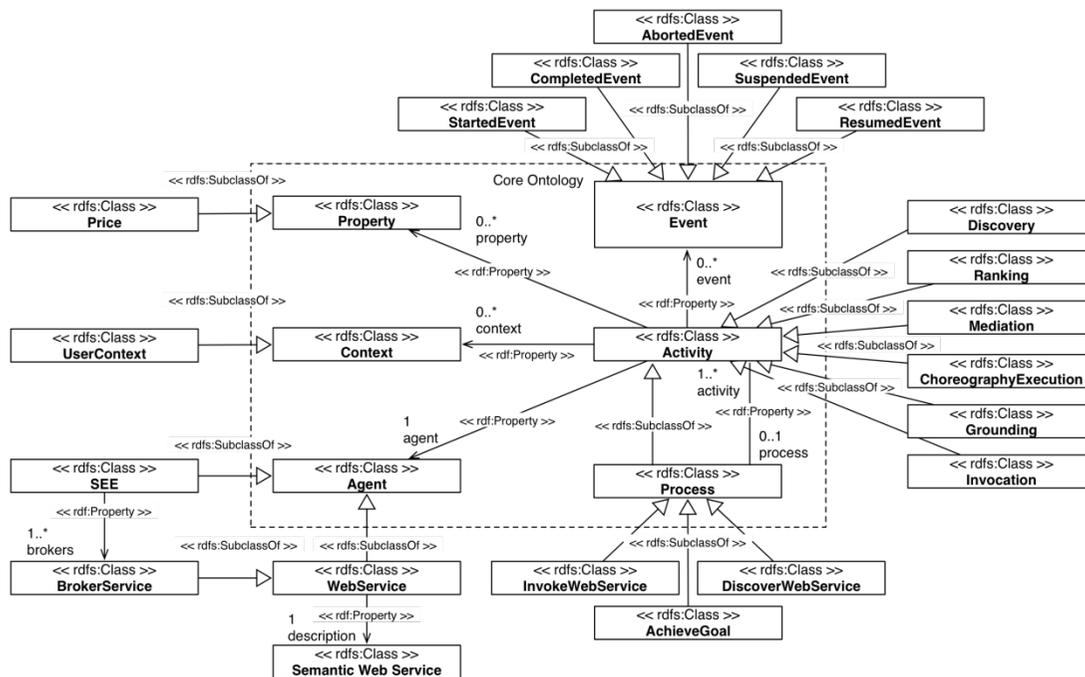
<sup>4</sup> Davide F. Barbieri, Daniele Braga, Stefano Ceri, Emanuele Della Vella, Michael Grossniklaus. C-SPARQL: A Continuous Query Language for RDF Data Streams. International Journal of Semantic Computing (IJSC), 2010

<sup>5</sup> Charles Forgy. "Rete: A Fast Algorithm for the Many Pattern/Many Object Pattern Match Problem", Artificial Intelligence, 19, pp. 17-37, 1982

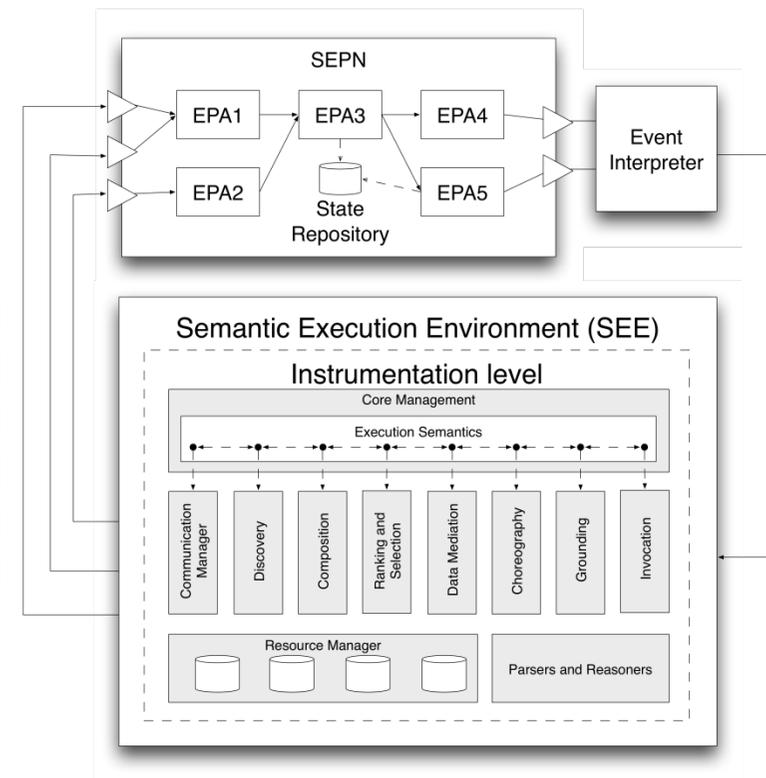
# *Evaluation*

- Study of the framework capabilities in the context of reference use-case
  - Fast flower delivery use-case by Etzion et al. (2010).
  - Expressiveness of the event pattern description language.
    - Time windowing, correlations of events, aggregation support.
  - Performance considerations.
- Event pattern detection engine performance evaluation
  - Measurements of throughput and scalability in the context of variable:
    - memory allocation,
    - frequency of arriving events, and
    - complexity of event pattern descriptions.

- Applying solution to the Semantic Execution Environment
  - Calculating aggregated statistics.
  - Communicating to external systems.



SEE Monitoring Ontology



SEE extended with SEPN

# *Future Work*

- Enriching  $sEPA$  functional capabilities
  - Declarative language with advanced concepts such as timing windows, geo-spatial-based event-relationships, supporting arbitrary event contexts, etc.
  - Accompanying Event Pattern Detection engine which supports those declarations.
- Design and implementation of formal EPN model
  - Investigating suitability and possible extensions to model EPN based on:
    - Current SOTA in Semantic Web Service frameworks.
    - Current SOTA in industrial efforts (W3C WS-Eventing, HP WS-Events).

# Questions?

