

Ensemble-oriented programming of self-adaptive systems with SCEL and jRESP

Michele Loreti



Dipartimento di Statistica, Informatica, Applicazioni
Università degli Studi di Firenze

AWASS 2013, Lucca, June 24-28, 2013

We aim at at developing linguistic supports for modelling (and programming) the behavior of service components and their ensembles, their interactions, their sensitivity and adaptivity to the environment

We aim at developing linguistic supports for modelling (and programming) the behavior of service components and their ensembles, their interactions, their sensitivity and adaptivity to the environment

SCEL

We aim at designing a specific language with

We aim at developing linguistic supports for modelling (and programming) the behavior of service components and their ensembles, their interactions, their sensitivity and adaptivity to the environment

SCEL

We aim at designing a specific language with

- **programming abstractions** necessary for
 - representing Knowledge, Behaviors and Aggregations according to specific Policies
 - programming interaction, adaptation and self- and context- awareness

We aim at developing linguistic supports for modelling (and programming) the behavior of service components and their ensembles, their interactions, their sensitivity and adaptivity to the environment

SCEL

We aim at designing a specific language with

- **programming abstractions** necessary for
 - representing Knowledge, Behaviors and Aggregations according to specific Policies
 - programming interaction, adaptation and self- and context- awareness
- linguistic primitives with **solid semantic grounds**
 - To develop logics, tools and methodologies for **formal reasoning** on systems behavior
 - To establish **qualitative and quantitative properties** of both the individual components and the ensembles

We need to enable programmers to model and describe the behavior of service components and their ensembles, their interactions, and their sensitivity and adaptivity to the environment they are working in.

We need to enable programmers to model and describe the behavior of service components and their ensembles, their interactions, and their sensitivity and adaptivity to the environment they are working in.

Notions to model

We need to enable programmers to model and describe the behavior of service components and their ensembles, their interactions, and their sensitivity and adaptivity to the environment they are working in.

Notions to model

- 1 The **behaviors** of components and their interactions

We need to enable programmers to model and describe the behavior of service components and their ensembles, their interactions, and their sensitivity and adaptivity to the environment they are working in.

Notions to model

- 1 The **behaviors** of components and their interactions
- 2 The **topology** of the network needed for interaction, taking into account resources, locations and visibility/reachability issues

We need to enable programmers to model and describe the behavior of service components and their ensembles, their interactions, and their sensitivity and adaptivity to the environment they are working in.

Notions to model

- 1 The **behaviors** of components and their interactions
- 2 The **topology** of the network needed for interaction, taking into account resources, locations and visibility/reachability issues
- 3 The **environment** where components operate and resource-negotiation takes place, taking into account open ended-ness and adaptation

We need to enable programmers to model and describe the behavior of service components and their ensembles, their interactions, and their sensitivity and adaptivity to the environment they are working in.

Notions to model

- 1 The **behaviors** of components and their interactions
- 2 The **topology** of the network needed for interaction, taking into account resources, locations and visibility/reachability issues
- 3 The **environment** where components operate and resource-negotiation takes place, taking into account open ended-ness and adaptation
- 4 The global **knowledge** of the systems and that of its components

We need to enable programmers to model and describe the behavior of service components and their ensembles, their interactions, and their sensitivity and adaptivity to the environment they are working in.

Notions to model

- 1 The **behaviors** of components and their interactions
- 2 The **topology** of the network needed for interaction, taking into account resources, locations and visibility/reachability issues
- 3 The **environment** where components operate and resource-negotiation takes place, taking into account open ended-ness and adaptation
- 4 The global **knowledge** of the systems and that of its components
- 5 The tasks to be accomplished, the properties to guarantee and the constraints to respect.

The Service-Component Ensemble Language (SCEL) currently provides primitives and constructs for dealing with:

The Service-Component Ensemble Language (SCEL) currently provides primitives and constructs for dealing with:

- 1 **Knowledge**: to describe how data, information and (local and global) knowledge is managed

The Service-Component Ensemble Language (SCEL) currently provides primitives and constructs for dealing with:

- ① **Knowledge**: to describe how data, information and (local and global) knowledge is managed
- ② **Behaviours**: to describe how systems of components progress

The Service-Component Ensemble Language (SCEL) currently provides primitives and constructs for dealing with:

- 1 **Knowledge**: to describe how data, information and (local and global) knowledge is managed
- 2 **Behaviours**: to describe how systems of components progress
- 3 **Aggregations**: to describe how different entities are brought together to form *components*, *systems* and *ensembles*

The Service-Component Ensemble Language (SCEL) currently provides primitives and constructs for dealing with:

- 1 **Knowledge**: to describe how data, information and (local and global) knowledge is managed
- 2 **Behaviours**: to describe how systems of components progress
- 3 **Aggregations**: to describe how different entities are brought together to form *components*, *systems* and *ensembles*
- 4 **Policies**: to model and enforce the wanted evolutions of computations.

1. Knowledge

SCEL is *parametric* wrt the means of managing knowledge that would depend on the specific class of application domains.

SCEL is *parametric* wrt the means of managing knowledge that would depend on the specific class of application domains.

Knowledge representation

- Tuples, Records
- Horn Clause Clauses,
- Concurrent Constraints,
- ...

SCEL is *parametric* wrt the means of managing knowledge that would depend on the specific class of application domains.

Knowledge handling mechanisms

- Pattern-matching, Reactive Tuple Spaces
- Data Bases Querying
- Resolution
- Constraint Solving
- ...

1. Knowledge (and Adaptation)

No definite stand is taken about the kind of knowledge that might depend on the application domain. To guarantee adaptivity, we, however, require there be some specific components.

No definite stand is taken about the kind of knowledge that might depend on the application domain. To guarantee adaptivity, we, however, require there be some specific components.

- Application data: Used for the progress of the computation.

No definite stand is taken about the kind of knowledge that might depend on the application domain. To guarantee adaptivity, we, however, require there be some specific components.

- Application data: Used for the progress of the computation.
- Control data: Providing information about the environment (e.g. data from sensors) and about the current status (e.g. its position or its battery level).

No definite stand is taken about the kind of knowledge that might depend on the application domain. To guarantee adaptivity, we, however, require there be some specific components.

- Application data: Used for the progress of the computation.
- Control data: Providing information about the environment (e.g. data from sensors) and about the current status (e.g. its position or its battery level).
- Knowledge handling mechanisms
 - **Add** information to a knowledge repository
 - **Retrieve** information from a knowledge repository
 - **Withdraw** information from a knowledge repository

Components behaviors are modeled as a process in the style of process calculi

Components behaviors are modeled as a process in the style of process calculi

- **Interaction** is obtained by allowing processes to access knowledge repositories, possibly of other components

Components behaviors are modeled as a process in the style of process calculi

- **Interaction** is obtained by allowing processes to access knowledge repositories, possibly of other components
- **Adaptation** is modeled by retrieving from the knowledge repositories
 - information about the changing environment and the component status
 - the code to execute for reacting to these changes.

Components behaviors are modeled as a process in the style of process calculi

- **Interaction** is obtained by allowing processes to access knowledge repositories, possibly of other components
- **Adaptation** is modeled by retrieving from the knowledge repositories
 - information about the changing environment and the component status
 - the code to execute for reacting to these changes.

Processes

nil | $a.P$ | $P_1 + P_2$ | $P_1[P_2]$ | X | $A(\bar{p})$ ($A(\bar{f}) \triangleq P$)

Components behaviors are modeled as a process in the style of process calculi

- **Interaction** is obtained by allowing processes to access knowledge repositories, possibly of other components
- **Adaptation** is modeled by retrieving from the knowledge repositories
 - information about the changing environment and the component status
 - the code to execute for reacting to these changes.

Processes

nil | $a.P$ | $P_1 + P_2$ | $P_1[P_2]$ | X | $A(\bar{p})$ ($A(\bar{f}) \triangleq P$)

The operators have the expected semantics. $P_1[P_2]$ (Controlled Composition) can be seen as a generalization of the many “parallel compositions” of process calculi. For the meaning of $a.-$, see next.

Actions

get(T)@ c | **qry**(T)@ c | **put**(t)@ c | **fresh**(n) | **new**($\mathcal{I}, \mathcal{K}, \Pi, P$)

1. Behaviours (and Actions)

Actions

get(T)@ c | **qry**(T)@ c | **put**(t)@ c | **fresh**(n) | **new**($\mathcal{I}, \mathcal{K}, \Pi, P$)

Action Targets

$c ::= n$ | x | self | P

Actions

get(T)@ c | **qry**(T)@ c | **put**(t)@ c | **fresh**(n) | **new**($\mathcal{I}, \mathcal{K}, \Pi, P$)

Action Targets

$c ::= n$ | x | self | P

Actions **manage knowledge** repositories by

- withdrawing information - **get**(T)@ c ,
- retrieving information - **qry**(T)@ c
- adding information - **put**(t)@ c

Actions operate on knowledge repository c and use T as a pattern to select knowledge items.

Aggregations

describe how different entities are brought together

- Model resource *allocation* and *distribution*
- Reflect the idea of *administrative domains*, i.e. the authority controlling a given set of resources and computing agents.
- are modelled by resorting to the notions of system, component and ensemble.

Aggregations

describe how different entities are brought together

- Model resource *allocation* and *distribution*
- Reflect the idea of *administrative domains*, i.e. the authority controlling a given set of resources and computing agents.
- are modelled by resorting to the notions of system, component and ensemble.

Systems

$$S ::= C \mid S_1 \parallel S_2 \mid ((\nu n))S$$

Components

$$C ::= \mathcal{I}[\mathcal{K}, \Pi, P]$$

Components

$$C ::= \mathcal{I}[\mathcal{K}, \Pi, P]$$

- An interface \mathcal{I} containing information about the component itself. In particular, each component C has attributes:
 - *id*: the name of the component C

Components

$$C ::= \mathcal{I}[\mathcal{K}, \Pi, P]$$

- An interface \mathcal{I} containing information about the component itself. In particular, each component C has attributes:
 - *id*: the name of the component C
- A knowledge manager \mathcal{K} providing control data (i.e. the local and (part of the) global knowledge) and application data; together with a specific knowledge handling mechanism

Components

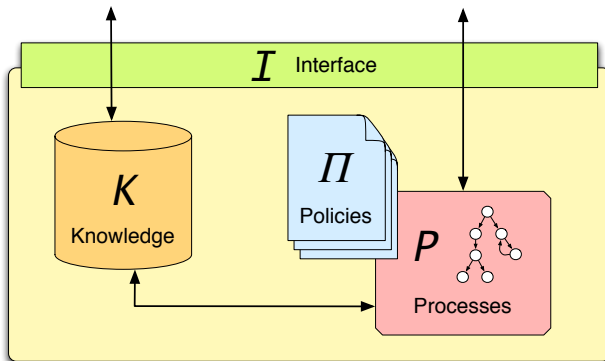
$$C ::= \mathcal{I}[\mathcal{K}, \Pi, P]$$

- An interface \mathcal{I} containing information about the component itself. In particular, each component C has attributes:
 - *id*: the name of the component C
- A knowledge manager \mathcal{K} providing control data (i.e. the local and (part of the) global knowledge) and application data; together with a specific knowledge handling mechanism
- A set of policies Π regulating inter-component and intra-component interactions

Components

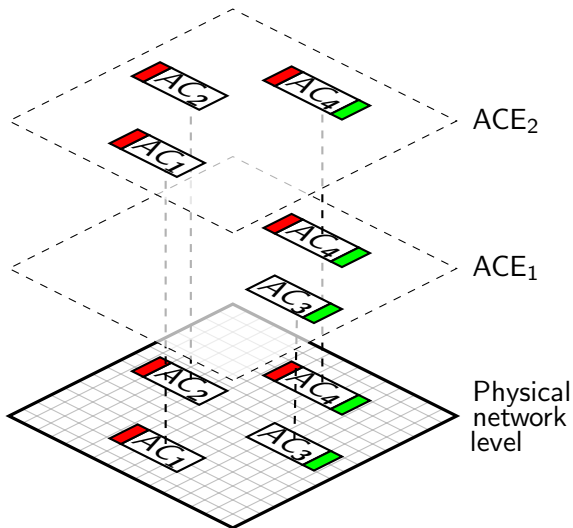
$$C ::= \mathcal{I}[\mathcal{K}, \Pi, P]$$

- An interface \mathcal{I} containing information about the component itself. In particular, each component C has attributes:
 - *id*: the name of the component C
- A knowledge manager \mathcal{K} providing control data (i.e. the local and (part of the) global knowledge) and application data; together with a specific knowledge handling mechanism
- A set of policies Π regulating inter-component and intra-component interactions
- A process term P that performs the local computation, coordinates their interaction with the knowledge repository and deals with adaptation and reconfiguration



Programming Abstractions

Important for improving code productivity



Policies deal with the way properties of computations are represented and enforced

- Interaction: interaction predicates, ...
 - Resource usage: accounting, leasing, ...
 - Security: access control, trust, reputation, ...
-
- SCEL is *parametric* wrt the actual language used to express policies.
 - Currently we (Pugliese, Tiezzi) are defining a specific language based on XACML.
 - When considering the operational semantics, we will see how policies are exploited to control components actions, their evolutions and their interactions.

SYSTEMS: $S ::= C \mid S_1 \parallel S_2 \mid ((\nu n))S$

COMPONENTS: $C ::= \mathcal{I}[\mathcal{K}, \Pi, P]$

KNOWLEDGE: $K ::= \dots$

PROCESSES: $P ::= \mathbf{nil} \mid a.P \mid P_1 + P_2 \mid P_1[P_2] \mid X \mid A(\bar{p}) \mid A(\bar{f})$

ACTIONS: $a ::= \mathbf{get}(T)@c \mid \mathbf{qry}(T)@c \mid \mathbf{put}(t)@c \mid \mathbf{fresh}(n) \mid \mathbf{new}(\mathcal{I}, \mathcal{K},$

TARGETS: $c ::= n \mid x \mid \mathbf{self}$

ITEMS: $t ::= \dots$ – for the moment just tuples

TEMPLATES: $T ::= \dots$ – for the moment tuples with variables

Basic design principles. . .

Basic design principles. . .

- 1 no centralized control

Basic design principles. . .

- ① no centralized control
- ② heavy use of *recurrent patterns* to simplify the development of specific
 - policies
 - knowledge
 - . . .

Basic design principles. . .

- ① no centralized control
- ② heavy use of *recurrent patterns* to simplify the development of specific
 - policies
 - knowledge
 - . . .
- ③ use of *open technologies* to support the integration with other tools/frameworks
 - Argos
 - DEECo
 - . . .

SYSTEMS:

$$S ::= C \mid S_1 \parallel S_2 \mid (\nu n)S$$

COMPONENTS:

$$C ::= \mathcal{I}[\mathcal{K}, \Pi, P]$$

PROCESSES:

$$P ::= \mathbf{nil} \mid a.P \mid P_1 + P_2 \mid P_1[P_2] \mid X \mid A(\bar{p}) \quad (A(\bar{f}) \triangleq P)$$

ACTIONS:

$$a ::= \mathbf{get}(T)@c \mid \mathbf{qry}(T)@c \mid \mathbf{put}(t)@c \mid \mathbf{exec}(P) \mid \mathbf{new}(\mathcal{I}, \mathcal{K}, \Gamma)$$

TARGETS:

$$c ::= n \mid x \mid \mathbf{self} \mid P$$

SYSTEMS:

$$S ::= C \mid S_1 \parallel S_2 \mid (\nu n)S$$

COMPONENTS:

$$C ::= \mathcal{I}[\mathcal{K}, \Pi, P]$$

PROCESSES:

$$P ::= \mathbf{nil} \mid a.P \mid P_1 + P_2 \mid P_1[P_2] \mid X \mid A(\bar{p}) \quad (A(\bar{f}) \triangleq P)$$

ACTIONS:

$$a ::= \mathbf{get}(T)@c \mid \mathbf{qry}(T)@c \mid \mathbf{put}(t)@c \mid \mathbf{exec}(P) \mid \mathbf{new}(\mathcal{I}, \mathcal{K}, \Gamma)$$

TARGETS:


$$c ::= n \mid x \mid \mathbf{self} \mid P$$

Hardware/Virtual Machine

Hardware/Virtual Machine


Networks


A vertical label "Networks" with two small pink rectangular boxes stacked vertically to its left, indicating a connection or interface.

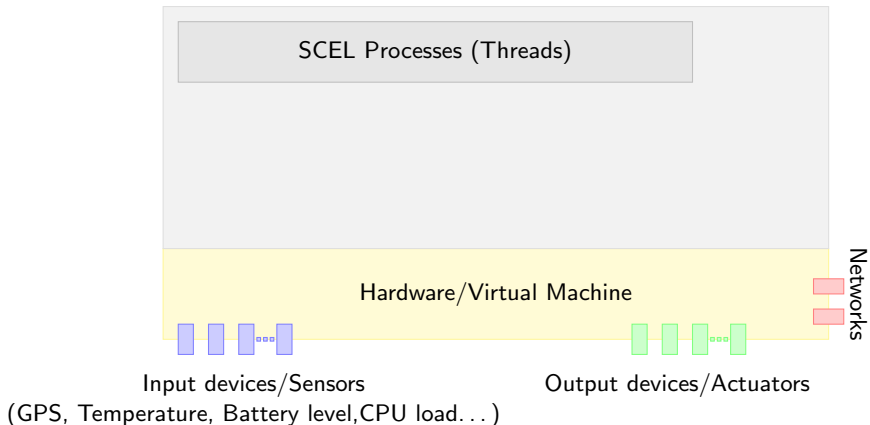


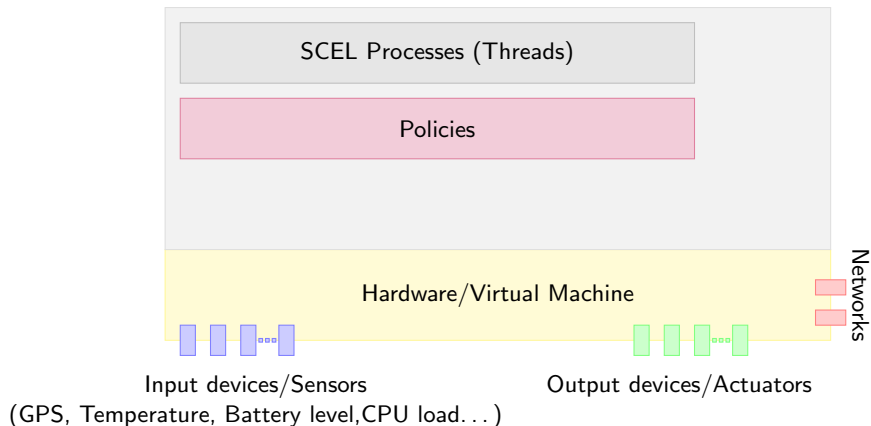
Input devices/Sensors

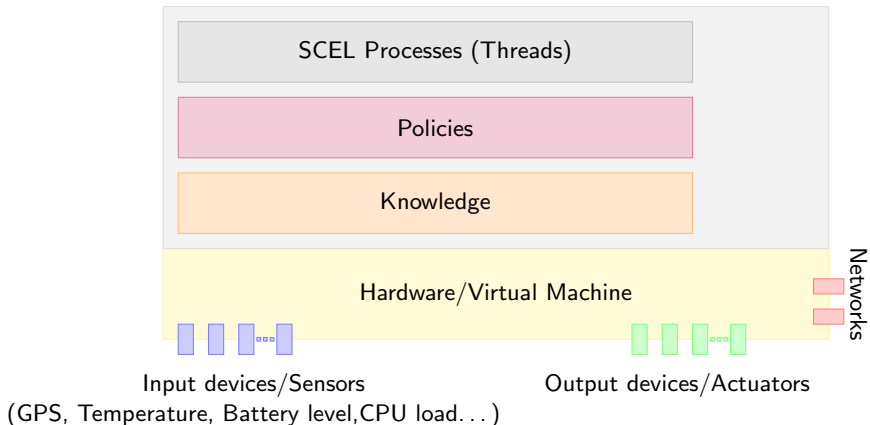
(GPS, Temperature, Battery level,CPU load...)

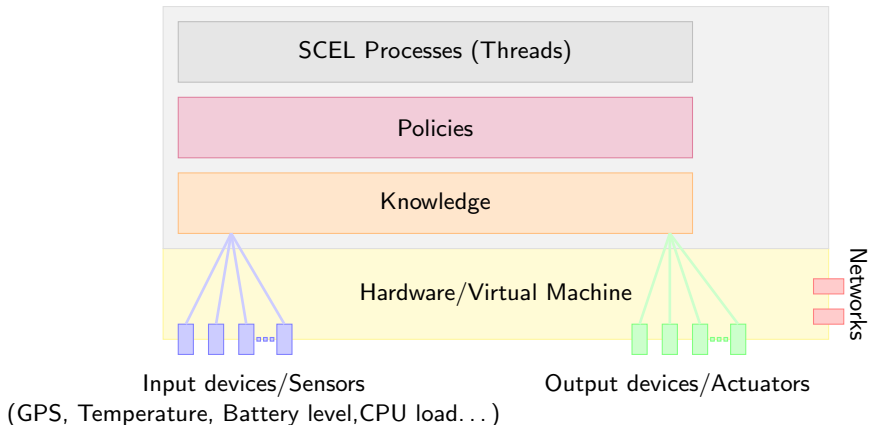


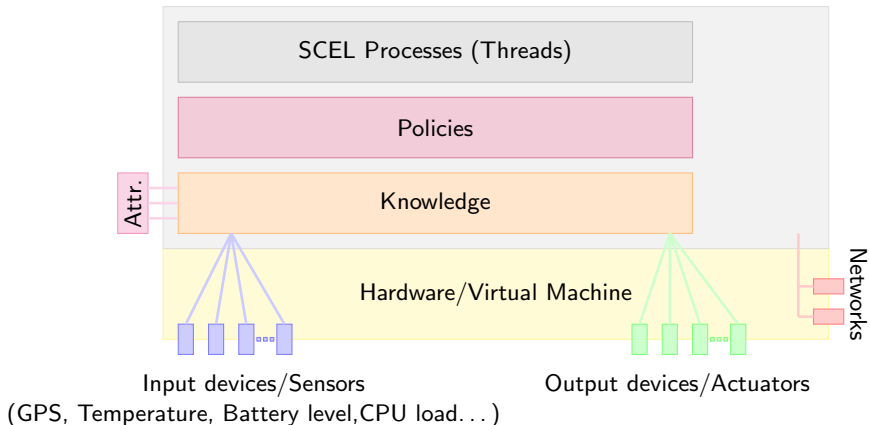


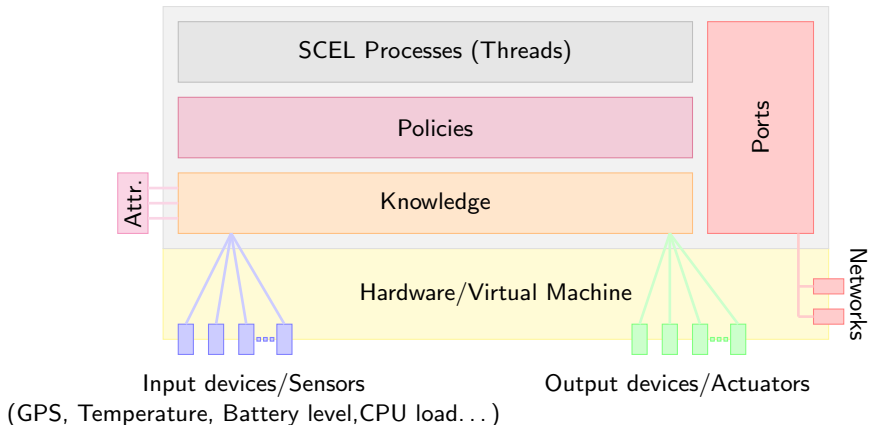












Components abstract from specific *Knowledge* implementations.

Components abstract from specific *Knowledge* implementations.

Indeed, *Knowledge* is a Java interface with the following methods:

- `put(Tuple t)`: boolean
- `get(Template t)`: Tuple
- `query(Template t)`: Tuple

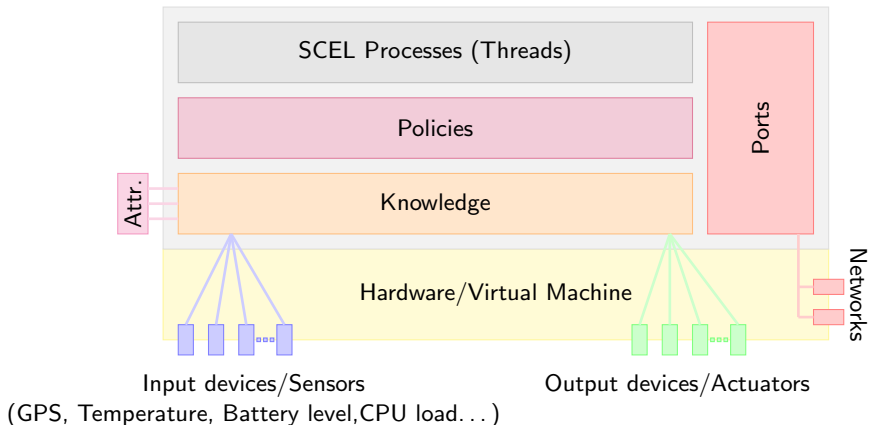
Components abstract from specific *Knowledge* implementations.

Indeed, *Knowledge* is a Java interface with the following methods:

- put(Tuple t): boolean
- get(Template t): Tuple
- query(Template t): Tuple

A single implementation of this interface is currently provided:

- class TupleSpace implements a *tuple space* (*à la klaim*)





Abstract class `Sensor` is used to identify a generic source of information:


- it can be associated to a logical/physical sensor
- values are exported as a *tuple*, each implementation has to define the structure of the tuple containing
 - (“GPS”, 45.8 , 37.2)
 - (“BATTERY” , %87)
- query actions are used to *retrieve* data from sensor


Abstract class Actuator is used to identify an external device that can be controlled by SCEL processes:

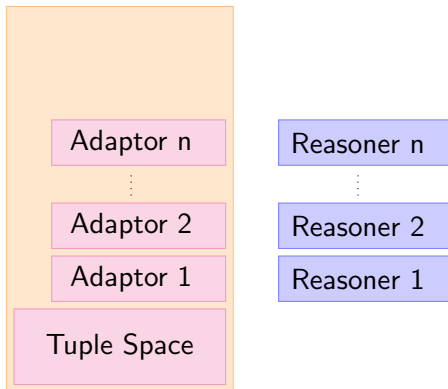
- it can be associated to a logical/physical actuator
- values are passed as a *tuple*, each implementation has to define the structure of the tuple containing
 - (“DIRECTION”, $\pi/3$)
- values are passed to a actuator via put actions

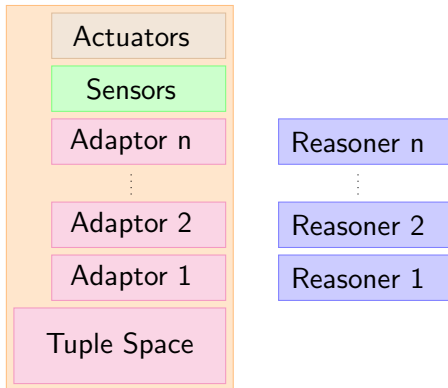













put(t)

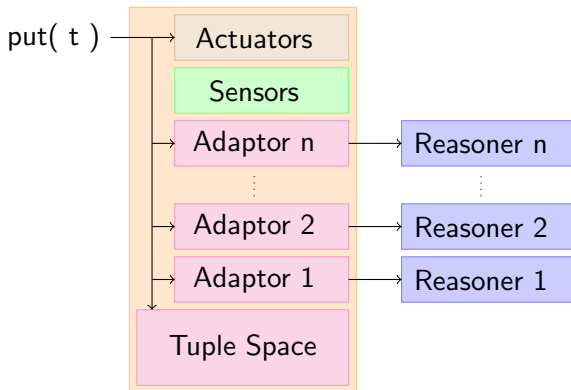


Reasoner n


⋮

Reasoner 2

Reasoner 1




When a `put(t)` is invoked, according to the structure of `t`, the *right* element is selected.



Adapters will use a *put* to add a knowledge item (e.g. a fact) represented by t in the corresponding reasoner.

query(T)

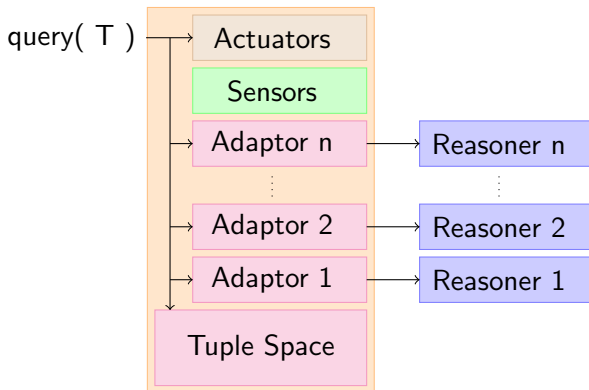


Reasoner n

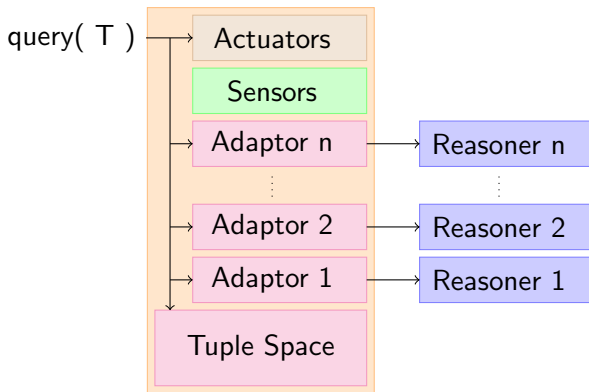
⋮

Reasoner 2

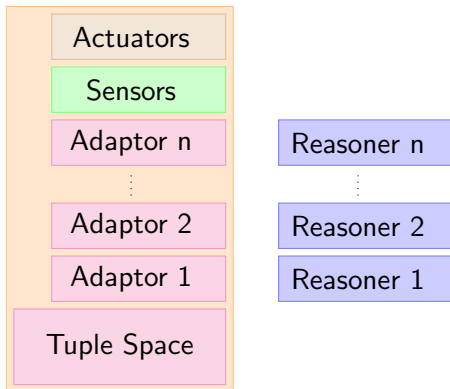
Reasoner 1



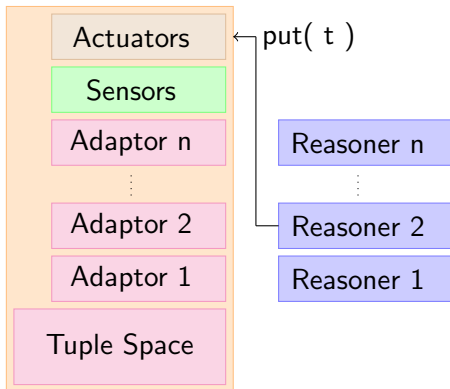
Similarly, a *get* will be *dispatched* according to template T .



Adapters will use a *query* to ask to the corresponding reasoner about a *query* described via *T*.



Reasoners can *autonomously* interact with the knowledge!



Reasoners can *autonomously* interact with the knowledge!

Action execution on each component is regulated by a *policy*.

Action execution on each component is regulated by a *policy*.

This is an interface that provides methods like:

- put(Agent a, Tuple t , Target l)
- acceptGet(Locality l , Template t)

each node *forwards* received request to its policy.

Action execution on each component is regulated by a *policy*.

This is an interface that provides methods like:

- put(Agent a, Tuple t , Target l)
- acceptGet(Locality l , Template t)

each node *forwards* received request to its policy.

Action execution on each component is regulated by a *policy*.

This is an interface that provides methods like:

- put(Agent a, Tuple t , Target l)
- acceptGet(Locality l , Template t)

each node *forwards* received request to its policy.

Policies are organized in a *stack*:

- the policy at one level relies on the one at the level below to actually execute SCEL actions
- the policy at the lower level is the one that allows any operation

Each node is equipped with a set of *ports* that are able to handle:

- *point-to-point* interactions
- *group* interactions (*ensemble* oriented)

Each node is equipped with a set of *ports* that are able to handle:

- *point-to-point* interactions
- *group* interactions (*ensemble* oriented)

Currently the following ports have been developed:

Each node is equipped with a set of *ports* that are able to handle:

- *point-to-point* interactions
- *group* interactions (*ensemble* oriented)

Currently the following ports have been developed:

- InetPort, this kind of ports uses TCP to *point-to-point* interactions and UDP for the *group* ones

Each node is equipped with a set of *ports* that are able to handle:

- *point-to-point* interactions
- *group* interactions (*ensemble* oriented)

Currently the following ports have been developed:

- InetPort, this kind of ports uses TCP to *point-to-point* interactions and UDP for the *group* ones
- ServerPort, in this case a centralized server is used to collect and dispatch nodes' actions

Each node is equipped with a set of *ports* that are able to handle:

- *point-to-point* interactions
- *group* interactions (*ensemble* oriented)

Currently the following ports have been developed:

- InetPort, this kind of ports uses TCP to *point-to-point* interactions and UDP for the *group* ones
- ServerPort, in this case a centralized server is used to collect and dispatch nodes' actions
- VirtualPort, this is used to *simulate* nodes running on *virtual* devices

Each node is equipped with a set of *ports* that are able to handle:

- *point-to-point* interactions
- *group* interactions (*ensemble* oriented)

Currently the following ports have been developed:

- InetPort, this kind of ports uses TCP to *point-to-point* interactions and UDP for the *group* ones
- ServerPort, in this case a centralized server is used to collect and dispatch nodes' actions
- VirtualPort, this is used to *simulate* nodes running on *virtual* devices

To simplify interactions with other framework/tools, even developed in languages that are different from Java, JSON format is used.

$\text{get}(T)@P$

`get(T)@P`

- 1 A request message containing template T and attribute names occurring in P is sent via a *multicast* channel

`get(T)@P`

- 1 A request message containing template T and attribute names occurring in P is sent via a *multicast* channel
- 2 All the components that receives the message reply with:
 - a tuple t matching template T
 - a list of requested attributes

$\text{get}(T)@P$

- 1 A request message containing template T and attribute names occurring in P is sent via a *multicast* channel
- 2 All the components that receives the message reply with:
 - a tuple t matching template T
 - a list of requested attributes
- 3 If the node originating the requests receives attributes that *satisfies* predicate P , action is executed. Otherwise, the request is sent again.

We have integrated a new port based on *FreePastry* framework.

- point-to-point interactions relies on DHT primitives (the key is the node name);
- group-oriented interactions are implemented via multicast.

We have integrated a new port based on *FreePastry* framework.

- point-to-point interactions relies on DHT primitives (the key is the node name);
- group-oriented interactions are implemented via multicast.

Group-oriented via multicast:

We have integrated a new port based on *FreePastry* framework.

- point-to-point interactions relies on DHT primitives (the key is the node name);
- group-oriented interactions are implemented via multicast.

Group-oriented via multicast:

- We associate a *topic* to each predicate;

We have integrated a new port based on *FreePastry* framework.

- point-to-point interactions relies on DHT primitives (the key is the node name);
- group-oriented interactions are implemented via multicast.

Group-oriented via multicast:

- We associate a *topic* to each predicate;
- When a predicate is satisfied/unsatisfied a component register/deregister for a *topic*;

We have integrated a new port based on *FreePastry* framework.

- point-to-point interactions relies on DHT primitives (the key is the node name);
- group-oriented interactions are implemented via multicast.

Group-oriented via multicast:

- We associate a *topic* to each predicate;
- When a predicate is satisfied/unsatisfied a component register/deregister for a *topic*;
- Operations on a predicate are then realised via a multicast on the corresponding topic;

We have integrated a new port based on *FreePastry* framework.

- point-to-point interactions relies on DHT primitives (the key is the node name);
- group-oriented interactions are implemented via multicast.

Group-oriented via multicast:

- We associate a *topic* to each predicate;
- When a predicate is satisfied/unsatisfied a component register/deregister for a *topic*;
- Operations on a predicate are then realised via a multicast on the corresponding topic;
- A *special* topic is used to coordinate activities.

Robots in the swarm are distributed over a physical area and have to reach different zones according an assigned task:

- each robot has to fulfil one of two different tasks ($task_1$ or $tasks_2$).

Robots in the swarm are distributed over a physical area and have to reach different zones according an assigned task:

- each robot has to fulfil one of two different tasks ($task_1$ or $tasks_2$).

Robots are unaware about the position of the two target zones:

Robots in the swarm are distributed over a physical area and have to reach different zones according an assigned task:

- each robot has to fulfil one of two different tasks ($task_1$ or $tasks_2$).

Robots are unaware about the position of the two target zones:

- to discover the location of the target, robots follow a *random walk*;

Robots in the swarm are distributed over a physical area and have to reach different zones according an assigned task:

- each robot has to fulfil one of two different tasks ($task_1$ or $tasks_2$).

Robots are unaware about the position of the two target zones:

- to discover the location of the target, robots follow a *random walk*;
- when a robot reaches the area, it 'publishes' its location in the local knowledge repository...

Robots in the swarm are distributed over a physical area and have to reach different zones according an assigned task:

- each robot has to fulfil one of two different tasks ($task_1$ or $tasks_2$).

Robots are unaware about the position of the two target zones:

- to discover the location of the target, robots follow a *random walk*;
- when a robot reaches the area, it 'publishes' its location in the local knowledge repository. . .
 - robots with the same *task* can get aware about position of target area.

Robots in the swarm are distributed over a physical area and have to reach different zones according an assigned task:

- each robot has to fulfil one of two different tasks ($task_1$ or $tasks_2$).

Robots are unaware about the position of the two target zones:

- to discover the location of the target, robots follow a *random walk*;
- when a robot reaches the area, it 'publishes' its location in the local knowledge repository...
 - robots with the same *task* can get aware about position of target area.
- robots stop moving when the level of their batteries goes under a given threshold.

Q: What are the “ensembles” in the considered scenario?

Q: What are the “ensembles” in the considered scenario?

A: The set of robots that can fulfil the same task.

Q: What are the “ensembles” in the considered scenario?

A: The set of robots that can fulfil the same task.

NB: There is no central coordinator!

Q: What are the “ensembles” in the considered scenario?

A: The set of robots that can fulfil the same task.

NB: There is no central coordinator!

Q: How can “ensembles” be identified?

Q: What are the “ensembles” in the considered scenario?

A: The set of robots that can fulfil the same task.

NB: There is no central coordinator!

Q: How can “ensembles” be identified?

A: Each robot publishes in its interface the task that it can fulfil.

Q: What are the “ensembles” in the considered scenario?

A: The set of robots that can fulfil the same task.

NB: There is no central coordinator!

Q: How can “ensembles” be identified?

A: Each robot publishes in its interface the task that it can fulfil.

By relying on *group-oriented* queries robots in the same ensemble can get aware about the position of the zone and then can move directly towards the target.

Each robot is a SCEL component with two main behaviours:

- *managed element (ME)*, that executes the next *control step* retrieved from local knowledge;
- *autonomic manager (AM)*, that depending on info received from the environment puts in the knowledge the next *control step*.

Each robot is a SCEL component with two main behaviours:

- *managed element (ME)*, that executes the next *control step* retrieved from local knowledge;
- *autonomic manager (AM)*, that depending on info received from the environment puts in the knowledge the next *control step*.

Info from the environment includes:

Each robot is a SCEL component with two main behaviours:

- *managed element (ME)*, that executes the next *control step* retrieved from local knowledge;
- *autonomic manager (AM)*, that depending on info received from the environment puts in the knowledge the next *control step*.

Info from the environment includes:

- values retrieved from robot's sensors (GPS sensor, target sensor,...);

Each robot is a SCEL component with two main behaviours:

- *managed element (ME)*, that executes the next *control step* retrieved from local knowledge;
- *autonomic manager (AM)*, that depending on info received from the environment puts in the knowledge the next *control step*.

Info from the environment includes:

- values retrieved from robot's sensors (GPS sensor, target sensor,...);
- messages collected/received from other robots.

Each robot is a SCEL component with two main behaviours:

- *managed element (ME)*, that executes the next *control step* retrieved from local knowledge;
- *autonomic manager (AM)*, that depending on info received from the environment puts in the knowledge the next *control step*.

Info from the environment includes:

- values retrieved from robot's sensors (GPS sensor, target sensor,...);
- messages collected/received from other robots.

Self-adaptation is realised by exploiting SCEL higher-order features:

- autonomic manager *AM* implements the adaptation logic by replacing the control step code in knowledge repository.

A component is created when an object of class Node is instantiated:

```
Node<TupleSpace> n =  
    new Node<TupleSpace>(  
        "Robot"+i,  
        new TupleSpace()  
    );  
n.addPort(vp);
```

A component is created when an object of class Node is instantiated:

```
Node<TupleSpace> n =  
    new Node<TupleSpace>(  
        "Robot"+i,  
        new TupleSpace()  
    );  
n.addPort(vp);
```

Class Node is parametrized with respect a class that implements the interface characterising a *knowledge manager*

A component is created when an object of class Node is instantiated:

```
Node<TupleSpace> n =  
    new Node<TupleSpace>(  
        "Robot"+i,  
        new TupleSpace()  
    );  
n.addPort(vp);
```

Class Node is parametrized with respect a class that implements the interface characterising a *knowledge manager*

- in the example above we use TupleSpace

A component is created when an object of class Node is instantiated:

```
Node<TupleSpace> n =  
    new Node<TupleSpace>(  
        "Robot"+i,  
        new TupleSpace()  
    );  
n.addPort(vp);
```

Class Node is parametrized with respect a class that implements the interface characterising a *knowledge manager*

- in the example above we use TupleSpace

New *knowledge managers*, **for instance the one based on KnowLang**, can be easily integrated!

The interface Knowledge identifies a generic knowledge repository and indicates the high-level primitives to manage pieces of relevant information coming from different sources.

The interface Knowledge identifies a generic knowledge repository and indicates the high-level primitives to manage pieces of relevant information coming from different sources.

This interface contains the methods for withdrawing/retrieving/adding piece of knowledge from/to a repository:

- `put(Template t)`: Tuple
- `get(Template t)`: Tuple
- `query(Template t)`: Tuple

The interface Knowledge identifies a generic knowledge repository and indicates the high-level primitives to manage pieces of relevant information coming from different sources.

This interface contains the methods for withdrawing/retrieving/adding piece of knowledge from/to a repository:

- `put(Template t)`: Tuple
- `get(Template t)`: Tuple
- `query(Template t)`: Tuple

Sensors and Actuators:

- external data can be collected into a knowledge repository via *sensors*;
- *actuators* can be used to forward data to external components.

```
n.addActuator(scenario.getDirectionActuator(i));  
n.addSensor(scenario.getLocationSensor(i));  
n.addActuator(scenario.getStopActuator(i));  
n.addSensor(scenario.getBatterySensor(i));  
n.addSensor(scenario.getTargetSensor(i));
```



```
n.addActuator(scenario.getDirectionActuator(i));  
n.addSensor(scenario.getLocationSensor(i));  
n.addActuator(scenario.getStopActuator(i));  
n.addSensor(scenario.getBatterySensor(i));  
n.addSensor(scenario.getTargetSensor(i));
```

Remark:

In the code above, *scenario* is the Java classes modelling the *physical environment* where robots work.

Attribute values are published on component interfaces via *attribute collectors*:

- when a request for an attribute is received, the corresponding collector is selected.
- node's knowledge is used to compute the actual attribute value.

Attribute values are published on component interfaces via *attribute collectors*:

- when a request for an attribute is received, the corresponding collector is selected.
- node's knowledge is used to compute the actual attribute value.

Each attribute collector is associated with a *name* and a *template*:

- when the attribute is evaluated a tuple matching the template is retrieved (via a predicative query action);
- the retrieved tuple is used to compute the actual attribute value.

```
n.addAttributeCollector(  
    new AttributeCollector(  
        "task",  
        new Template(  
            new ActualTemplateField( "task"),  
            new FormalTemplateField(Integer.class)  
        )  
    ) {  
        @Override  
        protected Attribute doEval(Tuple t) {  
            return new Attribute(  
                "task",  
                t.getElementAt(Integer.class, 1));  
        }  
    }  
);
```


SCEL processes are implemented as threads via the abstract class *Agent* which provides the methods for:

- executing another agent,
- for generating fresh names,
- for instantiating a new component
- and for withdrawing/retrieving/adding information items from/to shared knowledge repositories.

```
public class ManagedElement extends Agent {
    public ManagedElement() {
        super("ManagedElement");
    }
    @Override
    protected void doRun() throws Exception {
        while (true) {
            Tuple t = query(new Template(
                new ActualTemplateField("controlStep")
                ,
                new FormalTemplateField(Agent.class)) ,
                Self.SELF );
            Agent X = t.getElementAt(Agent.class, 1);
            call(X);
        }
    }
}
```

Autonomic manager (a fragment):

```
t = query( new Template(  
    new ActualTemplateField("informed") ,  
    new FormalTemplateField(Boolean.class)) ,  
    Self.SELF );  
boolean informed = t.getElementAt(Boolean.class, 1);  
if (informed) {  
    put( new Tuple( "controlStep" , new Informed() ) ,  
        Self.SELF );  
    get( new Template( new ActualTemplateField("seek") ) ,  
        Self.SELF );  
} else {  
    put( new Tuple( "controlStep" , new RandomWalk() ) ,  
        Self.SELF );  
    get( new Template( new ActualTemplateField("seek") ) ,  
        Self.SELF );  
}
```



We have briefly presented (first version of) a framework that permits developing and executing SCEL oriented applications in Java

We have briefly presented (first version of) a framework that permits developing and executing SCEL oriented applications in Java

Considered framework should be now *populated* with. . .

We have briefly presented (first version of) a framework that permits developing and executing SCEL oriented applications in Java

Considered framework should be now *populated* with. . .

- specific implementations for *policies*, *knowledge* the proposed solutions;

We have briefly presented (first version of) a framework that permits developing and executing SCEL oriented applications in Java

Considered framework should be now *populated* with. . .

- specific implementations for *policies*, *knowledge* the proposed solutions;
- significant examples that can also help assessing/improving the performance of runtime framework.

We have briefly presented (first version of) a framework that permits developing and executing SCEL oriented applications in Java

Considered framework should be now *populated* with. . .

- specific implementations for *policies*, *knowledge* the proposed solutions;
- significant examples that can also help assessing/improving the performance of runtime framework.

We are now working on a *top-level* programming language that, enriching SCEL with standard programming primitives, permits simplifying development of SCEL programs.

Good work!