SAPERE
Self-Aware Pervasive Service Ecosystems

Mirko Viroli
mirko.viroli@unibo.it

Alma Mater Studiorum—Università di Bologna, Cesena

Academic Year 2010/2011

What is SAPERE

A European Project

- European call for funding projects: FP7-ICT-2009.8.5
- Call name: Self-awareness in Autonomic Systems
- Period: 1/10/2010 - 30/9/2013
- Consortium: (i) Università di Modena e Reggio Emilia
  (ii) Università di Bologna
  (iii) University of Geneve, Switzerland
  (iv) University of St Andrews, Dublin, UK
  (v) Johannes Kepler Universitaet Linz, Austria

Impact on the Cesena site

- A "fulcrum" of challenging EU-wide research activities,
- An opportunity for students (projects, thesis, post-laurea)

Outline/Goals of this seminar

- Deepen the "Future Pervasive Computing" scenario
- Describe the SAPERE Project
- Show what is the structure of an European Project
- Present early research results/ideas
- Propose thesis/projects

Outline

Challenges of Future Pervasive Computing

1 — Increasing introduction of pervasive devices

- The world will be more and more populated of computing devices
    - GPS, PDAs, smart phones, tags, cameras, displays, sensors, actuators
- Internet, Telecom, TV networks will become integrated
    - Think at Skype, Digital TV
- The network will become a true virtual counterpart of the entire world (events, sociality, business, physical structure)

Challenges of Future Pervasive Computing

2 — Prosumption of large masses of data and services

- Devices will produce large masses of data
- Users will themselves inject their data in the system [5]
    - Think at Facebook, Youtube
- We will expect that any environment properly react to our preferences/situation
- Not just data, we are already facing universes of applications
    - E.g. Android Market may become a standard deployment tool
    - There is an increasing "long tail" in the market [1]
- Distinction between data and (atomic) services will blur

Challenges of Future Pervasive Computing

3 — Software will grow increasingly and be an "eternal beta"

- The standard notion of software is shifting
  - No longer a big monolithic engineering artifact
  - ..but rather a mash-up of services and data [2]
- similarly, the development cycle will become much different
  - No longer analysis-design-development-deployment
  - ..but rather it starts with initial services, and through years..
  - ..new simple services are injected
  - ..existing services are updated by small changes
  - ..even the availability of data and users will make the difference
- there is a feeling that a system is never in final state, but always in a "beta" state, that will improve next months
  - Open-source projects (even OS), social networks, ..

Challenges of Future Pervasive Computing

4 — Opennes, self-*, contex-awareness will be mandatory

- Opennes: we won't know which services, data, users, devices will be available soon, the infrastructure should work independently of this
- Self-* features should naturally emerge
  - Self-adaptiveness: tuning behaviour to ongoing changes
  - Self-organisation: find a better/new organisation (spatial/temporal)
  - Self-optimisation: be able to garbage services/data
  - Self-awareness: intrinsic identification of situations
- Context-awareness: data and services will be relative to the position/location in which they reside
  - centralisation of data and software will be abandoned
  - (or, will clouds take over?)

The challenge of Pervasive Computing

Oppurtunities for industry

- We are far from supporting what we have in mind
- Only specific solutions to specific problems so far (even by academia)

A good news for academia

A lot of work is going on in the following areas

- Research contexts: SOA, P2P, Grid, Cloud, Self-org, Coord
- Argue there are *two next big things* in Information Technology
  - Pervasive Computing
  - Bio-ICT convergence (e.g. nature-inspired computing)

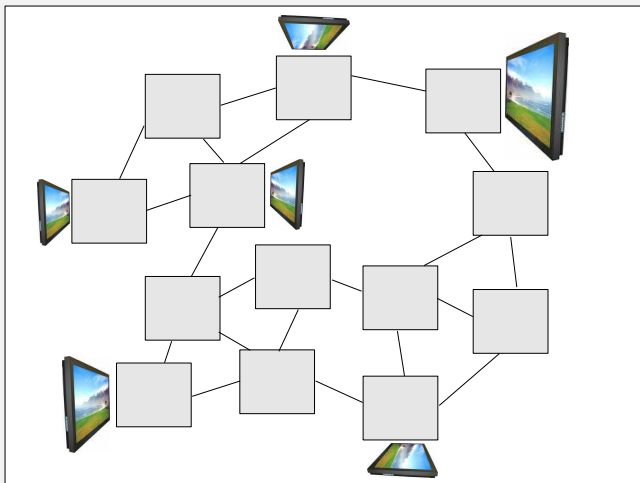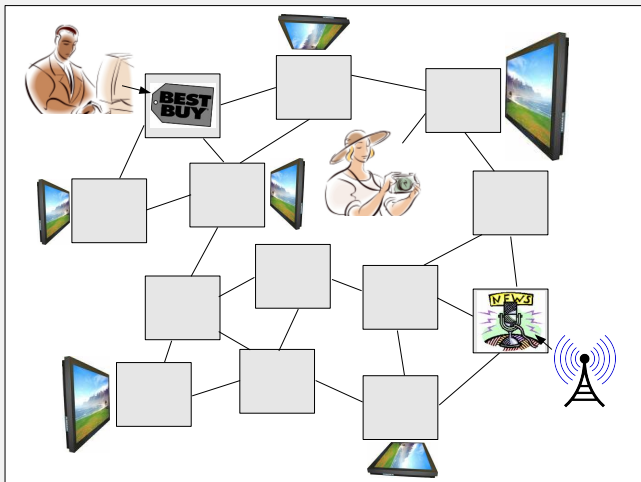# Outline

A scenario of pervasive displays

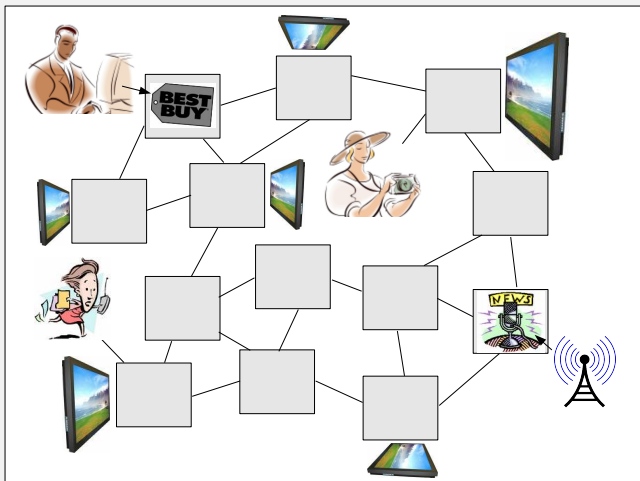A pervasive network with displays spread around

A scenario of pervasive displays

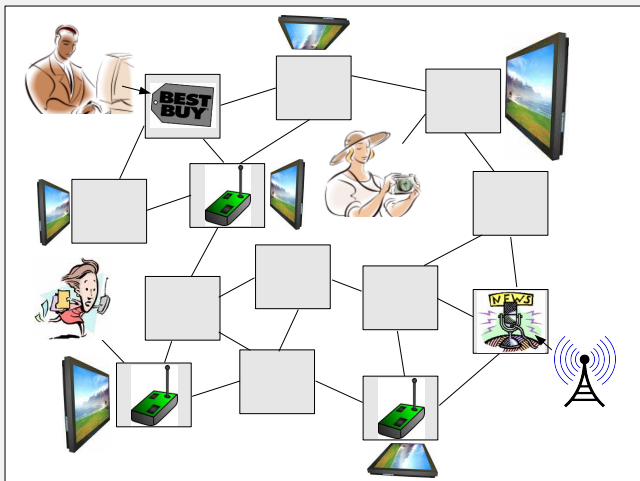## Visualisation services (news, ads, social data) get injected

A scenario of pervasive displays

Displays should adapt visualisation to users nearby

A scenario of pervasive displays

Sensors provide contextual-information to improve adaptation

Some interesting services and features

- Displays show information based on majority of people around
- Alerts are shown as a given person passes nearby
- Displays coordinate to avoid irritating users
- Displays coordinate to provide visualisation streams
- Adjacent displays show a common, bigger content
- Injection of a new display cause redirection there
- A display used as a shared map
- Using eye-glasses with eye-pointers for immersed interaction
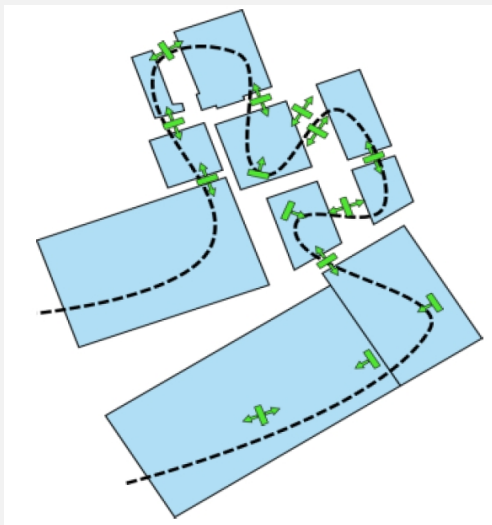
The crowd steering application scenario

A large scaled event area potentially consisting of multiple buildings which is populated by pervasive public displays
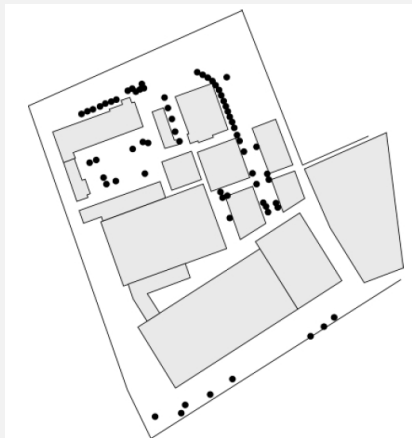
The crowd steering application scenario

Single User Steering via Public Displays

The crowd steering application scenario

Other cases

- Single User Steering via Public Displays
- Crowd Balancing
- Evacuation
- Intention Driven User Steering
- Adaptive Advertisement
- Collective Guidance

The crowd steering application scenario

Lower-level mechanisms

- Ad hoc Connectivity
- Component Deployment
- Presence Detection
- Coverage Adjustment
- User Counting
- Proximity Data Request
- Consensus Finding
- Distance and Orientation Estimation
- Entity Search
- User Path History

Abstract of SAPERE

The objective of SAPERE is the development of a highly-innovative theoretical and practical framework for the decentralized deployment and execution of self-aware and adaptive services for future and emerging pervasive network scenarios. The framework will be grounded on a foundational rethinking of current service models and of associated infrastructures and algorithms. In particular, getting inspiration from natural ecosystems, the project will demonstrate and experiment the possibility of modelling and deploying services as autonomous individuals in an ecosystem of other services, data sources, and pervasive devices, and of enforcing self-awareness and autonomic behaviours as inherent properties of the ecosystem, rather than as peculiar characteristics of its individuals only.

A standard, centralised SOA solution
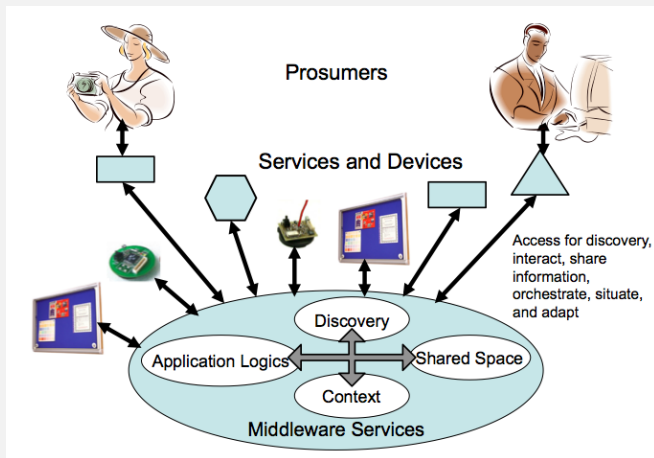
A standard, centralised SOA solution
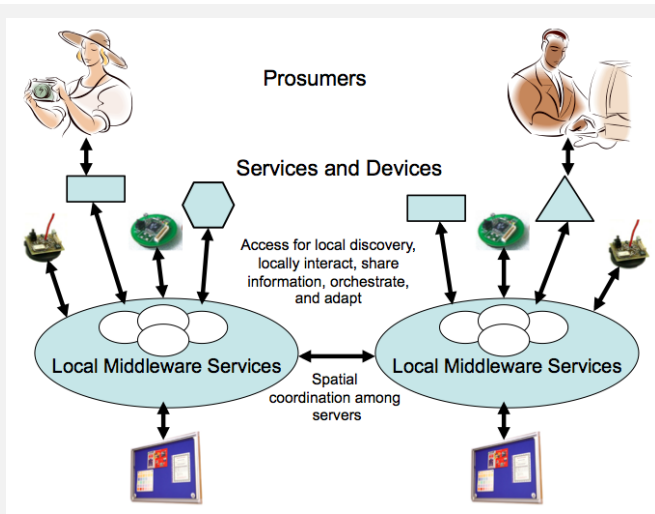
A centralised solution

One service for:

- Discovery: what components are available in the system?
- Context: where are components? (behaviour specialisation)
- Orchestration: coordinating components
- Shared space: depositing/retrieving local information
- Adaptation: reacting to contingencies

All components interact through such middleware services

De-centralising the SOA solution
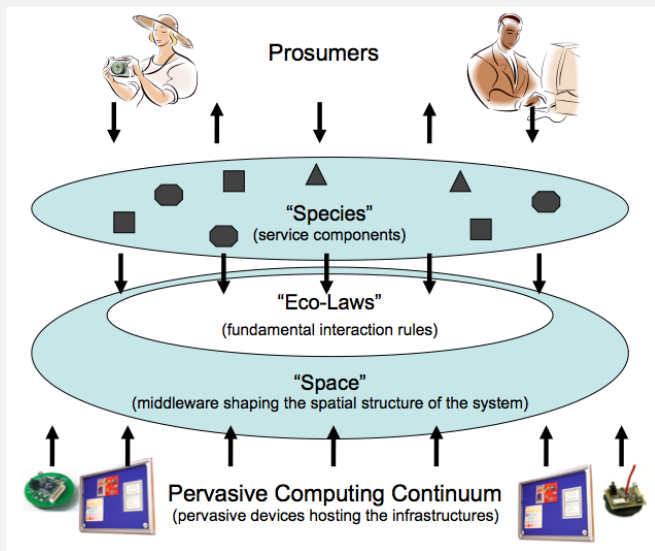
De-centralising the SOA solution

Duplicating middleware services in each location

- Better deals with situated components
- Simplifies contextualisation, discovery, and orchestration
- The role of shared spaces becomes more important
- Adaptation is still complex and crucial

Eco-inspired SOA solution

Eco-inspired SOA solution

Fully decentralised middleware services

- Locations become very small and form a huge dynamic set
- Contextualisation, discovery, and orchestration almost vanish
- Midd. service just as a single space
- In overall we have a network of spaces with service "tags"
- Adaptation is achieved by simple rules combining tags

Drawing a bridge with natural ecosystems

We have a set of spatially situated entities interacting according to well-defined set of natural laws enforced by the spatial environment in which they situate, and adaptively self-organizing their interaction dynamics according to their shape and structure

Architecture
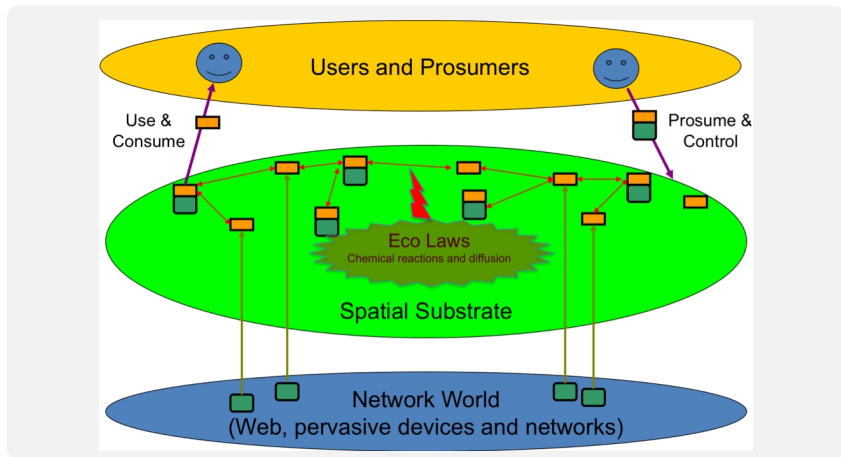
Pervasive Ecosystems

Example Patterns

Data/devices/services are added by injecting their tags, then:

- they could diminish until their population extinguishes
- they compete with other "species", and may survive
- they compose with patches injected to improve them
- they diffuse around
- they move where the context is favourable
- they aggregate with other copies, forming an overlay

Eco-laws and Live Semantic Annotations

Live Semantic Annotations (LSA)

- A unified description for devices, data, services
- Is about interface, status, and behaviour of a component
- It provides semantic information, and it is dynamic

Eco-Laws

- They resemble chemical reactions
- They take some reagent LSA, and provide some product LSA
- They can diffuse an LSA in the neighborhood
- They can aggregate LSAs like in chemical bonding
- They form a small & fixed set of natural eco-laws

Outline

What is SAPERE

A European Project

- European call for funding projects: FP7-ICT-2009.8.5
- Call name: Self-awareness in Autonomic Systems
- Call type: Future and Emerging Technologies (FET)
- Funding Scheme: STREP (Specific Targeted Research Project)
- Period: 1/10/2010 - 30/9/2013
- Overall research grant: $\approx$ 2.5MEuro
- Consortium
  - Università di Modena e Reggio Emilia – Franco Zambonelli
  - Università di Bologna – Mirko Viroli
  - University of Geneve – Giovanna di Marzo
  - University of St Andrews, Dublin – Simon Dobson
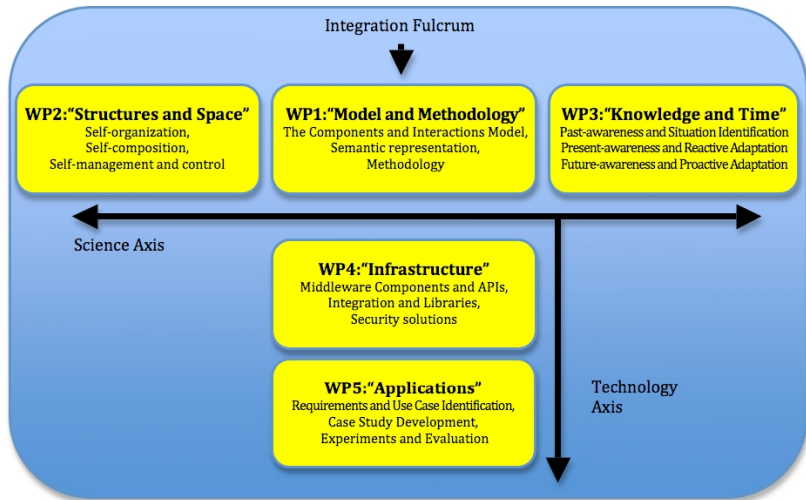 - Johannes Kepler Universitaet Linz – Alois Ferscha

Elements of a Project Document (in general)

- Abstract
- Objectives
- Novelty and Contribution w.r.t. existing works
- Workplan (division in tasks, timing, efforts)
- Technical descriptions
- Deliverables
- Measure success indicators
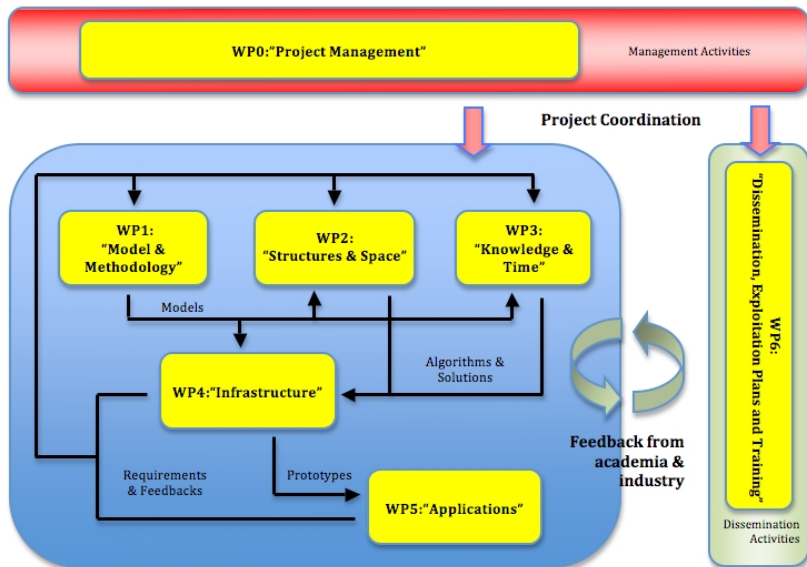- Dissemination activities
- Financial aspects

Strategy

Activities

WP1: Model & Methodology – Leader UNIBO

Task T1.1 – The Components and Interaction Models

- Abstract (representation independent) model of eco-laws
- Services structure and interactions
- Studying analysis tools for behaviour verification

Task T1.2 – Semantic representation

- Shape of LSAs
- Studying analysis tools for logic reasoning

Task T1.3 – Methodology

- Finding a SE methodology
- Conceiving tools for development/analysis

WP1: Deliverables

D1.1 – Early Operational Model (M12) – Editor: UNIBO

- First version of abstract computational model
- First version of live semantic annotation framework
- Early demonstration of the operational model at work

D1.2 – Complete Operational Model (M20) – Editor: UNIBO

- ...

D1.3 – Final Operational and Semantic Model (M24) – Editor: UNIBO

- ...

D1.4 – Early Report on Methodology (M24) – Editor: STA

- ...

D1.5 – Methodology and Analysis Suite (M32) – Editor: UNIBO

- ...

D1.6 – Final Report on Methodology and Suite (M36) – Editor: UNIBO

- Complete and refined engineering methodology documentation
- Assessment of methodology and of associated tools.

Efforts

Partic. no.	Partic. short name	WP0	WP1	WP2	WP3	WP4	WP5	WP6	Total person months
1	UNIMORE	**15**	12	0	22	**34**	16	2	101
2	BIRKBECK	1	15	**36**	0	21	0	1	74
3	STA	1	13	8	**36**	0	15	1	74
4	UNIBO	1	**36**	20	0	20	0	4	81
5	JKU	1	0	16	16	0	**40**	2	75
Total			19	76	80	74	75	71	10	405

UNIBO details

- Leader of WP1: "Model and Methodology"
- Leader of WP6: "Dissemination, Exploitation, .."
- Involved in WP2: "Structures & Space"
- Involved in WP4: "Infrastructure"

Some details on the Ecosystem Model

Syntax / semantics of eco-laws

How is an eco-law specified?

1. Chemical template: $X + Y \xrightarrow{r} Z$
2. Matching functions: how LSAs (l_x and l_y) match X and Y?
3. Bond functions: are l_x and l_y compatible/complementary?
4. Generation function: how is l_z produced from l_x and l_y?
5. Rate: how l_x and l_y affect r, i.e., velocity/probability?

Bio-ICT convergence

This can be viewed as either:

- a semantic-oriented chemical model
- a semantic-oriented population dynamics model

Some details on Infrastructure

Seen as a tuple-space infrastructure

- Agents (devices/services/data) inject their LSA as a tuple
- Agents monitor changes in their LSA, which become actions
- Users perceive the ecosystem by oberving population of LSA
- The infrastructure provide a networked set of LSA-spaces
- LSAs evolve/diffuse by eco-inspired coordination laws

Possible implementation framework

- TuCSoN as a basic middleware
- ReSpecT (or equivalent) as language to express eco-laws
- Need a semantic module, and a chemical module

Outline

A Proposal for the SAPERE "Abstract" Model

View at M6

- is nurtured by several discussions/thoughts
- admittedly suffers from current lack of reference case studies
- indeed takes some design choice
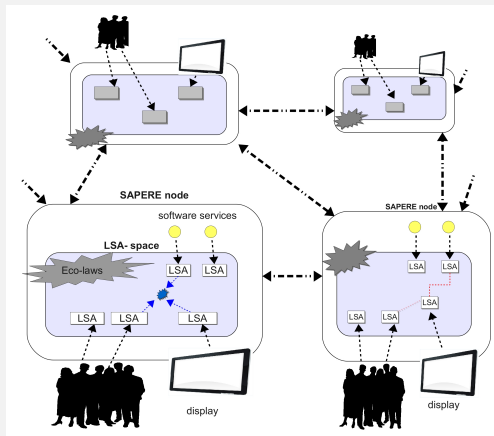- abstracts from incarnation technology

Outline

- the topological space
- LSA
- agent interaction
- eco-laws
- an abstract language for eco-laws

The topological space

The Mobile Network of LSA-Spaces

- Featuring SAPERE nodes, device nodes, LSA spaces, Agents, LSAs
- Can assume an eco-law is always local

Live Semantic Annotations

General features (mostly in line with technologies like RDF)

- a list of [property → value(s)] associations
- mandatory props: unique ID (uid), source ID (sid), a type (or concept)
- a property may expect primitive types, ids, channel ends, functions

Issues with LSAs

Simple vs Complex matching (i.e. semantics)

- Simple tag-based semantics (more lightweight & open)
- OWL-based reasoning (more standard & expressive)

Ontologies can play a key role in SAPERE

- they affect LSAs as well as eco-laws
- they can tackle general-concepts and application-oriented aspects
- not simple to understand how to manage their life-cycle in SAPERE

Injection of LSAs

## How are LSAs created and inserted (implicitly or explicitly)?

- Each agent has one individual LSA (iLSA), which identifies her existence
- An agent might also inject other LSAs (data, events, actions)
- (Some LSAs are also created by eco-laws)
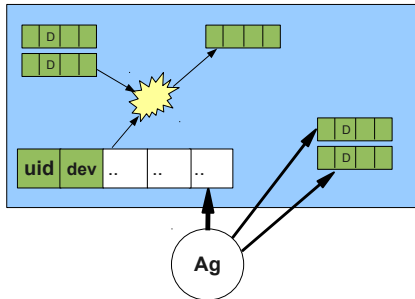
Action/Perception through the LSA

Agent interaction is "constrained" by her LSAs

- Actions by updating owned iLSA, or injecting new iLSAs
- Perception by reading observable LSAs
  (where the iLSA links to)
- (Querying for any LSA is prohibited)
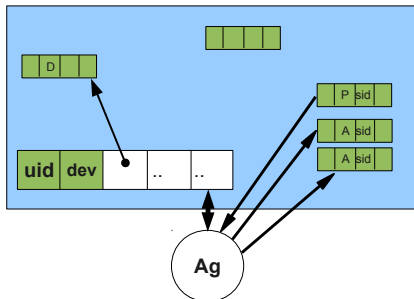
An Abstract Eco-law Framework

Features of an Eco-law

It is basically a tranformation pattern for a (small) set of LSAs

- it applies to a set R of reactant LSAs
- R is selected based on LSAs structure and mutual matching
- the choice of R has a "rate" (useful for ranking)
- possible effects of an eco-law:
  - it change values of properties in some LSAs of R
  - it removes some LSA of R
  - it creates new LSAs
  - it links some LSA of R

A Small Case I

A display, users, visualisation services scenario

- Display agent is a S/W either on the display or in the SAPERE node
- PDA agents represent users nearby the display (i.e. profiles)
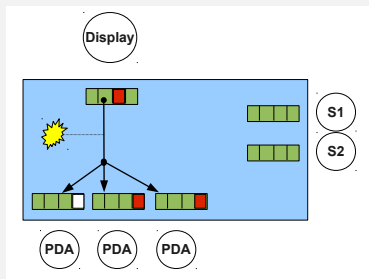- Visualisation services are e.g. ads
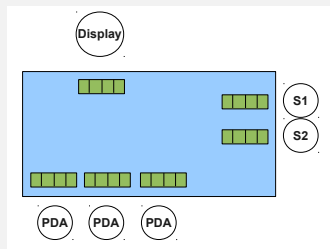
A Small Case II

"Gathering" eco-law

- An eco-law (either in one shot or by repeated application) links display's LSA with LSAs of users around
- A property of the display (red) reflects majority of preferences

A Small Case III

"Activation" eco-law

- An eco-law matches the display (the cumulative preference stored there) with a visualisation service
- A channel is created, by which the ad will be transmitted
- A property of the display (black) reflects activation of display

An Example Language of Eco-laws

Features

- Eco-laws structured as chemical reactions
- On left- and right- hand side we have patterns of LSAs
- Proper annotations are added to enhance expressiveness

Generality

- Chemical concentration might be added rather orthogonally
  (we won't speak about it here)
- We are assuming LSAs are as described before
- We also assume we have a fuzzy semantic matcher between
  primitive values
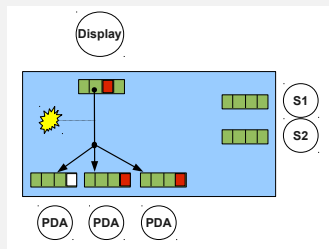 (e.g. strings, types, number intervals)

Example 1

"User Gathering" eco-law

```
{d}:[type=display, contextualizing=true]  +
{u}:[type=user]
-->[CTX-USR]
{d}:[context+={u}]  +  {u}
```
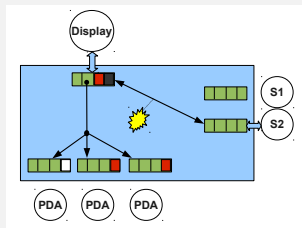
Example 2

"Display activation" eco-law

```
{d}:[type=display, status=ready, showService={s}]  +
{s}:[content={c}]  +
{t}:[type=#time, value={tt}]
-->[ACT]
{d}:[channel=#sink, status=showing, showContent={c}]  +
{s}:[channel+=#source]  +
{t}  +
{l}:[type=log, time={tt}, service={s}, display={d}]
```

A Candidate Set of Eco-laws for the Display Ecosystem

An open list (spatiality again neglected)

- A display gathering user information
- A display connecting with a service provider
- Action/perception laws to user-display interaction
- Composition/orchestration of services and displays
- Decay/dispose of some service

Remarking that..

It is only by case studies that we can address:

- Completeness: can we express all that we want?
- Expressiveness: is writing specification simple/natural enough?
- Coherency: are we considering the right abstractions?
- Generality: can we define application-indepent eco-laws?

Current Developments in UNIBO

An LSA-space prototype

- Status: an early prototype is ready
- Will evolve it to fully support eco-laws
- Will consider relying on W3C technologies (RDF,OWL)

A prototype simulation engine

- This is key when spatiality enters the picture
- Implemented in Java by known SSA
- Will move from chemistry to LSAs

Outline

SAPERE and students

Opportunities

- Small Projects (for this exam, for LMC)
- Master Thesis
- Master Thesis to a partner
- Post-degree collaboration
- Post-degree trip to a partner

Thesis so far

Completed

- Tosi 12/2010 – Architecture of LSA-space
- Virruso 12/2010 – Eco-laws and fuzzy matching
- Santarelli 03/2011 – LSA-space with eco-laws

Ongoing

- Desanti 10/2011 – Complete LSA-space Alpha
- Pronti 10/2011 – Complete Simulator Alpha
- Morgagni 07/2011 – Android Integration

Some available (see Montagna and Nardini's talks)

- Semantic Matching in eco-laws
- Self-organisation patterns and case studies

Outline

Chris Anderson.
The Long Tail: Why the Future of Business is Selling Less of More.
Hyperion, 2006.

Mehdi Jazayeri.
Species evolve, individuals age.
In *IWPSE '05: Proceedings of the Eighth International Workshop on Principles of Software Evolution*, pages 3–12, Washington, DC, USA, 2005. IEEE Computer Society.

Elena Nardini, Andrea Omicini, and Mirko Viroli.
Description spaces with fuzziness.
In Mathew J. Palakal, Chih-Cheng Hung, William Chu, and W. Eric Wong, editors, *26th Annual ACM Symposium on Applied Computing (SAC 2011)*, volume II: Artificial Intelligence & Agents, Information Systems, and Software Development, pages 869–876, Tunghai University, TaiChung, Taiwan, 21–25 March 2011. ACM.

Elena Nardini, Mirko Viroli, and Emanuele Panzavolta.
Coordination in open and dynamic environments with tucson semantic tuple centres.
In Sung Y. Shin, Sascha Ossowski, Michael Schumacher, Mathew Palakal, Chih-Cheng Hung, and Dongwan Shin, editors, *25th Annual ACM Symposium on Applied Computing (SAC 2010)*, volume III, pages 2037–2044, Sierre, Switzerland, 22–26 March 2010. ACM.
Awarded as Best Paper.

Raghu Ramakrishnan and Andrew Tomkins.
Toward a peopleweb.
Computer, 40(8):63–72, 2007.

Mirko Viroli and Franco Zambonelli.
A biochemical approach to adaptive service ecosystems.
Information Sciences, 180(10):1876–1892, 2010.