Towards Eternally Adaptive Pervasive Service Ecosystems

Franco Zambonelli[1] and Mirko Viroli[2]

[1] DISMI – Universita' di Modena e Reggio Emilia
42100 Reggio Emilia, Italy
franco.zambonelli@unimore.it
[2] DEIS – University of Bologna
47023 Cesena (FC), Italy
mirko.viroli@unibo.it

Summary. In this article, we first motivate the need for innovative open pervasive service frameworks that ensure capability of self-adaptability and long-lasting evolvability (i.e., eternity). On this basis, we discuss how such frameworks should get inspiration from natural ecosystems, by enabling modelling and deployment of services as autonomous individuals in an ecosystem of other services, data sources, and pervasive devices. A reference conceptual architecture is presented to clarify the concepts expressed, and the several possible approaches to realise the idea are surveyed and critically analyzed, also with the help of a simple case study.

1 Introduction

Pervasive sensing and actuating devices will densely populate our everyday environments and will be tightly integrated with the converging Telecom and Internet networks, paving the way the deployment of innovative pervasive services, i.e., context-aware services for interacting with the physical and social worlds around us [5, 10].

In this context of tight convergence and integration, a single innovative open software platform will have to be provided to host and orchestrate in an integrated and self-adaptive way the execution of general-purpose pervasive services, as well as the organization of large masses of contextual data that are typically involved in such scenarios. Also, such a platform should take into account the increasingly diverse and demanding needs of users (which will also seamlessly act as consumers and producers of data and services) [15], and must be able to flexibly tolerate evolutions over time without requiring significant re-engineering to incorporate innovations and changing needs.

Recently, a great deal of research activity has been devoted to produce solutions to match the emerging characteristics of future pervasive networks [7] and to solve problems related to, e.g., increasing dependability, reducing management efforts via self-* features, enforcing context-awareness and adaptability, tolerating evolution over time and eventually ensure that the overall service framework (if not all services within) can be highly adaptive and very long-lasting, even in the absence of explicit management actions. Unfortunately, most of the solutions so far are proposed in terms of "add-on", one-of solutions to be integrated in existing frameworks. The result of this process is often an increased complexity of current frameworks and the emergence of contrasting trade-off between different solutions.

In our opinion, there is need for tackling the problem at the foundation, by trying to answer the following question: *Is it possible to conceive a radically new way of modelling integrated network services and their execution environments, such that the apparently diverse issues of enabling pervasiveness, context-awareness, dependability, openness, diversity, flexible and robust evolution, can all be uniformly addressed once and for all?* The way towards the actual realization of eternally adaptive pervasive service environments can thus pass through a foundational re-thinking and deconstruction of current service architectures and models and of their supporting infrastructures.

As we will elaborate through the following of this paper, a promising solution is to get inspiration from natural systems, where situatedness, adaptivity and eternal evolvability are there because of the basic rules of the game, and adopt a natural metaphor to model and develop fully-fledged pervasive service frameworks and all the components within.

2 Scenarios and Motivations

2.1 Towards Pervasive Services (THIS COULD BECOME A SIDEBAR)

The ICT landscape, yet notably changed by the advent of ubiquitous wireless connectivity, will further re-shape due to the increasing deployment of pervasive computing technologies. RFID tags and alike will be attached to every, even small, objects and locations, carrying on digital information about themselves and enabling us to store in them and access additional information about. Wireless sensor networks, there included camera networks, will be distributed around in our cities, houses, and offices, to monitor specific physical phenomena, and possibly to perform actuating actions in the physical world. Devices such as mobile phones, PDAs, GPS and other localization devices, are increasingly able to produce and store notable amounts of data related to our personal, social and professional activities.

In addition, one must consider that the Web has rapidly evolved from a mere container of rather static information to a very dynamic shared information space [15]. Geographical information tools such as GoogleEarth can make available a variety of spatially situated information about facts and events occurring in the real-world [5, 10], whereas social Web platforms such as Flicker and Facebook are de facto becoming sort of real-time shared diaries.

The above scenario will overall lead to the emergence of a comprehensive, integrated, and very dense, decentralized shared infrastructure for the provisioning of general-purpose digital services. At the user level, the infrastructure can be used to access (e.g., via PDAs and on the move) innovative services for better perceiving/interacting with the surrounding physical world and the social activities occurring in it, and for acting on it (e.g., services for browsing the world [5]). It is also expected that users themselves will be able to personalize the infrastructure by deploying customized services over it (in other words, the overall pervasive infrastructure will be as open the same as the Web currently is). In addition, the infrastructure will be used as a way to enrich more traditional classes of digital services with the capability of dynamically and autonomously adapting their behavior to the context (spatial and/or social) in which they are invoked and exploited.

Indeed, we are already facing the commercial release of a variety of early examples of pervasive services, trying to exploit the possibilities opened by these new scenarios. These include, just to make a few examples, GPS navigator systems that are able to provide real-time traffic information and update routes accordingly, or smart phones that are able to inform us about the current positions of our friends. However, the road towards the effective development and execution of general-purpose services, able to exploit in full the potential of the emerging scenarios, is still a long one. In fact, this calls for a deep rethinking of current service models and service frameworks, in order to accommodate a number of important characteristics that are not accounted by any of the mentioned early examples of pervasive services.

2.2 Requirements

The key characteristics that we think should be exhibited by next generation pervasive services include:

- *Situatedness*, i.e., the need to naturally match the inherent situated nature of the services and their environment. On the one hand, services will have to deal with spatially situated and socially situated activities of users, and with their interacting with the surrounding physical and social worlds. On the other hand, the infrastructure will be embedded into the physical space, will have to deal with spatial concepts and spatial data, and its devices will typically have to interact based on spatial proximity (as induced by wireless communications).
- *Adaptivity*, i.e., the need to Inherently exhibit properties of self-organization, self-adaptation and self-management. Such properties are necessary in highly-decentralized and highly-dynamic scenarios (as the envisioned infrastructure is, due to its distributed nature, the unreliability of its components, and its openness to user contributions) to survive contingencies without human intervention and, thus, at limited management costs. In addition, they are necessary properties to enable personalization and contextualization of services behavior.
- *Prosumption and Diversity*, i.e., the capability of tolerating open models of service production and usage without limiting at all the number and classes of services provided. These features are altogether necessary both to accommodate the possibility for users to act as not only as consumers but also as producers of devices, data, and services (that is, to act as "prosumers"), and to capture the so called "long tail" of the market (that is, those non-negligeable percentage of

users that have peculiar needs and expectations, and that would therefore calls for peculiar services) .

- *Eternity*. i.e., the capability of tolerating evolutions of structure, components, and usage patterns over time. This is necessary to accommodate the increasingly diverse and changing needs of users as well as the very dynamic technological evolution, without forcing significant (and economically unbearable) re-engineering efforts to incorporate innovations and changes.

2.3 A Case Study Scenario

To clarify the concepts expressed, a simple case study (that will also act as running example through the paper) may be of help.

It is a matter of fact that we are increasingly surrounded by digital displays, wherever and at any time: from the small ones of our wearable devices and our domestic hardware, to the wide wall-mounted ones increasingly pervading our urban and working environments. Currently, in most of the cases, such displays are conceived to show information in a hardwired or manually-configured manner, being completely insensitive to the context in which they operate and to the specific needs of users. For instance, advertising displays typically cycles a number of pre-defined (and pre-paid) commercials independently of who's around.

For such a costly and intrusive hardware infrastructure to be highly effective and advantageous for both users and information providers, a totally different approach should be undertaken. Information should be displayed based on the current state of the physical environment (e.g., an advertiser would not be happy to pay for a hot chocolate commercial to be displayed in a warm day, when advertising about ice tea would lead to much more impact), possibly by exploiting information coming from surrounding sensors, and by coordinating actions among displays which are in physical proximity (i.e., to avoid users being irritated by seeing the same ad over and over as they pass by). Displays could adapt the information they show to the classes of users that are around (e.g., showing a soda commercial rather than a hot chocolate one in the presence of a teens' majority) and to what they have already seen in the recent past. The overall display infrastructure should be able to automatically re-distribute and re-shape the information displayed upon changing conditions (e.g., when new displays are deployed or some gets out-of-order). It should be possible for users to upload information to enrich the overall informative offer or simply to adapt the offer to their own needs . Finally, the overall infrastructure should tolerate in a seamless way both changes related to the technology of displays themselves, of the sensors around (e.g., the introduction of new classes of sensors that can be fuitfully exploited to improve the information offer), as well as the changes in the technological and business models according to which information and advertisements are displayed.

In a sentence, for the envisioned display infrastructure to become a ground for useful "pervasive information services", it must consider interactions with the various pervasive and personal devices around, must accordingly exhibit situatedness and adaptivity in its behavior, it has to be prosumption- and diversity-oriented and, to most extents, eternal.

2.4 Towards Nature-Inspired Service Ecosystems

To reach the goal, we should no longer conceive services and their interactions as it is usual made in standard SOA architectures [9]. There, services are simply considered as "loci" of functionalities (e.g., with reference to the case study, a display is a place where some information can be visualized), whose activities are orchestrated according to specific pre-defined patterns with the support of middleware services such as discovery, routing, and context services. As a consequence: either they miss in accounting situatedness and spatial concepts or do not elect them as primary abstractions; expressions of adaptivity such as self-organization, self-adaptability and self-management are not intrinsic properties of the system, but are typically enforced via ad-hoc one-of solutions, e.g., via the introduction of specific control tools as it is the concept of autonomic manager in the IMB's approach to autonomic computing [12]; long-term evolvability is simply not ensured, and most likely it can be achieved only at very high re-engineering costs.

Thus, the most promising direction is that of taking inspiration from natural systems, where spatial concepts, self-organization, self-management, and long-lasting evolvability are inherently there because of the basic "rules of the game". We are aware that nature-inspired solution have already been extensively exploited in the area of distributed computing (see e.g. [2, 13] for two recent extensive surveys). However, most of these proposals exploit the natural inspiration only for the effective implementation of specific algorithmic solutions or for the realization of specific distributed services. Similarly, although a number of initiatives recognize that the complexity and dynamics of modern ICT systems is comparable to that of natural systems and, thus, requires very flexible and innovative solutions (see e.g. the many initiatives worldwide named upon keywords such as digital, business, and/or services ecosystems [16, 8]) the idea that natural metaphors – other than a way to characterize the systems – can become the foundation on which to re-think the whole architecture of those systems is far from being metabolized.

Here we go in that direction, and argue that natural systems can act as the key metaphor around which to conceive, model, and develop, fully-fledged pervasive service frameworks and all the components within. You can think at natural systems at different abstraction levels, i.e., physical, chemical, biological, and ecological/social levels. Whatever the case, you can always recognize the following characteristics: above a spatial environmental substrate, individuals (or particles, or components) of different kinds interact, compete, and combine with each other according to complex patterns in respect of same basic laws of nature. That is, what you eventually perceive is always a sort of natural "ecosystem", independently of whether you are observing particles, chemicals, cells, or living organisms, and independently of whether such interactions are driven by physical laws, chemical laws, or laws of survival.

Accordingly, the shared pervasive infrastructure substrate will have to be conceived as the space in which bringing to life a sort of "ecosystem of services" , intended as spatially-situated individuals whose computational activities, while serving to own specific purposes, are subject to some (a limited set of) basic laws of the ecosystem. There, the idea is that the dynamics of the ecosystem (as determined by the enactment of its laws) will provide for spontaneously enforcing features of self-organization, self-adaptability, and eternal evolvability, being these inherent, endogenous properties of the ecosystem, rather than peculiar characteristics of its individuals [11]. Of course, it may be the case that the adoption of different modeling metaphors can impact on the way such features are enforced and on the overall complexity of modeling and controlling the service ecosystem.

3 A Reference Conceptual Architecture

Independently of any specific natural metaphor, a uniform reference conceptual architecture can be adopted for open pervasive service ecosystems (see Figure 1).

The lowest level is the concrete physical and digital ground on which the ecosystem will be deployed, i.e., a very dense infrastructure (ideally, a pervasive continuum) of networked computing devices and information sources. These includes all the hardware devices that are going to increasingly pervade all our everyday environments, such as PDAs, smart phones, tags, sensors, as well as actuators (e.g., the displays of our case study scenario). Such devices are typically interconnected with each other, and most of them are able to generate a large amount of information about the surrounding environment. In addition, such ground can also include the increasing amount of Web tools and data sources that already collect spatially-situated knowledge about nearly every aspect of the world.

At the top level, service developers, producers and consumers of services and data, access the open service framework for using/consuming data or services, as well as for producing and deploying in the framework new services and new data components or for making new devices available. With reference to our case study scenarios, actors include the users passing by, the display owners, and the advertising companies interested in buying commercial slots.

At both the lower and the top levels of the architecture openness and its dynamics stand: on the one hand, new devices can join/leave the system at any time; on the other hand, new users can interact with the framework and can deploy new services and data items on it, and possibly new devices too. As we have already anticipated, in our case study scenario it is possible to consider the integration at any time of new displays and new sensors, the presence of a continuous flow of users and commercial advertisers, some of which possibly interested in having their own devices integrated in the overall infrastructure and/or interested in contributing new ways to have information displayed.

In between these two levels, there are the abstract computational components of the pervasive ecosystem architecture.

The level of “Species” includes a variety of components. There, physical and virtual devices of the pervasive infrastructure, digital and network resources of any kind, persistent and temporary knowledge/data, contextual information, events and information requests, software service components, as well as software components/agents to represent users and prosumers, are all provided with a uniform abstract view, and represent the living entities (or the *individuals*) populating the ecosystem. Although such individuals are expected to be modelled and computationally rendered in a uniform way, they will have specific characteristics very different from each other, i.e., they will be of different "species". With reference to the case study, we will have different software species to represent displays and their displaying service, the various kinds of sensors distributed around the environment and the data they express, software agents to act on behalf of users, display owners, and advertisers. In such perspective, the distinction between active entities such as processes (e.g., advertising agents) and passive entities such as data items (e.g., users' personal preferences) is quite blurred, and pertaining only to the different behaviors that such different components may express, rather than to their modeling and to their overall role in the ecosystem.

In general terms, an ecosystem is expected to be populated with a set of individuals physically deployed in the environment (e.g., physical and network resources, static information, initialization data and services, and so on), and situated in some portion of the ecosystem space. Yet, the population of individuals is far from being static. In the short, term the set of individuals is subject to changes due to the already mentioned factors of dynamics and openness that typically affect the overall infrastructure. In the longer term, technological evolution and social evolution may notably affect the characteristics of the species that populate the ecosystem.

The "World" level provides the spatial fabric supporting individuals, their spatial activities and interactions, as well as their life-cycle. From a conceptual viewpoint, the "World" level gives shape to and defines the structure of the virtual world in which individuals live. Given the inherent spatial nature of pervasive services, it is clear that this level should consider that individuals exist in a specific portion of some metric space, and that their activities and interactions are directly dependent on their positions in space and on the shape of the surrounding space. What the actual structure and shape could be might depend on the specific abstractions adopted for modeling the ecosystem. With reference to the case study scenario, the specific structure of the physical environment as well as the positioning and characteristics of the deployed displays can overall contribute to the shaping of the ecosystem space.

From a more practical viewpoint, the spatial structure of the ecosystem will be implemented by means of some minimal middleware substrate, i.e., a software infrastructure deployed on top of the physical deployment context. Such middleware substrate will provide for supporting the execution and life cycle of individuals, and will enforce (rather than hiding, as most traditional middleware systems do) concepts of distance, locality, local interactions, and mobility, coherently to a specific structure of the space.

The way in which individuals live and interact (which may include how they produce and diffuse information, how they move in the environment, how they self-compose and/or self-aggregate with each others, aggregate, how they can spawn new individuals, and how they decay or die) is determined by the set of fundamental "Laws" regulating the eternal service ecosystems model. In our case study scenario, the eco-laws may provide for example at automatically and dynamically determining to display a specific information to display on a screen as a sort of automatic reaction to specific environmental conditions, or at having two displays aggregate and automatically synchronize with each other in showing specific advertisements.

The laws, or "eco-laws", are expected to act on the basis of spatial locality principles, as in real laws of nature: the enactment of the laws on individuals will typically affect and be affected by the local space around them and by the other individuals on. Although the eco-laws are expected to be always the same, they can have different effects on different species. Indeed, the enactment of the eco-laws may require the presence of some meaningful description (within the uniform modeling on individuals) of the information/service/structure/goals of each species and of their current context and state. These descriptions, together with proper "matching" criteria, define how the eco-laws apply to specific species in specific conditions of the space. With this regard, we emphasize that, in the proposed architecture, the concept of "semantic description" adopted in SOA architecture to facilitate service discovery turns into a concept of "alive semantic description" to rule the enactment of the eco-laws.

The proposed architecture – in its simplicity – represents a total deconstruction of the traditional perspective on service and network architecture. There is no such concept as vertical layering, which can be found both in service and in network architecture, and which indeed limits flexibility and adaptability. Rather, in between the network and the users there is a single universe of components, all of which underlying the same model and obeying the same eco-laws, such laws and the spatial structure in which they apply being the only concepts which are hardwired into the system. In our view, such a radical deconstruction is very important to ensure dependability, diversity, and evolution: nothing (i.e., no infrastructural components and no hardwired middleware or network services) is there to stay, everything can change and evolve, without undermining at the foundations the overall stability of the ecology.

The dynamics of the ecosystem will be overall determined by having individuals in the ecosystem act based on their own internal goals, yet being subject to the eco-laws for their actions and interactions. The fact that the way eco-laws apply may be affected by the presence and state of other individuals, provides for closing the feedback look which is a necessary pre-condition to enable self-* features. Indeed, the typical evolution patterns that can be driven by such laws may include forms of self-organization (e.g., service aggregation or service orchestration, where the eco-laws can play an active role in facilitating individuals to interact with each other and orchestrate their actions), self-adaptation (changing conditions will reflect in changes in the way individuals in a locality are affected by the eco-laws) and of decentralized self-management (the injection of new individuals can be used to modify the way eco-laws affect other individuals and, thus, to somehow control the evolution of the ecosystem dynamics form within the system).

As far as adaptation over time and long-term evolution are concerned, the very existence of the eco-laws can potentially make the overall ecosystem sort of eternal, and capable of tolerating dramatic changes in the structure and behavior of the species living in the ecosystem (i.e., the presence of brand new classes of services).

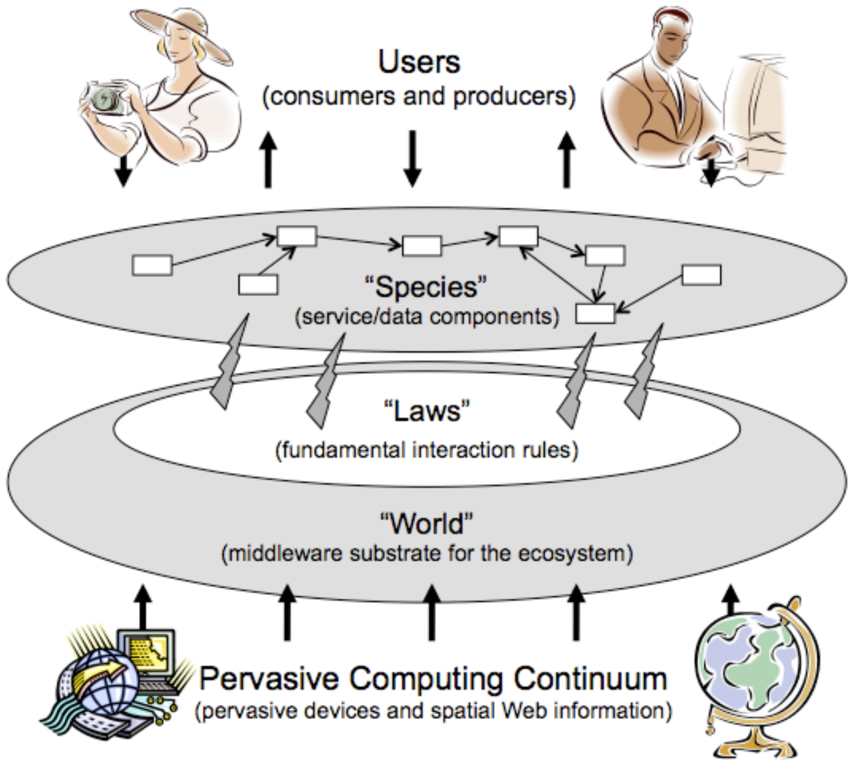

Fig. 1. A Conceptual Architecture for Pervasive Service Ecosystems

Simply said in ecological terms: while the basic laws of life (i.e., the basic infrastructure and its laws) are eternal and do not change (i.e., do not require re-engineering), the forms under which it manifests continuously evolve (i.e., the actual service and data species), naturally inducing new dynamics for the interactions between individuals and for the ecosystem as a whole.

4 Survey and Analysis of Possible Approaches

The key difference in the possible approaches that can be undertaken towards the realisation of eco-inspired service frameworks (as from the described reference architecture) stands in the metaphor adopted to model the ecosystem, its individuals, and its eco-laws. In particular – without excluding the existence of other useful natural metaphors or the possibility of conceiving interesting non-natural metaphors – the main metaphors that can be adopted and have been suggested are: physical metaphors [6, 14], chemical metaphors [3, 18], biological metaphors [2, 4], together with metaphors focussing on higher-level ecological/social metaphors (e.g., trophic networks [1, 17]).

We emphasise in any case that, so far, none of these metaphors has been actually adopted to extensively study and prototype an actual, open and general-purpose service framework: either the metaphor has been applied to specific application scenarios [14, 4] or its potentials for general adoption has been only envisioned [6, 1].

4.1 Metaphors

Let us now come to the distinguishing characteristics of each metaphor, a summary of which is in Figure 2.

Physical metaphors consider that the species of the ecosystem are sort of computational particles, living in a world of other particles and virtual computational fields, the latter acting as the basic interaction means. There, all activities of particles are driven by laws that determine both how particles should spread fields, how such fields propagate and gets re-shaped upon changing conditions, and how particles should be influenced by the local gradients and shape of the computational field (those whose description "matches" some criterion). Particles can change their status based on the perceived fields, and they can move or exchange data by navigating over such fields (i.e., by having particles that move following the gradient descent of a field, or by making them spread sort of data particles to be routed according to the shape of fields). The world in which such particles live and in which fields spread and diffuse can be either a simple (euclidean) metric world mapped in the physical space, or a network space mapped on the technological network, or it could also be a sort of relativistic world, in which shapes and distances in the environment are not "inherent" but are rather shaped by fields themselves (as in gravitational space-time).

With reference to our case study scenario, we can imagine displays (better, the corresponding display services) as masses that tend to emit sorts of gravitational fields (in the form of broadcast events or spanning trees). Such fields may also have different "flavours" (determined by, e.g., the species of users around and the specific environment conditions) and an intensity proportional to either their dimension or the available display slots. Advertiser agents can then simply behave as masses attracted by gravitational fields with specific flavor (on the basis of some matching criteria), to eventually get in touch with suitable displays for their commercials. Upon changing conditions, the structure and characteristics of diffused fields will change, providing for dynamically re-assigning commercials to displays. Similarly, we can easily conceive the definition of information routing services, or of load balancing services, or of whatever user-level services, by applying the simple basic concepts of the physical metaphor (e.g., by having data items and processes move in the network by following the gradients of specific computational fields).

Chemical metaphors consider that the species of the ecosystem are sorts of computational atoms/molecules. Their properties can be described by some sort of semantic descriptions, as the computational counterpart of the description of the bonding properties of physical atoms and molecules, yet made dynamic by the fact that such description reflect the current state and context of individuals. Accordingly, the laws that drive the overall behaviour of the ecosystem are sort of chemical laws, that dictates how chemical reactions and bonding between components take place (i.e., relying on some forms of pattern matching between the semantic description of components), and that can lead to the production of aggregates (e.g., of aggregated distributed components), of new components (e.g., of composite components or of components spawned by existing ones) and to the growth/decay of a set of components. The world in which components live is typically formed by a set of localities, intended as the "solutions" in which chemical reactions can occur.

With reference to the case study scenario, we can think of display services, of information services (concerning advertisement, news, and so on), as well as of users' and environmental personal data (e.g., users' preferences and position and status of the physical environments) as molecules. Displays can represents different localities (e.g., chemical solutions), in which components react. Chemical rules can dictate that when the preferences of a user U enters the locality of a display D and it matches an advertisement service S, then a new composite component is created which is in charge of actually displaying service S in D. Alternatively, one can think at activating such reaction only in the presence of some information I matching users' preferences, so as to avoid a user being overwhelmed by advertisements without receiving some interesting information too. Concurrently, one can think at having, in each locality, catalytic components in charge of re-enfocing the concentration of specific information or of specific advertisements in a locality, to reflect the current situation of users. Also, one can think at opening localities to enforce chemical bonds across displays such that, e.g., a high-activity of advertisement reactions on a display D can eventually propagate to neighbor displays.

Biological metaphors typically focus on biological systems at the small scale, i.e., at the scale of individual organisms (e.g., cells and their interactions) or of colonies of simple organisms (e.g. ant colonies). The species are therefore either simple cells or very simple animals, both of which can be seen as acting on the basis of very simple goal-oriented rules (e.g., move, find food, find mate). Similarly to physical systems, interactions takes place by means of signals of various flavors (i.e., chemical pheromones) spread by components in the environment and to which components can react. However, unlike in physical systems, components are not necessarily passively subjects to pheromones around: the reaction of components to the sensed pheromones as well as their own spreading of pheromones can depend on their current "mood" (which reflect its current state towards the achievement of a goal). Accordingly, within the biological metaphor, the laws of the ecosystem are only aimed at determining how such pheromones, depending on their specific flavors, should diffuse and/or evaporate, depending on their specific flavors. With this regard, we emphasize that pheromones, unlike fields, are typically more persistent than fields (i.e., they slowly evaporate and diffuse), which makes it easier for the infrastructure to support them, and which makes it also possible to exploit them as a sort of environmental memory, other than to reflect the current state of things in the environment. The world in which components live is typically a virtual computational landscape, either mapped on the network topology or on the physical space, that can influence the way pheromones diffuse and the way components can move over it.

With reference to the case study scenario, agents will be represented by simple agents roaming around spreading chemical signals with a flavor reflecting their personal interests. Displays can locally perceive such pheromones, and react by emitting some different pheromones to express their availability of showing some commercials of some specific nature, a choice which can also rely on the past history. Advertising agents, by their side, can sense the concentration of such pheromones and be attracted towards the displays where the concentration of users with specific interests is maximized. Both displays and advertising agents by their side, and depending on what advertisements they have displayed so far, can also emit additional flavors of pheromones, to store some memory of past events in the environments . The temporary persistence of the pheromones can also be effectively exploited by additional classes of components that, by moving around from display to display, can create proper pheromones trails as a basis for more global strategy (e.g., to create sorts of routing path for enabling advertiser agents to globally find the best displays to exploit).

Social metaphors focus on biological systems at the level of animal species and of their interactions, e.g., at the level of trophic interactions. There, components of the ecosystem are sort of goal-oriented animals (i.e., agents) belonging to a specific species (i.e., agent classes), that are in search of "food" resources to survive and prosper (e.g., specific resources or other components). Pure data items and resources can be abstracted as sorts of passive life-forms (i.e., vegetables). The laws of the ecosystem determine how the resulting "web of food" should be realized, that is, they determines how and in which conditions animals are allowed to search food, eat, and possibly produce and reproduce, thus influencing and ruling the overall dynamics of the ecosystem and the interaction among individuals of different species. Similarly to chemical systems, the shape of the world is typically organized around a set of localities, i.e., of ecological niches (think at a set of local pervasive computing environments), yet enabling interactions and diffusion of species across niches.

With reference to the case study scenario, and assuming that each display is associated to an ecological niche, we can imagine users as species of herbivors agents that roam around from niche to niche, possibly to eat those vegetable-like life-forms which is represent by information. Advertisers, on the other hands, can be sorts of carnivors (i.e., eating other active life-forms) in need both to find proper users' (e.g., users with matching interests) to survive. For both cases, the primary effect of an eating act is the feeding of displays with something to be displayed (information or commercials). Possible secondary effect of eating acts can be the reproduction and diffusion in the environment of the best-fit species (e.g., of most successful commercial). Those species that do not succeed in eating at a niche can either die or move to other niches to find proper food. Concurrently with the activities of the above species, one can think at the presence of background agents acting as sorts of bacteria that digest the activity logs at the various niches to enforce specific forms of distributed control over the whole system (e.g., by affecting the way information is propagated across niches).

	Species	*Laws*	*World*
Physical	Particles (computational components) and messages (computational fields)	Movements and activities driven by fields (gradient ascent/descent by components)	Space (network topology or physical space).
Chemical	Atoms and Molecules (semantic descriptions represents chemical properties)	Chemical Reactions (matching of semantic descriptions and bonding of components)	Localities (pervasive computing environments)
Biological	Cells or simple goal-oriented organisms (ants) and pheromones	Diffusion and evaporation of chemical pheromones (determining differentiation of behaviour and activities)	Space (network topology or physical space)
Social	Goal-oriented animals (Agents) of various species (Classes) and included passive life-forms (Resources and Data)	Trophic relations (eating), digest, produce, and reproduce	Niches (Pervasive computing environments)

Fig. 2. Metaphors for Pervasive Service Ecosystems

4.2 Space, Time, and Control

The analysis of which metaphor best suit the characteristics of modern pervasive computing scenarios cannot abstract from the identification of the key features that we want to achieve from it and that – consequently – we want our ecosystem to express.

We have stated earlier in this paper that the ecosystem should be situated, able to exhibit adaptatity, able to tolerate openness (prosumption), diversity, as well as eternal evolvability. The introduced conceptual reference architecture, by construction, considers an open world of situated components. Thus, within such frame, the suitability of the different metaphors should be analyzed in terms of:

- *(i)* the capability of the ecosystem of autonomously express various forms of self-organization among the distributed (i.e., spatial) activities of the components, so as to autonomously adapt the overall structure and behaviour of the system to specific contingencies;
- *(ii)* the capability of the ecosystem of tolerating changes over time, which includes the capabilities of adaptively accommodating new species or of surviving the extinction of species, as well as the capability of accommodating very diverse and composite behaviour with the same limited set of eco-laws.

In addition, since we should never forget that the service ecosystem is here to serve us, we cannot forget an additional important feature, that is,

- *(iii)* the need to allow humans (e.g., system administrator and users) to exert control over the behaviour of the ecosystem (or of some of its parts), i.e., of directing its activities and behaviour over space and time, both at design time and at run time.

All of these features, of course, should be enforced without paying the price of dramatically increasing the complexity of the ecosystem (i.e., the number and complexity of eco-laws, and the structure of its components and of the world in which they live).

Clearly, the analysis of the extent to which the presented metaphors are able to accommodate (and how easily and naturally) the above features is very complex, and would require much more room than the few pages of this paper and much more experience with real-world systems than it is actually available. Nevertheless, we can try at least to draw some considerations about, as summarised in Figure 3.

Physical metaphors have been extensively studied for their spatial self-organization features, and in particular for their capability of facilitating the achievement of coherent behaviours even in large scale system (e.g., for load balancing, data distribution, clustering and aggregation, differentiation of behaviors). Also, the conceptual tools available for controlling the spatial behaviour and the dynamics of such systems are well-developed, most of them related to properly acting on the way fields propagate and dynamically change. In fact, by acting on fields, it is possible to exert control over the overall behavior of the ecosystem, or of some of its part.

Unfortunately, the physical metaphor falls short in properly tolerating diversity, evolution and time adaptation. In fact, if one wants to preserve the simplicity of the eco-laws and of the supporting infrastructure (which is in charge of enforcing the eco-laws also by supporting field propagation and updates), one has to limit the amount of different fields in the system. However, to support very diverse components and behaviors (at a time and over time), the eco-laws must become complex enough to tolerate a wide range of different fields and propagation rules, with an overall increase in the conceptual complexity of the model and with an increase burden on the infrastructure, which has to support a large number of different laws for field diffusion and propagation.

Chemical metaphors can effectively lead to local self-organizing structures (e.g., local composite services and local aggregates), although it may have problems in dealing with more distributed self-organization (as explained below). Also, the same as real chemistry, chemical computational metaphors, can accommodate an incredible amount of different components and composites, yet with the same set of simple basic laws. In practice, this means that it can tolerate an increasingly diverse and evolving set of semantic description for components without affecting the basic eco-laws and without increasing the burden to the infrastructure. As far as control is concerned, one can think at using sort of catalyst or reagent components engineer and control (in a fully decentralized way) the dynamics and the behavior of a chemical ecosystem.

The key drawback of the chemical approach is that it typically relies on activities taking place within a locality, making it hard to naturally and easily enforce distributed self-organized behaviors. To this end, one should rely on complementing the chemical metaphor with physically- or biologically-inspired mechanisms to rule the diffusion of chemical substance across localities, and the creation of bonds between components in different localities. Clearly, this somewhat complicated the overall model and the supporting infrastructure.

Biological metaphors appear very flexible in enabling the spatial formation of both localized and (by exploiting diffusion of pheromones and movements of individuals) distributed morphological and activity patterns. This has been shown to have notable applications in a variety of applications to distributed systems [2]. As in physical metaphors, the number of patterns that can be enforced by the spread of chemical pheromones and by the reactions of simple individuals seem quite limited, which does not match with the need to accommodate diversity and time evolution. Also in this case, the problems can be somehow circumvented by increasing the variety of pheromone flavors and characteristics. This solution, for biological metaphors, can be more easily supported by the infrastructure, since pheromones (unlike fields) rely on local diffusion and slow evaporation dynamics.

A problem of the biological metaphor is the inherent difficulty in understanding how to properly engineer and control the overall behaviour of such systems. By thinking that, so far, the mechanisms of morphogenesis and biological self-organization in actual biological systems are not fully understood by scientists, one cannot hope to be able to properly enforce any desired form of control in complex biologically-inspired systems.

Social metaphors, such as the one based on trophic interaction that we have described, the same as chemical ones, promise to be very suitable for local forms of self-organization, e.g., think at self-organized equilibria of web food patterns in ecological niches. However, unlike chemical metaphors and more similarly to biological metaphors, the social metaphors can be effectively exploited also to enforce distributed forms of self-organization, by exploiting the capability of individuals of moving form locality to locality to achieve their goals (e.g., to find food and/or reproduce). In addition, social metaphors are particularly suited for modeling and tolerating evolution over time (think at how biodiversity has increased over the course of evolution, without ever mining the health existence of life in each and every place on earth).

However, since understanding how to properly control the overall behaviors, dynamics, and local/global equilibria of real ecological system it is very a difficult task. Accordingly, this would probably be very difficult also in their computational counterparts.

In summary, it is very difficult to assess once and for all which of the metaphors is the best for next generation of adaptive service ecosystems. Some exhibit suitable features for certain aspects, but fall short for others. The general lessons that we can learn from the performed analysis is that:

- For domain-specific scenario, a metaphor should be adopted depending on the specific characteristics of the domain and on the relative importance, in that domain, of the general requirements (i.e., distributed self-organization and adaptation, diversity and evolution, control). For instance, in the case study, one can evaluate that the capability of supporting adaptation and evolution locally at each display is more important than the capability of supporting distributed coordination of activities among display, thus making a chemical approach preferable over, e.g., a physical or biological one.

	Space (self-organization)	*Time (evolution and adaptation)*	*Control (decentralized management)*
Physical	+ Global self-organizing spatial structures	-- Few new components can be accommodated while keeping the laws simple	+ We know well how to build and control specific structures in physics
Chemical	-- Mostly local self-organizing structures	++ Several new components can be accommodated with the same basic laws	+ Reactants and catalysts can exert control over the dynamics and structure of reactions
Biological	+ Morphogenesis of local shapes, global patterns via movements and pheromones diffusion	+ Reasonable number of new individuals and pheromone flavours can be accommodated without increasing complexity too much	- Mechanisms and control of morphogenesis and biological self-organization not fully understood
Social	+ Local structures mostly, although sometimes leading to more global patterns and structures	++ Several new species can be accommodated with the same basic laws	-- Difficult to understand how to enforce control over ecosystems of many species

Fig. 3. Advantages and Limitations of the Different Metaphors

- For general scenarios, none of the presented metaphors is directly applicable as it is, each of them falling short in meeting some of the requirements. Some new synthesis, based on incorporating some features of a metaphor into a different one to produce a sort of hybrid metaphor, should be possibly studied to overcome this limitation.

By our side, we have extensively worked in the past with physically-inspired metaphors [14], and are currently shifting our attention to chemical [18] and social systems based on trophic networks [17], due to their specific suitability in tolerating diversity and evolution and the specific interest we have in this aspect. Yet, we are studying how to enriched the basic models with hybrid characteristics borrowed form biological and physical metaphors, to better support distributed self-organization and decentralized control.

5 Concluding Remarks

Nature-inspired metaphors appear very promising to act as a ground for re-thinking the foundation of modern and emerging pervasive service systems. Yet, the road towards the actual deployment of usable and effective ”eco-inspired” open service frameworks still requires answering to several challenging questions. Apart from the specific issues raised in this paper about the need to identify proper hybrid metaphors capable of inherently expressing all the desired features, additional remain to be analyzed relate to:

- *Modeling and Engineering*: given a specific nature-inspired metaphor, what actual shape will take the software individuals of the system, their semantic descriptions, the middleware infrastructure and its API? Also, how should the tools and the methodologies to develop services in such systems look ike?

- *Security*: how it is possible to accommodate and control proper security strategy within ecosystems of the kind we have envisioned? It is possible to make the capability of flexibly handling security and privacy an inherent feature of the ecosystem? Or it rather has to be something exogenous to it?
- *Legacy*: it is rather naif to hope that, all of a sudden, existing platforms and services will be abandoned to embrace an ecosystem-inspired approach. Thus, the issue of gradually moving towards such new approaches by incorporating the legacy of existing systems has to be addressed, the key question being how to achieve this without undermining the key properties of nature-inspired approaches and, rather, having them enrich legacy components.

The above, and possibly further issues we may have missed identifying, represent fascinating areas of research for anyone interested in future and emerging pervasive computing scenarios.

References

1. Gul Agha. Computing in pervasive cyberspace. *Commun. ACM*, 51(1):68–70, 2008.
2. Ozalp Babaoglu, Geoffrey Canright, Andreas Deutsch, Gianni A. Di Caro, Frederick Ducatelle, Luca M. Gambardella, Niloy Ganguly, Márk Jelasity, Roberto Montemanni, Alberto Montresor, and Tore Urnes. Design patterns from biology for distributed computing. *ACM Trans. Auton. Adapt. Syst.*, 1(1):26–66, 2006.
3. Alistair P. Barros and Marlon Dumas. The rise of web service ecosystems. *IT Professional*, 8(5):31–37, 2006.
4. Jacob Beal and Jonathan Bachrach. Infrastructure for engineered emergence on sensor/actuator networks. *IEEE Intelligent Systems*, 21(2):10–19, 2006.
5. Gabriella Castelli, Alberto Rosi, Marco Mamei, and Franco Zambonelli. A simple model and infrastructure for context-aware browsing of the world. *Pervasive Computing and Communications, 2007.*, pages 229–238, 19-23 March 2007.
6. Jon Crowcroft. Toward a network architecture that does everything. *Commun. ACM*, 51(1):74–77, 2008.
7. Simon Dobson, Spyros Denazis, Antonio Fernández, Dominique Gaïti, Erol Gelenbe, Fabio Massacci, Paddy Nixon, Fabrice Saffre, Nikita Schmidt, and Franco Zambonelli. A survey of autonomic communications. *ACM Trans. Auton. Adapt. Syst.*, 1(2):223–259, 2006.
8. Sebastian Herold, Holger Klus, Dirk Niebuhr, and Andreas Rausch. Engineering of it ecosystems: design of ultra-large-scale software-intensive systems. In *ULSSIS '08: Proceedings of the 2nd international workshop on Ultra-large-scale software-intensive systems*, pages 49–52, New York, NY, USA, 2008. ACM.
9. Michael N. Huhns and Munindar P. Singh. Service-oriented computing: Key concepts and principles. *IEEE Internet Computing*, 9(1):75–81, 2005.
10. Ramesh Jain. Eventweb: Developing a human-centered computing system. *Computer*, 41(2):42–50, 2008.
11. 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.

12. Jeffrey O. Kephart and David M. Chess. The vision of autonomic computing. *Computer*, 36(1):41–50, 2003.
13. Marco Mamei, Ronaldo Menezes, Robert Tolksdorf, and Franco Zambonelli. Case studies for self-organization in computer science. *Journal of Systems Architecture*, 52(8-9):443–460, 2006.
14. Marco Mamei and Franco Zambonelli. *Field-based Coordination for Pervasive Multiagent Systems*. Springer Verlag, 2006.
15. Raghu Ramakrishnan and Andrew Tomkins. Toward a peopleweb. *Computer*, 40(8):63–72, 2007.
16. Mihaela Ulieru and Steve Grobbelaar. Engineering industrial ecosystems in a networked world. In *5th IEEE International Conference on Industrial Informatics*, pages 1–7. IEEE Press, 23-27 June 2007.
17. Cynthia Villalba, Alberto Rosi, Mirko Viroli, and Franco Zambonelli. Nature-inspired spatial metaphors for pervasive service ecosystems. In *1st International SASO Workshop on Spatial Computing*, Venezia, Italy, 20–24October 2008. IEEE CS Press.
18. Mirko Viroli, Matteo Casadei, Sara Montagna, and Franco Zambonelli. A biochemical metaphor for developing eternally adaptive service ecosystems. In *24th Annual ACM Symposium on Applied Computing (SAC 2009)*, Honolulu, Hawai'i, USA, 2009. To appear.