The current issue and full text archive of this journal is available at
www.emeraldinsight.com/1742-7371.htm

Received 5 July 2011
Revised 15 July 2011
Accepted 15 July 2011

A survey on nature-inspired metaphors for pervasive service ecosystems

Franco Zambonelli
Dipartimento di Scienze e Metodi dell'Ingegneria,
Università di Modena e Reggio Emilia, Reggio Emilia, Italy, and
Mirko Viroli
Dipartimento di Elettronica Informatica e Sistemistica,
Università di Bologna, Cesena, Italy

Abstract

Purpose – Emerging pervasive computing scenarios require open service frameworks promoting situated and self-adaptive behaviors, and supporting diversity in services and long-term evolvability. This suggests adopting a nature-inspired approach, where pervasive services are modeled and deployed as autonomous individuals in an ecosystem of other services, data sources, and pervasive devices. However, there are many possibly nature-inspired metaphors that can be adopted, and choosing one may require a careful analysis of the pros and cons of the different metaphors. The purpose of this paper is to analyze the key requirements and desiderata for next generation pervasive computing services and associated infrastructures.

Design/methodology/approach – In this paper, the authors introduce and critically analyze a number of natural metaphors that can be adopted to realize these concepts and survey relevant proposals in the area.

Findings – The key result of this survey is that a uniform reference architecture can be a useful guide when framing the challenges involved in the design and implementation of future self-adaptive pervasive service ecosystems.

Originality/value – The survey in this paper, along with the proposed reference architecture, can be effective starting points towards the definition and implementation of general-purpose nature-inspired pervasive service ecosystems.

Keywords Metaphors, Pervasive computing, Context-aware computing, Self-organization, Nature-inspired approaches, Middleware, Ecosystems

Paper type General review

1. Introduction

The information and communication technologies landscape, yet notably changed by the advent of ubiquitous wireless connectivity, is further re-shaping due to the increasing deployment of pervasive computing technologies. Via radio frequency identification tags and alike, objects can carry on digital information of any sort. Wireless sensor networks and camera networks are being spread in our cities and buildings to monitor physical phenomena. Smart phones can increasingly sense and store data related to our personal and social activities, other than feeding and being fed by the Web with spatial and social real-time information (Campbell *et al.*, 2008).

International Journal of Pervasive Computing and Communications
Vol. 7 No. 3, 2011
pp. 186-204

1742-7371
DOI 10.1108/17427371111172997

Work supported by the SAPERE project (EU FP7-FET, Contract No. 256873).

This evolution is contributing to the emergence of integrated and dense infrastructures for the pervasive provisioning of general-purpose digital services. If all their components will be able to opportunistically connect with each other, such infrastructures can be systematically used to enrich existing services with the capability of autonomously adapting their behavior to the physical and social context of usage, and can also support innovative services for enhanced interactions with the surrounding physical and social world (Coleman, 2009). Users will play an active role, by contributing data and services, by making available their own sensing and actuating devices (Lane *et al.*, 2010; Reddy *et al.*, 2010), and possibly their own human capabilities (Yuen *et al.*, 2009). This will make pervasive computing infrastructures as participatory and capable of value co-creation as the Web (Spohrer *et al.*, 2007), eventually acting as globally shared substrates to externalize and enhance our physical and social intelligence, and make it become collective and more valuable.

We are already facing the release of early pervasive services trying to exploit the possibilities opened by these new scenarios: smart environmental displays reacting to users' presence; car navigation systems providing real-time traffic information; smart phones applications for interacting with close friends or to enrich what one see around with dynamically retrieved digital information. Pushing these scenarios forward, one can envision the future release of advanced services and applications like:

- *Pervasive and interactive adaptive display services*. Many digital displays already surround us, from wall-mounted public displays to personal displays of wearable devices and domestic hardware. These, along with sensors enabling to access information about users around and the state of the physical environment, will form the basis of a general user-centric infrastructure for the release of adaptive information services (Sippl *et al.*, 2010; Ferscha and Vogl, 2010; Rosi *et al.*, 2010). That is, users will finally take full advantage of such displays, being no longer passive subjects of what the displays show, but rather having the displayed information be specifically tuned to their own situations and needs.
- *Real-time traffic control*. By collectively exploiting sensors densely deployed over cities and streets, wireless inter-vehicle communications, and actuable traffic lights and traffic signs, many useful services will be possibly released. Other than those for sensing the current traffic situation to adapt recommended routes (Biem *et al.*, 2010), these could include services for detecting parking slots in real-time, supporting effective car and parking sharing, and actuating traffic lights and signs to steer traffic dynamics.
- *Social and augmented reality services*. The increasing communication and sensorial capability of smart phones, along with the mentioned diffusion of environmental sensors and pervasive displays, will make future wearable devices reach high levels of awareness about the surrounding social and physical situations. And, consequently, to transfer such awareness to users via some personal interfaces, as a sort of sixth sense to better interact with the surrounding world (Langlotz *et al.*, 2011), and, in it, to act, exchange information, socialize, or be involved in participatory coordinated activities (e.g. for contributing with sensing, computing, and physical capabilities (Yuen *et al.*, 2009; Zambonelli, 2011).

Yet, the road towards the effective future spread of these emerging application scenarios requires accommodating a number of emerging requirements and desiderata: the need for a situated and spatial computing model, the need to promote self-adaptivity in service composition and execution, the opportunity to support open production model, other than the need of tolerating long-term evolutions of services and technologies.

The fact is that such novel requirements cannot be simply accommodated by modifying existing paradigm, e.g. service-oriented architectures (SOA) (Huhns and Singh, 2005), but rather require a radical re-thinking of current service framework and architecture. In particular, nature-inspired approaches, being able to inherently accommodate spatiality, self-adaptivity, evolvability, propose themselves as a suitable solution.

Against this background, in this paper, we:

- Analyze the requirements and desiderata for next generation pervasive computing systems (Section 2).
- Discuss the limitations of current service-oriented approaches and argue for the need of innovative, nature-inspired, approaches (Section 3).
- Survey and critically analyze a number of different nature-inspired approaches, inspired by different metaphors, by which innovative pervasive service ecosystems can be modeled and architected (Section 4).
- Introduce a reference architecture to show how nature-inspired pervasive ecosystems can be actually implemented, and also to survey and discuss the critical implementation challenges and further general challenges (Section 5).
- Eventually, in Section 6, we conclude.

2. Requirements and desiderata

Let us now analyze, also with reference to the above-mentioned application scenarios, the key requirements and desiderata for next generation pervasive computing services and associated infrastructures.

2.1 Situatedness and spatiality

Pervasive services inherently deal with spatially – and socially situated activities, and should be able to interact with the surrounding physical and social world to adapt their behavior accordingly. The infrastructure itself, deeply embedded in the physical space, should effectively deal with spatial concepts and data (Werfel *et al.*, 2008; Mamei and Zambonelli, 2006).

In the pervasive display scenario, by exploiting information from surrounding sensors and from users' profile, an advertising service could recognize that it is a warm day and there are teenagers around, thus opting for displaying ice tea commercials rather than liquor ones. Tourists could interact with such displays to acquire information about the current state of nearby touristic attractions or restaurants. Also, coordinated visualization actions among adjacent displays could take place, to avoid irritating users with the same ads as they pass by, or to use adjacent displays as a single wide one to show complex multifaceted information.

Similarly, in traffic control systems, intelligent traffic lights can minimize waiting time by intercepting the situation about nearby cars and their spatial displacement, also to possibly give priority to certain classes of vehicles, e.g. car pools. In social and augmented reality services, it is by gathering information about the presence and preferences of people in a spatial location that, e.g. a social group can be dynamically created to share information or to be involved in some participatory activities.

2.2 Self-adaptivity

Given the inherent dynamics of their operational environment, pervasive services and infrastructures should exhibit properties of autonomous (i.e. self-adaptation and management), to dynamically and opportunistically interact with each other and to survive contingencies without human intervention, and at limited costs (Brazier *et al.*, 2009; Kephart and Chess, 2003; Mamei *et al.*, 2006).

In the display scenario, it would be desirable that when new displays are deployed or when new information is injected a spontaneous re-distribution of the overall displayed information could take place without explicit manual reconfiguration actions. For instance: the deployment of a big display in a room may suggest re-directing there all the ads previously forwarded to the personal displays of users; the injection of a new information service could induce aggregating it with the existing ones to provide a more complete yet uniform service.

For traffic control systems, adapting to contingencies might not only be desirable but compulsory. For instance, to adapt their schedule based on current traffic conditions, smart traffic lights must be able to tackle the unreachability of some sensors by opportunistically connecting with the vehicles passing by to collect information from them. Incorporating new components and providing effective services even upon unavailability of some components is desirable as well in social and augmented reality services. For instance, if a compass is no longer able to recognize the orientation of a tourist (required both for participatory sensing and for augmented reality services to produce digitally-annotated images), it can try to contact an image processing service to infer orientation based on the touristic attraction currently captured by its smart phone camera.

2.3 Prosumption and diversity

Infrastructures for pervasive services should tolerate open models of service production and usage without limiting the number and classes of services provided, and rather taking advantage of the injection of new services to improve and integrate the existing ones, as it is already happen in the Web services world (Barros and Dumas, 2006). That is, they should let users act as “prosumers” – both consumers and producers – of devices, data, and services. Not only this will facilitate meeting the specific needs of users (and capture the long tail of the market), but will also induce a process of value co-creation in the system and its services (Vargo *et al.*, 2008).

In the display scenario, the infrastructure should enable users – other than display owners – to upload information and services to enrich the offer or adapt it to their own needs. For instance, users could upload personal content (e.g. annotated pictures of the local environment) from own devices to the infrastructure, both for better visualization and for increasing the local information offer. Similarly, a group of friends could exploit a public display to upload software letting, it host a shared real-time map to visualize what is happening around. This would also require opportunistic access the existing environmental sensors or the available user-provided sensors, to make the map alive and rich in real-time information. Clearly, if the map accesses some sensors in unconventional ways to better detect situations around, this adds value both to such sensors and to all existing and future services requiring situation recognition.

In traffic control systems, a driver could upload in the infrastructure the whole history of events related to a recent trip to enhance the historical data available to the infrastructure, or to a new route recommender service tuned to his peculiar needs. In the area of mobile social services, the added value of exploiting user-generated services and content is already evident, due to modern smart phones and related application suites. Yet, the full potential for value co-creation is currently limited by the centralized model of data and service production (mostly relying on application and data stores), and by the impossibility of opportunistically making available and access – in a decentralized way – data, digital services, devices, not to mention human computation services.

2.4 Eternity

Beside short-term adaptation, the infrastructure should tolerate long-term evolutions of structure, components, and usage patterns, to accommodate the changing needs of users and technological evolution without forcing significant re-engineering efforts (Jazayeri, 2005).

In the long-term, any pervasive infrastructure will experience dramatic changes related to the technology being adopted, to the kinds of services being deployed, and to their usage patterns. Display infrastructures, real-time traffic control systems, and social and augmented reality services will somewhen integrate more sophisticated sensing means than today, new actuators via which to attract user attention and interact with them (e.g. personal projection systems to make any physical object become an interactive display or eyeglass displays for immersive perception), and will increasingly involve a variety of peculiar human computation capabilities.

While this will open up the way for brand new classes and generation of services to be conceived and deployed, it will require that such evolution – at both the level of devices and services – can be gradually accommodated without harming the existing infrastructure and services.

3. The natural inspiration

The above requirements can hardly be met by traditional approaches to service systems engineering, e.g. SOA (Huhns and Singh, 2005), and rather suggest heading towards nature-inspired approaches.

3.1 Limitations of service-oriented architectures

In SOA, services components are simply considered *loci* of functionalities whose activities are triggered and coordinated according to pre-defined interaction patterns, with the support of many centralized middleware services for discovery, routing, orchestration, and context-awareness. For instance, in the pervasive display scenario, one could imagine to set up a middleware server in which to host all the necessary infrastructural services to support the various components of the scenario, i.e. displays, information and advertising services, user-provided services, sensing devices, and personal devices.

Situatedness and spatiality are not primary abstractions, and can be accommodated only at the price of complicating service descriptions, discovery, and orchestration services, to account for the current location of components and their frequent updates. In fact, in dynamic scenarios, components would be forced to continuously access (or being notified by) the discovery service to preserve up-to-date information – a computational and communication wasting activity. Also, since discovery and interactions among components have to rely on spatial information (e.g. a display is interested only in the users and sensors in its proximity), this requires either sophisticated context-services to extract the necessary spatial information about components, or to embed spatial descriptions for each component into its discovery entry, again inducing frequent and costly updates to keep up with mobility.

As far as adaptivity is concerned, the typically static interaction patterns enforced by SOA tools make self-adaptability and self-management hard to be integrated in a system, calling either for adaptation tools and logics to be embedded in the middleware (as, e.g. proposed by autonomic computing approaches (Kephart and Chess, 2003; Brazier *et al.*, 2009) or for components that are able to recognize changes and autonomously adapt accordingly. However, such adaptation logics would have to be very complex and heavyweight to ensure capability of adapting to any foreseeable situation, and especially hard for long-term adaptivity.

To reduce the identified complexities and costs and better match the characteristics of the scenario, one could think at a more distributed solution, with a variety of middleware servers deployed in the infrastructure to serve, on a strictly local basis, only a limited portion of the overall infrastructure. Although this solution can somehow mitigate the identified problems, and in particular better match the spatial nature of the scenario, its does not eradicate the basic identified problems related to adaptivity and long-term evolvability.

3.2 Towards nature-inspired approaches

Many initiatives – like those named upon digital/business service ecosystems (Barros and Dumas, 2006; Ulieru and Grobbelaar, 2007) – recognize the above limitations of SOA, and suggest that the complexity of modern service systems (and of pervasive ones) is comparable to that of natural ecosystems. Yet, the idea that nature – other than a mean to metaphorically characterize their complexity – can become the source of inspiration for their actual modeling and implementation is far from being generally metabolized.

Nevertheless, an increasing number of experiences suggests that, rather than complicating SOA to integrate the emerging requirements, it could be worthwhile adopting natural systems as an inspiration to fully re-think their modeling and architectural assumptions (Agha, 2008). That is, modeling and architecting pervasive service systems in terms of nature-inspired decentralized ecosystems, so as to make spatiality, adaptability, openness, and long-lasting evolvability inherently there because of the basic "rules of the game," as it is in real-world natural systems.

In natural systems (and whether you think at the physical, chemical, biological, or social level), all the activities of the system components are inherently situated in space and driven by local interactions only. Such interactions are not ruled by pre-defined orchestrated patterns. Rather, interactions are simply subject to a limited set of natural laws (let us generically call them "eco-laws"), from which even complex patterns of interactions dynamically emerge via self-organization. In this way, adaptivity becomes an inherent characteristic deriving from the existence of self-organizing interactions patterns, whose structure can flexibly yet robustly re-shape in response to contingencies.

Accommodating new and diverse component species (to support both decentralized prosumption and long-term evolution) reduces to simply enabling components to enter the ecosystem in respect of its eco-laws, and to let the dynamics of the interactions evolve and re-shape in response to the appearance of such new species. This way, one can take advantage of the new interactional possibility of such new species and of the additional value they bring in, without requiring the individual components or the infrastructure itself (i.e. its eco-laws and structure) be re-engineered (Jazayeri, 2005).

Nature-inspired computing paradigms have indeed already been extensively investigated (Kari and Rozenberg, 2008). However, many proposals exploit the natural inspiration only for the definition of general computational models (Berry and Boudol, 1990; Paun, 2003) or – for distributed and pervasive computing – of adaptive algorithmic solutions in the context of assessed architectural frameworks (Babaoglu *et al.*, 2006; Mamei *et al.*, 2006). Rather, here, we argue on the opportunity to adopt natural metaphors as comprehensive conceptual and practical frameworks for architecting pervasive service systems and their supporting infrastructures.

4. Natural metaphors for pervasive service ecosystems

One can think at natural systems from many viewpoints, each of which can act as a guiding metaphor for modeling and implementing the service components that will form the pervasive service ecosystem, the space in which they will execute, and the eco-laws that will rule their interactions.

In particular, so far, four key metaphors are being extensively explored to this purpose (Table I): physical (Mamei and Zambonelli, 2006), chemical (Berry and Boudol, 1990), biological (Parunak, 1997), along with metaphors focusing on higher level social models (e.g. trophic networks (May, 1974)).

All metaphors, adhering to a natural inspiration, are by construction spatially situated, adaptive by self-organization, and open to host diverse and evolving species. However, when it comes to analyze such metaphors for exploitation in pervasive computing scenarios, different metaphors may tolerate with variable efficiency and complexity the enforcement of adaptive self-organization patterns and the support of diversity and evolution. Also, since a service ecosystem is here to ultimately serve us, it is necessary to analyze how and to which extent the metaphors facilitate engineering and exerting of decentralized control over the ecosystem behavior, in order to properly direct its self-organizing activities and not to lose control over it.

Ideally, a metaphor should be able to support these features while limiting the number and complexity of eco-laws and the complexity of individuals and their environment, thus also keeping the supporting infrastructure lightweight and the overall execution efficient.

4.1 Physical metaphor

Grounded on amorphous computing (Abelson *et al.*, 2000) and field-based coordination models (Mamei and Zambonelli, 2006), the physical metaphor considers service components as sort of computational particles living in a world of other particles and virtual computational force fields, and capable of perceiving and reacting to local force fields.

Fields act as basic interaction mean, and are modeled as distributed data structures spread in the network in the form of spanning trees, and with a location-dependent intensity value associated. Activities of particles are driven by eco-laws that determine how particles spread fields, how fields propagate and reshape upon changing

Table I.
An overview of possible natural metaphors for services ecosystems

Key characteristics			Analysis		
Service components	Eco-laws	Space	Adaptive self-organization	Diversity and evolution	Decentralized control
Physical					
Particles (computational components) and virtual computational fields	Movements and activities driven by fields (gradient ascent/descent by components)	Network topology or physical space	+ Local and global self-organizing spatial structures can be effectively accommodated	– – Few new components species can be accommodated while keeping the laws simple	+ We know well how to build and control specific structures in physics
Chemical					
Atoms and molecules (semantic descriptions representing chemical properties)	Chemical reactions (matching of semantic descriptions and bonding of components)	Localities (pervasive computing environments)	– – Mostly local self-organizing structures can be effectively accommodated	++ Several new components species can be accommodated with the same basic laws	+ Reactants and catalysts can exert control over the dynamics and structure of reactions
Biological					
Simple goal-oriented organisms (e.g. ants) and pheromones	Diffusion and evaporation of chemical pheromones (affecting behaviour and activities of components)	Network topology or physical space	+ 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					
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)	+ Local self-organizing structures can be mostly accommodated, 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

conditions, and how they influence particles. Particles change their status based on the locally perceived fields, by reacting to the sensing of some fields matching some type (or semantic pattern), and move or exchange data by navigating the gradient of such fields. For example, services can generate sorts of data particles that follow downhill some existing fields so as to reach specific locations. The space in which particles live and in which fields diffuse can be either a metric world mapped on the physical space, or a virtual/social space mapped on the technological network.

In the pervasive display scenario, for instance, we can imagine display services as masses emitting gravitational-like fields. Such fields have different "flavors" (i.e. different semantic descriptions, reflecting the characteristics of nearby users and the environmental conditions) and an intensity proportional to the available display slots. Information and advertiser agents can behave as masses attracted by fields with specific flavors, eventually getting in touch with suitable displays for their information and ads. Upon changing conditions, the structure and flavors of diffused fields will change, providing for dynamically re-assigning information and ads to different displays.

The physical metaphor has been adopted for its effectiveness in promoting self-organizing coherent spatial behaviors (e.g. for load balancing, clustering, and aggregation) even in large-scale system, as in the TOTA framework for mobile computing (Mamei and Zambonelli, 2009) and in the PROTO language for sensor networks (Beal and Bachrach, 2006).

Indeed, the conceptual tools available for controlling the spatial behavior and the dynamics of such systems (by acting on how fields propagate and dynamically change) are well-developed, making it possible to exert control over the overall system behavior. However, such metaphor hardly tolerates high diversity and evolution. In fact, to support very diverse species and behaviors (at a time and over time), eco-laws must become expressive and flexible enough to tolerate a wide range of different field types, structures, and propagation rules. This increases the complexity of the model and, consequently, of the supporting infrastructure in charge of propagating and continuously updating many different fields.

4.2 Chemical metaphor

Grounded on chemically inspired computational models (Berry and Boudol, 1990), the chemical metaphor considers service components as sorts of computational atoms/molecules capable of re-acting to each other and bonding (i.e. composing) accordingly.

Components have to have a semantic description associated, acting as the computational counterpart of the bonding properties of physical atoms/molecules, yet made dynamic to reflect the current state and context of individual components. Accordingly, the eco-laws driving the ecosystem behavior resemble chemical reactions, dictating how bonding between components take place, based on some forms of pattern matching between semantic descriptions, and leading to aggregated, composite, and new components, other than to growth/decay of species. The world where individuals live is typically formed by a set of spatially confining localities, associated to local pervasive computing environments and intended as the "solutions" in which chemical interactions occur and across which chemicals can possibly diffuse.

In the display scenario, we can think of display services, information services (concerning ads and news), and user and environmental data, as molecules.

Displays represent different localities in which components react. Chemical rules dictate that, when the preferences of a user entering the locality of a display match an information/advertisement service, then a new composite component is created in charge of displaying that service in the display. Concurrently, in each locality, catalytic components can be in charge of re-enforcing the concentration of specific information/advertisements, reflecting the current situation of users. Also, localities can be open to enforce chemical bonds across displays, so that a high advertisement activity on a display can propagate to neighbors.

The chemical metaphor can effectively lead to self-organizing structures like local composite services and local aggregates, and many frameworks are being proposed relying on chemistry as the basis for service composition, such as Chemical Tuple Spaces (Viroli and Casadei, 2009) and the Higher Order Chemical language (Banatre and Priol, 2009).

As in real chemistry, the chemical metaphor can accommodate an incredible amount of different components and composites with a single set of basic eco-laws. That is, it can tolerate an increasingly diverse and evolving set of semantic descriptions for components, as required by open service systems, without affecting eco-laws and without increasing the burden on the infrastructure. As far as control is concerned, one can use sorts of catalyst or reagent components to engineer (in a decentralized way) the dynamics of the ecosystem. A limitation of the chemical approach is that it typically relies on activities taking place within a locality or at least across neighboring ones via local diffusion (Paun, 2003), making it hard to naturally enforce distributed self-organized behaviors, like complex long-range coordination of components.

4.3 Biological metaphor

Grounded on the study of collective behaviors in multi-cellular organisms and insect colonies (Parunak, 1997), the biological metaphor focuses on biological systems at the scale of individual organisms, or of colonies of simple organisms like ants.

The service components are abstracted either as simple cells or as animals, again capable of perceiving the state of the environment around them, yet acting on the basis of simple goals like finding food and reproducing. Similarly to physical systems, interactions take place by means of diffused signals, i.e. chemical pheromones of different flavors (types or semantic descriptions) spread by individuals in the environment, and slowly diffusing and evaporating in it. However, individuals here are not necessarily passively subject to the sensed pheromones, but react to them depending on their current "mood" (e.g. their state towards the achievement of a goal). Consequently, eco-laws are only aimed at determining how such pheromones, depending on their specific flavors, should propagate and diffuse in the environment. The spatial environment is again a computational landscape either mapped on the network topology or on the physical space.

In the display scenario, users can be represented by simple agents roaming around and spreading chemical signals with a flavor reflecting their personal interests. Displays can locally perceive such pheromones, and react by emitting different pheromones to express the availability of ads and information. Advertising and information agents, by their side, sense the concentration of such pheromones and are attracted towards the displays where the concentration of users with specific interests is maximized. The persistence of pheromones – in contrast to the typical volatility of physical fields – can be exploited to enforce stateful global strategies, e.g. via additional components that, by moving from display to display, create pheromone trails acting as routing paths that advertiser agents can use to globally find the best displays.

The biological metaphor appear very flexible in enabling the spatial formation of both localized and distributed activity patterns, and has been exploited in many pervasive frameworks, e.g. for crowd mobility management (Werfel *et al.*, 2008), participatory sensing (Lee *et al.*, 2009), and ubiquitous service pro-visioning (Hossain *et al.*, 2009).

The problems of accommodating diversity and evolution that affect the physical metaphor are here smoothed, making the biological metaphor more suited for open environments. In fact, an increase in the variety of pheromone flavors (to support diversity and evolution) can be handled with less overhead by the infrastructure, since pheromones (unlike fields) rely on local diffusion and slow evaporation dynamics. On the negative side, since the mechanisms of morphogenesis and self-organization in actual biological systems are not fully understood yet, it can be consequently hard to understand how to enforce control in their computational counterparts too.

4.4 Social metaphor

Grounded on the study of the dynamics of natural ecosystems (May, 1974), social metaphors focuses on biological systems at the level of animal species and of their trophic interactions.

Service components are goal-oriented entities (sorts of life forms) belonging to specific species (i.e. classes), that are in search of “food” resources (i.e. other service components or data items) to survive and prosper, and that can represent in their turn food to others (both aspects reflecting in some semantic descriptions). Pure data items and resources can be abstracted as passive life-forms (i.e. vegetables). The eco-laws are aimed at enacting and ruling the resulting “food web,” by bringing animals in touch with locally matching food to promote interactions and compositions among individuals of different species, by ruling production and reproduction of components as well as their movements, and thus ruling the dynamics and evolution of the ecosystem. Similarly to chemical systems, the shape of the world is typically organized around a set of localities, i.e. ecological niches, yet enabling diffusion of species across niches.

In the display scenario, assuming that each display is associated to an ecological niche, we can imagine users as sorts of herbivor agents roaming from niche to niche, possibly to eat those vegetable-like individuals representing information. Advertisers can be sorts of carnivors (i.e. eating other active life-forms) in need to find users with matching profiles to survive. For both cases, the primary effect of an eating action is the feeding of displays with information or ads to show. A possible secondary effect is the reproduction and diffusion in the environment of the best-fit species, e.g. of most successful ads. Those species that do not succeed in eating can either die or move to other niches to find food. Concurrently with the activities of the above species, background agents acting as sorts of bacteria can 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).

The social metaphor (such as the one here described but without forgetting the related body of work in market-based multi-agent systems (Vytelingum *et al.*, 2010) and social norms (Salazar *et al.*, 2010), similarly to the chemical one, appears suitable for local forms of self-organization – think at self-organized equilibria of web food patterns in ecological niches. Indeed, frameworks inspired by the social metaphor have been effectively experienced in the context of ambient intelligence systems (Agha, 2008; Villalba and Zambonelli, 2011) and of heterogeneous multi-robot systems (Cakar and Muller-Schloer, 2009).

Unlike the chemical metaphor though, and similarly to the biological one, the social metaphor can be effectively exploited also to enforce distributed forms of self-organization, by exploiting the capability of individuals of moving around to find food and reproduce. In addition, social metaphors are suitable for tolerating diversity and evolution at no additional burden for the infrastructure – think at how biodiversity has increased over the course of evolution without requiring any change to the underlying infrastructure (i.e. the earth). However, understanding how to control the behavior and dynamics of complex social computational ecosystems can be as difficult as it is in real social systems and in natural ecosystems.

5. Reference architecture and challenges

The presented metaphors have each pros and cons, which makes it impossible to suggest one as a general-purpose solution, the choice being dependent on the specific characteristics of the application problem and its requirements. Accordingly, a first very general challenge that researches in the area have to face is:

- Experimenting with variations and contaminations of the presented metaphors, so as to possibly extract a new general-purpose synthesis out them, getting the best of the presented metaphors while leaving our their limitations.

However, when the viewpoint is that of re-thinking the whole modeling and architecting of service systems in nature-inspired terms, it is a useful exercise to frame the key aspects of the different metaphors into a unifying reference architecture and, with it:

- Reason on the different metaphor by their commonalities rather then by their differences, and thus.
- Uniformly address the practical and technological challenges involved in the implementation of nature-inspired pervasive service ecosystems, so as to sketch promising directions that can be independent of the specific metaphor adopted and of the specific application scenario.

All of which, by being aware that our proposed reference architecture is necessarily influenced by our own backgrounds and visions, and is not necessarily be only possible one (or the best one) to reason about nature-inspired pervasive service ecosystems. For instance, it is possible that the assumptions, we have made in defining it have contributed in emphasizing some specific challenges and directions while leaving others hidden.

5.1 The architecture and its implementation challenges

Getting on the specific of the reference architecture (Figure 1), it is clear that a pervasive service ecosystem, whatever the specific metaphor adopted, will be deployed and executed above a concrete physical and digital ground. This will be a dense infrastructure (ideally a continuum) of decentralized networked computing devices, yet without excluding the possibility of exploiting and integrating centralized data and service

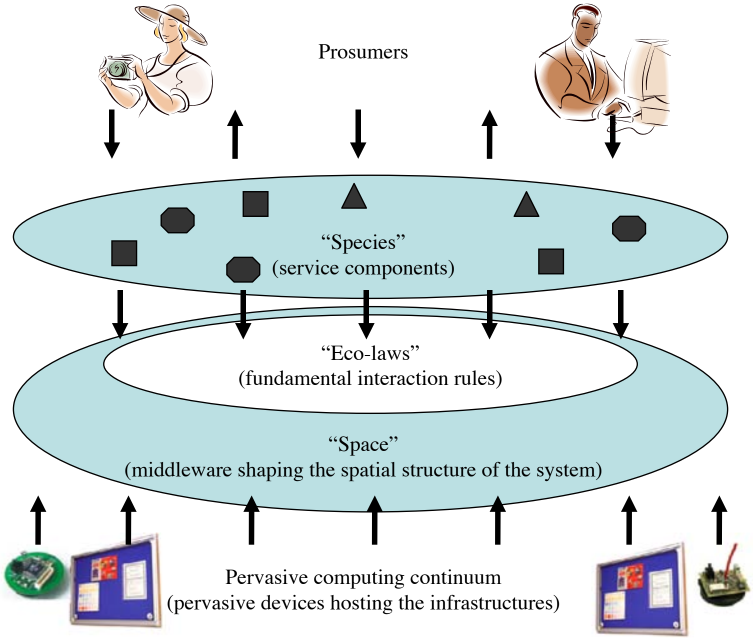

Figure 1.
A conceptual reference architecture for nature-inspired pervasive service ecosystems

centers within it. Prosumers access the framework for using/consuming data or services, for making available new services and new data components, other than to possibly make their own sensors, devices, and capabilities available to the infrastructure.

At both levels openness and dynamics arise: new devices can join/leave the system at any time, and new users can interact with the framework and can deploy new items on it. In the display scenario, for instance, we consider an infrastructure of displays, environmental sensors, and users' personal devices – with the possibility of integrating at any time new displays and new sensors – along with various classes of prosumers: display owners, advertisers, and simple users passing by and accessing the infrastructure to their own purposes.

In between these two levels, lay the abstract computational components of the pervasive ecosystem architecture and the challenges related to their implementation.

Species. This level includes a variety of components, belonging to different "species," representing the individuals populating the ecosystem. These may include, other than software service components *per se,* also the interfaces to the physical and virtual devices of the infrastructure, digital and network resources of any kind, persistent and temporary knowledge/data, contextual information, or personal user agents.

In the display scenario, for instance, we will have species to represent displays and their displaying service, the various sensors in the environment and the data they express, and software components to act on behalf of users, display owners, and advertisers.

Clearly, the adoption of a specific metaphor can lead to abstracting such components in terms of physical particles rather than of molecules or animals. However, in practical and technological terms, such components will always end up assuming the form of software agents (Muller, 1999), situated and executing in some portion of the ecosystem space, capable of perceiving the local situation around them, and possibly directed towards the achievement of some internal goals. The key challenge here is defining – either tuned for a specific metaphor or for some new synthesis – a suitable agent model, to properly balance the local degree of autonomous (possibly goal-oriented) behavior with the need of being subject to the reactive (eco-laws-driven) behavior as far as interactions are concerned.

Space. This level gives shape to the spatial fabric supporting individuals, their spatial activities and interactions, as well as their life-cycle. Given the spatial nature of pervasive services, this level has to situate individuals in space, so that their activities and interactions are directly dependent on their positions and on the shape of the surrounding space.

Practically, the spatial structure of the ecosystem will be reified by some minimal middleware substrate, deployed on top of the physical infrastructure, supporting the execution of individuals and their interactions. Although the metaphors, we have analyzed appear different in terms of spatial structure and interaction means, the middleware issues they involve are very similar. For individuals, the middleware will have to provide, via some application-programming interface, the following functionalities:

- the possibility of locally publishing themselves (e.g. for chemical and social metaphors, via their semantic description) and of locally spreading information (e.g. field and pheromones for the physical and biological metaphors); and
- the possibility of accessing information about their local spatial context (whether related to information about other individuals, or for accessing fields and pheromones), and detecting local events.

In addition, the middleware should provide for transparently absorbing dynamic changes and the arrival/dismissing of the supporting devices, without affecting the perception of the spatial environment by individuals.

Technologically, the above features do not necessarily require a complex suite of independent middleware services. For instance, a network of localized interaction spaces, spread on the pervasive substrate to store local information and notify events to local agents can be a good suit-able point (Eugster *et al.*, 2003). For instance, multiple tuple spaces accessible on a location-dependent basis, can uniformly act as repository of individuals' semantic descriptions and fields/pheromones and, thus, as a way for individuals to access and spread information and indirectly interact with each other. Indeed, recent proposals in the area of lightweight tuple-based middleware for pervasive and mobile devices, such as Linda in a Mobile Environment (Murphy *et al.*, 2006), can serve the purpose well.

Of course, mechanisms would be required to enforce (and dynamically re-shape upon contingencies) spatial relations among the network of tuple spaces, so as to enable coordinated spatial diffusion of information across them (as required for handling both fields/pheromones propagation and the movements of components). In the display scenario, for instance, one could assign one tuple space to each display, and have them dynamically self-configure (and re-configure) their spatial domain of competence accordingly to geographical and "line of sight" factors.

In this context, the algorithmic tools to achieve such dynamic spatial relations can be effectively inspired by recent works in the area of spatially structured overlay networks (Androutsellis-Theotokis and Spinellis, 2004). However, the challenge remains of defining proper spatial models, among the many possible ones, such as physical space models (mapping the space on physical coordinates), network space ones (conceiving the space as a graph), or logic space ones (shaping the space in terms of logical places such as rooms, streets, etc.). In addition, the challenge of having multiple space models co-exists and properly relate may appear in many application scenarios (e.g. pervasive services that have to manage both outdoor physical positions and indoor logic positions).

Eco-laws. The way in which individuals live and interact is determined by the set of fundamental "eco-laws." Enactment of eco-laws on individuals will typically affect and be affected by the local space around and by the other individuals around. In the display scenario, eco-laws might provide at automatically and dynamically determining to display a specific information on a screen as a sort of automatic reaction to specific environmental conditions, or at having two displays spontaneously aggregate and synchronize with each other in showing specific ads.

For the analyzed metaphors, the apparently different eco-laws all end up dealing with local pattern-matching processes and local spatial diffusion of information and events. This may include: for physical and biological metaphors, spreading fields and/or pheromones and have service components react to events related to local fields/pheromones of specific flavor (i.e. matching some criteria); for chemical metaphors, having components with matching descriptions properly bond with each other, possibly diffusing such process across localities; for social metaphors, having individuals gets properly in touch with matching "food" and move in space along with their semantic descriptions.

In general, the description of both service components and of data components such as fields and pheromones, along with the matching rules, define how the eco-laws apply to specific species in specific conditions of the space, and how local and spatial interactions among individuals have to take place. Thus, a key challenge here is to make the concept of "semantic description" of traditional SOA to facilitate discovery evolve into a concept of "alive semantic description." That is, dynamically changing as the context and state of components change, to enact eco-laws in dependence of current conditions. Equally challenging from the viewpoint of eco-laws is that individuals of different nature (service components, data components, sensors, and humans) should be uniformly treated, i.e. all of them should assume the form of alive semantic descriptions published and/or diffused in the tuple spaces network.

Practically, to be able to code and enact eco-laws, the middleware substrate should proactively mediate inter-component interactions. That is, it should act as an active space to store the continuously updating semantic descriptions, and adaptively support the matching process by triggering eco-laws in dependence of the current ecosystem conditions. Technologically, since modern tuple-based systems are enriched with the capability of reacting to events and of configuring the matching process (Murphy *et al.*, 2006; Viroli and Casadei, 2009), they could well act as starting point towards the definition of the support in which to locally embed and enact eco-laws.

5.2 Further general challenges

Within the proposed architecture, the overall dynamics of the ecosystem will be determined by individuals acting based on their own goals/attitudes, yet being subject to the eco-laws for their interactions with others. And accordingly drive – with an effectiveness that can depend on the specific metaphor and provided that there will exist tools to engineer and control such dynamics – most of the typical patterns of adaptive self-organization (e.g. spontaneous service aggregation or service orchestration, where the eco-laws play an active role in facilitating individuals to spontaneously interact and orchestrate with each other, also in dependence of current conditions) and adaptive evolution (changing conditions reflect in changes in the way individuals in a locality are affected by the eco-laws).

However, the capability of effectively developing pervasive service ecosystems capable of functioning in that way requires facing further research challenges.

Methodologies. There will be need of defining proper methodologies and tools for their development and engineering, as well as for the dynamic and decentralized control of the overall ecosystem dynamics. A key challenge here is be to understand the trade-offs between the power of top-down adaptation and bottom-up one, also by studying how the two approaches can co-exist (and they will indeed have to) and possibly conflict in future systems, and to contribute in smoothing the tension between the two.

Control and measurement. Methodologies should come along with tools to control the overall behavior of the pervasive systems and its sub-parts. Clearly, the challenging issue of defining such models and tools, come along with the issue of identifying means to measure such behaviors in order to understand if the control is effective. Defining sound measures for future pervasive scenarios is in itself challenging, given the size of the target scenario and the many – and often quite ill-defined – purposes it has to concurrently serve.

Security and privacy. To handle the extreme openness of the system, there is need to identify proper security mechanisms and policies. Specifically, it would be important to investigate the integration of security and privacy policies not in terms of external tools, but as part of the inherent self-adaptive dynamics of the system itself.

Legacy. For pervasive service ecosystems to hit the road, there will be need of proper strategies enabling a smooth transition to nature-inspired ecosystems, accounting for the legacy of current pervasive systems and infrastructures.

Scientific foundations. Finally, but not less important, the road towards pervasive service ecosystems will have to ground on new scientific foundations. A promising possibility is to ground on the emerging sciences of service systems (Spohrer *et al.*, 2007) and of participatory socio-technical systems (Olson *et al.*, 2010), and possibly be able to contribute to them.

6. Conclusions

Nature-inspired models and architectures appear the way to go for emerging and future pervasive service systems, in order to easily integrate spatial concepts, promote spontaneous adaptivity, support diversity and evolution in service components, and tolerate long-term evolutions at limited re-engineering costs.

Guided by the proposed reference architecture, we envision the possibility to adopt any of the natural metaphor, whenever suited to a specific application scenario, without always re-thinking modeling and architectural assumptions. In addition, the proposed architecture can be an effective starting point towards the definition and implementation of a unifying general-purpose synthesis of the presented metaphors.

Of course, the widespread deployment of nature-inspired pervasive service ecosystems also requires facing a number of fascinating research challenges, such as the ones we have identified and the many more that we may have missed.

References

Abelson, H., Allen, D., Coore, D., Hanson, C., Homsy, G., Knight, T.F. Jr, Nagpal, R., Rauch, E., Sussman, G.J. and Weiss, R. (2000), "Amorphous computing", *Communications of the ACM*, Vol. 43 No. 5, pp. 74-82.

Agha, G. (2008), "Computing in pervasive cyberspace", *Communications of the ACM*, Vol. 51 No. 1, pp. 68-70.

Androutsellis-Theotokis, S. and Spinellis, D. (2004), "A survey of peer-to-peer content distribution technologies", *ACM Computing Surveys*, Vol. 36 No. 4, pp. 335-71.

Babaoglu, O., Canright, G., Deutsch, A., Caro, G.A.D., Ducatelle, F., Gambardella, L.M., Ganguly, N., Jelasity, M., Montemanni, R., Montresor, A. and Urnes, T. (2006), "Design patterns from biology for distributed computing", *ACM Transactions on Autonomous and Adaptive Systems*, Vol. 1 No. 1, pp. 26-66.

Banatre, J.-P. and Priol, T. (2009), "Chemical programming of future service-oriented architectures", *Journal of Software*, Vol. 4 No. 7, pp. 738-46.

Barros, A.P. and Dumas, M. (2006), "The rise of web service ecosystems", *IT Professional*, Vol. 8 No. 5, pp. 31-7.

Beal, J. and Bachrach, J. (2006), "Infrastructure for engineered emergence on sensor/actuator networks", *IEEE Intelligent Systems*, Vol. 21 No. 2, pp. 10-19.

Berry, G. and Boudol, G. (1990), "The chemical abstract machine", *Proceedings of the 17th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ACM Press, New York, NY, pp. 81-94.

Biem, A., Bouillet, E., Feng, H., Ranganathan, A., Riabov, A., Verscheure, O., Koutsopoulos, H.N., Rahmani, M. and Guc, B. (2010), "Real-time traffic information management using stream computing", *IEEE Data Engineering Bullettin*, Vol. 33 No. 2, pp. 64-8.

Brazier, F.M., Kephart, J.O., Parunak, H.V.D. and Huhns, M.N. (2009), "Agents and service-oriented computing for autonomic computing: a research agenda", *IEEE Internet Computing*, Vol. 13 No. 3, pp. 82-7.

Cakar, E. and Muller-Schloer, C. (2009), "Self-organising interaction patterns of homogeneous and heterogeneous multi-agent populations", *Proceedings of the Third IEEE International Conference on Self-adaptive and Self-organizing Systems*, IEEE CS Press, Washington, DC, pp. 165-74.

Campbell, A.T., Eisenman, S.B., Lane, N.D., Miluzzo, E., Peterson, R.A., Lu, H., Zheng, X., Musolesi, M., Fodor, K. and Ahn, G.-S. (2008), "The rise of people-centric sensing", *IEEE Internet Computing*, Vol. 12 No. 4, pp. 12-21.

Coleman, B. (2009), "Using sensor inputs to affect virtual and real environments", *IEEE Pervasive Computing*, Vol. 8 No. 3, pp. 16-23.

Eugster, P.T., Felber, P.A., Guer-raoui, R. and Kermarrec, A.-M. (2003), "The many faces of publish/subscribe", *ACM Computing Surveys*, Vol. 35, pp. 114-31.

Ferscha, A. and Vogl, S. (2010), "Wearable displays – for everyone!", *IEEE Pervasive Computing*, Vol. 9 No. 1, pp. 7-10.

Hossain, M.S., Alamri, A. and El Saddik, A. (2009), "A biologically inspired framework for multimedia service management in a ubiquitous environment", *Concurrency and Computation: Practice and Experience*, Vol. 21 No. 11, pp. 1450-66.

Huhns, M.N. and Singh, M.P. (2005), "Service-oriented computing: key concepts and principles", *IEEE Internet Computing*, Vol. 9 No. 1, pp. 75-81.

Jazayeri, M. (2005), "Species evolve, individuals age", *Proceedings of the 8th IEEE International Workshop on Principles of Software Evolution*, ACM Press, Washington, DC, pp. 3-12.

Kari, L. and Rozenberg, G. (2008), "The many facets of natural computing", *Communications of the ACM*, Vol. 51 No. 10, pp. 72-83.

Kephart, J.O. and Chess, D.M. (2003), "The vision of autonomic computing", *IEEE Computer*, Vol. 36 No. 1, pp. 41-50.

Lane, N.D., Miluzzo, E., Lu, H., Peebles, D., Choudhury, T. and Campbell, A.T. (2010), "A survey of mobile phone sensing", *IEEE Communications Magazine*, Vol. 48 No. 9, pp. 140-50.

Langlotz, T., Wagner, D., Mulloni, A. and Schmalstieg, D. (2011), "Online creation of panoramic augmented reality annotations on mobile phones", *IEEE Pervasive Computing*, Vol. 99 No. 1, p. 10.

Lee, U., Magistretti, E., Gerla, M., Bellavista, P., Lio, P. and Lee, K.-W. (2009), "Bio-inspired multi-agent data harvesting in a proactive urban monitoring environment", *Ad Hoc Networks*, Vol. 7, June, pp. 725-41.

Mamei, M. and Zambonelli, F. (2006), *Field-based Coordination for Pervasive Multiagent Systems* (*Springer Series on Agent Technology*), Springer-Verlag New York, Secaucus, NJ.

Mamei, M. and Zambonelli, F. (2009), "Programming pervasive and mobile computing applications: the TOTA approach", *ACM Transactions on Software Engineering and Methodology*, Vol. 18 No. 4, pp. 1-56.

Mamei, M., Menezes, M., Tolksdorf, R., Zambonelli, R. and F. (2006), "Case studies for self-organization in computer science", *Journal of Systems Architecture*, Vol. 52 No. 8, pp. 443-60.

May, R.M. (1974), *Stability and Complexity in Model Ecosystems*, Princeton University Press, Princeton, NJ.

Muller, J.P. (1999), "Architectures and applications of intelligent agents: a survey", *Knowledge Engineering Review*, Vol. 13, pp. 353-80.

Murphy, A.L., Picco, G.P. and Roman, G.-C. (2006), "Lime: a coordination model and middleware supporting mobility of hosts and agents", *ACM Transactions on Software Engineering and Methodology*, Vol. 15 No. 3, pp. 279-328.

Olson, G., Mark, G., Churchill, E. and Rotman, D. (2010), "New missions for a sociotechnical infrastructure", *IEEE Computer*, Vol. 43, pp. 37-43.

Parunak, V. (1997), "Go to the ant: engineering principles from natural multi-agent systems", *Annals of Operations Research*, Vol. 75, pp. 69-101.

Paun, G. (2003), "Membrane computing: fundamentals of computation theory", Vol. 2751 of Lecture Notes in Computer Science, Springer, Berlin, pp. 177-220.

Reddy, S., Estrin, D. and Srivastava, M.B. (2010), "Recruitment framework for participatory sensing data collections", *Proceedings of the 8th International Conference on Pervasive Computing*, Vol. 6030 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, pp. 138-55.

Rosi, A., Codeluppi, A. and Zambonelli, F. (2010), "Adaptive pervasive advertisement: scenarios and strategies", *Journal of Pervasive Computing and Communication*, Vol. 6 No. 3, pp. 333-51.

Salazar, N., Rodrıguez-Aguilar, J.A. and Arcos, J.L. (2010), "Robust coordination in large convention spaces", *AI Communications*, Vol. 23 No. 4, pp. 357-72.

Sippl, A., Holzmann, C., Zachhuber, D. and Ferscha, A. (2010), "Real-time gaze tracking for public displays", *Proceedings of the First International Joint Conference on Ambient Intelligence*, Vol. 6439 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, pp. 167-76.

Spohrer, J.C., Maglio, P.P., Bailey, J.H. and Gruhl, D. (2007), "Steps toward a science of service systems", *IEEE Computer*, Vol. 40 No. 1, pp. 71-7.

Ulieru, M. and Grobbelaar, S. (2007), "Engineering industrial ecosystems in a networked world", *5th IEEE International Conference on Industrial Informatics*, IEEE Press, New York, NY, pp. 1-7.

Vargo, S.L., Maglio, P.P. and Akaka, M.A. (2008), "On value and value co-creation: a service systems and service logic perspective", *European Management Journal*, Vol. 26 No. 3, pp. 145-52.

Villalba, C. and Zambonelli, F. (2011), "Towards nature-inspired pervasive service ecosystems: concepts and simulation experiences", *Journal of Network and Computer Applications*, Vol. 34 No. 2, pp. 589-602.

Viroli, M. and Casadei, M. (2009), "Biochemical tuple spaces for self-organizing coordination", *Coordination Languages and Models*, Vol. 5521 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, pp. 143-62.

Vytelingum, P., Ramchurn, S.D., Voice, T.D., Rogers, A. and Jennings, N.R. (2010), "Trading agents for the smart electricity grid", *International Conference on Autonomous Agents and Multiagent Systems*, ACM Press, New York, NY, pp. 897-904.

Werfel, J., Bar-Yam, Y. and Ingber, D. (2008), "Bio-inspired environmental coordination in spatial computing systems", *International SASO Workshop on Spatial Computing*, IEEE CS Press, Washington, DC, pp. 338-43.

Yuen, M.-C., Chen, L.-J. and King, I. (2009), "A survey of human computation systems", *IEEE International Conference on Computational Science and Engineering*, IEEE CS Press, Washington, DC, pp. 723-8.

Zambonelli, F. (2011), "Pervasive urban crowdsourcing: visions and challenges", *5th International PerCom Workshop on Pervasive Life, Learning, and Leisure*, IEEE CS Press, Washington, DC, March, pp. 578-83.

About the authors

Franco Zambonelli is a Professor of Computer Science at the University of Modena and Reggio Emilia. He obtained his PhD in Computer Science and Engineering from the University of Bologna in 1997. His research interests include: pervasive computing, multi-agent systems, and self-adaptive and self-organizing computing. He has published over 60 papers in peer-reviewed journals, and has been invited speaker at many conferences and workshops. He is the Co-editor in Chief of the *ACM Transactions on Autonomous and Adaptive Systems*, and he is on the Editorial Board of the Elsevier *Journal of Pervasive and Mobile Computing*, of the *BCS Computer Journal* and of the *Journal of Agent-Oriented Software Engineering*. Also, he is on the Steering Committee of the IEEE SASO Conference. He has been scientific manager of the EU FP6 Project CASCADAS and is currently coordinator of the EU FP7 Project SAPERE. He is a senior member of ACM and IEEE. Franco Zambonelli is the corresponding author and can be contacted at: franco.zambonelli@unimore.it

Mirko Viroli is a Research Associate at DEIS, the Department of Electronics, Informatics and Systems of the Alma Mater Studiorum – Università di Bologna (UNIBO), Italy. He obtained his PhD in Computer Science and Engineering from the University of Bologna in 2004. He is an expert in computation models, in the areas of object-oriented programming languages, agent-based systems, coordination infrastructures, and self-organizing systems. He has written over 150 articles on such topics, more than 30 of which have appeared in international refereed journals. He currently leads the UNIBO team in the FP7 STREP project SAPERE.