Nature-inspired Spatial Metaphors for Pervasive Service Ecosystems

Cynthia Villalba[1], Alberto Rosi[1], Mirko Viroli[2], Franco Zambonelli[1]
1) University of Modena and Reggio Emilia – 42100 Reggio Emilia, Italy
2) University of Bologna – 40723 Cesena (FC), Italy
cynthia.villalba@unimore.it, alberto.rosi@unimore.it, mirko.viroli@unibo.it, franco.zambonelli@unimore.it

Abstract

Innovative paradigms and frameworks have to be identified to enable the effective deployment and execution of pervasive computing services. Such frameworks must be conceived so as to match the spatially-situated nature of pervasive services, and must be able to exhibit properties of self-organization and self-adaptability, self-management, and of long-lasting evolvability. This paper discusses how such frameworks should get inspiration from natural systems, by enabling modeling and deployment of services as autonomous individuals, spatially-situated in an ecosystem of other services, data sources, and pervasive devices, all of which acting, interacting, and evolving according to a limited set of spatial "eco-laws". In this context, this paper presents a reference architecture to uniformly frame ecosystem concepts, surveys and critically analyzes different nature-inspired spatial metaphors to realize the idea, and details our current research agenda concerning the development of service frameworks inspired to the ecological metaphor.

1 Introduction

Pervasive and mobile computing devices increasingly populate our everyday environments [7]. These, together with the increasing amount of Web tools that makes it possible to produce and access spatially-situated information about the physical world [5], will eventually define a comprehensive, integrated, and very dense, decentralized shared infrastructure for general-purpose usage. At the user level, the infrastructure can be used to access innovative services for better perceiving/interacting with the physical world and for acting on it. 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 in which they are invoked and exploited.

The inherent spatial nature of the above infrastructure and of all the services that will be deployed over it is very sharp. On the one hand, the infrastructure will be embedded into physical space, will have to deal with spatial concepts and spatial data, and its devices will typically interact based on spatial proximity (as induced by wireless communications). On the other hand, services will have to deal with spatially-situated activities of users, and with their interacting with the physical world.

The effective development and execution of services in the above infrastructure calls for a deep rethinking of current service models and for service frameworks, in order to:

- Naturally match the inherent spatial nature of the environment and of the services within.
- Inherently exhibit those properties of self-organization, self-adaptation and self-management that are necessarily required 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).
- Flexibly tolerate evolutions of structure and usage over time. This is necessary to account for increasingly diverse and demanding needs of users as well as for technological evolution, without forcing significant (and economically unbearable) re-engineering to incorporate innovations and changes.

To reach this 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, whose activities are orchestrated according to specific pre-defined patterns with the support of middleware services (such as discovery, routing, and context services) that either miss in accounting spatial concepts or do not elect them as primary abstractions; 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 [10]; 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 ecosystems [16, 8], 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, 11] 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. Here we go further, and argue that natural ecosystem can act as the key metaphor around which to conceive, model, and develop, fully-fledged pervasive service framework and all the components within.

You can think at physical systems, at chemical systems, at biological systems, as well as at the most properly called ecological systems. In all of them, you can always recognise the following characteristics: above a spatial environmental substrate, individuals of different kinds (or species) interact, compete, and combine with each other in respect of the basic laws of nature. Accordingly, in our scenario, the shared pervasive infrastructure substrate will have to be conceived as the space in which bringing to life an ecosystem of services, intended as individuals whose computational activities are subject to some basic laws of the ecosystem, and for which the dynamics of the ecosystem (as determined by the enactment of its laws) will provide for naturally enforcing features of self-organization, self-management, and evolvability.

In this context, the contributions of this paper are as follows:

- We introduce a unifying reference architecture for nature-inspired pervasive service ecosystems, to show how ecosystem concepts can be framed into a unifying conceptual scheme (Section 2).
- We survey the different metaphors that can be adopted for such ecosystems (Section 3), and discussed their advantages and limitations w.r.t. the capability of supporting self-organization and self-adaptation, self-management and decentralized control, and evolution over time (Section 4).
- We go into more details about the so called ecological metaphor, and sketch our current research work and our research agenda in that area (Section 5), and conclude (Section 6).

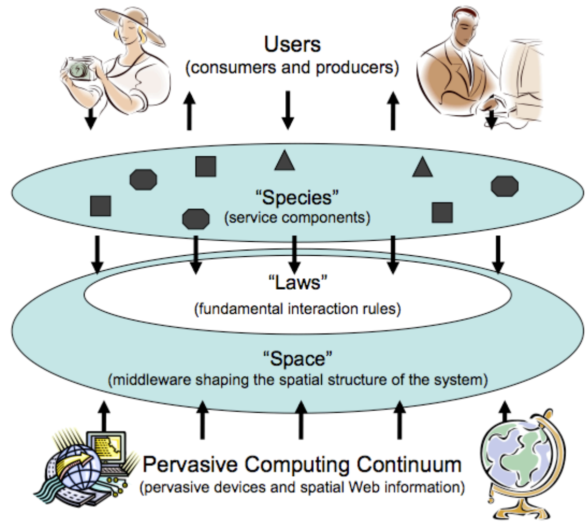

Figure 1. A Reference Architecture for Pervasive Service Ecosystems

2 A Reference Architecture for Pervasive Service Ecosystems

A unifying reference architecture can be identified around which to frame the key abstractions and the conceptual structure for spatial pervasive service ecosystems, independently of the specific metaphor adopted (see Figure 1).

At the lowest level is the physical 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 devices that are going to increasingly pervade all our everyday environments (e.g., PDAs, smart phones, sensors, tags), all interconnected with each other, and most of which generating 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 highest 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.

At both the bottom and the top levels, the architecture exhibits a high-degree of openness: new devices can join/leave the system at any time, and new users can interact with the framework and can deploy new services and data items on it. In between these two levels, there are the components of the pervasive ecosystem architecture.

The level of "Species" is the one in which physical and virtual devices of the pervasive system, digital and network resources of any kind, persistent and temporary knowledge/data, contextual information, events and information requests, and of course software service components, are all abstracted as "living entities" of the system (i.e., the *ecosystem individuals*) that populate the pervasive ecosystem space. 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".

In general terms, an ecosystem is expected to be populated with a set of individuals physically deployed in the environment (physical and network resources, contextual information, initialization data and services, and so on). Yet, the population of individuals is far from being static. First, the set of individuals is subject to changes (to tackle the physical system's mobility, faults, and evolution). Second, service developers and producers inject in the system new individuals at any time (they can insert new services and virtual devices, as well as data and knowledge). Third, producers and consumers can keep control and influence the behavior of (a limited set of) the individuals.

The "Space" level provides the spatial fabric supporting individuals, their spatial activities and interactions, as well as their life-cycle. From a conceptual viewpoint, the "Space" level gives shape to and defines the structure of the virtual world in which individual lives. 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 the modeling of the ecosystem.

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 concepts of 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. Such 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. The enactment of the eco-laws requires the presence of some meaningful description (within the uniform modeling on individuals) of the information/service/structure/goals of each species, and of proper "matching" criteria to define, based on such description, how the eco-laws apply to specific species in specific conditions of the space.

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 characteristic 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 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). 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.

3 Metaphors for Pervasive Service Ecosystems

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, the space in which they live, and its laws. 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 so far are: physical metaphors [6, 12], chemical metaphors [3, 14], biological metaphors [2, 4, 15], together with the most properly called ecological metaphors [1, 13].

	Species	Space	Laws
Physical	Particles (computational components) and messages (computational fields)	The Universe (a network), as shaped by waves and particles.	Navigation and activities driven by fields (gradient ascent by components)
Chemical	Atoms (semantically described) and Molecules (composed semantic descriptions)	Space (localities/bags of components)	Chemical Reactions (matching of semantic descriptions and bonding of components)
Biological	Cells (amorphous computing cells, modules of self-assembly components)	Space (Abstract computational landscapes, or physical landscapes)	Diffusion of chemical gradients and morphogens, differentiation of behaviour and activity
Ecological	Organisms (Agents) and Species (Classes) and Resources (Data)	Niches (Pervasive computing environments)	Survive (goal-orientation), eat, produce, and reproduce

Figure 2. Metaphors for Service Ecosystems

As far as we know, none of these metaphors has been so far adopted to extensively studying and prototyping an actual, open and general-purpose service framework: either the metaphor has been applied to specific application scenarios [12, 4, 15] or its potential general adoption has been only envisioned [6, 1].

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, which act as the basic interaction means. In fact, all activities of particles are driven by laws that determine how particles should be influenced by the local gradients and shape of some computational field (those whose description "matches" some criterion). In particular, they can change their status based on specific 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 space in which such particles live and in which fields spread and diffuse can be either a simple (euclidean) metric world, or it could 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).

Chemical metaphors consider that the species of the ecosystem are sorts of computational atoms/molecules, with properties described by some sort of semantic descriptions which are the computational counterpart of the description of the bonding properties of physical atoms and molecules. Indeed, the laws that drive the overall behaviour of the ecosystem are sort of chemical laws. They dictate 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 can lead to both the production of aggregates (e.g., of aggregated distributed components) or of new components (e.g., of composite components). In this case, the space in which components live is typically formed by a set of localities, intended as the "solution" in which chemical reactions can occur, altough of course it is intended that components can flow/diffuse across localities to ensure globality of interactions.

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 (unintelligent) animals, that act on the basis of very simple goal-oriented behaviours (e.g., move and eat) and that are influenced in their activities by the strength of specific chemical signals in their surroundings to which they are sensitive to (i.e., with which there is a match). Similarly to physical systems, in fact, components are expected (depending on their status) to be able to spread and diffuse (chemical) signals around, that can then influence the behaviour of other components. The laws of the ecosystem together with the shape of the spatial computational landscape in which individuals live determine how such signals should diffuse, and how they could influence the behaviour and characteristics of components.

Ecological metaphors focus on biological systems at the level of animal species and of their interactions. The 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 matching specific criteria). The laws of the ecosystem determine how the resulting "web of food" should be realised, that is, they determine 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 space 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.

4 Critical Analysis

As already stated in the introduction, a pervasive service ecosystem should be able to exhibit features of self-organization and self-adaptation (i.e., the capability of autonomously and adaptively self-organize and self-adapt the distributed spatial activities of the components) and self-management (here mostly intended as the possibility of exerting control and directing the behavior of the system form within the system itself, a fundamental feature not to lose control over the system and not to be forced to introduce complex management solutions), and should tolerate evolution and adaptation over time (i.e., should adaptively accommodate new species, should survive the extinction of species, and should be capable of accommodating very diverse and composite behaviour with the same limited set of eco-laws). 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 space in which they live.

The analysis of the extent to which the presented metaphors can be 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. Nevertheless, we can try at least to draw some considerations about this.

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 and data distribution), and the conceptual tools available for controlling the spatial behaviour and the dynamics of such systems are well-developed. However, the physical metaphor seems to fall short in evolution and time adaptation, in that it hardly tolerates the presence of very diverse components with very diverse behaviours (at least if we want to preserve the simplicity of the eco-laws).

Chemical metaphors, on the other hand, can effectively lead to local self-organizing structures (e.g., local composite services) and, to a more limited extent, to some sorts of global structures (e.g., networks of distributed homogeneous components, as in crystals). Real chemistry, and so chemical computational metaphors, can accommodate an incredible amount of different components and composites, yet with the same set of simple basic laws. This is an important pre-condition for facilitating evolution over time. As far as self-management is concerned, one can think at using sort of catalyst or reagent components to control the dynamics and the behaviour of a chemical ecosystem.

Biological metaphors appear very flexible in enabling the spatial formation of localised morphological and activity patterns, and this has been shown to have notable applications in a variety of applications to distributed systems. However, the number of patterns that can be enforced by the spread of chemical gradients and by the reactions of simple individuals seems (as it is in physical metaphors) quite limited, and this does not match with the need for time evolution and adaption. Moreover, it is quite difficult to understand how to properly control the overall behavior of such systems (just think at the fact that, so far, the mechanisms of morphogenesis are not fully understood by scientists).

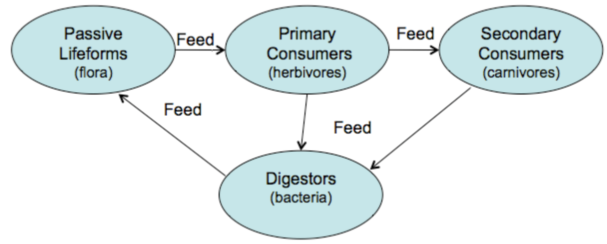

Figure 3. Key Elements for an Ecological System

Ecological metaphors, the same as chemical ones, promise to be very suitable for local forms of spatial self-organization (think at equilibria in ecological niches), and 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, unlike chemical systems, understanding how to properly control the local and global equilibria of real ecological system is a difficult task, and it 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.

Personally, and having already extensively studied in the physical metaphor in the past [12], we are now very interested in studying both the chemical metaphor (see for a preliminary study about [14]) and the ecological one (which is the specific subject of the PhD studies of the first author, and which is detailed in the next section), and in possibly ending up with a sound new "hybrid" metaphor, getting the best of the above.

5 Our Current Approach and Research Agenda

As from Section 3, the development of a pervasive service ecosystem inspired by the ecological metaphor should conceive the individuals within as the life forms of an ecosphere, each of which having the trivial ego-centric goal of surviving by finding the appropriate food and resources. The eco-laws thus reduce to simply ruling the dynamics of the food web (who eats who and when), and the spatial structure of the system (typically structured around spatially confined ecological niches) determines how life forms can find and look for food.

In general, an ecological system can consider the presence of different classes of living forms (see Figure 3). Passive life forms (i.e., the flora system) do not actively look for food, although their existence and survival must be supported by nutrients that have to be spread in space. Primary consumers (i.e., herbivors) need to eat vegetables to survive and prosper. Secondary consumers (i.e. carnivors) typically need to eat other animals to survive, though this does not exclude that can also act as primary consumers (eating vegatables too). The result of the metabolization of food by both primary and secondary consumers typically ends up in feeding lower-level "digestors" life forms (e.g., bacteria), densely spread in space, and that in their turn produce and diffuse necessary resources and nutrients for the flora.

Let us now translate the above concepts in computational terms. Passive life forms represent the data sources of the ecosystem, which are not to be considered proactive computational entities, i.e., they do not need to "do something" to exist and be used. Primary consumers represent those services that require to digest information to be of any use, and yet are computationally autonomous (they do not require external computational functionalities). Secondary consumers, instead, are those services that, to be of any use, need the support of other services (whether primary or secondary in their turn), other than possibly of information sources. Digestors can be generally assimilated to all those background computational services that are devote to monitor the overall activities of the system, and either produce new information about or influence the existing information.

To better clarify, let us present a simple case study we are currently in the process of developing.

Consider a scenario like a thematic park or an exhibition center, densely pervaded with digital screens where to display information, movies, advertisements, or whatever. We can consider each of these screens (i.e., the computational resources associated with each of them) as a spatially confined ecological niche. Different classes of visitors will watch these screens to look for different types of information (intended as passive life forms). Thus, we can think at sort of "user agents" executing on the users' PDAs that, once in the proximity of a screen (i.e., while finding themselves into that specific ecological niche) start looking for specific information to eat (i.e., to have it displayed). User agents would thus act as primary consumers. Concurrently, we can think at "advertising agents" that, acting on behalf of some advertising company, roam from screen to screen in search of specific classes of user agents (i.e. those interested in specific types of information), with the ultimate goal of displaying advertisements there where they could be more effective. Advertising agents would this act as secondary consumers. Background monitoring agents, executing on each ecological niche and possibly interacting with each other, can contribute replicating and spreading information there where it appears to be more appreciated, and can also contribute in supporting the spatial roaming of advertiser agents by directing them there where they could find more satisfaction. Thus, they would act as digestors.

The feedback loop that derives from the above activities can contribute the properly rule the overall dynamics of the screen ecosystem, by continuously self-organizing and self-adapting the way information flows in the system, as well as the way advertising agents move, act, and coordinate with each other. The possibility of exerting control over the dynamics of the system is ensured by the possibility of injecting in the system new classes of "digestor agents" that can radically influence the dynamics of information diffusion and the activities of advertising agents. The adaptation of the system over time is ensured by the fact that it is mostly irrelevant, for the overall functioning of the system, what specific classes of information user agents want, or what the specific goal of advertising agent is. In fact, independently of the specific species of life forms that will populate the system, the basic eco-laws will ensure that such life forms will either find their way of living and their role in the system (e.g., as it can be the case of useful information and of advertising agents that find appropriate users to which to display their ads), or will simply disappear (as it can be the case of useless information or of advertisements no users is interested in).

We personally believe that, within the above simple conceptual framework (mostly in line with that envisioned in [1]), we will be able to identify a simple yet usable general-purpose model for the design and development of specific data/service/control components, and will be able to develop a practical software framework for the execution of a wide class of distributed spatial services for pervasive environment. To this end, we are currently in the process of:

- Trying to identify a proper semantic representation of the needs and characteristics of each life forms (what food one agent class needs, and what kinds of nutrient it can represent to others);
- Define a simple agent-inspired computational model to have the action of agents, as well as their propagation and diffusion over space, driven by a simple set of "eat to survive" eco-laws;
- Implement a simple "middleware" infrastructure to test and put at work our ideas. Such middleware will typically rely on a spatially-structured network of nodes, each of which acting as the basic ecological niche for the local execution of components and their interactions, and interacting with close niches to enforce spatial diffusion and propagation of life forms;
- Put our ideas at work in a variety of application scenarios, to verify their generality and their extent of applicability;

- Evaluate how and to which extent to integrate and extend the sketched ecological approach with features and characteristics inspired by other metaphors, if at all needed.

6 Conclusions

In this paper, we have elaborated on the idea of getting inspiration from natural ecosystem to model and develop next generation pervasive service framework. That is, of conceiving future pervasive service frameworks as a spatial ecosystem in which services, data items, and resources are all modeled as autonomous individuals that spatially act and interact in accord to a simple set of well-defined "laws of nature". In this way, it is possible to deliver self-organization, self-adaptation, self-management, and long-lasting evolvability as inherent properties of the framework, rather than as complicated ad-hoc solutions.

The road towards the actual deployment of usable and effective pervasive service ecosystems, however, still requires answering to several challenging questions. Among the others: what actual metaphor is the best to be adopted among the possible ones? What should be the actual modeling of individual and of eco-laws? What should be the actual shape and properties of the space in which individual will live and interact? And how can we practically implement this? Finding at least some of these answers is the current goal of our research work.

References

[1] G. Agha. Computing in pervasive cyberspace. *Commun. ACM*, 51(1):68–70, 2008.

[2] O. Babaoglu, G. Canright, A. Deutsch, G. A. D. Caro, F. Ducatelle, L. M. Gambardella, N. Ganguly, M. Jelasity, R. Montemanni, A. Montresor, and T. Urnes. Design patterns from biology for distributed computing. *ACM Trans. Auton. Adapt. Syst.*, 1(1):26–66, 2006.

[3] A. P. Barros and M. Dumas. The rise of web service ecosystems. *IT Professional*, 8(5):31–37, 2006.

[4] J. Beal and J. Bachrach. Infrastructure for engineered emergence on sensor/actuator networks. *IEEE Intelligent Systems*, 21(2):10–19, 2006.

[5] G. Castelli, A. Rosi, M. Mamei, and F. 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] J. Crowcroft. Toward a network architecture that does everything. *Commun. ACM*, 51(1):74–77, 2008.

[7] D. Estrin, D. Culler, K. Pister, and G. Sukjatme. Connecting the physical world with pervasive networks. *IEEE Pervasive Computing*, 1(1):59 – 69, 2002.

[8] S. Herold, H. Klus, D. Niebuhr, and A. 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] M. N. Huhns and M. P. Singh. Service-oriented computing: Key concepts and principles. *IEEE Internet Computing*, 9(1):75–81, 2005.

[10] J. O. Kephart and D. M. Chess. The vision of autonomic computing. *Computer*, 36(1):41–50, 2003.

[11] M. Mamei, R. Menezes, R. Tolksdorf, and F. Zambonelli. Case studies for self-organization in computer science. *Journal of Systems Architecture*, 52(8-9):443–460, 2006.

[12] M. Mamei and F. Zambonelli. *Field-based Coordination for Pervasive Multiagent Systems*. Springer Verlag, 2006.

[13] M. D. Peysakhov, R. N. Lass, and W. C. Regli. Stability and control of agent ecosystems. In *AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 1143–1144. ACM, 2005.

[14] R. Quitadamo, F. Zambonelli, and G. Cabri. The service ecosystem: Dynamic self-aggregation of pervasive communication services. In *SEPCASE '07: Proceedings of the 1st International Workshop on Software Engineering for Pervasive Computing Applications, Systems, and Environments*, page 1, Washington, DC, USA, 2007. IEEE Computer Society.

[15] W.-M. Shen, P. Will, A. Galstyan, and C.-M. Chuong. Hormone-inspired self-organization and distributed control of robotic swarms. *Autonomous Robots*, 17(1):93–105, 2004.

[16] M. Ulieru and S. 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.