From Service-Oriented Architectures to Nature-Inspired Pervasive Service Ecosystems

Franco Zambonelli
Dipartimento di Scienze e Metodi dell'Ingegneria
Università di Modena e Reggio Emilia, Italy
Email: `franco.zambonelli@unimore.it`

Mirko Viroli
Dipartimento di Elettronica Informatica e Sistemistica
Alma Mater Studiorum – Università di Bologna, Italy
Email: `mirko.viroli@unibo.it`

***Abstract*—Emerging pervasive computing scenarios require open service frameworks promoting situated adaptive behaviors and supporting diversity in services and long-term evolvability. We argue that this naturally calls for a nature-inspired approach, in which pervasive services are modeled and deployed as autonomous individuals in an ecosystem of other services, data sources, and pervasive devices. As an evolution of standard service-oriented architectures, we present a general framework framing the concepts expressed, and discuss a number of natural metaphors that we can adopt to concretely incarnate the proposed framework and implement pervasive service ecosystems.**

I. INTRODUCTION

The ICT 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 RFID tags and alike, objects will carry on digital information of any sort. Wireless sensor networks and camera networks will be spread in our cities and buildings to monitor physical phenomena. Smart phones and alike will increasingly sense and store notable amounts of data related to our personal, social and professional activities, other than feeding (and being fed by) the Web with spatial and social real-time information [1].

This evolution is contributing to building 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 used to enrich existing services with the capability of autonomously adapting their behavior to the physical and social context in which they are invoked, and will also support innovative services for enhanced interactions with the surrounding physical and social worlds [2]. Users will play an active role, by contributing data and services and by making available their own sensing and actuating devices. This will make pervasive computing infrastructures as participatory and capable of value co-creation as the Web [3], 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 commercial release of a variety of early pervasive services trying to exploit the possibilities opened by these new scenarios: GPS navigator systems providing real-time traffic information and updating routes accordingly, cooperative smart phones that inform us about the current positions of our friends, and augmented reality services to enrich what we see around with dynamically retrieved digital information [4]. However, the road towards the effective and systematic exploitation of these emerging scenarios calls for a radical rethinking of current service models and frameworks.

II. CASE STUDY AND REQUIREMENTS

A simple case study – representative of a larger class of emerging pervasive scenarios – can help grounding our arguments and sketching the requirements of future pervasive services.

It is a matter of fact that we are increasingly surrounded by digital displays: from those of wearable devices to wide wall-mounted displays pervading urban and working environments. Currently, the latter are simply conceived as static information servers to show information in a manually-configured manner – e.g., cycling some pre-defined commercials or general interest news – independently of the context in which they operate and of the users nearby. However, such displays infrastructures can be made more effective and advantageous for both users and information/service providers by becoming general, open, and adaptable information service infrastructures.

First, information should be displayed based on the current state of the surrounding physical and social environment. For instance, by exploiting information from surrounding temperature sensors and from user profiles, an advertiser could have ice tea commercials – instead of liquor ones – being displayed in a warm day and in a location populated by teenagers. Also, actions could be coordinated among neighboring displays, e.g., 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. These examples express a general requirement for pervasive services:

- *Situatedness* — Pervasive services deal with spatially- and socially-situated activities of users, and should thus be able to interact with the surrounding physical and social world and adapt their behavior accordingly. The infrastructure itself, deeply embedded in the physical space, should effectively deal with spatial concepts and data.

Second, and complementary to the above, the display infrastructure and the services within should automatically adapt to their own modifications and contingencies in an automatic way without experiencing malfunctionings, and possibly taking advantage of such modifications. Namely, when new devices are deployed or when new information is injected, a spontaneous re-distribution and re-shaping of the overall displayed information should take place. For instance: the deployment of a big advertising 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. In terms of a general requirement for decentralized and dynamic scenarios:

- *Adaptivity* — Pervasive services and infrastructures should inherently exhibit properties of autonomous adaptation and management, to survive contingencies without human intervention and at limited costs.

Third, the display 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 may continuously upload personal content (e.g., pictures and annotations related to the local environment) from their own devices to the infrastructure, both for better visualization and for increasing the overall local information offer. Similarly, a group of friends can exploit a public display to upload software letting it host a shared real-time map to visualize what's happening around (which would also require opportunistic access to the existing environmental sensors and to any available user-provided sensors, to make the map alive and rich in real-time information). In general, one should enable users to act as "prosumers" – i.e., as both consumers and producers – of devices, data, and services. Not only this will make environments meet the specific needs of any specific user (and capture the long tail of the market), but will also induce a process of value co-creation increasing the overall intrinsic value of the system and of its services [5]. If the mentioned real-time 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 terms of a general requirement:

- *Prosumption and Diversity* — The infrastructure 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 by exploiting them to improve and integrate existing services whenever possible, and add further value to them.

Finally, beside short-term adaptation, in the longer-term any pervasive infrastructure will experience dramatic changes related to the technology being adopted, as well as in the kinds of services being deployed and in their patterns of usage. For instance, a display infrastructure will somewhen integrate more sophisticated sensing means than we have today and will possibly integrate enriched actuators via which to attract user attention and interact with them. This can be the case of personal projection systems to make any physical object become a display, or of eyeglass displays for immersive perception and action. While this can open up the way for brand new classes and generation of services to be conceived and deployed, it also requires that such evolution can be gradually accommodated without harming the existing infrastructure and services. As a general requirement:

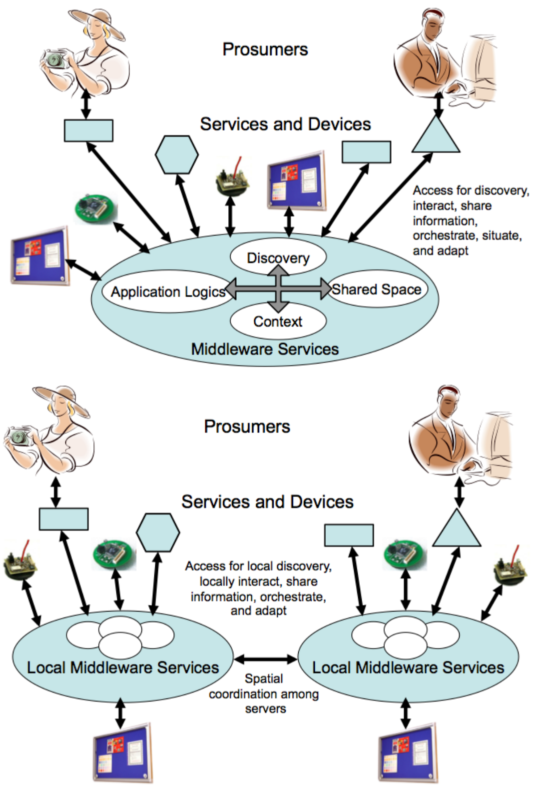

Fig. 1. Architecting pervasive service environments. Up: solution with a centralized middleware server. Bottom: solution relying on a distributed set of local middleware services.

- *Eternity* — 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 and expensive re-engineering efforts to incorporate innovations and changes.

III. FROM SERVICE-ORIENTED ARCHITECTURES TO NATURE-INSPIRED PERVASIVE SERVICE ECOSYSTEMS

Could the above requirements be met by architecting pervasive service environments around standard service-oriented architectures (SOA) [6]? Yes, to some extents, but the final result would be such a scramble of SOA to rather suggest re-thinking from scratch the architecture and its founding principles.

A. Centralized SOA Solution

In general, SOA consider the inter-related activities of service components to be managed by various infrastructural (middleware) services such as: discovery services to help components get to know each other; context services to help components situate their activities; orchestration services to coordinate interactions according to specific application logics; and shared dataspace services to support data-mediated interactions. To architect a pervasive display environment in such terms (Figure 1-up), one has to set up 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.

Such components access the discovery service to get aware of each other. However, in dynamic scenarios (users and devices coming and going), components are 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 (i.e., 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.

To adapt to situations and contingencies, components should be able to recognize relevant changes in their current environment and plan corrective actions in response to them, which again require notable communication and computational costs for all the components involved. Alternatively, or complementary, one could think at embedding adaptation logics into some specific server inside the middleware (e.g., in the form of autonomic control managers [7]). However, such 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.

B. *Decentralized SOA Solution*

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. For instance (Figure 1-bottom), one could install one middleware server for each of the available public displays. It will manage the local display and all local service components, thus simplifying local discovery and naturally enforcing spatial interactions. Adaptation to situations is made easier, thanks to the possibility of recognizing in a more confined way (and at reduced costs) local contingencies and events, and of acting locally upon them.

With the adoption of a distributed solution enforcing locality, the distinction between the logics and duties of the different infrastructural services fades: discovering local services and devices implies discovering something about the local context; the dynamics of the local scenarios, as reflecting in the local discovery tables, makes it possible to have components indirectly influence each other (being their actions possibly dependent of such tables), as in a sort of shared dataspace model. This also induces specific orchestration patterns for components, based on the local logics upon which the middleware relies to distribute information and events among components and to put components in touch with each other.

A problem of this distributed architecture is to require solutions both to tune it to the spatial characteristics of the scenario and to adaptively handle contingencies. The logics of allocation of middleware servers (i.e., one server per display) derives naturally only in a static scenario, but the arrival and dismissing of displays requires the middleware servers to react by re-shaping the spatial regions and the service components of pertinence of each of them, also correspondingly notifying components. To tackle this problem, the actual distribution of middleware servers should become transparent to service components—they should not worry about where servers are, but will simply act in their local space confident that there are servers to access. Moreover, the network of servers should be able to spontaneously re-organize its shape in autonomy, without directly affecting service components but simply adaptively inducing in them a re-organization of their interaction patterns.

C. *Nature-inspired Ecosystems*

Pushed towards a very dense and mobile network of nodes and pervasive devices, the architecture will end up being perceivable as a dense distributed environment above which a very dynamic set of spatially-situated components discover, interact, and orchestrate with each other. This is done in terms of a much simplified logics, embedded into the unique infrastructural service, subsuming the roles of discovery, context, dataspace, and orchestration services, and taking the form of a limited set of local rules embedded in the spatial substrate itself. That is, we would end up with something that notably resembles the architecture of natural ecosystems: a set of spatially situated entities interacting according to well-defined set of natural laws enforced by the spatial environment in which they situate, and adaptively self-organizing their interaction dynamics according to its the shape and structure.

Going further than architectural similarity, the natural metaphor can be adopted as the ground upon which to rely to inherently accommodate the requirements of pervasive service scenarios. Situatedness and spatiality are there by construction. Adaptivity can be achieved because of the basic rules of the game: the dynamics of the ecosystem, as determined by the enactment of laws and by the shape of the environment, can spontaneously induce forms of adaptive self-organization beside the characteristics of the individual components. Accommodating new and diverse component species, even towards a long-term evolution, is obtained by making components part to the game in respect of its rules, and by letting the ecosystem dynamics evolve and re-shape in response to the appearance of such new species. This way, we can take advantage of the new interactional possibility of such new services and of the additional value they bring in, without requiring the

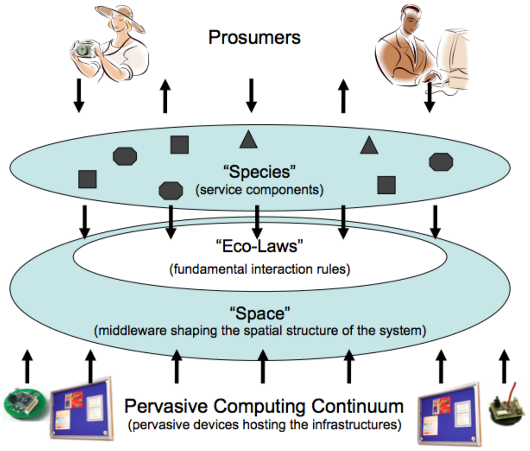

Fig. 2. A Conceptual Architecture for Pervasive Service Ecosystems

individual components or the infrastructure itself (i.e., its laws and structure) to be re-engineered [8].

Indeed, nature-inspired solutions have already been extensively exploited in distributed computing [9] for the implementation of specific adaptive algorithmic solutions or of specific adaptive services. Also, many initiatives – like those named upon digital/business service ecosystems [10] – recognize that the complexity of modern service systems is comparable to that of natural ones and requires innovative solutions also to effectively support diversity and value co-creation. Yet, the idea that natural metaphors can become the foundation on which to fully re-think the architecture of service systems is far from being metabolized.

IV. A REFERENCE CONCEPTUAL ARCHITECTURE

The above discussion leads to the identification of a reference conceptual architecture for nature-inspired pervasive service ecosystems (see Figure 2).

The lowest level is the concrete physical and digital ground on which the ecosystem will be deployed, i.e., a dense infrastructure (ideally a continuum) of networked computing devices and information sources. At the top level, prosumers access the open service framework for using/consuming data or services, as well as for producing and deploying in the framework new services and new data components or for making new devices available. In our case study, they include the users passing by, the display owners, and the advertising companies interested in buying commercial slots. At both levels openness and its dynamics arise: 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 our case study, we consider integration at any time of new displays and new sensors, and the presence of a continuous flow of new visualization services (e.g. commercial advertisers) and users, possibly having their own devices integrated in the overall infrastructure.

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

- **Species** — This level includes a variety of components, belonging to different "species" yet modeled and computationally rendered in a uniform way, representing the *individuals* populating the ecosystem: physical and virtual devices of the pervasive infrastructure, digital and network resources of any kind, persistent and temporary knowledge/data, contextual information, software service components, or personal user agents. In our case study, we will have different software species to represent displays and their displaying service, the various kinds of sensors distributed around the environment and the data they express, software agents to act on behalf of users, display owners, and advertisers. In general terms, an ecosystem is expected to be populated with a set of individuals physically deployed in the environment, situated in some portion of the ecosystem space, and dynamically joining/leaving it.
- **Space** — This level provides and 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 (as it is the case of information and advertising services in our case study), this level situates individuals in a specific portion of the 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 deployment context, supporting the execution and life cycle of individuals and their spatial interactions. From the viewpoint of such individuals, the middleware will have to provide them (via some API) with the possibility of advertising themselves, accessing information about their local spatial context (there included the other individuals around), and detecting local events. From the viewpoint of the underlying infrastructure, 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, this can be realized by a network of active data-oriented and event-oriented localized services (e.g., tuple spaces [11]), spread on the nodes of the pervasive substrate, and accessible on a location-dependent basis by individuals and devices. Indeed, recent proposals in the area of tuple-based coordination services for pervasive and mobile devices, such as TOTA [11] or LIME [12], can effectively candidate as the basic engine for reifying the space level, provided they are extended to automatically re-shape the spatial domain of competence of each node in response to contingencies (e.g., along the lines promoted in P2P computing by content-addressable networks [13]). In the case study, for instance, one could think at assigning one tuple space for each display, and have the various displays dynamically self-configure their spatial domain of competence accordingly to geographi-

cal and "line of sight" factors.

- **Eco-Laws** — The way in which individuals (whether services components, devices, or generic resources) live and interact is determined by the set of fundamental "eco-laws" regulating the ecosystem model. Enactment of eco-laws on individuals will typically affect and be affected by the local space around and by the other individuals around. In our case study, 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 advertisements.

  Although the set of eco-laws is expected to be always the same for a specific implementation, they will possibly have different effects on different species. Accordingly, their enactment may require the presence of some meaningful description (within the uniform modeling of individuals) of the information/service/structure/goals of each species and of their current context and state. These descriptions, together with proper "matching" criteria, define how the eco-laws apply to specific species in specific conditions of the space. We emphasize that, in the proposed architecture, the concept of "semantic description" of traditional SOA to facilitate discovery turns into a concept of "alive semantic description" (dynamically changing as the context and state of components change), to properly rule the dynamic enactment of eco-laws.

  Practically, to be able to code eco-laws and enact them, the middleware substrate should proactively mediate inter-component interactions, and act as an active space in which to store their continuously updating semantic descriptions, so as to adaptively support the matching process triggering eco-laws in dependence of the current conditions of the overall ecosystem. Technologically, since most tuple-based middleware systems are currently enriched with the capability of reacting to events and of configuring the matching process, they could well act as the interaction media in which to embed and enact eco-laws.

The proposed architecture represents a radically new perspective on modeling service systems and their infrastructures. The typically heavyweight and multifaceted layers of SOA are subsumed by an unlayered universe of components, all of which underlying the same model, living and interacting in the same spatial substrate, and obeying the same eco-laws—being the latter the only concept hardwired into the system.

This rethinking is very important to ensure adaptivity, diversity, and long-term evolution: no component, service or device is there to stay, everything can change and evolve, self-adapting over space and time, without undermining the overall structure and assumptions of the ecosystem. That is, by conceiving the middleware in terms of a simple spatial substrate in charge of enforcing only basic interaction rules, we have moved away from the infrastructure itself the need of adapting, and fully translated this as a property of the application level and of its dynamics.

The 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. Typical patterns that can be driven by such laws may include forms of adaptive self-organization (e.g., spontaneous service aggregation or service orchestration, where the eco-laws plays an active role in facilitating individuals to spontaneously interact and orchestrate with each other, also in dependence of current conditions), adaptive evolution (changing conditions reflect in changes in the way individuals in a locality are affected by the eco-laws) and of decentralized control (to affect the ecosystem behavior by injecting new components in it).

V. METAPHORS

Beside the above architectural guidelines, what actual shape can species, space, and eco-laws take in an actual implementation? Identifying and validating specific solutions in this direction will be a key challenge of pervasive computing in the next years. Yet, we argue that whatever solution will most likely get inspiration from one of the key natural metaphors already explored in the literature, or possibly extract specific desirable aspects from many of them towards a new synthesis. Key metaphors include physical [11], chemical [14] and biological ones [9], along with metaphors focussing on higher-level social models (e.g., trophic networks [15])—the key difference between them being in the way the species, the space, and the eco-laws are modeled and implemented (see Figure 3).

All the metaphors, by adhering to the proposed architecture, are by construction spatially situated, adaptive by self-organization, and open to host diverse and evolving species. However, when it comes to modeling and implementing, different metaphors may tolerate with variable efficiency and complexity the enforcement of adaptive self-organization patterns and the support of diversity and evolution. In addition, since the service ecosystem is here to ultimately serve us, it is necessary to analyze how and to which extent the metaphors facilitate exerting forms of decentralized control over the ecosystem behavior, in order to direct its self-organizing activities and behavior 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, the complexity of individuals and their environment, while keeping the infrastructure lightweight and the overall execution efficient.

A. Physical metaphors

These consider species as sort of computational particles, living in a world of other particles and virtual computational force fields, the latter acting as the basic interaction means. Activities of particles (to be practically modeled and implemented as reactive agents) are driven by laws that determine how particles spread fields, how fields propagate and reshape upon changing conditions, and how they influence particles

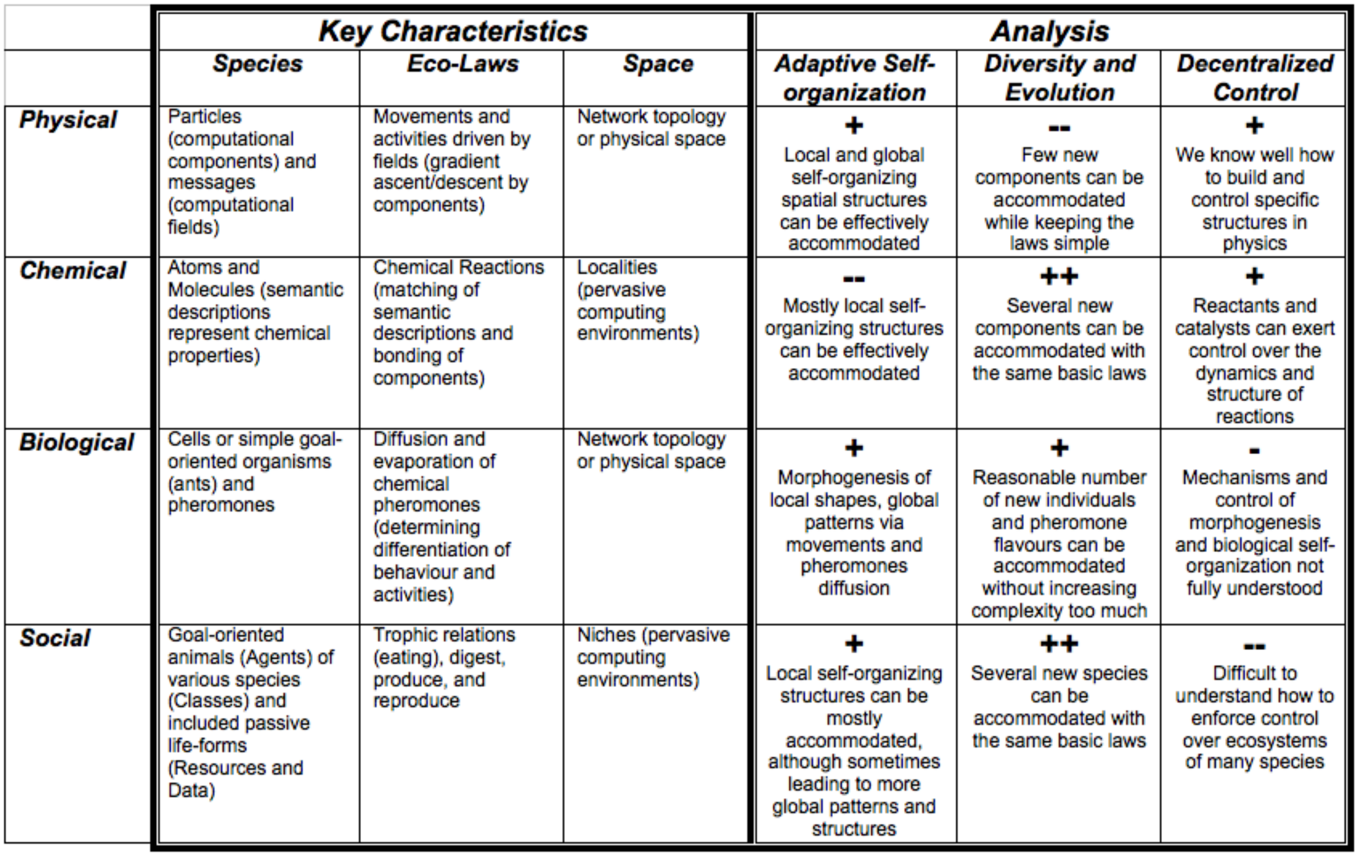

	Key Characteristics			Analysis		
	Species	**Eco-Laws**	**Space**	**Adaptive Self-organization**	**Diversity and Evolution**	**Decentralized Control**
Physical	Particles (computational components) and messages (computational fields)	Movements and activities driven by fields (gradient ascent/descent by components)	Network topology or physical space	+ Local and global self-organizing spatial structures can be effectively accommodated	-- Few new components 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 represent 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 can be accommodated with the same basic laws	+ Reactants and catalysts can exert control over the dynamics and structure of reactions
Biological	Cells or simple goal-oriented organisms (ants) and pheromones	Diffusion and evaporation of chemical pheromones (determining differentiation of behaviour and activities)	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

Fig. 3. Metaphors for Pervasive Service Ecosystems

(those whose semantic description "matches" some criterion). Particles change their status based on the perceived fields, and move or exchange data by navigating them (i.e., particles spread sort of data particles to be routed according to the shape of fields). The world in which such particles live and fields spread and diffuse can be either a simple (euclidean) metric world mapped in the physical space, or a virtual/social space mapped on the technological network. From the infrastructural viewpoint, a network of local tuple spaces will have to proactively support the storing of local field values, the propagation and continuous update of fields across the network, and the notifications about these changes to individuals. For instance, the TOTA middleware [11] can be adopted for the implementation and management of physically-inspired distributed field data structures.

In the case study, we can imagine display services as masses emitting gravitational-like fields (in the form of broadcast events or spanning trees over the network). Such fields have different "flavors" (i.e., different semantic descriptions, reflecting the characteristics of users around and the environment conditions) and an intensity proportional to either their dimension or the available display slots. Information and advertiser agents can behave as masses attracted by fields with specific flavor, 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.

Physical metaphors have been extensively studied for their spatial self-organization features, and for their effectiveness in facilitating the achievement of coherent behaviors even in large scale systems—for load balancing, data distribution, clustering, aggregation, and differentiation of behaviors. The conceptual tools available for controlling the spatial behavior and the dynamics of such systems are well-developed, most of them related to acting on how fields propagate and dynamically change—by which it is actually possible to exert control over the overall system behavior. On the other hand, such metaphors hardly tolerate high diversity and evolution. In fact, to support very diverse species and behaviors (at a time and over time), eco-laws must become complex enough to tolerate a wide range of different fields and propagation rules, with an increase in the complexity of the model and in burden on the infrastructure [11], that has to proactively support the propagation and continuous update of many field structures.

B. Chemical metaphors

These consider species as sorts of computational atoms/molecules (again modeled and implemented as reactive agents), enriched with semantic descriptions acting as the computational counterpart of the bonding properties of physical atoms/molecules (yet made dynamic to reflect the current state and context of individuals). Accordingly, the laws that drive the overall ecosystem behavior resemble chemical reactions, that dictate how chemical bonding between components take place (relying on some forms of pattern matching between semantic descriptions), and lead to aggregated, composite, and new components, and also to growth/decay of species. The world where individuals live is typically formed by a set of spatially confining localities, intended as the "solutions" in which chemical interactions occur and across which chemicals can eventually diffuse. The declarative tuple space model of TuCSoN coordination infrastructure can support an effective implementation for chemically-inspired interactions [14].

In the case study, we can think of display services, of information services (concerning ads and news), as well as of user and environmental data as molecules. Displays represent different localities (i.e., tuple spaces) 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 which is in charge of actually displaying that service in that display. Concurrently, in each locality, catalytic components can be in charge of re-enforcing the concentration of specific information or of specific information/advertisements, reflecting the current situation of users. Also, localities can be open to enforce chemical bonds across displays, so that high-activity of advertisement reactions on a display can eventually propagate to neighbors.

Chemical metaphors can effectively lead to self-organizing structures like local composite services and local aggregates. As in real chemistry, chemical computational metaphors can accommodate an incredible amount of different components and composites with a single set of basic laws. In practice, this means that it can tolerate an increasingly diverse and evolving set of semantic descriptions for components without affecting basic eco-laws and without increasing the burden to the infrastructure. As far as control is concerned, one can think at using sort of catalyst or reagent components to engineer and control (in a fully decentralized way) the dynamics and the behavior 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 [14], making it hard to naturally and easily enforce distributed self-organized behaviors, like creating a complex and distributed aggregation of components.

C. Biological metaphors

These focus on biological systems at the scale of individual organisms, or of colonies of organisms like ants. The species are therefore either simple cells or animals acting on the basis of simple goals like finding food and reproducing (to be modeled and implemented in terms of simple goal-oriented agents). As in physical systems, interactions take place by means of signals of various flavors (i.e., chemical pheromones) spread by individuals in the environment, and slowly diffusing and evaporating. The spatial environment is again a computational landscape either mapped on the network topology or on the physical space. Unlike physical systems, individuals here are not necessarily passively subject to the sensed pheromones, but they can 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. From the implementation viewpoint, any infrastructure that can support physical metaphors can also be adapted to support biological metaphors, by turning fields into persistent and slowly diffusing/evaporating data structures.

In the case study, users will 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 some different pheromones to express the availability of commercials 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. Displays, advertising and information agents by their side, and depending on what they have displayed so far, can also emit additional flavors of pheromones, to store memory of past events. The persistence of pheromones can also be exploited by additional components that, by moving from display to display, create pheromones trails as a basis for more global strategies—like identifying routing paths that advertiser agents use to globally find the best displays to exploit.

Biological metaphors appear very flexible in efficiently enabling the spatial formation of both localized and distributed activity patterns, and have a variety of applications [9]. The problems of accommodating diversity and evolution that affect physical metaphors are here notably smoothed. 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.

D. Social metaphors

These focus on biological systems at the level of animal species and of their interactions [15]. Individuals are sorts of goal-oriented animals (i.e., agents) belonging to a specific species, that are in search of "food" resources to survive and prosper, and that can represent in their turn food to others (both aspects reflecting in some proper semantic descriptions). Pure data items and resources can be abstracted as sorts of passive life-forms (i.e., vegetables). The eco-laws determine how the resulting "web of food" should be realized, namely, how animals search food, eat, and possibly produce and reproduce, thus influencing and ruling the overall dynamics of the ecosystem and the interaction among individuals of different species. Similarly to chemical systems, the shape of the world is typically organized around a set of localities, i.e., ecological niches, yet enabling interactions and diffusion of species across niches. From the implementation viewpoint, reactive tuple space models [12] can be effectively adopted towards the realization of a supporting infrastructure, where the possibility for the tuple space to enforce control over all interactions can be used as a mean to rule the food-web-based eco-laws.

In the case study, and assuming that each display is associated to an ecological niche, we can imagine users as species 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 proper users to survive (users with matching interests). For both cases, the primary effect of an eating action is the feeding of displays with information or commercials to show. A possible secondary effect of eating is the reproduction and diffusion in the environment of the best-fit species, e.g., of most successful commercials. Those species that do not succeed in eating at a niche 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).

Social metaphors, such as the one here described but without forgetting the large body of work in the area of social multi-agent systems, appear suitable for local forms of self-organization like chemical metaphors—think at self-organized equilibria of web food patterns in ecological niches. Unlike chemical metaphors, and more similarly to biological one, social metaphors can be effectively exploited also to enforce distributed forms of self-organization, by exploiting the capability of individuals of moving around (e.g., to find food and/or 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 social computational ecosystems can be as difficult as it is in real social systems.

VI. CONCLUSIONS

Nature-inspired models and architectures appear very promising for re-thinking the foundation of modern and emerging pervasive service systems. Yet, the general lesson that we can extract from our analysis is that:

- None of the natural metaphors being proposed so far is suited, as is, to act as a way to fully realize the eternally adaptive service ecosystem vision.
- Thus, along the lines of the proposed reference architecture and being guided by it, some new synthesis should be identified to incorporate the key features of the existing metaphors into a unifying general-purpose one.

In addition to the key problem of identifying suitable metaphors, the widespread deployment of nature-inspired pervasive service ecosystems requires methodologies and tools for their engineering, proper security mechanisms and policies, and means to integrate the approach with the legacy of current SOA systems. Lastly, the road towards pervasive service ecosystems will have to ground on the emerging science of service systems [3], and will possibly be able to contribute to it.

REFERENCES

[1] A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, R. A. Peterson, H. Lu, X. Zheng, M. Musolesi, K. Fodor, and G.-S. Ahn, "The rise of people-centric sensing," *IEEE Internet Computing*, vol. 12, no. 4, 2008.

[2] B. Coleman, "Using sensor inputs to affect virtual and real environments," *IEEE Pervasive Computing*, vol. 8, no. 3, pp. 16–23, 2009.

[3] J. C. Spohrer, P. P. Maglio, J. H. Bailey, and D. Gruhl, "Steps toward a science of service systems," *IEEE Computer*, vol. 40, no. 1, pp. 71–77, 2007.

[4] A. Ferscha and S. Vogl, "Wearable displays – for everyone!" *IEEE Pervasive Computing*, vol. 9, no. 1, pp. 7–10, 2010.

[5] S. L. Vargo, P. P. Maglio, and M. A. Akaka, "On value and value co-creation: a service systems and service logic perspective," *European Management Journal*, vol. 26, no. 3, pp. 145–152, 2008.

[6] M. N. Huhns and M. P. Singh, "Service-oriented computing: Key concepts and principles," *IEEE Internet Computing*, vol. 9, no. 1, pp. 75–81, 2005.

[7] J. O. Kephart and D. M. Chess, "The vision of autonomic computing," *IEEE Computer*, vol. 36, no. 1, pp. 41–50, 2003.

[8] M. Jazayeri, "Species evolve, individuals age," in *8th IEEE International Workshop on Principles of Software Evolution*, Washington, DC, 2005, pp. 3–12.

[9] 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.*, vol. 1, no. 1, pp. 26–66, 2006.

[10] M. Ulieru and S. Grobbelaar, "Engineering industrial ecosystems in a networked world," in *5th IEEE International Conference on Industrial Informatics*. IEEE Press, June 2007, pp. 1–7.

[11] M. Mamei and F. Zambonelli, "Programming pervasive and mobile computing applications: the tota approach," *ACM Trans. Software Engineering and Methodology*, vol. 18, no. 4, 2009.

[12] A. L. Murphy, G. P. Picco, and G.-C. Roman, "Lime: A coordination model and middleware supporting mobility of hosts and agents," *ACM Trans. on Software Engineering and Methodology*, vol. 15, no. 3, pp. 279–328, 2006.

[13] S. Androutsellis-Theotokis and D. Spinellis, "A survey of peer-to-peer content distribution technologies." *ACM Computing Surveys*, vol. 36, no. 4, pp. 335–371, 2004.

[14] M. Viroli and M. Casadei, "Biochemical tuple spaces for self-organizing coordination," in *Coordination Languages and Models*, ser. LNCS, vol. 5521. Springer-Verlag, June 2009, pp. 143–162.

[15] G. Agha, "Computing in pervasive cyberspace," *Commun. ACM*, vol. 51, no. 1, pp. 68–70, 2008.