

Simulation Experiences with an Ecological Approach for Pervasive Service Systems

Cynthia Villalba and Franco Zambonelli
DISMI – University of Modena and Reggio Emilia
Reggio Emilia, Italy
name.surname@unimore.it

ABSTRACT

Innovative frameworks have to be identified for the deployment and execution of pervasive service systems. These systems, composed by a massive number of components, should be able to exhibit properties of *self-organization* and *self-adaptability*, and of *long-lasting evolvability*. In this context, this paper discusses how such frameworks should get inspiration from natural systems. In particular, we focus on *ecological systems* that model and deploy services as autonomous individuals (i.e., *agents*), 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 “*laws of nature*”. The key characteristics of the proposed ecological approach are detailed with the help of a representative case study. Also, an extensive set of simulation experiments are reported to outline the peculiarities of its dynamic behavior and to show the potential effectiveness of the approach.

Keywords

Pervasive Ecosystems; Multi-Agent Systems; Self-* Systems; Eco-laws

1. INTRODUCTION

Pervasive and mobile computing devices increasingly populate our environments [9], and will eventually define a global-scale and very dense, decentralized computing and communication infrastructure for the provisioning of general-purpose pervasive services.

However, the effective development and execution of pervasive services calls for a deep rethinking of current service models and of service frameworks, in order to: (i) Naturally match the spatially-situated nature of the environment and of the services within, and rely on mostly localized spatial interaction to provide support for massive scalability [4, 26].

(ii) Inherently exhibit properties of self-organization, self-adaptation and self-management that are required in highly-decentralized and highly-dynamic scenarios [8]. (iii) Flexibly tolerate evolutions of structure and usage over time to incorporate innovations and changes [27].

To reach these goals, we should no longer conceive services and their interactions as in usual service-oriented architectures [5, 15], where services are simply functional entities statically orchestrated with the help of specific middleware services. Nor one can rely on one-of architectural solutions to achieve specific self-* features in existing systems, resulting in an increase of complexity [16]. Rather, the most promising direction is that of taking inspiration from natural systems [3, 18], where spatial concepts, self-organization, self-management, and long-lasting evolvability are inherently there because of the basic “rules of the game”.

Nature-inspired solutions have already been exploited in the area of distributed computing for the implementation of specific middleware solutions [3, 18]. Similarly, natural and ecological metaphors have been adopted to characterize the complexity of modern ICT and service systems [24]. Here we go further and argue that natural metaphors can act as reference architectural model to conceive, model, and develop a fully-fledged pervasive service framework and all the components within.

Although one can think at different classes of natural systems and from different perspectives (e.g., physical [19], chemical [25], biological [3], or social/ecological [1]) one can always recognize the following characteristics: above a spatial environmental substrate, autonomous individuals (i.e., agents) of different kinds 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 service agents, 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 paper provides the following contributions:

- We detail the ecological approach that we have started investigating. The approach, earlier presented in [?], abstracts the components of the ecosystem as sorts of goal-oriented organisms that, driven by laws of survival, interact with each other and self-organize their activities according to dynamic food-web relations (Sec-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

tion 2).

- We introduce a representative case study in the area of displays ecosystems [11] to clarify the concepts expressed (Section 3).
- We present a number of simulation experiments that we have performed to assess the effectiveness of our approach (Section 4). The experiments, simulating a spatial scenario of displays ecosystem, shows that our approach is capable of effectively exhibiting interesting self-organization features in a variety of different situations.
- We briefly discuss other nature-inspired approaches (Section 5).

Eventually, in Section 6, we conclude and outline open research directions.

2. THE ECOLOGICAL APPROACH

In this section, we go into more details about the key architecture of the proposed ecological metaphor and of its components.

2.1 The overall architecture

The overall architecture of our ecological approach is as follow (see Figure 1).

At the lowest level is the physical ground on which the ecosystem will be deployed, i.e., a very dense and widely populated infrastructure (ideally, a world-wide pervasive continuum) of networked computing devices (e.g., PDAs, tags) and information sources (Web 2.0 fragments). 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/deploying in the framework new services and new data components. Between these levels, we find the components of our ecological architecture.

The “Species” level is the one in which physical and virtual devices of the pervasive system, persistent and temporary knowledge/data, contextual information and of course software service components, are all abstracted as “living individuals” (or agents) of the system. Although such individuals are expected to be modeled in a uniform way, they will have specific characteristics very different from each other, i.e., they will be of different “species”. The dense population of devices and actors involved at the highest and lowest levels, together with their dynamics, reflect in the presence of a very massive and dynamically varying number of individuals and species.

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. From a practical viewpoint, the spatial structure of the ecosystem will be implemented by means of some minimal middleware substrate 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 or aggregate with each others, and

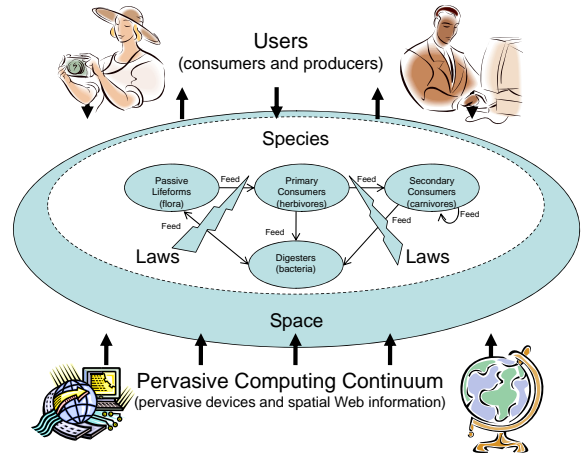


Figure 1: A reference architecture for pervasive service ecosystems.

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 (which is also what makes real ecosystems scalable): 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.

2.2 Key components

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). That is, individuals have the ego-centric goal of surviving by finding the appropriate food and resources.

The laws of the ecosystem thus determine how the resulting “food web” should be realized and ruled, 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.

As far real food-webs are concerned, we can consider the following key classes (Figure 1). 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 are in the space. Primary consumers (i.e., herbivores) need to eat vegetables to survive and prosper. Secondary consumers (i.e., carnivores) typically need to eat other animals to survive, though this does not exclude that can also act as primary consumers (eating vegetables too). The result of the metabolization of food by both primary and secondary consumers ends up in feeding lower-level “digesters” life forms (e.g., bacteria), densely spread in space, and that in their turn produce and diffuse resources and nutrients for the flora.

Translating the above concepts in computational terms, we can consider the individuals of the ecosystem as belonging to specific subclasses derived from four key abstract classes. Passive life forms represent the data sources and

the computational/memory/communication resources of the ecosystem, which are not to be considered proactive and autonomous computational entities (i.e., they are not expected to act). Primary consumers represent those services that require to digest information to be of any use (i.e., to reach their goals), and that are computationally autonomous in pursuing such goals. Secondary consumers, instead, are those services that, to reach their goals, need the support of other computational services/agents, other than possibly of information sources. Digesters components can be generally assimilated to all those background computational services that are devoted to monitor the overall activities of the system, and either produce new information about or influence the existing information and resources.

From Figure 1, it is clear that the existence of such a food-web chain (or, in computational terms, of such interdependencies in the activities of agents) provides for producing a close food-web loop. The existence of such loops is indeed a known basic regulatory ingredient to enable adaptation and self-organization in both natural and computational systems [18]. However, unlike in more structured approaches which consider the existence of a single (or of a limited set) of regulatory loops [16], the ecological approach can make a number of related loops co-exist, depending on the number of different classes of individuals and on their specific trophic relations.

As for the spatiality of the approach, the ecosystem space can be 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. Each locality will determine how the different species organize to live, will describe how individuals of each species respond to the distribution of resources and other species, and how they alter these factors. Clearly, such way of organizing spatiality around open ecological niches provides for the existence of spatially distributed regulatory loops that ensure adaptation and self-organization at the global level, rather than at the level of individual niches.

The ecological metaphor promises 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). In fact, while the basic infrastructure and its laws do not require re-engineering, the actual service and data species can dynamically change, naturally inducing new dynamics for the interactions between individuals and for the ecosystem as a whole.

2.3 Modeling of individuals

Let us now detail a bit further the modeling of the computational elements that populate the ecosystem (i.e., the ecosystem individuals) and of the mechanisms underlying their interactions (i.e., forming the basis of the eco-laws).

In our perspective, the key components of an individual include:

- A public semantic description of its own characteristics;
- A set of “needs” expressing which other individuals it needs to interact (i.e., “eat”) with;

- A “happiness” status acting as a main drive for the individual activities (i.e., as a sort of utility function for the individuals);
- An internal logic for the maximization of its happiness, based on its needs and its happiness status, and relying on a set of allowed actions.

Active individuals typically express all the above components, passive ones (i.e., pure data items and resources) typically express only a public semantic description.

We recall that, in our model, each active individual needs food and resources to survive. That is, it needs other components to “eat” (i.e., to interact with), whether passive or active ones. These needs are specific for each specific individual/species. Generalizing upon dynamic models of semantic service matching [2], the idea of our approach is that each individual expresses its own characteristics in a public semantic description. At the same time, each individual needs to eat some others, expresses its own needs in a sort of semantic “templates”, describing the characteristics of the needed resources or individuals.

Then, based on the needs of an individual, and on the semantic descriptions of the individuals around it, a discovery process based on semantic matching takes place. Such process of matching typically occurs within a niche, though without excluding the possibility for one description to diffuse in spatially close niches and enable cross-niches matching. The existence of a match of one individual with another needing it establishes a session of interaction between the two components, which we metaphorically abstract as an “eating” action by the individual that has found the right food. Pragmatically, we can think at a service component that has found the data/resources it needed or at a composite service that has found the needed computational partners required to fulfill its promises.

More formally, given the semantic description ind_a of an individual a and the needs ind_bneeds of another individual b , the process assumes the existence of a matching function

$$match(ind_a, ind_bneeds)$$

that returns some value expressing the degree of match between individual. The eco-laws of the ecosystem have the duty of automatically triggering possible matches in a niche (or across close niches) and of establishing a connection between matching individuals.

We do not go into details here on the specifics of such semantic representations and of the matching function. What we want to emphasize here is that both the semantic descriptions of individuals and their expressed needs are dynamic entities that can change over time depending on the current state and context of individuals. For instance, an individual that has already fully satisfied some of its needs, will have to change its expressed needs to avoid keep on searching for something it does no longer need.

Upon occurrence of a match, the involved individuals can, according to their own specific behaviors, start interacting with each other (e.g. a primary consumer having found matching information, can decide to absorb or consume it). The specific capability (i.e., the set of actions allowed) within individuals will determine the course of such interactions. In general, though, the internal actions of an individual are mostly driven by the need to maximize its happiness.

The happiness value of an individual expresses the satisfaction degree of each individual, in other words it expresses how well an individual satisfies himself. The happiness value of an individual is generally a function of time, and it is affected by the previous happiness value and by the current state of the individual.

The happiness increases when the individual finds the proper resources to eat (i.e., a lot of matching individuals), because this implies it is effectively achieving the goals it has been conceived for. However, the happiness of an individual can also be affected by different situations in its niche (or in close niches), and specifically by the happiness of the other individuals in the niche.

Thus, we can distinguish between the instantaneous happiness of an individual, which is merely a function of the current matchings, i.e., for an individual i in a niche with n individuals:

$$H_i^{inst}(t) = f(\text{match}(\text{ind}_j, \text{ind}_i \text{need}), j = 1, \dots, n). \quad (1)$$

and the overall happiness of an individual, which is a function of the current happiness of the individual, of the previous happiness, and of the overall happiness of the other individuals in the niche, i.e., for an individual i :

$$H_i^{overall}(t) = f(H_i^{inst}(t), H_i^{overall}(t-1), H_j^{overall}(t-1)), j = 1, \dots, n). \quad (2)$$

The set of allowed actions (usually aimed at egoistically increase the happiness of the individual) are the final elements that define an individual. Beside specific internal computational actions and beside interactions with other individuals, additional action capabilities of an individual can include reproducing itself, moving and/or replicating across niches, or even dying. For example, an individual that does not find suitable matches in the current niche can decide to move to another niche to search for matches that can increase its happiness.

3. A CASE STUDY

Let us introduce a representative case study in the area of adaptive display systems [11, 12].

Consider a scenario with different kinds of pervasive devices spread on it, like a thematic park densely pervaded with digital wall-mounted screens where to display information for visitors, 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. In the niches we find visitors looking for information, advertisers that want to display commercials on the screens, and the displays themselves, that have the goal of maximizing their own exploitation.

Visitors will watch these screens to look for different information. 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 with the goal to find the information required by users and maximizing users’ happiness.

Concurrently, we can think at “advertising agents” that, acting on behalf of some advertising company, roam from




Figure 2: Food web for the ecosystem of displays case study.

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 where they could be more effective. Advertising agents would then act as secondary consumers.

Background monitoring agents, executing on each ecological niche, can contribute replicating and spreading information where it appears to be more appreciated, and can also contribute in supporting the spatial roaming of advertiser agents by directing them where they could find more satisfaction. Thus, they would act as digesters associated to each display. Specifically, here we can conceive the existence of one “display agent” associated to each display.

In addition, we can think at several “information agents” in charge of controlling the spreading and diffusion of information across displays. We emphasize that, although we conceptually prefer to think at information as something passive that is externally managed by sorts of digester agents, for the sake of modeling it is also possible to think at information as something “active”, that has its own “information agents” associated and with the goal of properly spreading information around.

The resulting patterns of interactions among agents can be defined according to the food web that we have already introduced in Figure 1, instantiated for the introduced case study as from Figure 2. Information (or, which is the same, information agents) feed user agents, and these feed the advertiser agents. Display agents close the feedback loop in the system (between information agents and user agents; and user agents and advertiser agents) by affecting how the other agents move (and there included how information is spread by information agents). This is necessary for the ecosystem to self-regulate (as a whole). It is expected that the system continuously evolve in time as a result of the interactions and actions occurring along such food web.

As for the specific actions of agents, we can think at user agents moving from place to place and thus being in the view range of different displays (and thus looking specific information on them and being exposed to different advertisement), at advertising agents moving and replicating over the spatially distributed displays so as to capture the more suitable target audience, at display agents that provide to update the display, and at information agents that distribute and replicate information on need.

Let us now detail the happiness functions for each class of the case study.

Users are happy when they see what they are interested in. The displays near to an user can satisfy (or not) his interests and consequently influence his happiness in a positive or a negative way. In other words, if the users watch what they

want/need in the displays around, they will feel happy. The function that gives us the instantaneous happiness of an user agent is defined by the following:

$$H_{user}^{inst}(t) = \frac{1}{n} \sum_{k=0}^n match(dT_k(t), user). \quad (3)$$

Where $H_{user}^{inst}(t_0) = 0$.

n is defined as the number of displays that are around the user in a specific time t . The *match* function expresses the degree of satisfaction of the user when he watches the display k in a given time t . A match takes place when what it is displayed (dT - displayed task) it is in agreement with the interests of the user.

The overall happiness for one user agent is:

$$H_{user}^{overall}(t) = H_{user}^{inst}(t) + \sum_{k=0}^{t-1} H_{user}^{overall}(k) - \Delta. \quad (4)$$

Where $H_{user}^{overall}(t_0) = 0$.

The function is the result of the sum of the whole values that the overall happiness had ($\sum_{k=0}^{t-1} H_{user}^{overall}(k)$), and also the instantaneous/current happiness ($H_{user}^{inst}(t)$). Δ is a constant value that stand for the unhappiness degree of the user, it makes his happiness value decrease over time. This models the fact that the longer a user does not see something interesting, the more he gets unhappy. Eventually, if the user never sees something interesting, it happiness will return to 0.

Advertiser agents are happy when they display their clips in an environment with many of happy users interested in their advertisements. The idea here is that a commercial clip is well received by users that are interested in it and that are "happy", i.e., that have not being bored by uninteresting clips so far. The function that evaluates the instantaneous happiness of advertiser agents is thus given by:

$$H_{adv}^{inst}(t) = \frac{1}{n} \sum_{k=0}^n match(dT_k(t), adv) \left(\frac{1}{m} \sum_{k=0}^m H_{user}^{overall}(t-1) \right). \quad (5)$$

Where $H_{adv}^{inst}(t_0) = 0$.

In this function n and m are defined as the number of displays and users respectively, which are in the near environment of the advertiser agent in a specific moment t and then, influencing in his happiness value. This happiness is influenced by what the displays around show and by the average of the overall happiness of the users around ($\frac{1}{m} \sum_{k=0}^m H_{user}^{overall}(t-1)$). The happiness of users influences in the happiness of advertisers by the definition of the final goal of the last: to show their advertisements to happy and interested users. From this goal we can deduce that: what increase user happiness also increase advertiser happiness. By other hand, the match function returns the degree of happiness when the advertiser agent watches the display k . A match takes place when what it is displayed (dT - displayed task) it is in agreement with the interests of the users around and of the advertiser agent, which can be: his own advertisement or an interested information for users. In a practical example: if the display k shows information

(but this information keeps users happy), then the match function returns a positive value.

The overall happiness for one advertiser agent is given by:

$$H_{adv}^{overall}(t) = H_{adv}^{inst}(t) + \sum_{k=0}^{t-1} H_{adv}^{overall}(k) - \Delta. \quad (6)$$

Where $H_{adv}^{overall}(t_0) = 0$.

Similar to the overall happiness of users, this function includes the values of the overall happiness already happened ($\sum_{k=0}^{t-1} H_{adv}^{overall}(k)$) and the current/instantaneous one ($H_{adv}^{inst}(t)$). Δ is the constant value that decreases the overall happiness of advertisers over time, expressing the degree of unhappiness.

The happiness of information agents is very simple. Information agent just want that the displays around them be interested in their information and thus, show them. The instantaneous happiness for an information agent is similar to the others instantaneous happiness described before and is given by:

$$H_{inf}^{inst}(t) = \frac{1}{n} \sum_{k=0}^n match(dT_k(t), inf). \quad (7)$$

Where $H_{inf}^{inst}(t_0) = 0$.

n is defined as the number of displays around the information agent, i.e., the number of displays that the information agent can see in a given instant t . The match function returns the degree of happiness of information agents when they see a display showing the task $dT_k(t)$ in a given time t .

The overall happiness for one information agent is obtained as the overall happiness described before, and is given by:

$$H_{inf}^{overall}(t) = H_{inf}^{inst}(t) + \sum_{k=0}^{t-1} H_{inf}^{overall}(k) - \Delta. \quad (8)$$

Where $H_{inf}^{overall}(t_0) = 0$.

The final goal of display agents is to show advertisements because they pay for advertise on it. Consequently, their happiness is proportional to the number of advertisements showed and the degree of advertisers' happiness, since displays want to keep happy their advertiser clients.

The function that evaluates the happiness of the displays is thus given by:

$$H_{dis}^{inst}(t) = \frac{1}{n} \sum_{k=0}^n H_{adv}^{overall}(k). \quad (9)$$

Where $H_{dis}^{inst}(t_0) = 0$.

n is defined as the number of advertisers that are around the display. The display's happiness is given by the average of the overall happiness of the advertiser agents that are in the niche.

The overall happiness for one display is given by:

$$H_{dis}^{overall}(t) = H_{dis}^{inst}(t) + \sum_{k=0}^{t-1} H_{dis}^{overall}(k) - \Delta. \quad (10)$$

Where $H_{dis}^{overall}(t_0) = 0$.

This function is defined as the that previous overall happiness.

4. SIMULATION RESULTS

In this section, we describe the simulation environment that we have realized to evaluate the effectiveness of the proposed approach in the introduced case study, and then present several results from the simulation experiments assessing the approach in terms of dynamic behaviour, stability and scalability.

4.1 The simulation environment

The simulation environment has been realized above the REPAST Toolkit (Repast), to simulate an environment as the one of our case study. In particular, it considers a 2D space (i.e., the area of a thematic park) populated with displays. There, a number of spatially situated agents belonging to different classes (users, advertisers, display, and information agents) act, interact and evolve according to a simple set of eco-laws (driving their matchings and interactions) and each with the egocentric goal of maximizing their individual happiness.

Displays (and thus display agents over them) have a fixed position and each represent a niche. Displays are assumed to be able both to display information (some data they own and that can be useful to users) and advertisements (based on the actions of advertiser agents). Display agents have indeed the goal of showing information and advertisements, and their happiness is proportional to the overall happiness of advertiser agents and user agents that live in their niche, as from previous section.

Users (i.e., user agents) randomly move in the simulation space, that is, be moving close to specific displays, they move from niche to niche. Each user agent is assumed to have an individual profile, expressing its preferences in terms of information and advertisements he would like to see. According to these preferences a matching function is defined for the instantaneous happiness of a user watching a display, that measures the degree of satisfaction of the user when watching information and/or advertisements.

The matching function returns a value between 0 and 1, depending on the degree of match between its preferences and the information/advertisements currently being shown in the display in which he is situated. In any case, we emphasize that in the simulation user agents do proactively pursue the goal of increasing their happiness, but are simply exposed to information and advertisements around.

Advertiser agents reside on niches have the goal – that they proactively pursue – of showing a specific advertisement to a well-targeted audience. Again, a matching function varying between 0 and 1 is defined that affect the happiness of advertisers depending on how well the overall audience (i.e., the users currently in that display niche) matches the advertisements. Clearly, the currently displayed advertisement is selected as the one best matching the current audience. Advertiser agents can move to a different display in the case no suitable matching is found in the current one.

Information agents reside on niches and act manage specific types of information. That is, by monitoring the overall situation on one niche, can affect the information stored (and displayed) on that niche. That is, by increasing in one niche that information that results being of use to the classes of users in that niche, and by deleting useless information instead. Also information agents can move from niche to niche to properly distribute on need the information they hold.

Despite some limitation of the current simulation envi-

ronment (that is: the number of agents is static since agents cannot die or reproduce; information and advertiser agents are not capable of dynamically adapting their behaviour and interacting with each other, but rather follow a set of individual static rules), experimenting over it enables to get useful insight on the overall behaviour of our proposed ecosystem approach.

4.2 Overall ecosystem behavior

We primarily use the following simulation scenario: 25 niches, 500 user agents, 100 advertiser agents, 100 information agents.

Let us first analyze the overall behavior of the ecosystem in terms of how the average happiness of the agents involved evolve over time.

Figure 3 shows the evolution of happiness for the scenario. After a short transitory period, the overall average happiness stabilizes to a nearly steady value for all the agents involved. That is, the ecosystem reaches a sort of global equilibrium in which its global parameters are stable, independently of the specific parameters of the simulation. Indeed, such stability is reached not only in this scenario, but also in other scenarios with different number of agents and niches involved.




Figure 3: Average happiness of agents for the whole ecosystem.

In any case, such apparent equilibrium is by no means the expression of a static situation. Rather, as the simulation continues executing and the agents in it continue acting and moving, it is an expression of the capability of the ecosystem as a whole to globally reach in a self-organizing way a coherent dynamic behavior.

To better analyze these results, let us see what happens at the level of individual niches, where we will discover many complex behaviour that are hidden behind the above discussed global stability.

4.3 Niche-level behavior

Let us see how things evolve in a single niche over time. Figures 4 and 5 show the evolution of the happiness and of the number of agents, respectively, in one of the 25 niches of the scenario (niche 11). However, the patterns of evolution are very similar for all niches.

What we see here is the high-degree of dynamics inside a niche, where the average happiness of the various classes of agents tends to vary with time, along with the number of agents themselves. Despite this, the stable balance already discussed at the global level is achieved.

Similar test of agents' happiness in the others niches were done. The results show us that the average happiness of user agents is well balanced in each of the niches. That




Figure 4: Happiness of agents over time in one niche.




Figure 5: Number of Agents over time in one niche.

is, the specific dynamics of the ecosystem, as enforced by the happiness functions of its various agents and by their interaction, have the effect of trying to make users as happy as possible, which results in users being nearly equally happy in every niche. Such balance of user happiness goes at the price of somehow unbalancing the happiness of the agents of other classes.

The results concerning the average number of agent in the others niches show us that the number of advertiser and information agents notably varies from niche to niche, to due to the fact that they are forced to be mobile in order to satisfy their own happiness.

4.4 Scalability

To evaluate the capability of the ecosystem to preserve its behavior in large scenarios, we have evaluated it under different scenarios with a larger number of users and a larger number of niches.




Figure 6: Average happiness of agents with different number of users (100 information agents, 100 advertiser agents and 25 niches).

Figure 6 shows the average happiness of agents when increasing the number of users involved. Interestingly, one can see that there are no substantial changes in the happiness of

the various agent classes involved in the simulation. In other words, the overall behavior of the system keeps stable even by changing the number of users, showing that its overall properties are maintained independently of the actual number of agents involved in the scenario.




Figure 7: Average happiness of agents with different number of niches (with 25 niches the number of user agents is that of 500, and then it is increased proportionally with the increase in the number of niches).

Figure 7, on the other hand, shows the average happiness of agents when increasing the number of niches (while preserving a proportional number of user agents). Specifically we consider 25, 36, 49, 64, 81 and 100 niches with 500, 720, 980, 1280, 1620 and 2000 of user agents respectively. Also in this case, no substantial changes can be perceived as the number of niches (and, thus, the size of the system) increases. This prove that our ecosystem approach, being based on local interactions only, is capable of preserving its properties even in the large.

In summary, the above two figures testify that the proposed ecological approach is suitable for large-scale systems with a massive number of service components.

4.5 Influence of other parameters

To further assess the robustness of our approach, we have evaluated how its behavior changes by tuning both the parameters involved in the matching functions and by varying the number of tasks that a display can concurrently perform (i.e., the number of information and advertisements it can concurrently show).

Varying the parameters of the matching functions implies modifying the rate of happiness increase (or decrease) in agents. For example, increasing the average return value of the matching between users and advertisements implies making users more tolerant to see advertisement.

Figure 8 shows exactly the effect of increasing the average return value of the matching function between users and advertisement. If users are more tolerant to see advertisements, which means that the match function when users see advertisements will return a higher value, then users and advertisers happiness increase, since the system shows more advertisements because. Information happiness decreases because the times to show its information also decrease (Figure 8).

From the definition of the happiness of advertisers, we know that if in a niche the display shows interesting information that increases the happiness of users, this will also increase the happiness of advertisers according to their degree of tolerance on seeing information. In Figure 9 we

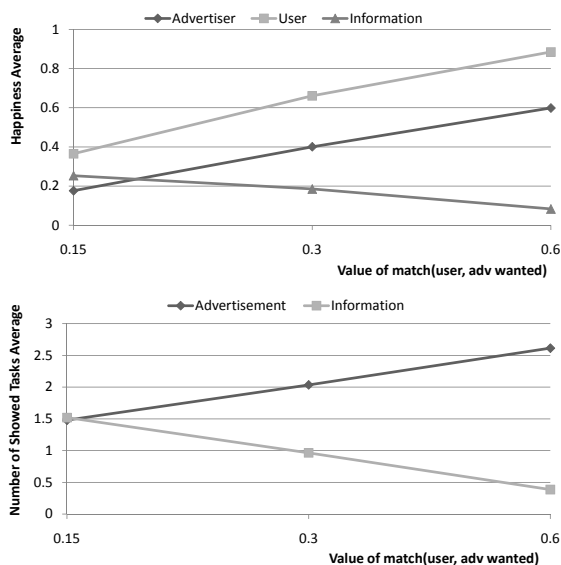


Figure 8: Average happiness for agent classes (up) and number of showed tasks (bottom) when varying degree of happiness of users when they see advertisements.

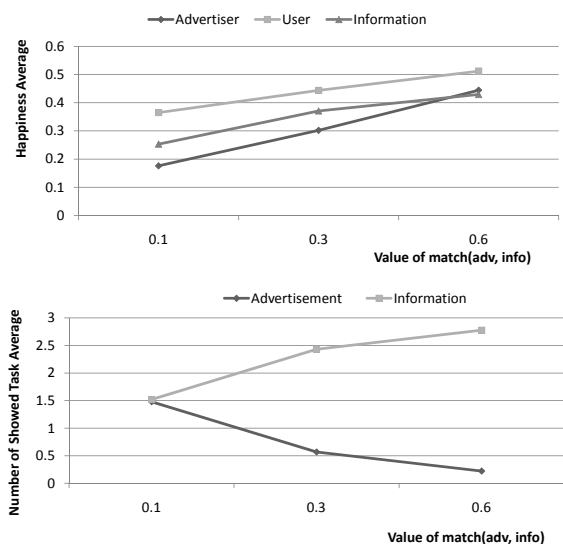


Figure 9: Average happiness for agent classes (up) and number of showed tasks (bottom) when varying degree of happiness of advertisers when they see information.

increase the degree of happiness when advertisers watch information. In other words, advertisers are more tolerant to see information, which means that the match function between an advertiser and a displayed information will return a higher value.

As expected, the happiness of agents increases, also the happiness of advertisers (since their level of tolerance also increases), even though, the number of advertisers showed decreases and the number of information showed increases.

As a final test, we modified the simulation by varying the




Figure 10: Average happiness for agent classes when varying the number of task that each display can show simultaneously.




Figure 11: Percentage of information and advertisement showed by increasing the number of task that each display can show simultaneously.

number of tasks (information and/or advertisement) that each display can concurrently show. Figure 10 shows how the average happiness of agents is proportional to number of tasks that each display can show simultaneously. This is simply due the fact that each display can accommodate at a time the needs of a larger number of users and advertisers. However, this does not change the overall behaviour of the ecosystem. For instance, Figure 11 shows that the relative percentage of information and advertisements shown on the displays does not substantially change. As more information and advertisements can be concurrently shown, the percentage of information tend to grow over advertisement. However, the overall balance between advertisement and information is preserved, due to the inherent dynamics of the regulatory feedback loops which are part of the system.

4.6 Summary

The above presented simulation experiments, although performed within a single case study and still exhibiting some limitations, shows very interesting properties. In particular:

- In the proposed ecological approach, the components of the system are able to reach stable self-organized global system behaviours, despite the inherent dynamics of the system itself;
- The approach promises to be scalable to very large-scale systems with a massive number of agents;
- The approach appears to exhibit very little sensitivity to variations in the parameters of agents and of the ecosystems.

5. RELATED APPROACHES

The key difference between our and other possible approaches that can be undertaken towards the realization of eco-inspired service frameworks stands in the metaphor adopted to model the ecosystem and its individuals.

Physical metaphors consider that the species of the ecosystem are sort of computational particles, living in a metric space 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 fields should diffuse and how particles should be influenced by the local gradients and shape of some computational field. Physical metaphors have been proposed to deal with several specific middleware-level aspects in dynamic network scenarios [4, 19] or in the area of amorphous computing [22]. However, they appear not suitable for general adoption in large-scale pervasive service ecosystems, due to the fact that they hardly tolerate diversity and evolution: the number of behaviors and interactions that can be enforced via fields is limited, unless one wants to sacrifice simplicity by increasing the number of different fields and the complexity of eco-laws.

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 reactive behaviors and that are influenced in their activities by the strength of specific chemical gradients (i.e., pheromones) in their surroundings. Biological metaphors are very similar to physical ones. Indeed, they too have been proposed to solve specific algorithmic problems in distributed network scenarios [3, 23]. However, the number of patterns that can be enforced by the spread of chemical gradients and by the reactions of simple individuals seem (as it is in the physical metaphor) quite limited, and this does not match with the need for time evolution and adaptation.

Chemical metaphors consider that the species of the ecosystem are sorts of computational atoms/molecules, living in localized solutions, and 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. The laws that drive the overall behavior of the ecosystem are sort of chemical laws that dictate how chemical reactions and bonding between components take place to realize self-organizing patterns and aggregations of components. Chemical metaphors have been proposed to facilitate dynamic service composition in pervasive and distributed computing systems [25], and they appear to well tolerate diversity and evolution (because they can accommodate a large number of different components with the same simple set of laws). However, a metaphor strictly adhering to the idea of a single chemical solution would hardly address the aspects of spatial distribution that characterize large-scale pervasive systems, unless properly coupled with some bio-inspired aspects, as we have recently proposed in [25].

Economic and market-based models have been already extensively studied in recent years [7, 14, 21] and they are well suited for large-scale Internet scenarios. However, they are possibly too complex and elaborated for pervasive computing scenarios, where simplicity of both individuals and interaction rules is a primary requirements (due to the presence

of resource constrained computation and communication devices).

As a final note, we emphasize that – as far as the modeling of agents in terms of happiness and actions – similar proposals have been made in the area of goal-oriented and utility-oriented agent systems [10, 17, 20]. The peculiarity of our approach is to couple this into with the overall natural inspiration applied to pervasive computing scenario and to exploit a flexible model of dynamic composition.

Ecological models like the one we have presented in this paper are largely unexplored, e.g., in [1] the adoption of a trophic-based interaction model is argued but not put at work.

6. CONCLUSIONS

In this paper, we have elaborated on the idea of modeling and developing next generation of pervasive service framework by getting inspiration from ecological systems. That is, by conceiving future pervasive service frameworks as a spatial ecosystem in which services, data items, and resources are all modeled as autonomous individuals (agents) that locally act and interact according to a simple set of well-defined “eco-laws”.

Beside the current limitations of our simulation environment, the presented results suggest that the proposed ecological approach has the potential to act as an effective general-purpose framework for spatially-situated and adaptive pervasive service ecosystems.

Beside the promising results, there are still a lot of open issues to exploit to turn our idea into a practically usable one. First, more simulations on a larger set of case studies are needed. Second, the proposed modeling for individuals and interactions needs to be better formalized and its properties more formally analyzed. Third, the security threats of our approach have to be identified. Finally, of course, there is need to put the approach at work in a real implementation and to test it on the field.

7. REFERENCES

- [1] G. Agha. Computing in pervasive cyberspace. *Commun. ACM*, 51(1):68–70, 2008.
- [2] M. Autili, P. Benedetto, and P. Inverardi. Context-aware adaptive services: The plastic approach. In *FASE '09: Proceedings of the 12th International Conference on Fundamental Approaches to Software Engineering*, pages 124–139, Berlin, Heidelberg, 2009. Springer-Verlag.
- [3] 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.
- [4] J. Beal and J. Bachrach. Infrastructure for engineered emergence on sensor/actuator networks. *IEEE Intelligent Systems*, 21(2):10–19, 2006.
- [5] F. M. Brazier, J. O. Kephart, H. V. D. Parunak, and M. N. Huhns. Agents and service-oriented computing for autonomic computing: A research agenda. *IEEE Internet Computing*, 13(3):82–87, 2009.
- [6] D. Cornforth, M. Kirley, and T. Bossomaier. Agent heterogeneity and coalition formation: Investigating

- market-based cooperative problem solving. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 556–563, Washington, DC, USA, 2004. IEEE Computer Society.
- [7] S. Dobson, S. Denazis, A. Fernández, D. Gaïti, E. Gelenbe, F. Massacci, P. Nixon, F. Saffre, N. Schmidt, and F. Zambonelli. A survey of autonomic communications. *ACM Trans. Auton. Adapt. Syst.*, 1(2):223–259, 2006.
- [8] D. Estrin, D. Culler, K. Pister, and G. Sukhatme. Connecting the physical world with pervasive networks. *IEEE Pervasive Computing*, 1(1):59–69, 2002.
- [9] D. T. Etienne de Sevin. An affective model of action selection for virtual humans. In *AISB '05: Proceedings of Agents that Want and Like: Motivational and Emotional Roots of Cognition and Action symposium at the Artificial Intelligence and Social Behaviors*, 2005.
- [10] A. Ferscha. Informative art display metaphors. In *4th International Conference on Universal Access in Human-Computer Interaction*, volume 4555 of *Lecture Notes in Computer Science*, pages 82–92. Springer, 2007.
- [11] A. Ferscha, A. Riener, M. Hechinger, and H. Schmitzberger. Building pervasive display landscapes with stick-on interfaces. In *Workshop "Information Visualization and Interaction Techniques", associated with CHI'06 International Conference, Quebec, Canada*, page 9, April 2006.
- [12] N. Haque, N. R. Jennings, and L. Moreau. Resource allocation in communication networks using market-based agents. *Knowledge-Based Systems*, 18(4-5):163–170, 2005.
- [13] M. N. Huhns and M. P. Singh. Service-oriented computing: Key concepts and principles. *IEEE Internet Computing*, 9(1):75–81, 2005.
- [14] J. O. Kephart and D. M. Chess. The vision of autonomic computing. *Computer*, 36(1):41–50, 2003.
- [15] P. Maes. The agent network architecture (ana). *SIGART Bull.*, 2(4):115–120, 1991.
- [16] M. Mamei, R. Menezes, R. Tolksdorf, and F. Zambonelli. Case studies for self-organization in computer science. *Journal of Systems Architecture*, 52(8):443–460, 2006.
- [17] M. Mamei and F. Zambonelli. *Field-Based Coordination for Pervasive Multiagent Systems (Springer Series on Agent Technology)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005.
- [18] J. F. Oliver Simonin. Modeling self satisfaction and altruism to handle action selection and reactive cooperation. In *SAB '00: Simulation of Adaptive Behaviors*, pages 314–323, 2000.
- [19] S. D. Ramchurn, C. Sierra, L. Godo, and N. R. Jennings. Negotiating using rewards. *Artificial Intelligence*, 171(10-15):805–837, 2007.
- [20] D. Servat and A. Drogoul. Combining amorphous computing and reactive agent-based systems: a paradigm for pervasive intelligence? In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 441–448, New York, NY, USA, 2002. ACM.
- [21] W.-M. Shen, P. M. Will, A. Galstyan, and C.-M. Chuong. Hormone-inspired self-organization and distributed control of robotic swarms. *Autonomous Robots*, 17(1):93–105, 2004.
- [22] 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, 2007.
- [23] C. Villalba and F. Zambonelli. An ecological approach for large scale pervasive multiagent systems. In *Proceedings of the 4th AAMAS Workshop on Massive Multiagent Systems*, Budapest, HU, 2009.
- [24] M. Viroli and F. Zambonelli. A biochemical approach to adaptive service ecosystems. *Information Sciences*, 180, 2009. To appear.
- [25] F. Zambonelli and M. Mamei. Spatial computing: An emerging paradigm for autonomic computing and communication. In *1st International Workshop on Autonomic Communications*, volume 3457 of *Lecture Notes in Computer Science*, pages 44–57. Springer, 2005.
- [26] F. Zambonelli and M. Viroli. Architecture and metaphors for eternally adaptive service ecosystems. In *Intelligent Distributed Computing, Systems and Applications, Proceedings of the 2nd International Symposium on Intelligent Distributed Computing*, volume 162 of *Studies in Computational Intelligence*, pages 23–32. Springer, 2008.