MyAds: a System for Adaptive Pervasive Advertisements

Antonio Di Ferdinando [a], Alberto Rosi [b], Ricardo Lent [a] ,
Antonio Manzalini [c] and Franco Zambonelli [b].

[a] *Imperial College London, Department of Electrical and Electronic Engineering, Exhibition Road, London SW7 2BT.*

[b] *Università di Modena e Reggio Emilia, DISMI, Via Amendola 2, 42100 Reggio Emilia, Italy.*

[c] *Telecom Italia Lab, Via Reiss Romoli 274, 10148 Torino, Italy.*

Abstract

In this paper we show how pervasive technologies can be employed on a public-display advertisement scenario to enable behavioral self-adaptation of content. We show this through *MyAds*, a system capable of exploiting pervasive technologies to autonomously adapt the advertisement process to the trends of interests detected among the audience in a venue. After describing the rationale, the architecture and the prototype of MyAds, we describe the advantages brought by the use of such a system, in terms of impact on the audience and economic efficiency. The comparison of MyAds performances with different advertisement selection techniques confirms the validity of our advertisement model, and our prototype in particular, as a mean for maximizing product awareness in an audience and for enhancing economic efficiency.

Key words: Pervasive computing, self-adaptation, autonomic communication, auctions.

Email addresses: `A.Di-Ferdinando@imperial.ac.uk` (Antonio Di Ferdinando), `Alberto.Rosi@unimore.it` (Alberto Rosi), `R.Lent@imperial.ac.uk` (Ricardo Lent), `Antonio.Manzalini@telecomitalia.it` (Antonio Manzalini), `Franco.Zambonelli@unimore.it` (Franco Zambonelli).

1 Introduction

Advertising has an extreme importance in today's society, and business companies invest constantly growing parts of their budgets in engineering ways to attract higher numbers of customers towards the products and services they offer. Historically, advertisement has always followed the main technological trends exploiting, from time to time, the new channels that technological development has been capable of providing [20]: newspapers, billboards, radio, television and, more recently, the web and the email.

The traditional advertisement strategy consists in taking a generic audience to achieve product awareness through a massive bombing of advertisements. However recent marketing studies are drawing attention on the fact that this strategy, besides of being economically onerous for companies, is becoming less effective as the bombing strategy sometimes challenges the audience in reaching peaks of irritability [19], for instance through email spam.

Given the above situation, new advertisement techniques and strategies need to be studied. Advertisement through the web [15] [16], for instance in the form of so called *web ads*, represents a first attempt in the direction of refining the advertisement policy, where user potential needs and preferences are detected, on the basis of past activities, to the end of suggesting links to products of potential interest. Several enterprises in the Internet market arena, such as *Google* and *Amazon*, adopted the strategy of suggesting links to products or services on the basis of user information (i.e. past navigation history, purchases or preferences). Behind this process lies the business strategy of refining advertisement in a way to suggest products and services only to potentially interested users, leaving uninterested ones aside (thus avoiding negative impact).

Web ads represent only a first step towards the design of adaptive pervasive advertisement systems. Given web ads success, it is natural to imagine the impact of a fully pervasive environment [5] [25], in which several distinct pervasive devices concur in building a snapshot of the audience at a specific venue: whereas traditional advertisement systems rely on companies pushing contents to users with the hope to meet a latent need, pervasive advertisement explores the way of providing contents in a revolutionary sort of effective publish/subscribe-like [10] way: users have the ability to express (implicitly or explicitly) their interest in a product, in an event, or a pattern of events, and are subsequently notified of any event, generated by a publisher, which matches their registered interest. Given these considerations, pervasive advertisement appears to have all the characteristics for a fast and effective commercial diffusion.

Large advertisement displays already pervade our cities and everyday environment. These are strategically placed in locations where the audience turnover is high, thus shopping malls, university cafeterias, metro or railway stations just to mention a few. As of today, these displays are typically governed by policies whereby advertisements are displayed sequentially in a static cyclical way regardless of the audience the display is exposed to. In this paper we focus on the idea of achieving adaptive, i.e. *audience-sensitive*, advertisement conjugating the rationale of adapting ads to users' characteristics with the display-based advertisement technique. In more detail, we present the architecture and a prototype for an innovative advertisement system called *MyAds*, whereby the exposure of advertisements to a potentially interested audience can be maximized. MyAds features a modular and adaptive architecture for capturing and processing user profiles and an advertisement allocation policy based on a real-time auction. The latter enhances the competition among advertising companies in such a way to allow allocation of an advertisement slot to the advertiser who values it the most. Potential advertising companies are thus given real-time information they can employ in the process of evaluating (whereas not predicting) the impact of the exposure of an own advertisement on the basis of dominant interests in the current audience. This, in turn, enables them to refine the investment, from time to time, as this information is disclosed to them.

Summarizing, our paper makes at least 3 contributions: (i) it presents an innovative paradigm of pervasive advertisement for the detection of users through short-range wireless technologies enabling the display logic to "sense and understand" the characteristics of its potential audience; (ii) it proposes a distributed prototype platform that shows how adaptive pervasive advertising systems are potentially easy to realize and deploy while yet effective from a commercial point of view; (iii) it proposes and evaluates a content selection paradigm based on auctions and experimentally studies its impact by comparison with other possible strategies.

The remainder of this paper is organized as follows: in Section 2 we introduce scenario and approaches to adaptive pervasive advertisement. In Section 3 we describe the general architecture of the MyAds system. In Section 4 we describe the way components have been developed in the prototype while, in Section 5, we discuss, study and evaluate strategies for advertisement selection. Finally, Section 6 draws some conclusion and identify areas for future research.

2 Adaptive Pervasive Advertisement: Scenarios and Approaches

Adaptive pervasive advertisement is a relatively new advertisement paradigm, whose nature is strongly influenced by the peculiar characteristics of the application scenario. In the remainder of this section, we detail the scenario we consider in our work and introduce related research and industrial approaches.

2.1 Scenarios

The scenario considers an environment equipped with sensorial and computational capabilities as a modern exhibition centre, a museum or a stadium, in which it is realistic to assume the presence of a pervasive infrastructure of embedded devices such as sensors of various types, Wi-Fi connections, RFID tags and other location systems. Without loss of generality, we assume visitors to carry portable devices, such as PDAs or smartphones, and/or be provided with RFID-based tickets storing a sort of personal profile embedding implicit heterogeneous information about the user (e.g. the nature of the visit, visitor category, or eventual fees paid). This profile could be furthermore enriched with explicit personal information (e.g. sport practiced, favorite music, hobbies, etc...) provided, on a voluntary basis, during a registration/subscription (to events or services) phase from the user itself.

We consider the presence, in the venue, of a number of advertising displays that show information about the centre itself, events thereafter hosted and organized, as well as third-party commercial advertisements. Public point of interest may afford the costs of deploying such infrastructure if this enables the provision of good services to visitors (and accordingly attract a higher number of persons) and an increase in the revenues.

Moreover, potential advertising companies can be interested in delivering their commercials on the basis of real sampling of dominant interests in the current audience, and refining their investments, from time to time, as this information is disclosed to them.

Therefore, we mainly recognize three typologies (groups) of people that have direct or indirect stake in such a scenario because they can affect, or be affected by, a pervasive advertisement system:

- the *Audience*, or i.e. people attending to broadcasted ads;
- the *Display Owners*, i.e., those who sustain the cost for displays' purchase and maintenance;
- the *Advertisers*, i.e., industrial companies that want to maximizing exposure of products to interested customers taking advantage of information about audience's interests that a system for adaptive pervasive advertisement could make available;

In our vision, the stakeholders at game are not antagonist and don't pose any joint optimization problem because the stakes these three groups share are not conflicting; to the contrary, a system for adaptive pervasive advertisement reveals to enforce and support each group, each providing a valid instrument to face competitors in their relative domain of competition (advertisers between advertisers, display owner between display owner and so on). Display owners featuring such a system, compared to the one which they don't, are in condition of selling valued-added advertisement slots because ads transmitted have a good probability of hitting the target. Advertisers which trust in an adaptive pervasive advertisement system know that the ratio between cost per ad transmitted and number of people interested by their content will be certainly lower. In fact, it is realistic to consider the possibility to advertise a product to an audience known (and guaranteed) to be potentially interested as a valuable business asset. The enforced efficiency on costs that comes from represents for sure a strong competitive edge towards competitors. Finally, people in the audience should consider a two fold advantages: from one side they are reached by ads not boring them because closer to personal interests, from another side if they agree to provide a personal information, they could be rewarded with benefits of different natures (as reduction on the ticket cost, etc.).

In such a scenario, our contribution is not therefore limited to (i) developing a way to exploit the pervasive infrastructure and (ii) selecting appropriate contents to be shown dynamically according to the way this information changes. When we consider the commercial aspect for this scenario, a third challenge arises: it consists in (iii) employing a content selection technique capable of enhancing the economic efficiency for each of the display under consideration (see Section 3.2).

2.2 Related Approaches

After decades of technological improvements, it became clear that the use of pervasive technologies can bring tremendous innovation to the advertising industry. Having pervasive advertisement a two fold impacts on research and business interests, in the last few years researchers and industries have deeply investigated the scenario from different perspectives.

Great emphasis has grown on the investigation of ways to embed pervasive technologies in a wide range of everyday activities and contexts [18]. A number of initiatives have industrial origin, where the employment of, in particular, RFID tags in advertisement and communication applications is investigated as a mean for providing users with new services. This is the case of the *Tokyo Ubiquitous Network Project* initiative [7] [13], which foresees the equipment of the Ginza area in Tokyo with RFID tags, infrared and wireless transmitters to enhance a pervasive interaction with shoppers aimed at, among other things, advertising various offers from shops in the area, downloading local maps and accessing visitor information.

A slightly different extent lies behind the *MINI Motorby* [17] initiative started by Mini USA, which attempts to enhance the direct communication with consumers through the use of personalized "talking" billboards. Specifically, users are provided, after online registration, with a key fob containing built-in RFID tags that, when in range, will interact with strategically placed *MINI Motorboards* to let funny greeting messages appear on the board. Message content is generated from a number of personal factors (e.g. gender, status, profession) and operative (e.g. geographic position, current time, direction) nature. The RFID technology employed allows users to be identified at a very wide range (i.e. up to 500 feet) while driving or even simply carrying their key. At the time of writing this paper, both initiatives described above are in an experimental evaluation phase.

Industrial initiatives, such as the ones outlined above, are matched by initiatives of academic origin [24], from which our work takes inspiration. With this respect, our work shares the vision of the research conducted by Ranganathan et al. [20], which traces a new horizon for the advertisement where the path for reaching customers, and delivering the right advertisements to the right people, is achieved in a novel fashion. However, although this work highlights the reason why pervasive advertisement might potentially be an effective instrument, it does not provide information nor directions on how to actually build such system from an architectural point of view.

The study carried out by Rashid et al. [21] in 2005 focuses on ways of extracting information from an environment enriched with pervasive technologies. In more detail, users' needs and desires are inferred by exploiting location-based information over Bluetooth and Wi-Fi connections. Although the use of this approach can lead to the generation of valid context-based profiles for an audience, we found it inappropriate for our purposes in a way to introduce an interdependence between the accuracy of profiles generated and the environment surrounding the application. As a consequence, the range of interests of potential inferences results inevitably narrowed proportionally to the activities possible in the venue.

Finally, we found in the approach of Rogers at al. [22], named *Bluescreen*, the one that shares our vision the most. Similarly to our work, in fact, this approach presents an environment where a public-display advertisement scenario is governed by an auction-based allocation system. Even more similarly, the auction scenario is populated by bidding agents that access users' profiles and exploit information acquired to decide the worthiness of submitting an offer. However, despite these similarities our approach differs in at least two fundamental aspects. The first is related to the main system objective;

whereas, in fact, Bluescreen aims at keeping to the minimum the number of interactions needed to show advertisements to the totality of people expected to transit through the system, our idea is to select content in such a way to maximize the impact among an unpredictable audience (see Section 5.1). The second aspect is, on the other hand, related to the way Bluescreen exploits the information acquired through profiles in an audience. Whereas the former employs a probabilistic model, which considers past transits to predict future ones, our approach is more experimental and involves the provision of statistical audience information on a real-time and continuous basis.

3 ***MyAds*: Rationale and Architecture**

Our approach towards the realization of pervasive adaptive advertisement is in a system that we named *MyAds*. The system enriches public-display advertisement scenarios with autonomous and dynamic adaptation of the content to behavioral patterns detected in the audience in proximity of a display.

In this section, we shortly discuss (*i*) the rational behind the design of MyAds and (*ii*) the key characteristics of MyAds components of its overall architecture.

3.1 Rationale

The key guidelines that have driven the design of the MyAds architecture have been modularity and adaptivity. In particular:

- Modularity stems from the need to clearly identify the main functional blocks of our architecture and realize them in terms of different interacting macro-components, according to well-defined interaction patterns. Enforcing modularity in MyAds implies making it possible to dynamically change and/or optimize the internal structure of components without affecting the whole system, e.g., to transparently support different pervasive devices than RFID to acquire users profiles, or to transparently change the policy for allocating content to displays transparently.
- Adaptivity stems from the need to make it possible to dynamically instantiate new components within the overall architecture, even abstracting from their actual allocation, without affecting the actual working of the system and its functional behavior. Enforcing adaptivity in MyAds implies making it possible, e.g., to dynamically tolerate the instantiation or replication of new display components, to tolerate changing loads and usage patterns, or to have the selection of content to take place in a different geographical

location than that of the actual displays.

3.2 System Components

MyAds can be structurally regarded as composed by three classes of macro-components, namely the *Content Display* component, in charge of managing the displaying of content on a display and the acquisition of local contextual information (the profiles of users close to the display and possibly additional local environmental information); the *Information Processing* component, in charge of managing and elaborating users' profiles and possibly additional contextual information; and the *Content Allocation* component, in charge of evaluating which content (e.g., which advertisements) to actually allocate on displays.

In line with the adopted rationale, these three components have been designed in a modular way and to support their flexible instantiation within a system. A minimal instantiation of the MyAds system has to consider one instance of each component (see Figure 1). However, in the presence of multiple displays, one has to instantiate multiple instances of the Content Display component, and possibly (depending on the actual computational load on this components) of multiple Information Processing and Content Allocation components. As far as the geographical allocation of this components is concerned, MyAds tolerates the possibility for each instance of this components to be allocated wherever, depending on needs and on resources availability. For example, one can think at a scenario in which an exhibition in Reggio Emilia locally instantiates one Content Display component for each of its displays, exploits the Information Processing component available at Telecom, and outsources the actual decisions on advertisements selection to a company based in London, to a there located Content Allocation component. Indeed, we have personally tested our system in a similar configurations, by distributing the various components across different institutions in Europe.

In addition, it is worth emphasizing that the above components can be instantiated and put to work also to serve different or additional application purposes other than adaptive advertisement. For instance, one can think at exploiting the Content Display component and the Information Processing one to select other classes of information (e.g., information about events) other than (or in addition to) advertisements. As another example, one can think at putting the capabilities of the Information Processing and of the Content Allocation components to the service of different pervasive environments (e.g., active actuating environments or immersive games) than pervasive displays. However, a more deeper investigation of the potential application possibilities of MyAds is beyond the scope of this paper.

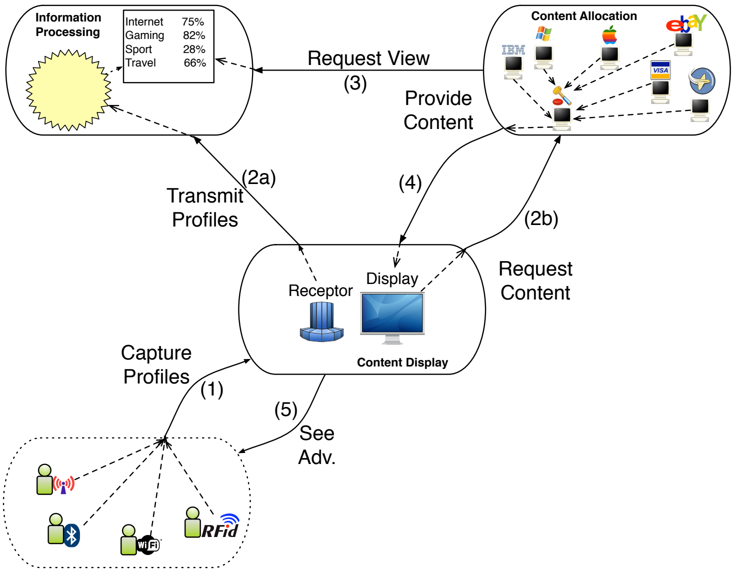

Fig. 1. The MyAds architecture.

Let us now detail the actual characteristics of the three introduced macro components.

Content Display Component. The Content Display component represents the interface through which MyAds manages a physical display, and is concerned with the two main tasks of requesting content allocation, to be later shown on the display, and interacting with pervasive devices in order to acquire user data and possibly environmental data around the display.

The first task involves the component to adapt the logic of the physical display to the characteristics of the content that, from time to time, will be selected. To this extent, the service time is divided into *advertisement time slots* (*slots* for short), and allocation of content for each of these is delegated to the Content Allocation Component, soon to be described. The length of a slot is decided dynamically on the basis of the length of the previous slot. This allows to define the timeliness requirements for the slot currently under consideration. Thus, for instance, when showing the content allocated for the slot s_1, known to have length l_1, the system will consider l_1 as the timeliness requirement for the allocation of the slot s_2. Consequently, it will adapt the allocation mechanism for s_2 to comply with this requirement.

The second task for the component starts from the consideration that the enhancement of pervasive features will involve coupling displays with hardware pervasive receptors. Consequently, this component is the ideal starting point for the gathering of pervasive data, and is thus equipped with means for gathering audience profiles and forwarding them to the Information Processing component, next to be described, where they can be processed. It's worth nothing that not only people profiles could be of interest, but also sensors in the environment (e.g. cameras, mics, sensor network and so on) could participate in completing and expanding available information about environment, visitor behaviors and more in general about the real local contexts.

The management of the Content Display Component only considers the local context of the physical display. By doing so, a higher degree of accuracy for the content exposure can be achieved, as the local Content Display component can select the best allocation technique on the basis of the evolution of its local context. Clearly, this requires the presence of multiple Content Display components for the management of multiple displays.

In talking about the user profile collection process, we take the assumption that users build their own profile and give their consent to share it. This assumption is not trivial, and triggers serious privacy concerns. However, without loss of generality, we claim that this assumption does not restrict the applicability scope of our work, as many ways of inducing users to share their profile, mainly based on incentives, can be observed in everyday life, as we already discussed in Section 2.1. In addition, privacy is further supported and reinforced in the profile processing step, where the Information Processing component can make it available to third parties in an aggregated and/or anonymous way.

On the other hand, further concerns can be raised about the "validity" of profiles, intended as the relevance the profile shows with respect of the owner's real interests. With this respect, a sociological dissertation should be carried out about the motivations that lead a person to frequent a place, which is beyond our interest. However, by observing the growth of tools with similar social nature, such as for instance mailing lists and SMS notifications, we are confident about the validity of the assumption that the owner of a profile is interested in receiving information (i.e. news, offers, events) concerning areas of personal interest.

Information Processing Component. The Information Processing component has the task of receiving audience profiles (and possibly additional information about the local context) and processing them in order to generate *Knowledge Views* built by organizing information extracted from profiles generating "snapshots of local reality", i.e., compact, meaningful and usable contextual knowledge. This allows the construction of aggregate views of local reality from the raw information collected, thus easing retrieval and extraction. These snapshots are made available on an on-demand basis, so that other components can acquire information on the context (i.e. become "context-aware") by simply requesting this information on a continuous basis.

The scope of such knowledge views is typically local to the context sourcing the profiles knowledge refers to, however knowledge views of larger scope may be built by reifying processing on local views on a larger, eventually system-wide, scope. This might be the case, for instance, of a scenario where possible advertisers require guarantees on the statistical relevance of audience flowing in the venue with the business topic they want to expose.

The typical scenario of MyAds will see the instantiation of one Information Processing component only. However, as knowledge views are built on a per-display basis and have a scope local to the display sourcing profiles, more complex scenarios might involve the need of different processing techniques for different displays or even balancing a workload that has become excessive for the processing by one sole component. In this case, other instances of this component may be used or, alternatively, the same instance may be enhanced to provide more than one processing algorithm (or more knowledge views) in the former case.

Content Allocation Component. The Content Allocation Component has the non-trivial task of selecting the (advertisement) content to be shown on the display at a specific time. The actual selection process relies on a *Content Allocation Policy* (a *policy* for short) to provide the logic, parameters and priorities that allow to individuate the right content provider on the basis of the peculiarities of the scenario. Regardless of these factors, ideal allocation policies need to meet three fundamental characteristics for any scenario where high levels of dynamism need to be achieved and sustained:

- *Flexibility:* The policy accommodates the needs of both the slot seller and buyers, eventually meeting changes in expectations at both sides, in front of a audience that can change rapidly.
- *Openness:* Whereas possible, the policy must accommodate for an arena of potential advertisers that grows and shrinks dynamically over time, eventually with even high frequency.
- *Optimality:* The policy should guarantee optimality in the content selection process on the basis of the arena of advertisers *and* contextual information available at any given moment.

The use of policies allows isolating content selection in a self-contained object, thus making it easy to test distinct ones by switching their use within the component. This feature was useful in the development of our prototype (see Section 4.4), where we developed three different allocation policies that met the above characteristics at various degrees. We employed each of them through experimental evaluations (see Section 5) and compared the results: between them, the policy based on auctions confirmed to reach the best outputs.

The way the components described above interact is depicted in Figure 1, which shows a typical basic scenario where one display is managed and one instance of each components is present.

The logic starts with the Content Display component interacting with the pervasive devices carried by the audience to the extent of acquiring profiles from the audience, as interaction (1) in Figure 1 shows. In detail, the pervasive (the figure mentions Bluetooth, RFID and WiFi) receptors included in the component detect the devices in the audience as these appear in the range, and asks each of them for the person profile they are carrying. The receptor simply forwards each profile received to the Information Processing component, as interaction (2a) shows.

At the same time, the reception of new profiles declares the need for content to be shown to this audience. For this reason, concurrently to interaction (2a) the Content Display component requests the Content Allocation component for contents to fill the next available slot (interaction (2b)), as the present one has already started. This latter component applies the policy to achieve a decision on the advertiser to allocate the slot to. In order for this selection to respect the rules of audience-sensitivity mentioned in Section 2.1, the decision needs to consider and evaluate the audience context local to the display. To this extent, the Content Allocation component interacts with the Information Processing component to acquire relevant contextual information in the form of aggregate views referring to the relevant display (interaction (3)).

Once the advertiser for the slot has been selected, the Content Allocation component returns the requesting component with the content to be displayed (interaction (4)), provided by the advertiser, so that it can be exposed to the audience (interaction(5)).

4 Prototype Implementation

The system described above has been implemented in a prototype platform, which has been used to test our ideas on a real distributed environment. Components developed for the platform conform to the specifications of the *Autonomic Communication Element* (ACE). The remainder of this section details the ACE component model and the way they have been employed to build the actual prototype.

4.1 Autonomic Communication Elements

The CASCADAS project (www.cascadas-project.org) aims at defining a general-purpose component-based paradigm for autonomic and situation-aware services for next generation network infrastructures and for pervasive computing scenarios.

The key concept at the root of CASCADAS is the ACE (Autonomic Communication Element), intended as a unifying software engineering abstraction for the development of component-based distributed services. The ACE model has been implemented in an associated ACE toolkit released by the CASCADAS consortium and is available in open source. The ACE model takes inspiration and leverages from existing autonomic component-models and adaptive agent-based models with features conceived to facilitate the design and development of complex, self-adaptive [4,2] and self-organizing network services running on a wide range of heterogeneous devices [3]. In particular:

- ACEs are able to run both on high-end computers as well as on tiny devices like sensors, due to their internal light and modular structure. Moreover, ACEs are planned to be able to relocate themselves dynamically to different devices at run time by making use of mobile code techniques.
- The internal functioning of ACEs relies on specific behaviors (i.e., roughly speaking, goal-oriented functions) that can be associated to individual ACEs or ACEs classes, and on an internal control loop that can enable self-monitoring and self-adaptation. More in detail, the service is co-modeled by a plan providing an explicit and machine processable representation of the actions the ACE will undertake (concretely, this is a XML file encoding a finite-state automaton with the actions to be performed), and by a set of functionalities that can be dynamically invoked while the ACE executes its plan. The benefits of this separation are numerous: (i) plan generation or modification is possible without intervening at the code level, but only writing the plan representation. (ii) The plan is supervisable since it describes the operations to be performed from a high-level point of view (as a finite state automaton). (iii) Code (functionalities) is structured by means of plug-and-play individual activities.
- The ACE components are provided with two communication mechanisms: (i) a distributed publish-subscribe mechanism to advertise and look for services offered by other ACEs using semantic descriptions. This is called the GN/GA (Goal Needed/Goal Achievable) protocol in which ACEs publish the goals they can achieve and look for goals they need for their task. The suitability of this kind of service discovery protocols is widely recognized in several pervasive and autonomic computing approaches. (ii) A direct message-passing mechanism allowing to flexibly contracting service usage. This mechanism supports bilateral and multilateral communication along

previously defined connection partners and allows to implement properties, such as encryption or fixed-number-of-participants constraint.

The combined use of all the above mechanisms allows ACEs to dynamically and adaptively connect with each other to provide advanced autonomic services. This characteristic is intended to fulfill the self-organization scientific aspect described in the introduction. Using the discovery mechanism ACEs can find interaction partners in open and dynamic environments where both the other ACEs around are unknown and they come and go at any time. Using the contracting mechanism ACEs can autonomously organize with other ACEs into contract chains to create advanced and complex services. Moreover, a desirable feature within the ACE abstraction allows exposing groups of self-organized ACEs as standard ACEs that export a service in reality composed out of a mix of remote services and local functionalities. These design principles work together to support the development of autonomic communication services. The adoption of ACEs, and the exploitation of its distinguish features, has enabled us to develop the pervasive adaptive advertisement application as a truly distributed system. In particular:

- Their self-organization characteristics allowed to easily develop the system components described in Section 3.2 as a composition of ACEs each providing part of the service and becoming and ACE in it selfs.
- The characteristics of the component model made the prototype system capable of sustaining the high levels of dynamism a targeted message delivery application is required. Pervasive advertisement components can flexibly self-configure their behavior (e.g., modify the sampling rate in a sensor) by acting on the self model at run time.
- The ACE GN-GA discovery mechanisms of the toolkit support components to discover and interact each other in a flexible decentralized way and make it possible for new instances of components to dynamically enter into play.

MyAds has been developed through one or more ACEs providing a number of local functionalities, which were combined with other functionalities located remotely on other ACEs in order to form more sophisticated services seamlessly to the application logic. For instance, the information processing component was developed as a number of ACEs that engage in a coordinated interaction in order to provide a coherent processing logic. These, in turn, were further combined with other services in order to deliver the whole adaptive advertisement application. These concepts are better described in the next section.

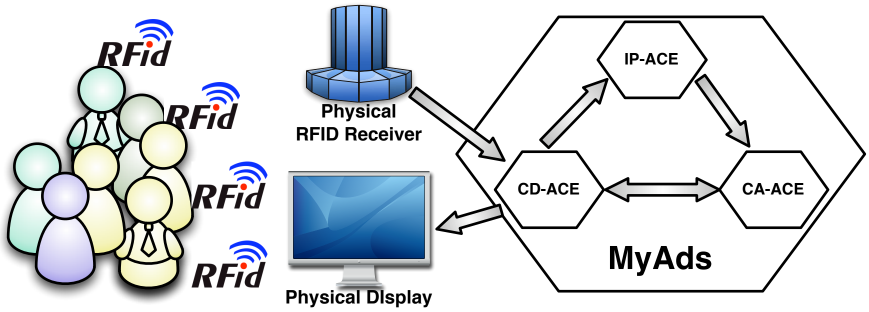

Fig. 2. ACE self-organization in MyAds prototype.

4.2 Prototype Components Implementation

The Content Display, Information Processing and Content Allocation components described in Section 3.2 have all been developed as sets of ACEs, each of which offers one or more peculiar services.

The Content Display ACE (*CD-ACE* for short) makes available services for gathering the pervasive data and displaying the advertisement. The Information Processing ACE (*IP-ACE*) enacts the information processing service, while the Content Allocation ACE (*CA-ACE*) serves content allocation. The combination of such services, resulting from the self-organization of these ACEs as shown in Figure 2, allows to meet the requirements of the pervasive adaptive advertisement service MyAds is concerned with. Upon system startup the CD-ACE *contracts* the IP-ACE, for acquiring the capability of processing profiles acquired into aggregate views, and the CA-ACE, for enhancing content selection. This latter component also contracts, in turn, the IP-ACE for accessing the aggregate views related to the display, which constitute the basis for the audience-sensitive content selection.

As already mentioned, the ACE abstraction allows exposing groups of self-organized ACEs as standard ACE. The prototype so built is itself an ACE exporting the pervasive adaptive advertisement service, and as such it might be employed as part of other systems anywhere such a service might be considered of interest. This feature is also clearly visible in the way the ACEs for the system components are structured. In fact, each component ACE results from the self-organization of "smaller" ACEs providing local functionalities. Let us now detail the specific implementation of each component.

Content Display component. The Content Display component has been developed as three ACEs, with the first (the *receptor ACE*) interfacing the pervasive receptor and thus managing the profiles acquisition task, the second (the *display ACE*) managing the physical display in order to actually show the content allocated, and the third, the *DB ACE*, providing information storage and retrieval functionalities for the multimedia contents the component will

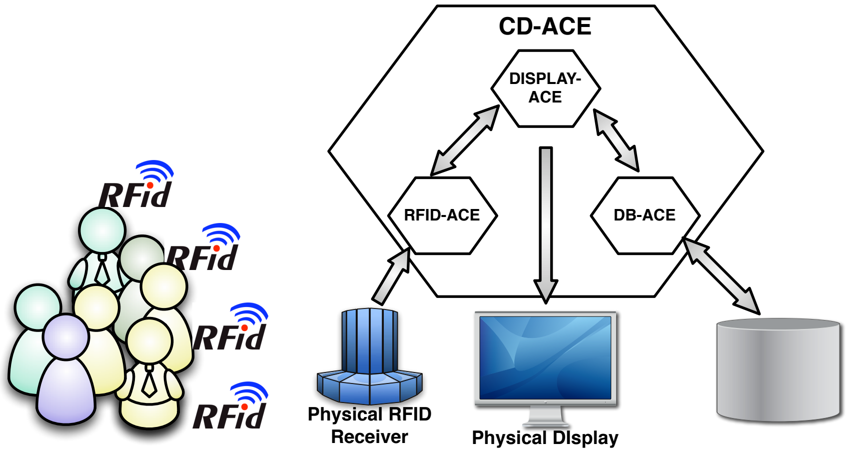

Fig. 3. CD-ACE in the prototype system.

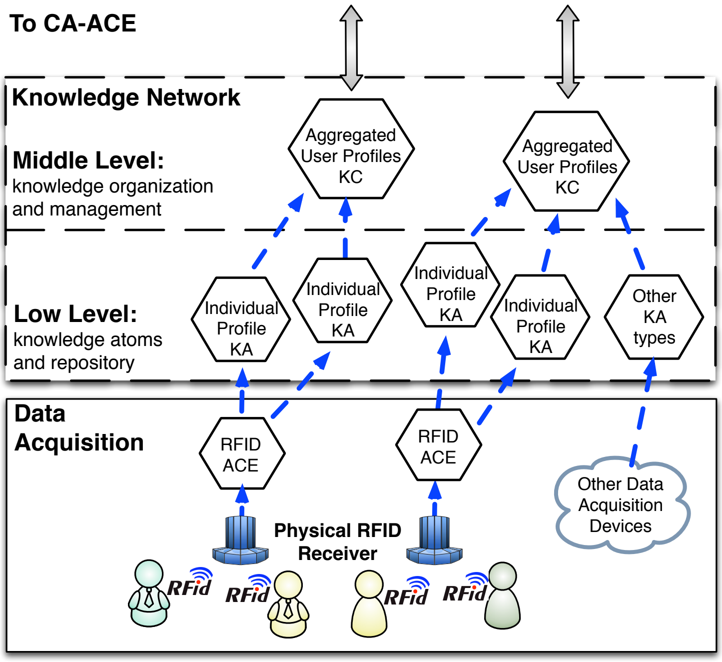

Fig. 4. Knowledge Network's layer structure.

receive as a result of content allocation. The structure for this component is shown in Figure 3. In the initialization phase the display ACE self-organizes with both receptor and DB ACEs . The effect of doing so is tying the screen to the audience in its physical closeness and storing multimedia content received, respectively. In particular, the acquisition of this latter functionality is made to the extent of possibly reusing multimedia content received in a past allocation, whenever possible, in such a way to reduce the impact on network traffic.

Information Processing component. The Information Processing component has been developed as a collection of ACEs that conform to the logic of *Knowledge Networks* [6] (KNs), that is a collection of *Knowledge Atom ACEs* (KAs), see Figure 4. Users profiles are provided to the KN through a layer of data acquisition (in our implementation via a RFID antenna) that transforms fetched profiles in KAs (once for each individual profile that have been acquired). A KA represents the atomic unit of knowledge, and is typically connected to a data source. A knowledge atom provides a uniform abstraction to access contextual information independently of its type, size or context. This is required to provide generic access to knowledge from within the knowledge network as well as from services and components that are outside, independently of the specific characteristics of the data source (whether, e.g. a sensor, a tag, or a Web atom). On the basis of the requests made by upper level application components, the *Knowledge Container* (KCs) components can organize knowledge in the form of "knowledge views", by making it possible to enforce and reify structural and behavioral relations between knowledge atoms and between other knowledge containers, in order to access such structured knowledge as if it were atomic information. Also, other than organizing knowledge, they can encapsulate algorithms and methods to manipulate knowledge, e.g., for analyzing, aggregating, pruning or transforming it.

Besides, it's worth noticing that Knowledge Networks improve overall system privacy because their snapshots of people profiles are aggregated and anonymous views of people interests. In a way, they act as a filter between people profiles and advertisers. In addition, the KN could embed any arbitrary privacy management policies.

Content Allocation component. The Content Allocation component was developed along similar lines as a number of ACEs that engage in continuous interactions to the extent of realizing the allocation logic. However, differently from the other components, which are likely to keep the number of ACEs constant with respect to a scenario, the ACEs in this scenario are variable proportionally to the number of advertisers present. These latter are thus central to the components, and their presence drives the allocation logic. As a consequence, the structure of this component changes substantially in accordance to the allocation logic, and the corresponding actors, employed. This concept can be observed in Figure 5, and will be further clarified in Section 4.4, where three distinct allocation policies will be presented.

4.3 Pervasive Profiles

As mentioned at the beginning of this section, the goal of the prototype developed was that of testing our ideas experimentally. For this reason, we decided to keep its structure to a basic set. This also applies to the way we modeled profiles, in which we decided to include only informations which enabled us to extract the data retained of interest. In detail, we structured the profiles in an object containing the following data:

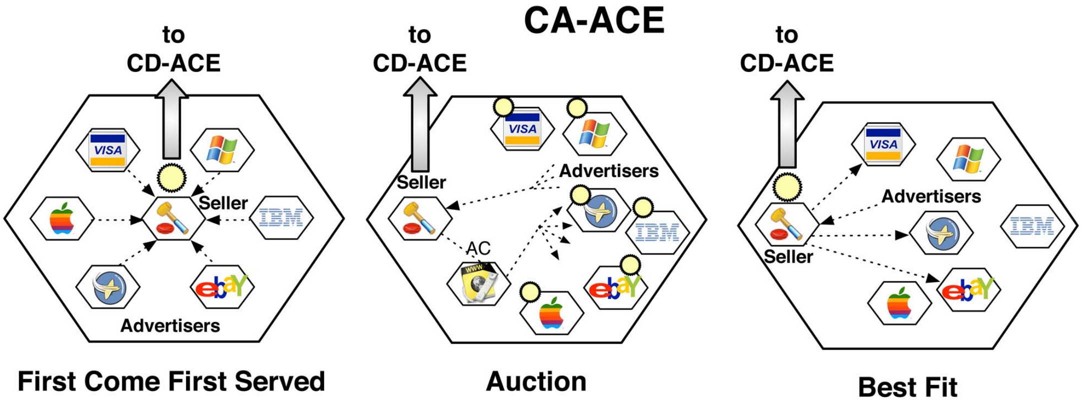

Fig. 5. CA-ACE in the prototype system.

- *ID*: a unique identification code that allows the identification of the profile of a person appearing in the range of the antenna univocally. The use of a unique ID is important in the application to guarantee that a person is counted exactly once in the processing of profiles.
- *Age slot*: it defines the age slot for the profile owner. This characterization is made upon creation of the profile owner, and the categorization is based on the person being *teenager*, *adult*, *senior* and *mature*. This information helps advertisers within the Content Allocation component in deciding on the worthiness of the slot under allocation.
- *Gender*: it defines gender of the person belongs to. Similarly for the age slot above, knowledge of this information allows the advertisers to bind the appropriateness of their own product more accurately, and is assigned upon profile generation.
- *Interests*: bind the profile to a list of subjects the owner declares an interest in. Interests considered within the scope of our prototype contain *reading*, *sport*, *internet*, *travel*, *cinema*, *music* and *gaming*. This information, assigned at profile generation time, constitutes the building block for the decision an advertiser takes about investing on the particular slot under consideration.

As already mentioned, our prototype acquired profiles by means of badges embedding RFID tags which are powered by the reader by radiated energy. With large crowds this could presents some issues with tag collisions, and the number of individuals who could be sensed at any given point would be limited. To overcome this limitation we provided our system of a software buffer in charge of managing tag collisions. The building of every single profiles snapshot is composed by 5 seconds of continuos sampling by the RFID antenna. Given the sampling frequency of 100 scans for seconds, every snapshot is built on a total of 500 RFID tag scans. With this workaround, even if tag collisions occur, sampling the tags more times at second lowers considerably the probability that a single profile couldn't be accessed.

For the prototype described in this section, three distinct allocation policies have been developed, namely *First Come First Served* (FCFS), *Best Fit*(BF) and *Auction*. Each of the policies triggers a different structure for the CA-ACE, as can be observed in Figure 5. A Seller ACE is always present to take the final decision on allocation, based on different interaction patterns with the set of advertising ACEs. The usage of each policy within the CA-ACE naturally excludes the use of all other ones, and the switch from one to the other requires explicit human intervention in the current prototype. The rationale behind this choice couples with the main motivation behind the development of three policies, which is the one of comparing the results obtained with each of them so as to individuate the best one. The remainder of this section describes each policy in more detail.

First Come First Served (FCFS) Policy: advertisements are selected with a round-robin policy in a paradigm that governs the vast majority of advertisement displays currently observable in public spaces. Typically, advertisers pay a fixed time-based fee (e.g. weekly/monthly) in order to equally share the rights of advertisement on the display. However, policies involving advertisers to pay the actual display of an own advertisement on a singular basis can also be found.

This approach does not meet the flexibility, openness or optimality characteristics mentioned in Section 3.2. In particular, it is not audience-sensitive in a way that the sequential order with which content is displayed remains unaltered throughout the whole service provision time. However, it has been included in order to provide comparison with an allocation model that is not only employed in real-world scenarios,but also widely spread.

With respect to our system, the use of this policy is realized having seller to let advertisers notify their interest in the slot of time. The order with which bidders register at the seller is used as a basis for building the actual advertisement sequence and following the Round Robin principle, it resolves conflicts assigning time slices to each bidders in equal portions and in circular order, handling all processes without priority.

Best-Fit (BF) Policy: advertisements are selected on the basis of the suitability for the audience considered at allocation request time. The calculation on the best fit is done on the relevance of the interests detected in the audience with the ones owned by advertisers, which makes the BF policy audience-sensitive.

Problems arising from conflicts, where more advertisers are interested on the same slot, can be solved in a variety of ways. In our case, a sub-policy aiming at balancing the number of times advertisements are shown is employed. Advertisers are assumed to pay a price proportional to the advertisement value when this is exposed to a certain audience, in a paradigm close to the one employed for allocating TV commercials. With respect to our prototype, the best-fit allocation policy is realized with advertisers notifying the seller their interest in advertising in the slot of time, and the latter handling contextual information for selection purposes.

This policy meets the openness and optimality characteristics, as no assumption is made on the size of the arena of advertisers and the selection is made optimal from the actual best best fit logic. However, the flexibility requirement is only partially met, as the dynamics naturally rests the allocation responsibility to the seller, which might limit the fulfillment of advertisers' needs.

Auction Policy: this policy takes from the BF policy above and simply employes a conflict resolution mechanism that gives advertisers the freedom of proposing their own price for the slot under allocation. In other words, the auction allocation policy, soon to be described, is a sub-set of the BF allocation policy.

The auction scenario complies to the model defined in [11] and [12], implemented in [8]. The display is represented within the auction environment by a *seller ACE*, which offers the rights of advertisement on a specific slot of time under auction. This offer is advertised through an *Auction Centre* (AC) ACE, which the advertisers in the arena consult in order to come acquainted of slots currently under allocation. Then, advertisers eventually interested in the slot act as *bidder ACEs* and submit offers. The AC ACE has a passive role in our scenario, as in our evaluation we focused on the management of one screen only, however on more complex scenario it would provide a single point of reference for the allocation of slots belonging to distinct displays.

The auction model chosen in our scenario conforms to the rules of an *English* auction, where an initially low price is increased by a number of iterative bids. The length of auctions are specified by the seller ACE, and bound by the duration of the advertisement slot under auction.

Once acquainted of new auctions through the AC ACE, advertisers query the Information Processing component for knowledge views. Then, according to relevance of the dominant interests there contained with own business, they ponder the submission of an offer. This process is reiterated by each of the bidder ACEs, in an independent autonomous fashion, for the whole duration of the auction. Upon termination, the current highest bidder is declared winner, and provides the seller with the content chosen. This, in turn, is forwarded to the Content Display as the right content to show.

The Auction policy fully meets all three characteristics for the ideal policy described in Section 3.2. In fact, it inherits the openness and optimality from the BF policy above and enhances flexibility by ensuring that both slot seller and advertisers have an appropriate reward even in front of challenges due to the audience turnover. In particular, this is done through the level of competition intrinsically injected in the allocation, which leaves leadership of the allocation to advertisers while guaranteeing the slot seller with an appropriate final reward. Even more, this policy couples the need for incentives typical of auctions with context-awareness features of the system to form a synergy whose result is the audience-sensitive content allocation policy; advertisers are provided with contextual information about the audience in proximity of displays, so that they can autonomously ponder the worthiness of their own potential involvement in the relevant auctions.

We will validate these considerations in Section 5, where we show results from experimental evaluation.

5 System Evaluation

Evaluation of the system aimed at determining the performances in the two aspects of (i) timeliness (intended as the capability of guaranteeing that the content under display is still of interest for the audience by the time it is actually shown), and (ii) added value of audience-sensitive advertising, with corresponding economic effectiveness issues.

The experimental setting for our evaluation was built by deploying a distributed prototype of components on a number of machines. The Information Processing component (where all KA ACEs were executed) and the Content Display component, both located in Reggio Emilia, IT, consisted in one machine each. The Content Display component employed a hardware Alien Technology RFID antenna, with a range of about 1.5 meters, as receptor for pervasive data, and 15 passive UHF RFID tags, placed on small badges that people were asked to carry when passing by the display. For the sake of simplicity, RFID tags within badges were statically associated to profiles built in advance. Profiles were structured according to the profile structure described in Section 4.3, so as to simplify their processing into aggregated views in the Information Processing component. For the sake of our setting, these latter views provided a statistical characterization of interests amongst the audience. In addition we have built a simulation environment for testing the system with a wider audience (as detailed in Section 5.2).

The Content Allocation component was located in London, UK, where it was distributed on a number of machines organized in the following way: an arena with a maximum of 15 ACEs, simulating advertisers representing business enterprises contending the slot was simulated. In addition to these, an ACE representing the display owner (i.e. the seller) and, in the case of the auction allocation policy soon to be described, an ACE representing an Auction Centre were developed. Each of these ACEs resided on a dedicated machine, and therefore the Content Allocation component consisted of a maximum of 17 ACEs deployed on as many hosting machines.

The total number of machines utilized for the whole application was therefore 20 machines, distributed across Europe. All machines involved were desktop PCs with a standard equipment of single Intel Pentium III processor 1Ghz and 2Gb RAM. Of these, the ones local to the same domain were connected through a 1Gbit LAN connection, while the ones on different domain communicated via the Internet through a 100Mbit Ethernet adapter.

5.1 Analysis of Timeliness Guarantees

The real-time nature of the system requires the interaction model to allow consumption of the information as soon as this becomes available. Therefore, timeliness has a paramount importance in MyAds.

The interaction model previously described highlights how the advertisement process relies on information that, for its own nature, may change rapidly in time. This information needs to be gathered, interpreted, aggregated and consumed in the smallest possible time in order to be meaningful, because excessive delays might result in invalidating the context described. Therefore, an analysis on the life cycle of the information is required to better understand whether MyAds, and our prototype in particular, fulfills the timeliness requirements of such a highly dynamic scenario.

The life cycle of the contextual information can be broadly divided into two phases, as Figure 6 shows. The *Generation* phase, contained within the Content Display component, accounts for the time necessary for generating the information in the system. The *Consumption* phase, on the other hand, refers to the time necessary for the consumption of such information in the system and embraces both the Information Processing and the Content Allocation components.

The time needed for the generation of the snapshot will essentially be given by the sum of the time for the user profiles to be read by the pervasive receptor, indicated as t_1 in Figure 6, and the time needed to interpret such information and format it according to the specifications of the snapshot, denoted as t_2. t_1 is fundamentally bound by the specifications of the pervasive technology employed. In our prototype, we used an RFID receptor capable of reaching passive UHF RFID tags in a range of about 1.5 meters. This choice is supported by

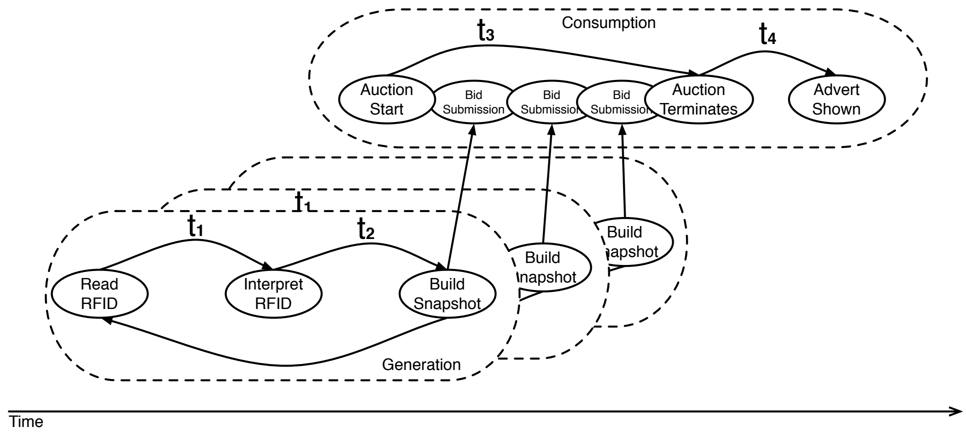

Fig. 6. Life cycle for the Contextual Information.

the wide range of successful commercial solutions, available nowadays, which allow to put emphasis on the simplicity of embedding RFID tags in badges, tickets or more general service access enablers. A trustworthy bound for the time t_1 can therefore provided on the basis of our experience with the use of the RFID system, which individuates in around 100 the maximum number of RFID tags that can be read per second in order to be considered reliable. This basically means that in very dynamic environments, which somewhat represent a worst case for us, the interpretation of a single RFID tag takes on average 10 milliseconds. Thus, we can approximate $t_1 \approx 10$ milliseconds.

The length of time t_2 is, on the other hand, proportional to the actual computational power of the machine hosting the ACE for the Information Processing component. An estimation is hard to obtain, but we can approximate its length by taking the experimental setting used for the tests shown above, where the component was instructed to generate the information every second, a time which we consider proper for assuring a fresh view to Content Allocation component. Thus, the condition to satisfy in order to be able to respect timeliness requirement is $t_1 + t_2 < 1$ seconds. This allows to bind the time for building snapshots to $t_2 < (1 - t_1)$ seconds. Bearing in mind the average above for t_1, of negligible amount, we can approximate an upper bound for t_2 as $t_2 \approx 1$ second, which is large enough to allow the component to produce a snapshot even for the least powered machine. This estimation, as well as the others to follow, does not consider factors such as network delays, as these are beyond our control. In any case, there is the possibility to tolerate network delays given t_1 and t_2, both < 1 second. The Consumption phase is bound by two intervals of time t_3 and t_4, which account for the duration of the auction and the gap time between termination of the latter and the actual display of the advertisement respectively. In the experimental setting used for tests, the first was a system parameter, and thus configurable, while the latter was a "technical" time driven by the communication delay towards the display. Within t_3, the contextual information is used in the process of deciding the submission of

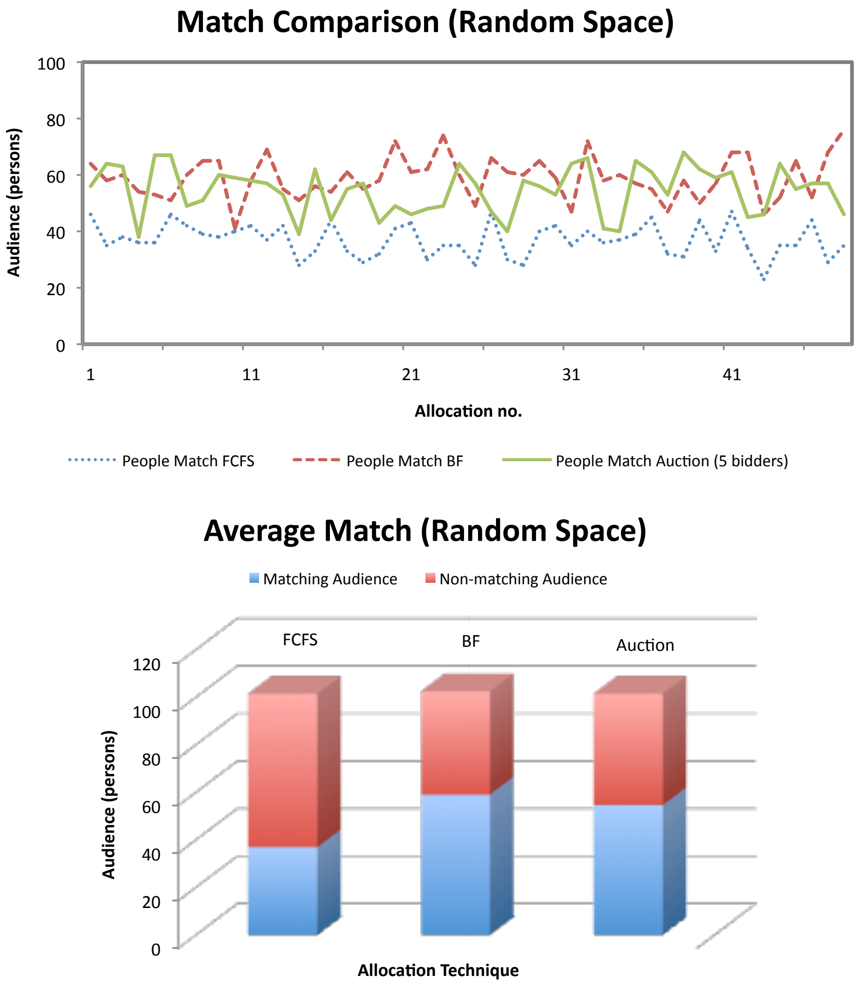

Fig. 7. Experiments on the random space: matches on subsequent allocations (top) and matching average values (bottom).

an offer. Thus, considering that t_4 is a technical time, hence not customizable, the key for exploiting timeliness in the system is to encourage offers to be submitted towards the end of the auction. This requirement couples perfectly with the rationale behind the most widely used bidding strategies, which aim at keeping the price low until the last instants of the auction, when all bids can be concentrated in an attempt to reduce the final expense.

The analysis above allows us to conclude that MyAds, and in particular the prototype presented in this document, is capable of fulfilling the timeliness requirements needed by the scenario. The system employs a logic capable of meeting the high dynamism requirements of the scenario, and in our prototype this is joined to a robust autonomic component architecture that brings the added value of robustness and reliability.

We have performed performance measurement focused at quantifying the impact of the advertisement on the current audience and characterizing the economic efficiency of the investment.

The main evaluation parameter considered when evaluating the impact of exposing an own advertisement to visitors is the *audience matching*, calculated through an analysis of the composition of the audience in the moment the advertisement is shown in the display and whose impact is assumed to be proportional to the value so calculated (i.e., the bigger the number, the better the impact).

Characterization of the economic efficiency of the investment, the second aspect of our performance evaluation, has been carried out by relating the matching quantified for an advertisement to the price paid for showing the advertisement itself. This will generate a value for the price paid for each of the persons in front of the display, and thus will give us a measure to define the goodness of the investment. Clearly, the smaller the price per-person (pp), the better the investment.

We can define an advertisement as *matching* a person if the subject of the former is of interest for the latter, as per the information specified in the profile. For the sake of simplicity, we consider this as a binary metric (i.e. 1 if the person is interested, 0 otherwise), and use it to determine the percentage of relevant audience in front of the display at the moment the advertisement is shown. The percentage of audience matching a displayed advertisement is the metric we used for estimating system's performance.

The value of advertising in an audience-sensitive fashion, and in particular through the use of an auction-based policy, was assessed by evaluating the matching of the auction-based policy against those already described of the FCFS (audience-insensitive) and of the BF (audience-sensitive) policies.

The evaluation was carried out by means of indicators through which comparing the performance of our prototype when using different allocation policies. Specifically, indicators concerned the impact, on the audience, of the advertisement chosen with different policies, and the economic efficiency triggered by each of the policies employed.

When large audiences needed to be simulated, a software profile generator was combined to the RFID tags. The generator simulated an audience of average size of around 100 persons, and where new persons arrive every 200 milliseconds on average and depart from the display every 20 seconds on average. Both inter-arrival and departure rates were exponentially distributed.

Persons were generated with a list of interests among the set mentioned in Section 4.3. To each of these was assigned a generation probability, so as to be able to "drive" the audience generation towards predictable scenarios. Thus, for instance, by giving higher probabilities of generation to a specified interest, i.e. *books*, we were able to predict the dominant interest among the audience flow generated in this way. This feature gave us the possibility of building a number of ad-hoc scenarios for the performance evaluation of MyAds, and we executed our prototype in two scenarios that we believe are the most interesting. In the first, we decided to simulate a venue populated by an audience that does not necessarily have a common interest, whose characterization suits venues such as metro and train stations, cafeterias or even off-the-street displays. We achieved this by uniformly distributing interest generation probabilities among the audience. Consequently, all interests had the same probability of being generated, and the nature of dominant interests was thus unstable. In the results presented below, we named this the *random* scenario.

The second scenario, on the other hand, simulates a situation at the opposite end of the spectrum; the profile generation process was configured in a way to prioritize the generation of specific interests over the remaining ones. This led to the clear emerging of a dominant interest among the audience, which suits scenarios where the display is located in a venue with a clear character attracting the audience, that is electronics fairs, thematic shopping malls, museums, etc. In the results below, we named this the *specific* scenario.

In the remainder of this section, we detail the policies employed in the comparison and describe the experimental setting before showing the results.

5.2.1 Results

Results shown in the remainder of this subsection have been gathered through sampling of subsequent executions on the application's continuous execution. Then, results are both shown as an excerpt of the results on 50 consecutive allocations, in graphs based on a temporal line, and as an aggregate, averaged over 1000 continuous allocations.

Impact on the audience: The impact of an advertisement on the audience of a heterogeneous public space is compared in the graphs shown in Figure 8. The graphs measure the impact of the advertisement on the audience in terms of match. The graph in the upper part of the figure shows a temporal representation, while the one in the bottom part delivers an aggregate view. Clearly, the matching of audience-sensitive strategies (i.e. BF and auction) makes the comparison with the FCFS strategy trivial on audiences similar size. Even though it may appear natural, however, this result shows that exploiting the spontaneous pervasive infrastructure in a public-display advertisement

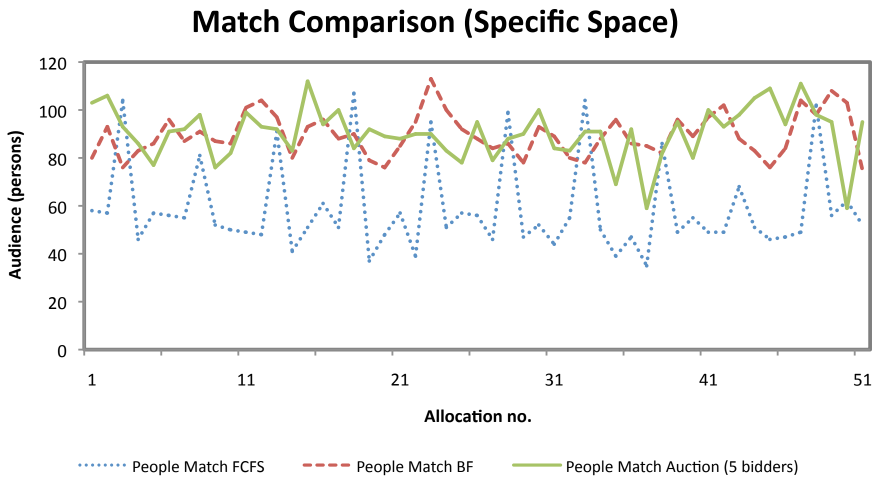

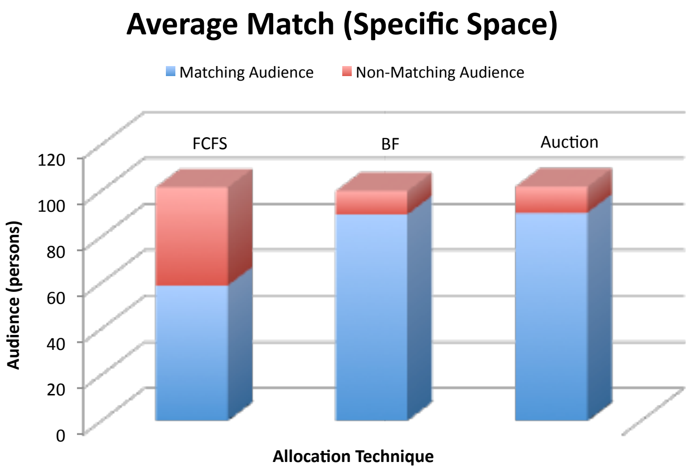

Fig. 8. Experiments on the specific space: matches on subsequent allocations (top) and average matching values (bottom).

scenario does bring an added value to the advertisement process.

On another note it is possible to see that the match obtained through the auction strategy is often slightly smaller that the one obtained through BF. This result can be explained by recalling our considerations about the influence of the level of competition in the auction, and considering that the auction setting in the graph involved a reduced number of advertisers, fixed to 5. As a consequence, it appears clear that the auction allocation policy needs to be exploited appropriately in order to maximize its efficiency.

When the simulated public space has specific characteristics, and thus the trend of interests in the audience tends to prioritize certain ones, the situation changes as Figure 8 from the second set of execution shows. As for the case of the heterogeneous environment, BF and auction strategies are shown to have a similar impact on the audience, with the FCFS strategy suffering of the lack of audience-sensitivity. However, the first thing that captures our attention in this figure is the difference in match level; while, in fact, in the previous case the matching reached percentages close to 60% with auction and BF strategies, in this case the matching for the same strategies reaches much sensitively higher peeks. In the context of the auction strategy, the influence of competition can be observed in Figure 9, where the number of competing advertisers was let grow from 5, to 10 and 15. It is possible to note how the match percentage slightly increases as the number of bidders increases. This is due, as already mentioned previously, to the enhanced competition level. In fact, as competition becomes tighter, with a higher number of advertisers present, more bids are submitted. This increases the rate at which bids are received by the auction seller, thus reducing the interval of time between the submission of the bid that will lately result the winning one and the auction expiration. As a result, the audience to which the actual advertisement will be shown will be more likely to be the same as the one upon which the decision of submitting was taken.

Economic Efficiency: The indicator we used to evaluate the economic efficiency was created by relating the cost of the investment to the singular person in front of the display for each of the allocation policies. This calculation led to determining the cost of the investment *per person* (pp).

Results from the first set of experiments (i.e. the heterogeneous environment) are shown in Figure 10. The average cost for the investment, intended as the final cost an advertiser sustains to advertise an own product, is calculated for all strategies and compared. In the FCFS and BF case, the cost has been calculated in the following way: bearing in mind that these two schemes imply the advertiser to pay a fixed price, the actual price model has been calculated by considering the cost per person obtained in the auction case as the actual *market price*. Starting from this, then, the affective price in the FCFS and BF cases was built by relating this value with the match obtained in each of these strategies.

Results in Figure 10 follow the predictable trend of "higher the matching, higher the price", which seems reasonable considering the results emphasized previously in terms of matching. However, it is worth noting that the auction strategy allows the best trade-off between economic efficiency and matching. This same trend can also be found in the results from the second set of experiments, which we omit for the sake of space, referring to the display being placed on a space with well-defined characteristics. Predictably, the sole difference in this latter scenario lies in the average cost per person, which results smaller in the BF and auction strategies, this latter with a bidders arena of size 5. This is obviously a direct consequence of the enhanced match obtained in this scenario through audience-sensitive strategies. Similarly, for the same reason the price for the FCFS strategy is seen to increase according to the

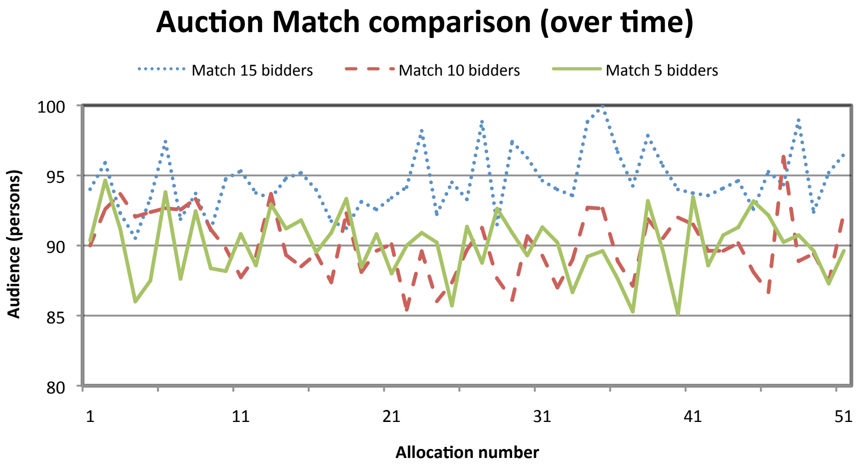

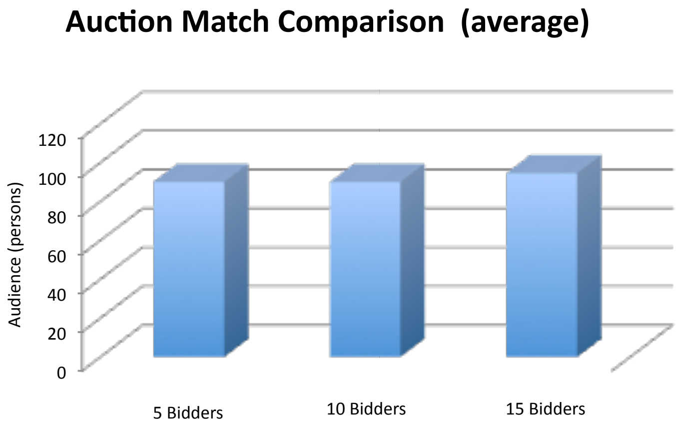

Fig. 9. Match comparison for the Auction policy with 5, 10 and 15 advertisers: subsequent allocations (top) and as average matching value (bottom).

advertisement price-building model described previously. When the arena of bidders is enlarged to the size of 10 and 15 bidders, whose graph we likewise omit for the sake of space, the cost of the advertisement clearly reflects the enhanced competition, whereby the more advertisers compete, the higher the final cost for the investment. This lets us argue that, even though the auction allocation strategy is better exploited in scenarios characterized by competition among advertisers, in certain situations of particularly high competition, the price to pay for the high dynamism of the scenario, and the high impact, might reflect on the advertisers' economic efficiency.

5.3 Summary of Results

The discussion in Subsection 5.1 allows us to conclude that MyAds, and in particular the prototype presented in this document, is capable of fulfilling the timeliness requirements needed by the scenario.

The graphs presented and discussed in the Subsection 5.2 clearly show that the BF and auction strategies outperform the classical, cyclical, allocation strategy, here implemented through a simple FCFS strategy. While this is clear in environments that attract audience with heterogeneous interests, as witnessed by graphs from the first set of executions, it results utterly evident in specialized spaces that attract an audience with clearly marked interests, such as for instance specialized fairs or museums simulated in the second set of executions.

The trade-off of this enhanced effectiveness lies in the average price advertisers find themselves to pay. The small price in the case of the FCFS strategy is, in fact, increased when the BF and auction strategies are employed. However, a more detailed study aimed at analyzing the price of the advertisement with respect to the match on the audience, shows that the average cost per person sustained is smaller in the case of employment of BF and auction.

Within the range of audience-sensitive strategies, BF and auction show similar levels of effectiveness, in terms of match on the audience, with the former appearing to perform slightly better than the latter, at first glance, when the levels of competition for the auction strategy are low. As a consequence, increasing this latter factor enhances the impact on an audience on average, but it also determines an increase of the cost for the investment that, in some cases, might contribute to reduce its overall worthiness. This leads to the final consideration that the auction strategy needs to be exploited properly, e.g. balancing the level of competition among advertisers, in order to optimize its outcome.

6 Conclusions and Future Work

In this paper, we have outlined the advantages of audience-sensitive advertisement techniques in a public-display advertisement scenario. The possibility to maximize the impact of the exposition of an advertisement is probably the main advantage, but it is indeed not the only one. Among the others worth mentioning there is, for instance, the consequent reduction of the so-called "bombing" effect, whereby the audience is unduly exposed to unwanted or inappropriate contents. Our approach to audience-sensitive advertisement in-

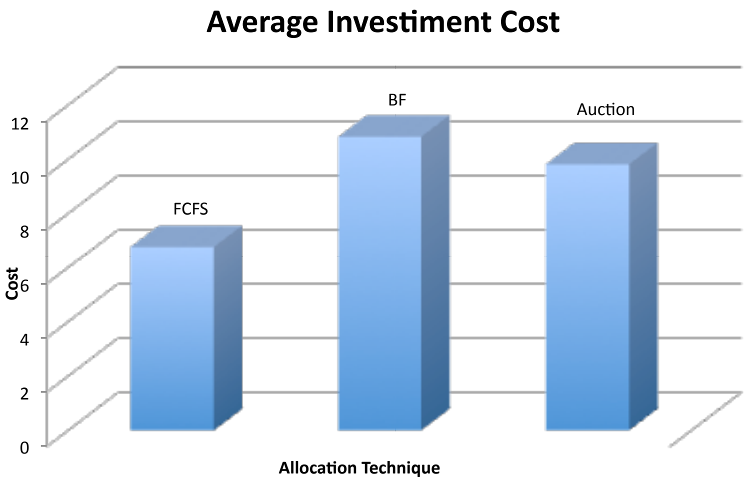

Fig. 10. Average investment cost.

volves making available real-time information about the audience currently by the display to content providers. Such information is used by content owners to ponder worthiness of the investment required, giving them a high degree of freedom on the actual possibility of getting involved in the allocation. This high degree of freedom is made explicit through the use of an auction as content allocation policy.

We have presented a system that we named MyAds, realizing the above audience-sensitive advertisement technique to autonomously adapt the contents to the evolution of the interests among an audience, and validated the ideas behind it through experimental execution of a prototype platform on a test-bed of distributed machines.

Results confirm the advantages of the use of audience-sensitive advertisement techniques, whereby the impact of the advertisement can be optimized towards a bigger part of the audience. Furthermore, they show that a proper exploitation of the auction-based allocation paradigm leads to an enhancement of the economic efficiency at both sides of the transaction, in particular at advertisers' side where the ratio between the cost of the investment and the final impact results optimized in a way to increase the segment of interested audience while reducing the sustained cost per person.

Future research can be carried out in several directions, and for the sake of space we will only cite the most interesting ones. First and foremost, more sophisticated strategies (for content allocation, but also for provision of situation-aware information) can be studied and evaluated for performance through comparison with the following ones. In addition, as emphasized in Subsection 3.1, the described components can be instantiated and put to work also to serve different or additional application purposes other than adaptive advertisement. On another direction, techniques for integrating sensors in the platform setting can be exploited in order to widen the range, and enhance the accuracy, of the provision of contextual information. Finally, new and more sophisticated ways of exploiting the pervasive infrastructure might also be a direction worth taking, possibly with the integration of algorithms and techniques borrowed from Opportunistic and Delay Tolerant networking but taking advantage of the high degree of connectivity the infrastructure is expected to have.

References

[1] The ACE Autonomic Toolkit. Available at the Sourceforge Website at: http://sourceforge.net/projects/acetoolkit/.

[2] L. Baresi, A. Di Ferdinando, A. Manzalini and F. Zambonelli, The CASCADAS Framework For Autonomic Communications, to appear in A. Vasilakos, M. Parashar, S. Karnouskos, eds., Autonomic Communication (Springer, Heidelberg and New York, 2009).

[3] E. Hoefig, B. Wuest, B. Benko, A. Mannella, M. Mamei, E. Di Nitto, On Concepts for Autonomic Communication Elements, IEEE International Workshop on Modeling Autonomic Communications Environments, Dublin, Ireland, 2006.

[4] The CASCADAS Project: http://www.cascadas-project.org.

[5] G. Castelli, A. Rosi, M. Mamei, F. Zambonelli, A Simple Model and Infrastructure for Context-Aware Browsing of the World, Pervasive Computing and Communications, 2007. Fifth Annual IEEE International Conference on , pp.229-238, 19-23 March 2007

[6] G. Castelli M. Mamei and F. Zambonelli, Engineering Contextual Knowledge for Pervasive Autonomic Services ,to appear in International Journal of Information and Software Technology 50 (2008) 36-50.

[7] G. Chhabra, Shoppers to Be Attracted Via RFID Advertising, in The RFID Weblog: http://www.rfid-weblog.com/50226711/shoppers_to_be_attracted_via_rfid_advertising.php, Jan. 2007.

[8] A. Di Ferdinando, R. Lent, E. and Gelenbe, A framework for autonomic networked auctions. In Proceedings of the 2007 Workshop on innovative Service Technologies (Rome, Italy, October 28 - 30, 2007). INSERTech, vol. 392, ICST, Brussels, Belgium, 1-10.

[9] A. Di Ferdinando, A. Rosi, R. M. Lent, F. Zambonelli and E. Gelenbe, A Platform for Pervasive Combinatorial Trading With Opportunistic Self-Aggregation, in Proc. WoWMoM AOC 2008 1-6.

[10] P. Th. Eugster, P. A. Felber, R. Guerraoui and A.-M. Kermarrec, The many faces of publish/subscribe. ACM Comput. Surv. 35, 2 (Jun. 2003), 114-131

[11] E. Gelenbe, Analysis of automated auctions, in Proc. ISCIS 2006, Lecture Notes in Computer Science, Vol. 4263 (Springer, Berlin, 2006) 1-12.

[12] E. Gelenbe, Analysis of single and networked auctions, to appear in ACM Trans. on Internet Technology, 2008.

[13] M. Fitzpatrick, Tagging Tokyo's streets with no name, `http://www.guardian.co.uk/technology/2007/may/10/japan.guardianweeklytechnologysection` May 10 2007.

[14] M. Ghorbel, M. Mokhtari and S. Renouard, A distributed approach for assistive service provision in pervasive environment. In Proceedings of the 4th international Workshop on Wireless Mobile Applications and Services on WLAN Hotspots (Los Angeles, CA, USA, September 29 - 29, 2006). WMASH '06. ACM, New York, NY, 91-100.

[15] S. Kumar, V. Jacob, and C. Sriskandarajah, Scheduling advertising at a web site, in Proc: WITS 2000 37-42.

[16] M. Langheinrich, A. Nakamura, N. Abe, T. Kamba, and Y. Koseki, Unintrusive customization techniques for Web advertising. Comput. Netw. 31, 11-16 (May. 1999), 1259-1272

[17] Mini Cooper and RFID Technology, in Technology News Daily. `http://www.technologynewsdaily.com/node/5962`.

[18] New Advertising Mission Statement, in Center for Interactive Advertising (ciAd) `http://www.ciadvertising.org/sa/spring_06/adv391k/905allison/RFID.html`.

[19] N. K. Fritz, Claim Recall and Irritation in Television Commercials: An Advertising Effectiveness Study, Journal of the Academy of Marketing Science 7(1979) 1-13.

[20] A. Ranganathan and R. Campbell, Advertising in a pervasive computing environment. In Proceedings of the 2nd international Workshop on Mobile Commerce (Atlanta, Georgia, USA, September 28 - 28, 2002). WMC '02. ACM, New York, NY, 10-14.

[21] O. Rashid, P. Coulton and R. Edwards, Implementing location based information/advertising for existing mobile phone users in indoor/urban environments, Mobile Business, 2005. ICMB 2005. International Conference on, pp. 377-383, 11-13 July 2005

[22] A. Rogers, E. David, T. R. Payne and N. R. Jennings, An advanced bidding agent for advertisement selection on public displays. In Proceedings of the 6th international Joint Conference on Autonomous Agents and Multiagent Systems (Honolulu, Hawaii, May 14 - 18, 2007). AAMAS '07. ACM, New York, NY, 1-8.

[23] M. Sipser, Introduction to the Theory of Computation. SIGACT News 27, 1 (Mar. 1996), 27-29.

[24] C. Narayanaswam, D. Coffman, M. C. Lee , Y. S. Moon, J. H. Han, H. K. Jang, S. McFaddin, Y. S. Paik, J. H. Kim, J. K Lee, J. W. Park, D. Soroker, Pervasive symbiotic advertising. In Proceedings of the 9th Workshop on Mobile Computing Systems and Applications (Napa Valley, California, February 25 - 26, 2008). HotMobile '08. ACM, New York, NY, 80-85.

[25] Essex, D. 2009. Matchmaker, matchmaker. Commun. ACM 52, 5 (May. 2009), 16-17.

Acknowledgements

The material here presented is based on a research funded by the CASCADAS (*Component-ware for Autonomic Situation-aware Communications, and Dynamically Adaptable Services*) project (EU FP6-027807). Authors would like also to thank the Fraunhofer Institute for Open Communication Systems FOKUS for the collaboration in the experimental phase.