

CONTEXT-AWARE COORDINATION IN THE SENSORS' CONTINUUM

Nicola Biccocchi, Gabriella Castelli, Alberto Rosi, Franco Zambonelli

DISMI – Università di Modena e Reggio Emilia – Via Amendola 2 – Reggio Emilia – ITALY
name.surname@unimore.it

ABSTRACT

Pervasive computing technologies such as sensor networks and RFID tags will soon densely populate our everyday environments. These, together with the increasing diffusion of geospatial Web 2.0 tools such as GoogleEarth, will soon form the basis of a shared distributed information space capable of producing and storing data about the physical and social worlds and their processes. This opens up the possibility of exploiting such an information space as a general-purpose coordination infrastructure to facilitate users in gathering information about the world, interact with it in a context-aware way, and coordinate with each other via the mediation of that infrastructure. However, the extremely distributed and heterogeneous nature of such infrastructure and the potentially incredible density of the information produced within (at the very extreme, a spatio-temporal continuum of information), introduces several issues related to the management of such infrastructure, i.e., the need for properly aggregating data to abstract from its actual density and to enable multilevel views, and the need for representing data in a simple, uniform, and easy to be managed way. In this paper, after having sketched our general vision for such future coordination infrastructure, we analyse and discuss the key research challenges to data aggregation and data representation, and present our current research and experimental activity in these areas.

Keywords: Coordination, Pervasive Computing, Context-awareness, Sensor networks.

1 INTRODUCTION

Pervasive computing systems will increasingly populate our everyday environments [15]. RFID tags and alike will soon be attached to every, even small, objects and locations, carrying on digital information about such objects/location and enabling us to store in them additional information about [35]. Wireless sensor networks will be distributed around in our cities, houses, and offices, to monitor specific physical phenomena [10]. Devices such as mobile phones, PDAs, GPS and other localization devices [19], will be able to produce and store notable amounts of data related to our personal, social and professional activities.

In addition, one must also consider that the Web has rapidly evolved from a mere container of rather static information to a very dynamic shared information space. Web 2.0 technologies, and in particular geospatial tools such as GoogleEarth, makes it possible to dynamically make available a variety of information about facts and events occurring in the real-world [3, 8].

Both the above trends clearly contribute in building a sort of distributed “digital world model”, representing and making available a real-time

representation of our world and of its processes [28]. Even more important, the dense deployment of pervasive devices together with the availability of geographic information systems can form the basis of a globally distributed, *shared infrastructure*, for general-purpose context-aware coordination activities.

Let us assume that users and services can be given access to pervasive devices (i.e., computer-based sensors and RFID tags) around them for both reading and storing data, and that they are similarly given the possibility of accessing Web-based geospatial information tools (e.g., GoogleEarth). Then, this will make it possible to let users and services:

- *Act* in the world with improved sensorial capabilities, i.e., being able to perceive facts and events of their current context beyond those allowed by their inherent capabilities.
- *Interact* with the world, i.e., engage in actuating actions affecting the work around and store, other than simply consume, contextual-information (i.e., by storing them in nearby pervasive devices and/or by geo-locating it before storing on the Web).

- *Coordinate* with each other through the world, i.e., by exploiting the overall shared infrastructure as a sort of shared coordination medium for mediated (*à la Linda* [1]) data-oriented interactions, and in particular for context-aware, spatially-situated, interactions.

One important aspect of such a sort of shared distributed coordination infrastructure is that one cannot expect it to have a regular, definable, structure. On the one hand, one can envision that the mass deployment of sensor networks and tags will make such infrastructure so densely embedded in the physical world to make it worth abstracting it as a sort of “*sensing continuum*” able to produce (and store) infinitely dense and infinitely detailed amounts of information in space and time. On the other hand, one should consider that the actual density of such infrastructure can be varying and its structure very heterogeneous (including both different classes of situated pervasive devices and geospatial Web information spaces). Nonetheless, users and services should be enabled to interact with and coordinate through it according to their own specific needs and without being functionally affected by the actual characteristics of the infrastructure.

To tackle the above issue, we think that three key challenges have to be faced:

- The need of defining proper mechanisms for self-organized data aggregation, enabling users and services to access and exploit the infrastructure abstracting from its actual density and structure, and rather enabling multiple levels of observation and action.
- The need of proper models for data representation, supporting the above forms of aggregation and multi-level perception, and also enabling a uniform representation of contextual data coming from a variety of devices and geospatial tools.
- The need of verifying whether such mechanisms and models can actually translate into a usable, general-purpose, distributed framework for the support of context-aware coordination activities.

In this paper, after better having detailed our general vision of the pervasive coordination infrastructure and framed its key assumptions and abstractions in Section 2, we analyse in detail the above three research challenges and present our current research activity in that direction. In particular: in Section 3 we discuss the need of innovative approach to data aggregation, and present our current experience in that area; in Section 4, we discuss the need for innovative data models and present our proposal in this directions; in Section 5, we present the prototypical infrastructure we have developed to experience with the proposed approaches and some exemplary coordination services developed within it. Section 6 concludes the paper.

2 COORDINATION IN THE SENSORS' CONTINUUM: VISION AND CHALLENGES

2.1 Towards a Pervasive Coordination Infrastructure

From the operational viewpoint, our overall vision for the shared distributed coordination infrastructure is sketched in Figure 1.

Users are situated in some location of the physical world, densely embedded with pervasive devices such as sensors and tags. All such devices produce very detailed information about the physical world and its processes (at the very extreme, one could imagine a future when there will be nearly a *continuum* of information available for each point of the physical space). Such an information, possibly integrated with additional geo-located information that can be available on the Web, makes available to users a sort of local, contextual, digital real-time world model, which may also include information about the users themselves.

From their current locations, users with some portable computing devices such as PDAs (or which is the same, software services/agents executing on such devices) can be given access to such local view of the world model. Typically, this can occur by proper wireless technology mounted on the PDA that enables it to access and interact with in-range sensors and RFID tags, and that also enables accessing geo-located Web resources related to their current location via WiFi or other telecommunication means.




Figure 1: The shared coordination infrastructure.

Beside simply exploiting the infrastructure for accessing existing information about the world, which is important to enable them to *act* in the world with improved sensorial capabilities, users can also be enabled to *interact* with it in a more participatory way. This includes making it possible for users to dynamically configure such devices (and, if actuator devices exist, to enforce specific actuating actions) as well as to exploit the overall infrastructure as a sort of distributed shared data space.

Indeed, other than enabling reading information, devices such as computer-based sensors and RFID tags enable storing information in their memories. Thus, users and services can be allowed to

dynamically contribute to the overall building of the digital world model, by storing information on the distributed infrastructure, the same as they can post geo-located tags on existing geospatial web tools.

The possibility of considering the overall distributed infrastructure as a sort of shared dataspace also opens up the possibility of exploiting it as a coordination medium, i.e., as a sort of globally distributed Linda-like tuple space [1]. Via such space, users and services can be enabled to exchange information, synchronize their activities, and more in general to indirectly *coordinate* with each other. It is worth mentioning that such coordination can be either explicit (a user/service inserts information with the goal of informing another user/service) or implicit (the actions and processes of a user/service reflect in a change of the world situation and, thus, implicitly affect the contextual activities of other users/agents).

2.2 Issues and Challenges

The above description of the pervasive shared coordination infrastructure, although apparently very simple, indeed hides a number of inherent complex issues to be tackled. In particular, the key challenges all relate to translating the wide availability of information and devices into a usable and manageable world model and shared dataspace, facilitating its practical exploitation as a general-purpose coordination infrastructure by users and services.

As stated in the introduction, the overall infrastructure can be formed by a variety of diverse heterogeneous devices, deployed in an environment at variable and typically unpredictable density, and can also include geo-spatial information located at some centralized servers on the Internet. To better understand how such characteristics impact on the exploit ability of the overall infrastructure, let us refer to Figure 2.

At the lower level, we can assume that the physical environment is filled with a very dense, virtually continuum, pervasive network of sensors' devices. Also, we can assume that such a continuum of sensors will be able to continuously produce infinitely dense amounts of spatial information. Of course, no actual pervasive network will ever be continuous, and the amount of information produced (and possibly stored) will always be discrete in time and space. Nevertheless, users and services must be ready to exploit such infrastructure without knowing in advance what the density of the network and the amount of produced information are. Accordingly, the sensors' continuum is a useful abstraction to represent the extreme situation that should be accounted for.

Users and services, in the real world, are necessarily of limited rationality: the amount of information they can handle and produce is limited both in time and space. So, for the pervasive

infrastructure to be actually usable, one should consider limiting the amount of information available to users/services to that actually usable by them, whatever the actual density of the pervasive network and of the information produced. The upward arrows in Figure 2 represent this concept: starting from the continuum of data produced, some process of aggregation/pruning of data must be enforced to produce a level of granularity suitable to the actual characteristics of users/agents, i.e., to abstract the perception of the infrastructure and of its data via a sort of information overlay. Depending on the needs and characteristics of users and services, overlays at various levels of granularity can be produced.

Interestingly, such overlays could also not only act as a mere space of aggregation/pruning of information, but it could also be exploited as a place where to enrich information, as produced by the lower levels, and make it more semantically-expressive and knowledgeable by services. That is, such overlay could become a sort of knowledge overlay network [3].

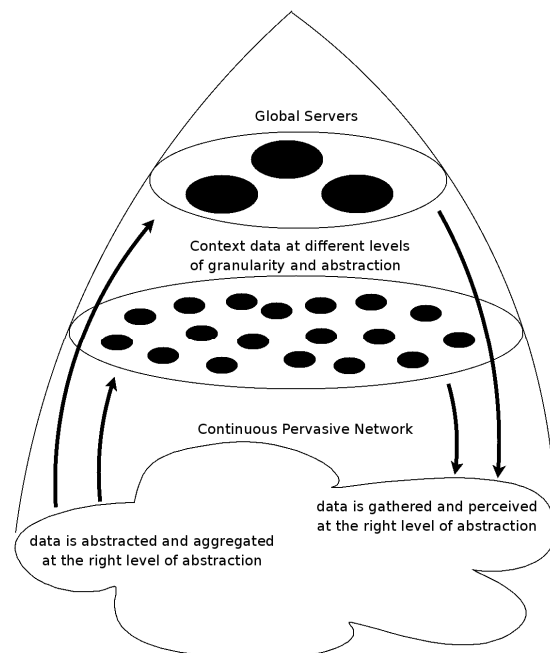


Figure 2: The Continuum Abstraction.

The downward arrows of Figure 2 are instead meant to represent a different concept. In the overall infrastructure, one should account for the presence of a variety of heterogeneous devices and tools, all of which can contribute to the overall building of the world model (and all of which possibly playing an active role in the coordination infrastructure). Just to make an example, data coming from a Web cam and geo-tags on Google Earth could both be fruitfully exploited to complement data coming from pervasive sensors and could contribute in reaching a better understanding of a situation. This integration, could be for example practically realized by visualizing in

an integrated way on a web-enabled application such as Google Earth environmental data coming from a sensor network, data coming from a camera, as well as Google Earth place marks.

The above process of integration can be of great use to make available more data (and, with reference to the shared infrastructure, more memory resources) than actually available from low-level sensors. In fact, beside the sensors' continuum vision, one should also account that there can be regions of an environment very scarcely populated by sensors, and that this should by no means affect the actual behavior of users/services, that should be able to proceed with the context-aware coordinated activities without being influenced by the actual structure and density of the surrounding infrastructure. Also, such process of integration can support the process of transforming data into knowledge, and in building proper knowledge overlay networks, other than mere information overlays.

All of these said, the three key challenges already outlined in the introduction can be better specified as follows:

- There is need of proper algorithms and approaches to data aggregation that, abstracting from the actual density of a sensing infrastructure, can be used to building expressive information overlays at different levels of granularity. Of course, due to the fully distributed and decentralized nature of the scenario, self-organized unsupervised approaches are a necessary choice.
- There is the need of identifying a uniform model for representing contextual data at different levels of granularity and coming from different devices and tools. Such data model should also be suitable to support expressive, semantically enriched, knowledge representations, and should be suitable to enforce mediated (tuple-based) coordination activities.
- There is need of experiencing in practice with the above concepts, to verify whether our general vision for the shared coordination infrastructure can translate into a usable, general-purpose, coordination framework, also supporting the proper accommodation of data aggregation and data integration approaches.

In the following of the paper, we will go into details about the three above challenges, by shortly over-viewing related activities in such areas and by presenting our current research in these directions.

3 DATA AGGREGATION

The main goal of a sensor network (and of the majority of pervasive computing systems) is to collect data from the environment and to suitably present them to application services and components. As already said, we believe that the number of sensors and devices embedded in our everyday

environments will quickly reach unforeseen densities. Accordingly, to avoid being overwhelmed with an unmanageable amount of data, techniques for in-network filtering, processing and aggregation of data must be provided.

3.1 Related Work

There are a lot of approaches trying to limit and simplify the data produced by sensors and increase network's lifetime.

The most common research line is focused on protocols that, due to their reactive nature, try to respond on demand to queries injected by nodes. If the answer is found in some region of the network, it is routed directly to the issuer node [23, 16, 17, 6, 31]. Due to aggregation and filtering, this kind of in-network data aggregation, can reduce both size and complexity of the data streams produced by the sensors.

On the other side, a new research line is starting to investigate mechanisms to allow direct access to sensor data by multiple and mobile users.

In most of the applications, users are not interested in punctual sensor readings, but in the overall status of an area. Accordingly, diffusive algorithms [11] and gossip-based aggregation algorithms [21, 14] have been proposed as simple yet very effective approaches to compute and make available at each node globally-aggregated values. However, in large scale scenarios, these values may lose significance and may not properly represent the status of the user region.

Due to this problem, the issue of recognizing regions characterised by similar data patterns is considered every day more important. In our approach, we borrow from [11, 21, 14] by exploiting aggregation algorithms that have a mixed diffusive-based and gossip-based inspiration. However, other than for computing global network values, we use them for computing regional aggregated values.

Indeed, some algorithms for self-organizing region partitioning similar to our have been proposed [9, 26]. The key differences with it are that: (i) these algorithms require a priori information about the typical patterns exhibited by the environment, while our approach does not and it is fully self-organizing; (ii) these algorithms are not conceived for other goals than recognizing regions, while our approach goes further, by exploiting regions as a basis for aggregation and for making regions act as sorts of virtual macro sensors.

3.2 Region Identification

Let us assume to sense, over a certain area covered by a sensor network, a particular property (see e.g. Figure 3a), and that each sensor is capable of measuring the local value v of such property. For instance, v could be a temperature, a light level, or whatever a sensor can measure. The goal of the region formation algorithm [4] is to have sensors

organize themselves into a disjoint set of regions each characterized by “similar” properties (e.g. see Figure 3b). Organization in regions occurs through a process of building an overlay of virtual weighted links between neighbor nodes, such that nodes belonging to the same region have strong links, while neighbor nodes belonging to different regions have weak (or null) links.




Figure 3: The scalar field of a physical property (a) and a sensor network organized over it (b).

Basically, our algorithm works as follows. Let s_i and s_j be two sensors. They can be considered neighbors if they are within the wireless radio range r . Define the values sensed by s_i and s_j as $v(s_i)$ and $v(s_j)$, and let us assume that generic distance function D can be defined for couples of v values. Region formation is then based on iteratively computing the value of the logical link $l(s_i, s_j)$ for each and every node of the system, as in the following *Update_link* procedure:

```
Update_link(other_node_data):
  if  $D(v(s_i), v(s_j)) < T$  {
     $l(s_i, s_j) = \min(l(s_i, s_j) + \delta, 1)$ 
  } else {
     $l(s_i, s_j) = \max(l(s_i, s_j) - \delta, 0)$ 
  }
```

Table 1: *Update_link* procedure. Two nodes increase the weight of their link if the distance D between their values is lower than a threshold T .

Where: T is a threshold that determines whether the measured values are close enough for $l(s_i, s_j)$ to be reinforced or, otherwise, weakened, and δ is a value affecting the reactivity of the algorithm in updating link.

The distributed execution of the algorithm is based on a sort of gossip scheme [24]: each node periodically wakes up, broadcasts a message to its neighbors, exchange with them the needed data, and then executes the *Update_link* procedure for each of them (See Table 1). After some iterations, nodes belonging to the same region will share a link with a value close to 1. The regions formed can then be used to aggregate data coming from nodes belonging to each of them. In the next paragraph this aspect will be over viewed.

Considering that the number of nodes executing the algorithm is known and that the frequency of execution is parameterizable, it is rather clear that our algorithm tends to impose an uniform load, on

the system. What really matters is that users, either human or digital, can select the most suitable trade-off between accuracy of the data and the energy consumption of the network producing them.

3.3 Per-Region Aggregation

As already introduced, the local availability of globally aggregated data may be useful independently of regions. However, in large scale scenarios, these values give very few details on the status of the network and are definitely of little use for users wishing to acquire information about the environment around themselves. For this reason, our approach also exploits per-region aggregation algorithms.

By considering the situation in which regions are already formed, computing aggregation functions is simply realized by executing the gossip-based algorithm proposed in [21], not only on the whole network, but also between nodes that are in the same region.

```
Main:
Do_forever:
  Sleep( $t$ );
   $neigh[] = Send\_Data\_Exchange\_Request();$ 
  For each( $neigh[]$ )
     $link\_data, global\_data, region\_data =$ 
      Exchange_data();
    Update_link( $link\_data$ );
    Aggregation( $global\_data$ ); // if needed
    /* connected is true is  $l(s_i, s_j) > 0.8$  */
    If( $connected$ ) Aggregation( $region\_data$ );
Done

Aggregation(other_node_data)
  if( $Min_i > Min_j$ )  $Min_i = Min_j$ 
  if( $Max_i < Max_j$ )  $Max_i = Max_j$ 
  /* other distributed aggregation functions */
```

Table 2: Algorithm *main body* and *Aggregation* function.

In this section we introduced and described a possible approach to make manageable the huge amount of data that we believe will be produced in the near future. Starting from a sensor continuum (see lower level of Figure 2), we provide users and services with a simplified view of the underlying sensor network in terms of a cluster of virtual macro sensors (see middle level of Figure 2). defined algorithms and tools, it is mandatory to define a general data model for representing contextual data.

4 DATA MODELS

An effective data model for contextual information should enable to uniformly represent very diverse facts and situations about the world, generated by a variety of heterogeneous sources and at different levels of granularity. Given the intrinsically

distributed and decentralized nature of contextual information that can be provided by the pervasive infrastructure, access to such data/knowledge by services and users, there include generation of new knowledge by users and services (as required to enforce data-mediated coordination) should abstract from the actual distribution and allocation of data and knowledge. Also, of course, such data model should enable representing data at different levels of granularity, and should accommodate the integration with data aggregation algorithms, as described in the previous section.

In general terms, the identification of a data model should consider that data items should form a sort of “live” overlay where new information and knowledge is continuously generated, analyzed, aggregated, related with existing information, and makes it possible to extract meaningful knowledge from the available data in multiple knowledge views.

4.1 Related Work

In the past few years, a lot of research works in the area of pervasive and ubiquitous computing has focused on identifying proper models for context-aware information.

Early works in this area [30, 13], concentrate on the issue of acquiring and making available context data, but lack in focusing on the identification of a uniform model and a common semantic. Other researchers [29, 18] have focused more on the issue of providing rich data models for contextual information and of facilitating querying by services. Such approaches aim at developing a general and flexible data model, however they result in a long list of elements/variables, and lack that uniform representation and that ease of manipulation that is instead needed for the envisioned scenario.

Some recent proposals [36, 7] suggest simple models for contextual information, structured around well-defined data fields to properly represent specific aspects of the context. While we consider these as very interesting approaches, the proposed structure appears too much application-specific to be adopted as general-purpose solutions.

An increasing number of research works get inspiration from tuple space coordination models and suggest that coordination in pervasive environment could take advantage from tuple-based coordination and from representing contextual information in the form of unstructured tuples [22, 20]. Interestingly, such proposals also accounts for the possibility of providing application-specific views.

4.2 Data Representation

None of the analyzed proposals effectively tackle all the needs for an effective context model. Our proposal of the W4 contextual data model tries to couple the advantages of structured contextual data model and of tuple-based coordination models.

The W4 model considers that most world facts can be expressed by means of a simple yet

expressive 4-fields tuple (Who, What, Where, When): “someone or something (Who) does/did some activity (What) in a certain place (Where) at a specific time (When)”. We also call such tuples knowledge atoms, as they are atomic units of factual knowledge.

The four-fields (Who, What, Where, When) of the W4 data model each describes a different aspect of a fact.

The Who field associates a subject - human person or an unanimated part of the context - to a fact. It is represented by a type-value pair, in the form of a string, with an associated namespace that defines the “type” of the entity that is represented. E.g., “person:Gabriella”, “tag:tag#567”.

The What field describes the activity performed by the subject. This information can either come directly from the data source (e.g., a sensor), or be inferred from other context parameters or it can be explicitly supplied by the user. This field is represented as a string containing a predicate:complement statement. E.g.: “work:pervasive group”, “read:temperature=23”.

The Where field associates a location to the fact. The location may be a physical point represented by its (lat, lng), a geographic region, or a logical place. In addition, context-dependent spatial expressions like “here” can be used for context-aware querying.

The When field associates a time or a time range to a fact. Also in this case, context-dependent expressions can be defined (e.g., “now”, “today”) and can be used for context-dependent querying.

The W4 model can provide a uniform representation of data coming from different devices, such as sensors, geo-located tags, cameras, etc., and is capable of expressing different level of observation (e.g., a single light sensor Vs the whole sensor network).

In the rest of Section 4, we describe the W4 Model from a theoretical perspective. Implementation details are discussed in Section 5.2.

4.3 Data Generation

The W4 model relies on the reasonable assumption that software agents are associated with data sources and are in charge of creating W4 tuples and inserting them in some sort of shared data spaces. W4 agents have the goal of collecting all the necessary information to produce a W4 tuple which is as accurate and complete as possible. This occurs by sensing and inferring information from all the devices and sources available (e.g., sensors, RFID tags, GPS devices, Web services), and by combining them in a W4 tuple. Let us make two simple examples to clarify this concept.

Let us assume Gabriella is walking in the campus park. An agent running on her GPS-equipped PDA can periodically (e.g., every n seconds) create the following tuple:

```
Who: user:Gabriella
What: walk:4kmh
Where: lonY, latX
When: 2006/10/17:10.53am
```

There, the Who is entered implicitly by the user at the login, What and Where can be derived by the GPS (e.g., the speed of Gabriella can be used to deduce that she is walking), When can be provided both by the PDA or by the GPS.

Similarly, the algorithm presented in Section 3 can produce W4 tuples describing the current situation of an identified spatial region:

```
Who: macrosensor: #123
What: temperature: 25 °C
Where: region: #XYZ
When: 2006/10/17:10.53am
```

4.4 Data Manipulation

Knowledge atoms are stored as W4 tuples in a shared information space (or in multiple information spaces). To interact with them, we take inspiration from tuple space approaches [1] and define the following API:

```
void inject(KnowledgeAtom a);
KnowledgeAtom[] read(KnowledgeAtom a);
```

The inject operation is trivial: a service accesses the shared data space to store a W4 tuple there.

The read operation is used to retrieve tuples from the data space via querying. A query is represented in its turn as a W4 tuple with some unspecified or only partly specified values (i.e., a template tuple). Pattern matching operations work rather differently from the traditional tuple space model: the pattern matching mechanism allows for matches between incomplete information, and we exploit differentiated pattern matching mechanisms for the various fields. Current mechanisms work as follows:

- Who and What. Pattern-matching operations in these fields are based on string-based regular expressions. E.g., “user:*” will match any user.
- Where. We currently adopt the PostGIS language.
- When. In this case, the template defines a time interval. Everything that happened within that interval matches the template. Concise time descriptions as well as contextual ones (e.g., “now” or “before”) are converted into actual time intervals before pattern matching.

Let us assume that a user wants to know the average temperature of a region, identified by the aggregation algorithm described in section 3 running over a sensor network. He will ask (via a read operation):

```
Who: macrosensor: *
What: temperature: *
Where: region #XYZ
When: now
```

Then he will get in return a tuple representing the average temperature currently in region #XYZ.

4.5 W4 Knowledge Network

Although pattern-matching techniques proved rather flexible to retrieve context information, our idea is to exploit the W4 structure to access the context repository in a semantically enriched fashion. More specifically, we are currently working towards integrating in our proposal general-purpose mechanisms and policies to link together knowledge atoms, and thus forms W4 knowledge networks in which it will be possible to navigate from a W4 tuple to the others. Our idea is that, by querying such a knowledge network instead of a flat W4 tuple space, one can obtain much higher knowledge, as resulting from the analysis, manipulation and inference upon the link structure of the knowledge network.

To this end, we adopt an algorithmic approach which relies on a two-phase process. The first step is the identification of all possible correlations between knowledge atoms, and the creation of link between W4 atoms. The second step is the generation of new knowledge atoms, by analyzing which of the identified link can lead to a new W4 atom as a process of merging related atoms. The idea is to support self-aggregation algorithms similar to that of Section 3. While those are specific for spatial regions, in Knowledge Networks we try to use the same basic approach to aggregate and organize W4 tuples based on each of the W field.

5 INFRASTRUCTURES AND SERVICES

The algorithms and models proposed so far needs to be make running by a proper hardware/software infrastructure to show that the overall vision makes sense and can translate into usable user services.

5.1 Related Work

Our vision for a pervasive shared coordination for supporting context-aware activities is, of course, shared with the research community on pervasive computing. Consequently, other proposals exist with similar objectives.

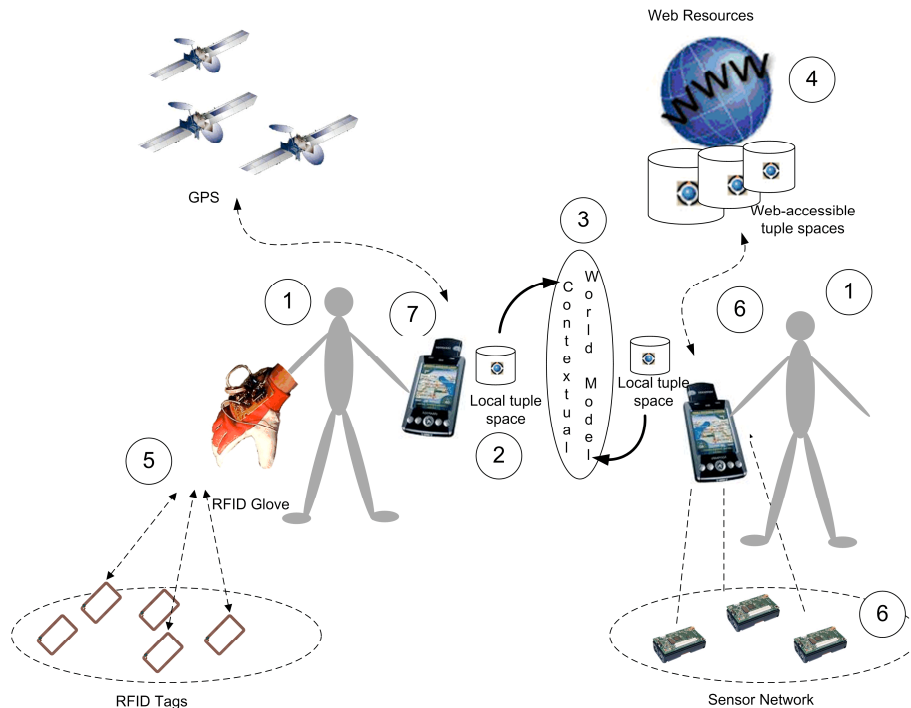


Figure 4: Schematic System Overview.

A large number of proposals consider the possibility of exploiting user-generated geo-located information to facilitate context-aware interactions among users. For example, MapWiki [34] defines a Web-based collaborative environment for spreading by users contents on a shared map in a ubiquitous and location-dependent way. In this way, users in the same location can share comments and impressions on that location. Unlike our approach, MapWiki (and the many similar systems being frequently proposed in this period) relies on a centralized server-based approach on which users can upload their own data and geo-located annotations. They do not consider localized and server-less interactions. Also, they do not consider exploiting the presence of embedded pervasive computing devices to gather and generate information about the surrounding environment.

Other proposals, instead, put the focus on exploiting pervasive devices to facilitate the interactions of a user with its surrounding environments. This translates in proposals for novel user interfaces [33], novel mechanisms to “query the world” around [8], and novel devices to improve such interactions with the physical world [32]. However, these proposals lack in identifying the possibility of exploiting such pervasive infrastructure as a general-purpose coordination infrastructure.

There exist some approaches that focus on tuple-based infrastructures to facilitate users in gathering information from an environment enriched with pervasive computing devices and coordinating through them. TinyLime [12] proposes a tuple-based model to access information from surrounding sensors and supporting ad-hoc sharing of information

between close users. In our research group, we have previously implemented a pervasive tuple-based distributed memory realized by making use of RFID tags and to be accessed by mobile users [27]. Both these proposals go in the direction of exploiting pervasive devices as a sort of shared coordination infrastructure. However, they focus on specific classes of technologies, rather than considering general pervasive devices. Also, they are not based on any well-structured context model as W4 is, and do not consider the exploiting information from sources such as the Web.

5.2 The Infrastructure

The general hardware/software infrastructure we have prototyped enables accessing to a distributed space of contextual information and coordinating through it, and it includes tools for acquiring, storing, organizing, and aggregating data in form of W4 tuples.

The implemented infrastructure is organized as follows (see also Figure 4):

1. Putting humans at the centre, our architecture considers users with portable computing devices (i.e., laptops or PDAs), integrating localization devices (i.e., GPS), devices to acquire information from the world (i.e., RFID readers and sensors), and an Internet connection (via WiFi).
2. Contextual information about the world, including user data, data coming from pervasive devices, and more generally data available on the Web, is represented by means of W4 tuples and stored into a local tuple space (locally

accessible from laptop and PDA) to be later accessed by application services executing on the PDA or by other users in the proximity. The implementation is based upon the LightTS tuple space [2] that is a Java light-weight implementation of the Linda Tuple Space, particularly suitable for context-aware application and easily extensible. Currently, data in each W field are represented as simple strings, and basic pattern matching operations are defined to retrieve tuples.

3. The content of local tuple spaces for user in close proximity (as determined by the range of their WiFi ad-hoc connectivity) is the same for all such users. A content replicator system (as described in [5]) ensure that all W4 tuple dynamically gathered/stored by a user is replicated on the tuple spaces of the connected users. In this way, users can exploit their local tuple spaces as a shared media for coordinating with the other users around, as in the shared "Contextual View of the World Model" of Figure 1. We are perfectly aware that full replication is not the optimal solution, and we are planning to experience with more effective solutions for sharing information (i.e., P2P overlays). Nevertheless, the actual solution is enough to test with our concepts.
4. A number of additional remote Web-accessible tuple spaces can be also used to store/retrieve W4 information. Thus, services can dynamically connect to the Web to retrieve additional information to integrate with that coming from the W4 local tuple spaces. Also, each space could host specific information and algorithms for the management and aggregation of data.
5. A RFID reader (in the form of a wearable glove) connected to the laptop or to a generic PDA via a serial interface can be used to collect information from RFID tags dispersed in the environment. This information, enriched with the physical location where it has been collected (provided by the GPS device) is stored in the local tuple space.
6. Data coming from sensor network nodes can be directly accessed and stored in the local tuple space and converted in the form of W4 tuples possibly pre aggregated, based on an in-network aggregation algorithms as those described in Section 3.
7. Specific services can be realized by means of application agents running locally on a PDA and accessing, via the W4 model, both local and remote tuple spaces. Also, application can interface themselves with a graphical interface (Google Earth or Google Maps) to make data immediately understandable.

The current implementation is merely a proof-of-concept and rather simple. Despite a number of limitations it has been of great use to experience with

our concepts and to give a concrete meaning to our vision, though in laboratory experiments and under low data load. We are currently performing extensive experiments to test the effectiveness and robustness of the infrastructure under high data loads and in actual distributed environment.

5.3 Services

To test our infrastructure and evaluate its strength and expressiveness, in a setting characterized by specific conditions, we designed and implemented the following services. They show how easy we obtain context-aware services exploiting our infrastructure.

A first application we have experimented has the goal of providing real-time and historical information to a user equipped with a GPS device and a RFID reader. In particular, we focused on the scenario in which a tourist wants to automatically build and maintain a diary of his journey.

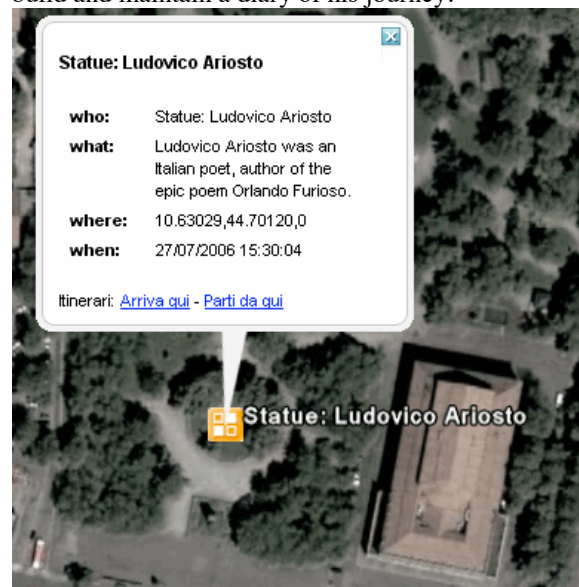


Figure 5: A monument embedding a RFID tag.

The proposed service allows keeping track of all user movements and have them displayed on the map of the visited places. Simply, application agents can periodically (e.g., each minute) store in the tuple space the position of a user. Later on, the user can simply retrieve the history of his journey by having the agent querying the tuple space. Furthermore, the support for RFID allows tourists to access tourist information stored in RFID tags attached to monuments (Figure 5) and art-pieces.

As a second example, an user equipped with a GPS device can decide to share his location with other users and, analogously, he may wish to be aware of the location of others users.

This can happen either by uploading W4 tuples to a Web-accessible repository, or by exchanging them in ad-hoc way and storing them in the local tuple space only. Either way, collected W4 tuples can be used to display users' locations on a real-time

map (which can also highlight other interesting Web-retrieved information for the group, such as museums or hotels, depending on the specific interests of the group). In addition, starting from a people map, we can implement coordination services that, working on reciprocal user positions, can produce coordinated movements, acts and behaviours.

Regarding the coordination capabilities of our approach, we can imagine a system to let mobile users to coordinate themselves into crowded places in order to avoid bottlenecks (see Figure 6). Mobile users are perceived (using Bluetooth, Wi-Fi, RFID tags, cameras) by the deployed sensor network. Sensors aggregate their values (containing the number of perceived users in their sensing range) and provide users with a shared tuple space containing those data. By this way, mobile devices are made aware of over crowded rooms and are able to suggest to the users alternative routes. This is only a brief example to highlight some coordination results achievable with our approach.



Figure 6: In black very crowded area, in grey low crowded area, in white available space.

6 CONCLUSIONS

Pervasive computing technologies, together with the increasing diffusion of geospatial Web tools, can contribute forming the basis of a shared distributed coordination infrastructure to support the context-aware activities of users and services.

In this paper, we have identified the key challenges that may arise towards the practical exploitation of such infrastructure, mostly related to its distributed and heterogeneous nature and to the potentially incredible density of the information produced within (at the very extreme, a spatio-temporal continuum of information). In particular, such characteristics introduces issues related: (i) the need for properly aggregating data to abstract from

its actual density and to enable multilevel views; (ii) the need for representing data in a simple, uniform, and easy to be managed way; (iii) the need for translating the abstract architectural vision into a usable prototype infrastructure.

Our current research work focused on facing the above three challenges, as presented in this paper, is far from being complete. First, we are continuing our research on self-organized data aggregation algorithms, with the goal of extending them to make them more general and adaptive. Second, we are currently studying how to integrate ontology and ontological reasoning into the W4 data model. As our current implementation is merely a proof-of-concept, we are using a sort of “informal ontology”, but we are aware that for a full implementation the exploitation of ontology is mandatory. We plan to use ontology to describe data in each W field, in our opinion the exploitation of the W structure can be beneficial to improve the typical semantic reasoning mechanism of ontology. We are also evaluating the adoption of a standard language to describe data, such as RDF ([25]). Third, we are continuing experiencing with our prototype infrastructure, in particular with regard to the integration of further devices and to the study of proper tuple distribution strategies.

Acknowledgements: Work supported by the integrated project CASCADAS (IST-027807), funded under the FET Program of the European Commission.

7 REFERENCES

- [1] S. Ahuja, N. Carriero, D. Gelernter, “Linda and Friends”, IEEE Computer, 19(8):26-34, 1986.
- [2] D. Balzarotti, P. Costa, G.P. Picco, “The LighTS Tuple Space Framework and its Customization for Context-Aware Applications”, International Journal on Web Intelligence and Agent Systems, 5(2): 215-231, 2007.
- [3] M. Baumgarten, N. Biccocchi, R. Kusber, M. Mulvenna, F. Zambonelli, “Self-organizing Knowledge Networks for Pervasive Autonomic Communication Services”, IEEE Conference on Systems, Man, and Cybernetics, Montreal (CA), October 2007.
- [4] N. Biccocchi, M. Mamei, F. Zambonelli, “Self-organizing Spatial Regions for Sensor Network Infrastructures”, 21st IEEE Conference on Advanced Information Networking and Applications, 2007.
- [5] M. Bortenschlager, G. Castelli, A. Rosi, and F. Zambonelli, “A Context-sensitive Infrastructure for Coordinating Agents in Ubiquitous Environments”, Special Issue on Engineering Environments for Multiagent Systems of International Journal on Multiagent and Grid Systems, Eds. D. Weyns and A. Omicini, 2008.
- [6] A. Boulis, S. Ganerival, and M. B. Srivastava,

- “Aggregation in sensor network: An Energy-Accuracy Trade-off”, Proceedings of IEEE SANPA, May 2003.
- [7] J. Bravo, R. Hervás, I. Sánchez, G. Chavira, S. Nava, “Visualization Services in a Conference Context: An approach by RFID Technology”, *Journal of Universal Computer Science*, 12(3): 270-283, 2006.
- [8] G. Castelli, A. Rosi, M. Mamei, F. Zambonelli, “A Simple Model and Infrastructure for Context-aware Browsing of the World”, 5th IEEE Conference on Pervasive Computing and Communication, New York (NY), March 2007.
- [9] E. Catterall, K. Laerhoven, M. Strohbach, “Self-Organization in Ad-Hoc Sensor Networks: An Empirical Study”, *International Conference on Simulation and Synthesis of Living Systems*, Sydney (AU), 2002.
- [10] C.-Y. Chong, S. Kumar, “Sensor Networks: Evolution, Opportunities, Challenges”, *Proc. of the IEEE*, 91(8), 2003.
- [11] A. Corradi, L. Leonardi, F. Zambonelli “Diffusive Load Balancing Policies for Dynamic Applications”, *IEEE Concurrency*, 7(11):22-31, 1999.
- [12] C. Curino, M. Giani, M. Giorgetta, A. Giusti, A.L. Murphy, G. Picco, “Mobile Data Collection in Sensor Networks: The TinyLIME Middleware”; *Journal of Pervasive and Mobile Computing*, 1(4):446:469, 2005.
- [13] A. K. Dey, G. D. Abowd, D. Salber, “A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-aware Applications”, *Human-Computer Interaction*, Taylor and Francis Ltd, 16(2-4):97-166, 2001.
- [14] A. G. Dimakis, A. D. Sarwate, M. J. Wainwright, “Geographic Gossip: Efficient Aggregation for Sensor Networks”, *Proceedings of the International Conference on Information Processing in Sensor Networks*, Nashville (TN), ACM Press, April 2006.
- [15] D. Estrin, D. Culler, K. Pister, G. Sukjatme, “Connecting the Physical World with Pervasive Networks”, *IEEE Pervasive Computing*, 1(1), 2002.
- [16] Johannes Gehrke, Samuel Madden, “Query Processing In Sensor Networks”, *IEEE Pervasive Computer Journal*, April 2004.
- [17] S. Goel and T. Imielinski, “Prediction-based Monitoring in Sensor Networks: Taking Lessons from MPEG”, *ACM Computer Communication Review*, Vol. 31, No. 5, October 2001.
- [18] K. Henriksen, J. Indulska, A. Rakotonirainy, “Developing Context-aware Pervasive Computing Applications: Models and Approach”, *Journal of Pervasive and Mobile Computing*, Elsevier, 2(1):37-64, 2006.
- [19] J. Hightower, G. Borriello, “Location Systems for Ubiquitous Computing”, *IEEE Computer*, 34(8), 2001.
- [20] J. Hong., “The Context Fabric: An Infrastructure for Context-Aware Computing”, *Conference on Computer Human Interaction*, Minneapolis (MN), USA, 2002.
- [21] M. Jelasity, A. Montresor, O. Babaoglu, “Gossip-based Aggregation in Large Dynamic Networks”, *ACM Transactions on Computer Systems*, 23(3): 219-252, 2005.
- [22] C. Julien, G. Roman, “EgoSpaces: Facilitating Rapid Development of Context-aware Mobile Applications”, *IEEE Transactions on Software Engineering*, IEEE CS Press, 32(5):281-298, 2006.
- [23] S. Madden, J. M. Hellerstein, “Distributing queries over low-power wireless sensor networks”, *Proceedings of the 2002 ACM SIGMOD International Conference on Management of Data*, June 2002.
- [24] A. Montresor, M. Jelasity, and O. Babaoglu, “Gossip-based Aggregation in Large Dynamic Networks”, *ACM Transactions on Computer Systems*, 23(3): 219-252, 2005.
- [25] L. Nixon, J. B., R. Tolksdorf; A. Wood, R. Menezes, “Coordinating Knowledge in Pervasive Environments”, *16th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises*, WETICE 2007., Paris (FR), 2007.
- [26] A. Panangadan, G. S. Sukhatme, “Data Segmentation for Region Detection in a Sensor Network”, *Center for Robotics and Embedded Systems*, University of Southern California, Technical Report 05-005, 2005.
- [27] M. Mamei, R. Quagliari, F. Zambonelli, “Making Tuple Spaces Physical with RFID Tags”, *21st ACM Symposium on Applied Computing*, Dijon (F), 2006.
- [28] D. Roush, “Second Earth”, *MIT Technology Review*, July/August 2007.
- [29] B. Schilit, N. Adams, and R. Want. “Context-Aware Computing Applications”, *Workshop on Mobile Computing Systems and Applications*, Lake District, United Kingdom, 1994.
- [30] A. Schmidt, K. A. Aidoo, A. Takaluoma, U. Tuomela, K. Van Laerhoven, W. Van de Velde, “Advanced Interaction in Context”, *Symposium on Handheld and Ubiquitous Computing*, Karlsruhe, Germany, 1999.
- [31] T. Schoellhammer, B. Greenstein, D. Estrin, “Hyper: A Routing Protocol To Support Mobile Users of Sensor Networks”, *Proceedings of the MobiHoc 2006 Conference*, Florence (I), May 2006.
- [32] R. Siagri, “Pervasive Computers and the GRID: the Birth of a Computational Exoskeleton for Augmented Reality”, *International Conference on the Foundations of Software Engineering*, ACM Press, Dubrovnik (CRO), Sept. 2007.

- [33] R. Simon, P. Forhlich, "A Mobile Application Framework for the Geospatial Web", The 2007 International World Wide Web Conference, ACM Press, Banff (CA), May 2007.
- [34] Y. Teranishi, J. Kamahara, S. Shimojo, "MapWiki: A Ubiquitous Collaboration Environment on Shared Maps", 6th International Symposium on Applications and the Internet Workshops, Phoenix (AZ), 2006.
- [35] R. Want, "An Introduction to RFID Technology", IEEE Pervasive Computing, 5(1):25-33, 2006.
- [36] C. Xu, S. C. Cheung, "Inconsistency Detection and Resolution for Context-aware Middleware Support", Symposium on the Foundations of Software Engineering, Lisbon, Portugal, 2005.