

A Self-Organizing Approach for Building and Maintaining Knowledge Networks

Gabriella Castelli, Marco Mamei, Franco Zambonelli

DISMI - Dipartimento di Scienze e Metodi dell'Ingegneria
University of Modena and Reggio Emilia
Via Amendola 2 , 42100 Reggio Emilia, Italy
`{gabriella.castelli,marco.mamei,franco.zambonelli}@unimore.it`

Abstract. Pervasive and mobile devices can generate huge amounts of contextual data, from which knowledge about situations occurring in the world can be inferred for the use of pervasive services. Due to the overwhelming amount of data and the distributed and dynamic nature of pervasive systems, this may be not a trivial task, indeed the management of contextual data should be run by a dedicate middleware layer, i.e., knowledge networks in charge of organizing and aggregate such data to facilitate its exploitation by pervasive services. In this paper we introduce a unsupervised, distributed and self-organizing approach to build and maintain such a layer based on simple agents that organize and extract useful information from the data space. We also also present a java-based implementation of the approach and discuss experimental results.

Key words: Distributed Middleware, Knowledge Networks, Context Awareness

1 Introduction

Pervasive and mobile devices and Web 2.0 are already able to generate an overwhelming amount of data about users and context, from which knowledge about situations and facts occurring in the world should be inferred for the use of pervasive and mobile services. A service in need of understanding what is happening around can access the produced pieces of data and analyze them to understand what is the current situation of its context. Nevertheless there are a number of complexities inherent in this process, such as the communication efforts to retrieve the useful knowledge out of an overwhelming amount of data, and the computational efforts to analyze, relate and aggregate such information.

Accordingly, a real challenge for future pervasive services is the investigation of principles, algorithms, and middleware infrastructures, via which this growing amount of distributed information can be properly represented, organized, aggregated, and made more meaningful, so as to facilitate the successful retrieval by pervasive services [7, 3]. Many approaches [10, 4] are currently going in the direction of adopting specific middleware layer, i.e., a knowledge network layer embedding data and knowledge network algorithms and providing effective access to such data by services.

Unlike traditional approaches (e.g., centralized and/or deterministic) to data organization and aggregation, we propose a distributed self-organized approach to organize, link, and aggregate, related items of contextual information. This choice better suit the decentralized, dynamic, and unpredictable nature of pervasive and mobile systems. In particular, in this paper:

1. We present a middleware architecture and prototype to store and manage contextual data coming from diverse pervasive devices structured according the W4 Data Model [9]
2. We present an original self-organized algorithmic approach to perform knowledge networking over a massive amount of distributed pieces of knowledge stored in the above middleware
3. We show a several experiments we performed to test the system and discuss evaluation results

The remainder of this paper is organized as follows. Section 2 briefly summarizes the W4 Data Model that is used to represent contextual data provided by pervasive devices. In Section 3 introduce the W4 Knowledge Networks idea, describe the algorithmic approach to organize isolated and distributed pieces of paper into networks of correlated data items and then present the overall middleware architecture. Section 4 presents the implemented prototype and the evaluation results. Section 5 discusses related work, and Section 6 concludes.

2 The W4 Approach

We adopt the W4 Data Model to represent and structure the data to to illustrate our self-organizing approach for building and maintaining knowledge networks. The W4 Data Model has been firstly described in [9]. Here we shortly summarize its key features and give an overall architectural view of the W4 Middleware.

2.1 The W4 Data Model

The proposed w4 model starts from the consideration that any elementary data atoms as well as any higher-level piece of contextual knowledge, in the end, represents a fact which has occurred. Such facts can be expressed by means of a simple yet expressive four-fields tuples (Who, What, Where, When): “*someone or something (Who) does/did some activity (What) in a certain place (Where) at a specific time (When)*”.

More in particular the four-fields of the W4 data model each describes a different aspect of a contextual fact:

- The *Who* field associates a subject to a fact. The *Who* field is represented by a type-value pair.
- The *What* field describes the activity performed by the subject. It is represented as a string containing a predicate:complement statement.

- The *Where* field associates a location to the fact. In our model the location may be a physical point represented by its coordinates, a geographic region, or it can also be a place label.
- The *When* field associates a time or a time range to a fact. This may be an exact time/time range or a context-dependent expression, e.g., *now*.

The way it structures and organizes information makes the W4 data model able to represent data coming from very heterogeneous sources and simple enough to promote ease of management and processing (although we are perfectly aware that it cannot capture each and every aspect of context, as freshness of data, reliability, access control, etc).

2.2 The W4 API and Architecture

In the W4 model, we rely on the reasonable assumption that software drivers are associated with data sources and are in charge of creating W4 tuples and inserting them in some sorts of shared data spaces.

Knowledge atoms are stored in the form of W4 tuples in distributed data spaces. The interface to access the W4 middleware took inspiration from tuple-space approaches [1] and consists in two basic operation:

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

The *inject* operation is equivalent to a tuple space "out" operation: an agent 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 template tuple. Upon invocation, the read operation triggers a pattern-matching procedure between the template and the W4 tuples that already populate the data space. Pattern-matching operations work rather differently from the traditional tuple space model and may exploit differentiated mechanisms for the various W4 fields. In [7] we provide several examples of knowledge representation and knowledge generation using the W4 Data Model.

Figure 1 depicts the overall architecture of a W4 system. At the bottom there are the diverse data sources that produce data formatted according the W4 Data Model and feed a number of W4 tuple spaces. The W4 tuple spaces are part of the W4 system, in the system there are a variety of agents that can be in execution:

- *Spiders*: that access the tuples in the W4 tuple spaces and are able to jump from a tuple space to another and link tuples that are correlated
- *Browsers*: that can browse a knowledge network to solve a query and to infer new tuples

Those agents (i.e., spiders and browsers) and the algorithms to create and manage knowledge networks will be discussed in Section 3. Many W4 knowledge networks can be realized and coexist in the W4 system, each realizing a specific

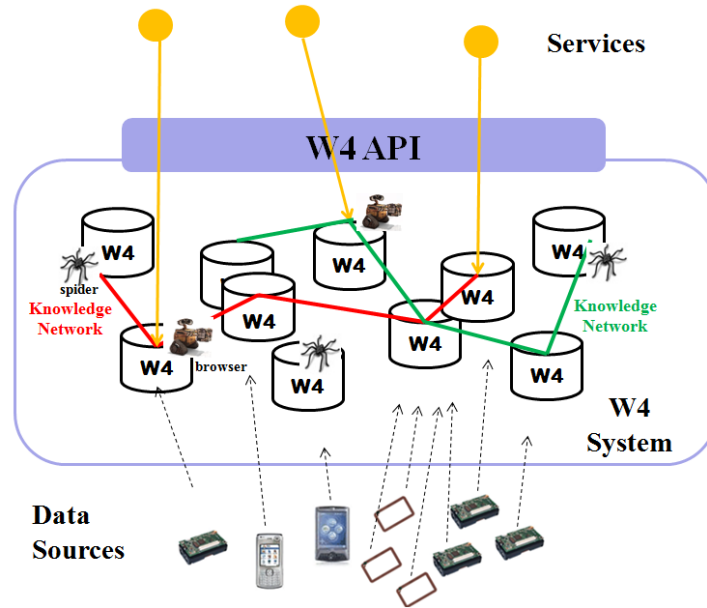


Fig. 1. The W4 System Architecture

view over the data. Eventually, at the top there are the various services that access the W4 system to submit query and retrieve data, to whom the internal W4 system and data location are completely transparent. Indeed they can act over the system via the W4 API.

We developed a prototype implementation of the described architecture in a small pervasive computing testbed by extending the LighTS Tuple Space [2], a light weight tuple space implementation particularly suitable for context-aware application, and by realizing spiders and browsers as simple Java agents. The implemented W4 middleware runs on laptops and on PDAs equipped with wireless interface and J2ME-CDC (Personal Profile) Java virtual machine.

3 W4 Knowledge Networks

Although the W4 data model proved to be rather flexible to manage contextual data, the idea is to exploit the W4 structure to access and exploit distributed contextual data in a more sophisticated and effective way. More specifically we propose general-purpose mechanisms and policies to link together knowledge atoms, and thus form W4 knowledge networks in which it will be possible to navigate from a W4 tuple to the others. Moreover, new information could be produced combine and aggregating existing tuples while navigating the space of W4 tuples.

This paper focuses on the realization of W4 knowledge networks, whose basic idea has been anticipated also in [9] and [7], but that only in this paper is eventually realized and evaluated.

3.1 The W4 Knowledge Networks Idea

The W4 knowledge networks approach is based on the consideration that a relationship between knowledge atoms can be detected by a relationship (a pattern-matching) between the information contained in the atoms fields. In particular, for the W4 model, we can identify two types of pattern matching correlations between knowledge atoms:

- *Same value – same field*: We can link together W4 tuples belonging to the same user, about the same place, activity or time (or, more in general, those W4 tuples in which the values in the same field match according to some pattern-matching function). Matching two or more *same value – same field* relationships, we can render complex concepts related to groups of W4 tuples, e.g. *All students (same subject) who are attending a class (same activity) at the same room (same location)*.
- *Same value – different field*: We can link atoms in which the same information appears in different fields. This kind of pattern matching can be used for augmenting the expressive level of the information contained in the W4 tuples. For example, a knowledge atom having *When: 18/09/2008* can be linked with another atom like *Who: Fall Class Begin , When: 18/09/2008* to add semantic information to that date.

Exploiting those correlations it is possible to find the relationships between one particular W4 tuple with other tuples in the data space, which may then be used to create a web of linked information both to more effectively navigate in the space of information (e.g., for effectively gathering information correlated to a specific context) and as a basis for more elaborated inference and reasoning (e.g., for representing in a comprehensive and expressive way complex situations).

For instance (see also Fig. 2), suppose that Gabriella's PDA, at a certain time, creates the following tuple: (student:Gabriella, - , room:room.35, 01/09/2008 10:05 am), where - means an empty field. A correlation can be found with the following tuple: (class: Computer Foundation, - , room: room.35, 01/09/2008 10:00-12:00 am). Then a new tuple carrying higher level logical information may be created: (student:Gabriella, attending: Computer Foundation,room:room.35, 01/09/2008 10:05 am).

3.2 The W4 Knowledge Networks Algorithm

An unsupervised, distributed and self-organizing approach to generate and maintain the knowledge networks' layer is clearly required by the decentralized nature of pervasive computing systems and the overwhelming amount of generated data,

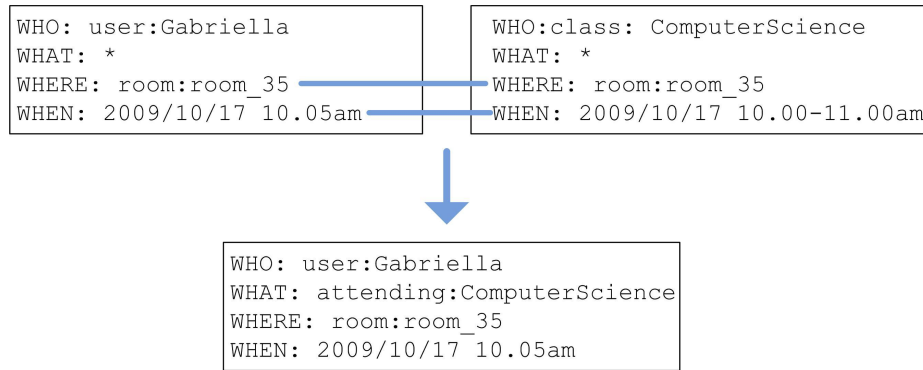


Fig. 2. W4 knowledge network data inference

which prevent the use of a centralized process for data management. To this end, we adopt a swarm-based approach which relies on a two-phase process.

The first phase is the identification of all possible correlations between knowledge atoms, and the creation of links between W4 atoms. This can be done by a number of simple agents (which we call *spiders* as they weave their webs between correlated tuples), in charge of identifying relationships between tuples. Each spider is associated with a pattern matching function that takes a w4 tuple and matches it against a w4 tuple used as a template. The function returns a boolean value meaning whether the matching is successful or not, and accordingly suggesting creating a link or not. Obviously the simpler the matching function is (i.e., few fields of a w4 tuple are involved), the more the resulting net of links can be reusable. So each spider is responsible for a single relationship over W fields. E.g., A1 is in charge of linking together all the tuples with corresponding *who* fields, for instance all the tuples whose *who* field corresponds to “user:Gabriella”, while another agent can search tuples with corresponding *where* field.

Spiders continuously surf W4 tuple spaces in order to retrieve tuples that fulfill the specific relationship, those tuples are virtually linked together thus creating a W4 knowledge network for the given relationship. Obviously spiders must be capable of analyzing W4 tuples stored in different tuple spaces and building correlation networks that extends over distributed tuple spaces. The spiders’ algorithm follows:

```

define:
    rel; //the relation to be satisfied
    knet; //the Knowledge network reference

Main:
    Do forever:
        TupleSpace ts = random();
        jumTo (ts);
        tuple t[] = ts.read(rel);
        knet.add(t);

```

```
Done;
```

The spider chooses a random tuple space and checks if any tuple in the tuple space fulfills the relationship. If it is positive the tuples are added to the knowledge network *knet* by adding a reference to the last tuple space that was found earlier, i.e. drawing a link between the last tuple space added to the knowledge network *rel* and the current one. This process continuously repeats. In this way, a single knowledge network of links between correlated tuples is generated. More spiders can work concurrently, building the knowledge networks layer in a self-organizing fashion.

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. This is performed by another class of agents, called *browsers*. Browser agent surf the knowledge networks trying to generate new W4 atoms. Each browser is capable of inferencing a specific type of relationship. The browsers' algorithm follows:

```
define:
    rel;      // the relation that the browser is capable to infer

Main:
    Do forever:
        TupleSpace ts = random();
        tuple t = ts.random();
        ts.add (GenerateNewKnowledge(t));
    Done;
```

The browser chooses a random tuple *t* in the system, and locates all the knowledge networks in which the tuple *t* is involved. Then the browser start to browse each of the found knowledge networks. For each tuple \hat{t}_i found in a related knowledge networks, the browser checks if he is able of generate a new w4 atom carrying higher knowledge. If positive, the new atom is generated and added to the current tuple space. The issue of bounding the amount of knowledge generation has been discussed in [8].

3.3 Using W4 Knowledge Networks

The idea at the base of the w4 knowledge networks approach is that spiders and browsers continuously surf, analyze, correlate and infer new knowledge. In this way

Although the knowledge networks can be used as the basis for knowledge reasoning and can lead to the generation of new knowledge, even when new data are not generated, the web of links between atoms can be fruitfully used during querying to access and retrieve contextual information more effectually. When a query is submitted to the W4 tuple space system, a query-solving agent capable

of browsing knowledge networks, i.e. a browser, analyze the query template and determine one or more knowledge networks to which the matching tuples should belong. Then the browser agent choose a random W4 tuple space in the system and scans it until he finds an entry point for one of the identified knowledge networks, i.e. a tuple belonging to one of those knowledge networks. When the entry point is found, the agent starts to jump from the entry point tuple to the other tuples in the identified knowledge network, checking if they matches the template and finally returns the retrieved tuples. This is beneficial for services because fewer read operations have to be performed when exploiting knowledge networks instead of a set of data spaces in which information is not related to each other.

4 Result Evaluation

To complete the W4 Approach presentation, we report several experimental results we performed to evaluate the effectiveness and feasibility of the proposed approach and compare it to other tuple space middle-ware solutions.

To extensively test the approach we developed a simulated environment based on the Repast framework [<http://repast.sourceforge.net>] and integrated with the LighTS technology. We represented a location with a number of users each moving in the environment. The virtual environment is split in 100 zones, each of them holds a private W4 tuple space that stores all the tuples generated in it. Periodically a W4 tuple for each user is generated based on the current position and the current time. In this scenario many tuples are stored in the W4 tuple spaces, and services may find difficult to access those data.

4.1 Efficiency

The first set of experiments aims to measure the efficiency of the W4 system in retrieving information and comparing it with an exhaustive search in tuple spaces and with an hash based approaches based on the performance of the above systems when a non destructive query is submitted to the system.

The exhaustive search is performed on a tuple space that embeds the W4 data model facilities but not the W4 knowledge networks mechanisms. When a query is submitted to this simplified W4 tuple space system, a query agent chose a random tuple space in the system and scans it seeking for the W4 tuples that fulfill the query template. Then a random tuple space is chosen again, until the whole system is scanned.

Hash based tuple space is a well known and popular technique for data indexing in distributed environment. Here we follow an approach similar to [16] in which a single field of the tuple structure is used for the hashing operation and indexing purpose. When a tuple is injected in the distributed system, the hashing operation is performed over the designated field and the tuple is then stored on the resulting tuple space. Then when a query is submitted to the system,

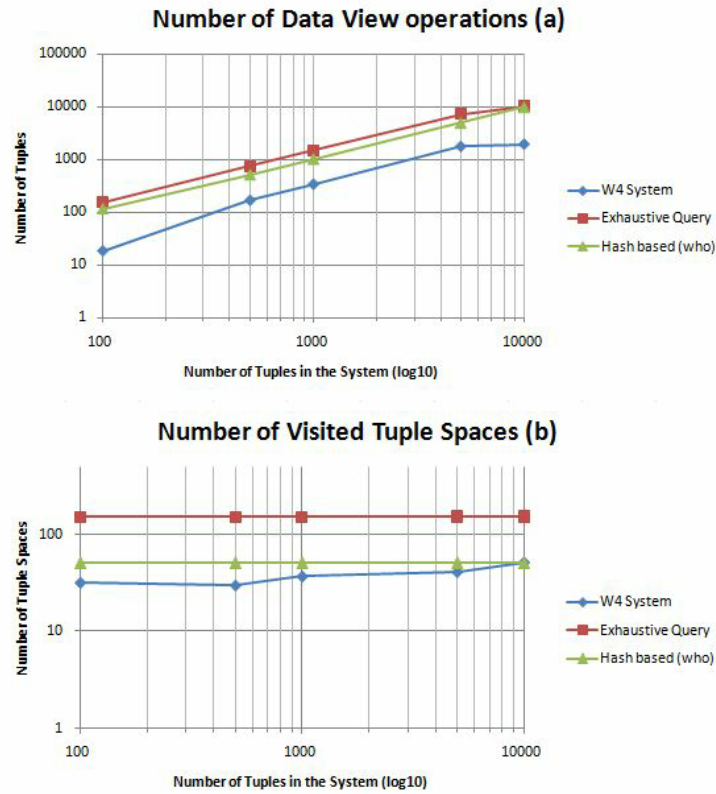


Fig. 3. Efficiency: (a) Number of view operations done by query-solving agents. (b) Number of view operations done by query agents

the same hashing operation is performed on the query template and the result indicates the tuple space to scan for results. In this simulation we considered the hash performed on the who field. Similar results are achieved considering the other fields of a W4 tuple. Of course other hash based approaches that use tuple spaces exist. In particular, there are approaches that consider to hash more than one W field at a time. Those approaches typically rely on pointers to tuples, and must deal with data duplication, update issues and so on. We intentionally didn't consider these approaches, they may obviously have better performances but, in our opinion, the issues that arise make the approaches not comparable to the W4 one as the system management becomes more complex.

The experiment works as follows: we fed the W4 system with a fixed amount of tuples and made them organized in W4 knowledge networks. Then, for the sake of experiments, we stopped the data sources and submitted to the system the following complex query: "Retrieve all the users that were near agent A5 was, on time 500". For a W4 System this means the following two queries should be subsequently resolved:

```

WHO      user:A5
WHAT     *
WHERE    ?var1
WHEN     500

```

Please note that the *when* field is automatically transformed in a time interval. The second query looks as follow:

```

WHO      *
WHAT     *
WHERE    var1
WHEN     500

```

Here both the *where* is automatically considered as a bounding box and the *when* field is automatically transformed in a time interval. To solve such a query to knowledge networks must be investigated, one relating all tuples from *user:A5*, and one relating all tuples that refers a specific region of the space.

We run the simulations 15 times and depicted the average values. Figure 3 (a) shows the number of tuple spaces visited by the query-solving agent in the considered systems. The W4 tuple space systems performs better than the other considered approaches. Indeed, in the medium case, the exhaustive search has to query half the number of tuple spaces in the system to solve the first sub-query and the whole number of the tuple spaces to solve the second ones. The hash approach works better than the exhaustive query because one of the sub-query is solved thanks to the hashing operation, nevertheless the other sub-query have to be solved traditional as in the case of the exhaustive Search. However exploiting the W4 knowledge networks is even better because the number of accessed tuple spaces is determined by the number of tuple spaces involved in the knowledge networks of interest.

Figure 3 (b) shows the number of read operation performed. Also in this case the W4 tuple space system performs better then the other systems. As in the previous case, the exhaustive search have to access half the number of tuples in the systems to solve the first sub-query, and all the tuples in the system to solve the second one. Here the performance of the hash based systems are significantly different depending if the hashing is performed on the who field or on the where field. However the W4 system performs better because all the fields are considered equally important when building knowledge networks.

4.2 Effectiveness

Provided that the W4 approach exhibits a good behavior in accessing contextual data (i.e. the access costs are lower than the other considered approaches), we run a second set of experiments to test the effectiveness of the knowledge networks' approach, in terms of accuracy of provided results when the knowledge networks algorithms are running, i.e. the fraction of the documents that are relevant to the query that are successfully retrieved (also called recall in information retrieval).

In order to test the effectiveness of the approach, we started feeding the system with W4 tuples generated by the simulated environment and let the

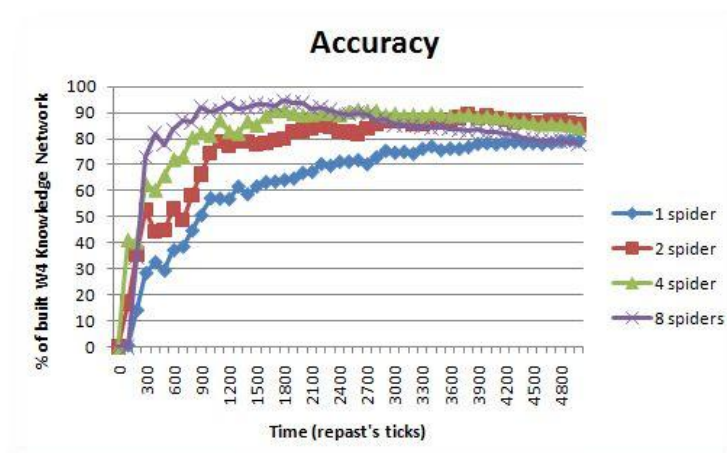


Fig. 4. Accuracy of the indexation over time.

knowledge networks keep organizing. Periodically we checked the content of a specific knowledge networks respect to the content of the whole W4 system to measure the percentage of tuples that has been indexed.

Of course the indexing works as quicker as more spiders are involved in, to this end Figure 4 compares the results obtained when a different number of spiders is running. The obvious result is that the more spiders are working, the quicker the knowledge network reaches its indexation level. We can see that it takes a certain amount of time for the knowledge network to reach its stable value of indexation that is in the satisfactory range of 80-90% depending on the number of spiders run. This suggests that the W4 system could be improved by taking into account the W4 tuples' injection rate in order to autonomously determine the right number of spiders running.

Accordingly to this observation, we performed a second set of experiments varying the dynamism (i.e. the tuples injection rate) of the system respect to the number of spiders running, the results is quite interesting because they give an idea of the number of spiders that should run simultaneously in relation to the dynamism of the system. As we expect, Figure 5 shows that as the dynamism of the system increases, i.e. tuples are injected in the w4 system more quickly, the percentage of indexation decrease. That is, when the tuple injection rate increase, it may be needed to run more spiders and browsers to keep the good level of indexation.

4.3 Scalability

Another key factor in distributed systems is their ability to scale. The first set of experiments about the system' efficiency shows that the system scale at same extent when increasing the number of tuples stored in the whole system. To test further the system scalability we performed another set of experiments fixing the

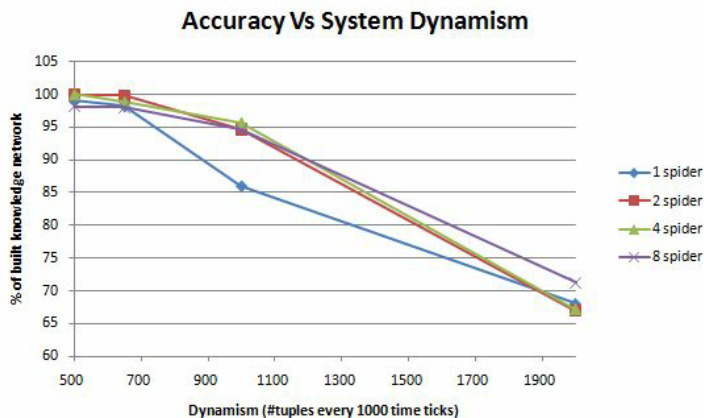


Fig. 5. Accuracy of the indexation Vs dynamism of the system.

number of tuples per tuple space, and varying the number of tuple space in the w4 system. We performed the measurements as described for Efficacy experiments, and measured the percentage of tuple spaces accessed to solve the query (the number of tuples accessed is not represented because it is highly correlated, as shown in 3). Figure 6 depicts results. We can see that the performances improve when the number of tuple spaces in the system increase. This is due to the fact that the more the system is wide and distributed, the more selective the knowledge networks can be. Indeed the W4 approach make sense in a distributed environment rather than in a centralized and static one.

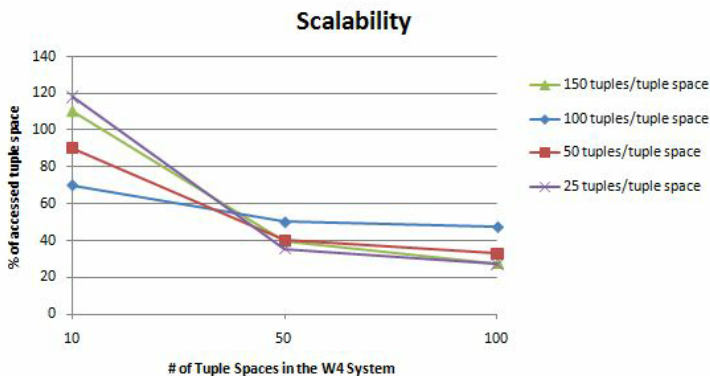


Fig. 6. Scalability of the w4 systems respect to the number of tuple spaces in the system (number of tuples per tuple space fixed).

4.4 Overhead

To summarize experimental results show that accessing pre-organized W4 knowledge networks instead of a flat or hash-based tuple space system greatly improves the system performance in terms access costs while the accuracy of provided results remains satisfactory. However, the generation of overlay networks never comes without overhead costs. The idea behind W4 knowledge networks is to pre-organize data in a fashion that will be useful for a number of services and will decrease delay experiences by services. In order to keep knowledge networks feasible and manageable one should find a good trade-off between the number of knowledge networks to be build in the system and overhead costs associated, in general only knowledge networks that might be useful (i.e., accessed by agents) should be built and maintained.

5 Related Work

Context is a very fluid notion and although several researchers claim that it is very hard to abstract it in terms of variables and data models [12], it is also a widespread opinion that a more pragmatic perspective should be adopted. Early works in this area, as from Schmidt et al. [25] and Dey et al. [11], concentrates on the issue of acquiring context data from sensors and of processing such data but they generally miss in identifying a uniform model to describe the data and analyzing the issues at the middleware level. A different thread of researchers [14] focuses more on the issue of providing rich data models for contextual information and of facilitating querying by services. However, this does not go toward simplicity and generality, which we instead feel should both be goals also to make the middleware implementation more light-weight. Some recent proposals, such as [26, 5] focus on providing models for contextual data that adopt a uniform well-defined structure. Indeed, our W4 proposal accounts for a very similar structuring for contextual information, and enriches it further with a well-defined API, and with the possibility of linking data atoms and of providing application-specific views to services.

An increasing number of research works get inspiration from tuple space middleware models [1] and propose representing and storing contextual information in the form of tuples to be stored in distributed tuple spaces. Egospaces [17] adopts this perspective, without committing to a specific pre-defined structure for context tuples, which can make it difficult for services to uniformly deal with tuples represented in different formats. Other proposal, such as The Context Fabric model [15] rely on well-structured context tuples. Recent proposals focusing on sensor networks, suggest exploiting a tuple-based approach to provide application-specific views on sensorial data [20]. In general we consider tuple-based approaches very suitable for organizing and accessing contextual information, but we also think that there is need of more structuring and flexibility than those exhibited by the existing approaches.

In the above described work, the issue of relating contextual data atoms with each other and of providing different views to different applications is not generally addressed. More recently, other proposals have adopted a similar endeavor but have considered the issue of adopting specific ontologies to model context information and enable other than efficient querying also efficient context-reasoning [23, 18]. Although such approaches tend to be too application-specific, they attribute the importance of linking independent atoms of contextual information (with ontological relations) and of reasoning not only on individual data items but also on their relations, an idea which is fully shared by our knowledge network vision.

Some recent proposals go beyond context modeling and representing and start consider also context reasoning, i.e. considering related piece of information for generating new knowledge as W4 knowledge networks do. Many works, such as and [22], are focused on situation learning and situation relationships in smart environment. Other works, such as [21] propose predicate logic as an effective language for context-aware reasoning. The Knowledge Networks approach we propose aim to be more general and propose an approach different from traditional ones, considering self-organizing agents. Campbell et al. [6] consider the possibility of extracting higher-level knowledge from raw sensed data merging feature vectors in an opportunistic fashion for people-centric application. The idea of merging and considering data coming from diverse sources is shared with the W4 knowledge networks approach. However in the W4 Approach we go further considering multiple knowledge views that can be accessed by multiple services.

In [24] the context management processes is performed by BDI agents that generate and administrate the context artifacts at run time. We share with [24] the idea of a middleware responsible for contextual data, but we also argue overwhelming amounts of data require novel approaches for data management, such as bio-inspired and self-organizing ones.

Also other areas of research contributed towards the realization of our knowledge networks vision, in particular data mining and pattern discovery and granular computing. See [7] for a critical survey. Data mining [13] always come to play when dealing with an overwhelming amount of data to be analyzed, indeed the analysis task performed by the knowledge networks layer can be considered as a sort of data mining process. Also Granular Computing [27] is of interest, indeed W4 knowledge atoms can be seen as information granules, providing knowledge at different scales.

6 Conclusion and Future Works

Despite the promising results achieved so far in the study of the W4 self-organized knowledge networks algorithms, some research issues still have to be faced. More experiments should be done to evaluate properly the overhead costs, as introduced in Section 4. Moreover, in the current implementation of the W4 system, the number of tuples stored in the system is constantly increasing as

new data are injected in the system. There is the need for a "garbage collection" solution and we plan to experiment with a concept of knowledge tuple fading as introduced in [19]. Finally, security and privacy issues need to be analyzed w.r.t. accessing W4 tuples and their realtions.

References

1. S. Ahuja, N. Carriero, and D. Gelernter. Linda and friends. *Computer*, 19(8-9):26–34, 1986.
2. D. Balzarotti, P. Costa, and G. P. Picco. The LighTS Tuple Space Framework and its Customization for Context-Aware Applications. *International Journal on Web Intelligence and Agent Systems*, 50(1-2):36–50, 2007.
3. Bettini, Brdiczka, Henricksen, Indulska, Nicklas, Ranganathan, and Riboni. A survey of context modelling and reasoning techniques. *Pervasive and Mobile Computing*, in press.
4. N. Bicochi, G. Castelli, M. Mamei, A. Rosi, F. Zambonelli, M. Baumgarten, and M. Mulvenna. Knowledge networks for pervasive services. In *ICPS '09: Proceedings of the 2009 international conference on Pervasive services*, pages 103–112, New York, NY, USA, 2009. ACM.
5. J. Bravo, R. Hervs, I. Snchez, G. Chavira, and S. Nava. Visualization services in a conference context: An approach by rfid technology. *Journal of Universal Computer Science*, 12(3):270–283, 2006.
6. A. Campbell, S. Eisenman, N. Lane, E. Miluzzo, R. Peterson, H. Lu, X. Zheng, M. Musolesi, K. Fodor, and G.-S. Ahn. The rise of people-centric sensing. *Internet Computing, IEEE*, 12(4):12–21, July-Aug. 2008.
7. G. Castelli, M. Mamei, and F. Zambonelli. Engineering contextual knowledge for autonomic pervasive services. *International Journal of Information and Software Technology*, 52(8-9):443–460, 2008.
8. G. Castelli, R. Menezes, and F. Zambonelli. Self-organized control of knowledge generation in pervasive computing systems. *ACM Symposium on Applied Computing, 2009*, 8-12 March 2009.
9. G. Castelli, A. Rosi, M. Mamei, and F. Zambonelli. A simple model and infrastructure for context-aware browsing of the world. In *PERCOM '07: Proceedings of the Fifth IEEE International Conference on Pervasive Computing and Communications*, pages 229–238, Washington, DC, USA, 2007. IEEE Computer Society.
10. D. D. Clark, C. Partridge, J. C. Ramming, and J. T. Wroclawski. A knowledge plane for the internet. In *SIGCOMM '03: Proceedings of the 2003 conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pages 3–10, New York, NY, USA, 2003. ACM.
11. A. K. Dey, G. D. Abowd, and D. Salber. A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human Computer Interaction*, 16(2):97–166, 2001.
12. P. Dourish. What we talk about when we talk about context. *Personal Ubiquitous Computing*, 8(1):19–30, 2004.
13. J. Galloway and S. J. Simoff. Network data mining: methods and techniques for discovering deep linkage between attributes. In *APCCM '06: Proceedings of the 3rd Asia-Pacific conference on Conceptual modelling*, pages 21–32, Darlinghurst, Australia, Australia, 2006. Australian Computer Society, Inc.

14. K. Henriksen and J. Indulska. Developing context-aware pervasive computing applications: models and approach. *Pervasive and Mobile Computing*, 2, 2005.
15. J. I. Hong. The context fabric: an infrastructure for context-aware computing. *CHI '02 extended abstracts on Human factors in computing systems*, pages 554–555, 2002.
16. Y. Jiang, G. Xue, Z. Jia, and J. You. Dtuples: A distributed hash table based tuple space service for distributed coordination. *Grid and Cooperative Computing, 2006*, pages 101–106, October 2006.
17. C. Julien and G.-C. Roman. Egospaces: facilitating rapid development of context-aware mobile applications. *Software Engineering, IEEE Transactions on*, 32(5):281–298, 2006.
18. D. Lee and R. Meier. Primary-context model and ontology: A combined approach for pervasive transportation services. *Pervasive Computing and Communications Workshops, 2007. PerCom Workshops '07. Fifth Annual IEEE International Conference on*, pages 419–424, 2007.
19. R. Menezes and A. Wood. The fading concept in tuple-space systems. In *Proceedings of the 2006 ACM Symposium on Applied Computing*, pages 440–444, Dijon, France, 2006. ACM, ACM Press.
20. L. Mottola and G. P. Picco. Logical neighborhoods: A programming abstraction for wireless sensor networks. In *Proc. of the 2nd Int. Conf. on Distributed Computing on Sensor Systems (DCOSS)*, 2006.
21. A. Ranganathan and R. H. Campbell. An infrastructure for context-awareness based on first order logic. *Personal Ubiquitous Comput.*, 7(6):353–364, December 2003.
22. P. Reignier, O. Brdiczka, D. Vaufreydaz, J. L. Crowley, and J. Maisonnasse. Context-aware environments: from specification to implementation. *Expert Systems: The Journal of Knowledge Engineering*, 24(5):305–320, November 2007.
23. I. Roussaki, M. Strimpakou, N. Kalatzis, M. Anagnostou, and C. Pils. Hybrid context modeling: A location-based scheme using ontologies. *Pervasive Computing and Communications Workshops, IEEE International Conference on*, (1):2–7, 2006.
24. I. Salomie, T. Cioara, I. Anghel, and M. Dinsoreanu. Rap - a basic context awareness model. *Intelligent Computer Communication and Processing, 2008. ICCP 2008. 4th International Conference on*, pages 315–318, 2008.
25. A. Schmidt, K. A. Aidoo, A. Takaluoma, U. Tuomela, K. V. Laerhoven, and W. V. d. Velde. Advanced interaction in context. *Proceedings of the 1st International Symposium on Handheld and Ubiquitous Computing*, 8, 1999.
26. C. Xu and S. C. Cheung. Inconsistency detection and resolution for context-aware middleware support. *Proceedings of the 10th European Software Engineering Conference Held jointly with 13th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pages 336–345, 2005.
27. Y. Yao. Three perspectives of granular computing. *Journal of Nanchang Institute of Technology*, 25(2):16–21, 2006.