Future Generation Computer Systems 56 (2016) 675–683

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Scalable semantic aware context storage

Mário Antunes *, Diogo Gomes, Rui L. Aguiar

Instituto de Telecomunicações, Universidade de Aveiro, Aveiro, Portugal

HIGHLIGHTS

- We analyse efficient ways to deal with unstructured information.
- Context information should be organized through a bottom-up characterization.
- We discussed and analyse four possible context storage solutions, and two context organization schemes.
- 1-dimension model leads to poor scalability and semantic extraction.

ARTICLE INFO

Article history:
Received 15 November 2014
Received in revised form
5 September 2015
Accepted 9 September 2015
Available online 28 September 2015

Keywords:
IoT
M2M
Context information

ABSTRACT

The number of connected devices collecting and distributing real-world information through various systems, is expected to soar in the coming years. As the number of such connected devices grows, it becomes increasingly difficult to store and share all these new sources of information. Several context representation schemes try to standardize this information, but none of them have been widely adopted. In previous work we addressed this challenge, however our solution had some drawbacks: poor semantic extraction and scalability. In this paper we discuss ways to efficiently deal with representation schemes' diversity and propose a novel *d*-dimension organization model. Our evaluation shows that *d*-dimension model improves scalability and semantic extraction.

1. Introduction

When we think about the Internet we mostly consider servers, laptops, routers and fixed broadband that have penetrated almost every household. But the fact is that the Internet is diversifying as we speak. Everyday new kinds of devices (from mobile phones to environmental sensors networks) connect to the Internet, and share massive amounts of data. According to the ICT Knowledge Transfer Network, the number of mobile devices is expected to increase worldwide from 4.5 billion in 2011 to 50 billions by 2020 [1].

In M2M (machine to machine) scenarios, an entity's context can be used to provide added value: improve efficiency, optimize resources and detect anomalies. The following examples illustrate the importance of context information in M2M scenarios. Fusing data from several sensors makes it possible to predict a driver's ideal parking spot [2,3]. Projects such as Pothole Patrol [4] and Nericell [5] use vehicular accelerations to monitor road conditions and detect potholes. TIME (Transport Information Monitoring Environment) project [6] combines data from mobile and fixed sensors in order to evaluate road congestion in real time.

These projects provide valuable insight about context information potential in advanced context-aware applications. However, many of these projects follow a vertical approach. This has hindered interoperability and the realization of even more powerful IoT scenarios. Another important issue is the need felt for a new way to manage, store and process such diverse machine made context information; unconstrained and without limiting structures.

Common definitions of context information [7–9] do not provide any insight about its structure. In fact, each device can share context information with a different structure. e.g. sensory and location information can be used to characterize an entity context, yet the two can have different structures. One important objective of context representation research [10–12] is to standardize the process of sharing (with different platforms) and understanding context information. Context-aware platforms strongly benefit from a uniform environment: the storage process is easier (the information follows a known structure) and the analysis of the information becomes simpler. Standard context management platforms commonly store context information in relational databases. We can devise a mapping process only if there is a common context representation. Multiple context representations have been proposed, such as ContextML [13], SensorML [14]

* Corresponding author.
E-mail addresses: mario.antunes@av.it.pt (M. Antunes), dgomes@av.it.pt (D. Gomes), ruilaa@av.it.pt (R.L. Aguiar).

http://dx.doi.org/10.1016/j.future.2015.09.008

COBRA-Ont [15]. All these representations try to solve the same problem, but each representation is quite different and incompatible with the other. None of the above mentioned representations have been widely accepted either by the academia or the industry. Usually, each context-aware platform defines its own context representation based on the platform scenarios. This breaks compatibility between platforms and limits the quantity of context information that can be used in M2M applications, impairing future developments.

It is possible (but unlikely) that in the future a context representation standard will be widely adopted. Until then, context-aware platforms have to deal with multiple context representations. The work presented in this paper addresses this problem and analyses possible representation schemes independent of storage solutions. In previous work we addressed this challenge, however our solution had some drawbacks: poor semantic extraction and scalability. We propose a novel *d*-dimension organization model, that improves scalability and semantic extraction.

The remainder of the paper is organized as follows. In Section 2 we analyse how context information can be organized and define the basic requirements for context storage solutions. We discuss the impact of time constrains in storage solution designed for M2M scenarios in Section 3. Two context organization models are proposed and analysed in Section 4. We studied the organization models' impact on context information solutions in Section 5. The spatial requirements of both models is estimated in Section 6. Section 7 contains implementation details of our context storage solution. The results of the organization models evaluation are presented in Section 8. Finally, the discussion and conclusions are presented in Section 9.

2. Context organization

Context information is an enabler for further data analysis, potentially exploring the integration of an increasing number of information sources. As previously mentioned, nowadays no widely accepted context representation scheme exists; instead there are several approaches to deal with context information. These approaches can be divided into three categories: (i) adopt/create a new context representation, (ii) normalize the storing process through ontologies or (iii) accept the diversity of context representations.

Previous works have defined a new optimized context representation [16,17]. However, this approach imposes limits to the quantity of information that can be shared with other context-aware platforms. On later works [17] the authors recognized that the usage of a single context representation limits the information expressiveness.

Another possibility would be employing ontologies to normalize the storage process. Each context representation scheme is mapped into the internal data model through an ontology [18]. This type of platform supports several context representations, yet it is necessary to define a new ontology for each new representation. Defining a new ontology is a tedious task that requires human intervention. The scale of M2M scenarios makes this task very difficult.

Finally, we can accept the diversity of context representation as a consequence of economic pressures, and develop methods/techniques do deal with it. According to the authors [19–21], the best solution to classify context information is through bottom-up characterization. Bottom-up characterization is massively dimensional, and there is no global consistency imposed by current practice.

Based on this analysis we defined the basic requirements for context storage solution [22]: ability to scale, generalize storing process and discriminative retrieval. The number of connected devices is increasing, as such the quantity of context information is also increasing, and a context storage must cope with this increase. The two last requirements complement each other. In other words, the ideal context database must store and accurately pinpoint any piece of information.

Although context information is not manually tagged by users, we can model bottom-up characterization as an information retrieval problem. Organizing documents based on its content is one of the major objectives of information retrieval (IR) research: information retrieval informs on the existence (or non-existence) and whereabouts of documents related with user's query (similar to a web search engine). There are several methods that provide discriminative retrievals such as relational models, semantic web, ontologies/taxonomies among others. However, these methods either require knowledge about the context structure (relational model) or manually defined relations amongst entities (semantic web, ontologies/taxonomies).

Information retrieval systems use discriminative terms to index documents. Keywords/discriminative terms are sequences of terms that provide a compact representation of a document's content. Ideally, keywords represent in condensed form the essential content of a document. Moreover, it is possible to enrich an information retrieval system with semantic information. Some of the most popular semantic methods are based on latent analysis [23–25]. These well known methods analyse the co-occurrences of terms in a corpus of documents in order to find hidden/ latent variables, regarded as topics or concepts. Since the number of concepts is usually greatly inferior to the number of words and it is not necessary to know the document categories/ classes, these methods are thus unsupervised dimensionality reduction techniques.

Other popular semantic methods are based on estimating distance between two units of language. Semantic distance [26] is a measure of how close or distant two units of language are, in terms of their meaning. For example, the nouns *banana* and *fruit* are closer in meaning than the nouns *banana* and *car*. With this addition, and a new similarity metric that takes into account the semantic value of the relevant terms, users can search for concepts instead of simple words. Hence, discriminative retrieval can be achieved with information retrieval techniques.

3. Storage time constrains

As previously mentioned storing and processing context information yields significant advantages for M2M scenarios. New kinds of M2M applications will greatly benefit from analytical algorithms, however different applications will require different analytical techniques.

Analytical techniques can be divided into two different categories: batch and stream processing. Batch processing requires the whole dataset and can process it multiple times. On the other hand, in stream processing each element is processed once when it arrives. These two kinds of analytical algorithms require very different access to information: batch requires random access while stream requires sequential access. Random and sequential access can be achieved through request/reply and publish/subscribe interfaces respectively. The access to information greatly conditions the underlying storage solution.

One of our objectives is developing a storage solution capable of supporting several M2M applications and facilitate interoperability between applications. Furthermore, a complex M2M application may require both batch and stream algorithms. As an example, let us considers the LITES [27] project. This project delivers an intelligent public street lighting service using solid-state lights LED in order to drastically reduce energy consumption. The core element of the solution is the dimming of the lamp based on the environment. In order to optimize the power consumption it is necessary