Proceedings of the

3rd International Workshop on Information Management for Mobile Applications

Message from the Workshop Chairs

The increasing functionality and capacity of mobile devices have enabled new mobile applications which require new approaches for data management. Users want to have a seamless integration of their data on their mobile with other devices, which can be either classical devices such as a desktop PC or other mobile devices. Although the capabilities of mobile devices are growing, their limitations have to be taken into account when designing efficient and effective mobile applications. For example, constraints on energy, CPU power, storage, display size, communication bandwidth, and real-time capabilities have to be considered.

Information management in mobile applications is a complex problem space which requires the consideration of the aforementioned constraints. In addition, mobile data can have various forms, such as sensor data, user profiles & user context, spatial data, and multimedia data. Smartphones, mobile and wearable sensors, and other portable systems are used in various applications to collect, process, and exchange an increasing amount of data. Applications can run on several devices (mobile, PCs, multimedia), but the exchange, the integration, and the querying of data between these devices remains a challenging problem.

The International Workshop on Information Management for Mobile Applications (IMMoA'13) is a continuation of the successful IMMoA and HIMoA workshops in the previous years. It aims at a broad range of mobile application fields: Business, (Serious) Games, Leisure, and Transport. IMMoA'13 was co-located with VLDB 2013 in Riva del Garda, Italy, and provided a forum for discussion about technologies and mechanisms, which support the management of mobile, complex, integrated, distributed, and heterogeneous data-focused applications.

We received about ten high quality submissions of which we could accept six as full papers. The papers have been peer-reviewed by three to four reviewers each. In addition, we had three invited papers and a keynote presentation by Maria Luisa Damiani from the University of Milan on "Moving objects beyond raw and semantic trajectories". The presentations covered topics such as pattern recognition, models for spatio-temporal data, mobile data management in health applications, and videogames to evaluate information management in vehicular networks.

We would like to thank all authors, presenters, and the reviewers for their good work that resulted in a successful workshop. Furthermore, we would like to thank the DFG Research Cluster Ultra High-Speed Mobile Information and Communication (UMIC, http://www.umic.rwth-aachen.de) at RWTH Aachen University, Germany, for their support in organizing this event.

Thierry Delot
University of Valenciennes & Inria Lille, France

Sandra Geisler
RWTH Aachen University, Germany

Sergio Ilarri
University of Zaragoza, Spain

Christoph Quix
Fraunhofer FIT, Germany

Workshop Chairs and Program Committee

Workshop Chairs

- Thierry Delot, Dept. of Computer Science, University of Valenciennes & Inria Lille, France
- Sandra Geisler, Information Systems, RWTH Aachen University, Germany
- Sergio Ilarri, University of Zaragoza, Spain
- Christoph Quix, Fraunhofer FIT, St. Augustin, Germany

Program Committee

Hyung-Ju Cho, Ajou University, Korea
Christine Collet, Grenoble INP, France
Francesco Guerra, University of Modena and Reggio Emilia, Italy
Marwan Hassani, RWTH Aachen University, Germany
Arantza Illarramendi, University of the Basque Country, Spain
David Kensche, Thinking Networks, Germany
Dejan Kovachev, RWTH Aachen University, Germany
Sanjay Madria, Missouri University of Science and Technology, USA
Jochen Meyer, OFFIS Institute for Informatics, Germany
Nathalie Mitton, INRIA Lille, France
Filip Perich, University of Maryland Baltimore County, USA
Florence Sedes, University Paul Sabatier Toulouse, France
Satish Narayana Srirama, University of Tartu, Estonia
Dragan Stojanovic, University of Nis, Serbia
Masaaki Tanizaki, Hitachi Central Research Laboratory, Japan
Goce Trajcevski, Northwestern University, USA
Raquel Trillo, University of Zaragoza, Spain
Upkar Varshney, Georgia State University, USA
Jari Veijalainen, University of Jyväskylä, Finland
José Luis Zechinelli Martini, Universidad de las Amricas, Puebla, Mexico

Table of Contents

Moving objects beyond raw and semantic trajectories

Maria Luisa Damiani
University of Milan, I
damiani@di.unimi.it

Ralf Hartmut Güting
FernUniversität Hagen, D
rhg@fernuni-hagen.de

Fabio Valdés
FernUniversität Hagen, D
fabio.valdes@fernuni-hagen.de

Hamza Issa
University of Milan, I
issa@di.unimi.it

ABSTRACT

Mobile applications, for example for road traffic monitoring, mobile health and animal data ecology, call for methods enabling rich and expressive representation of moving objects. This demand motivates the increasing concern for the paradigm of semantic trajectories. In this paper, I overview related research, focusing in particular on the novel data model of *symbolic trajectories* proposed for the efficient and flexible handling of semantics-aware trajectories through a Moving Object DBMS.

1. INTRODUCTION

Semantic trajectories is a relatively recent paradigm developed to provide applications with knowledge about the movement of moving entities. The key idea is to supplement the raw mobility data (i.e. raw trajectories in the following) - typically sequences of GPS points - with contextual data [4]. For example, semantic trajectories can be used to describe the sequence of points of interest visited by tourists in a city, or the sequence of transportation means used by an individual to reach the working place from home. Basically a semantic trajectory consists of a raw trajectory augmented with annotations regarding the whole trajectory or parts of it. Probably because of its simplicity and naturalness, the concept of semantic trajectory has attracted the interest of numerous researchers over the last years. Current research develops along diverse streams including: ontology/conceptual modeling, mobility pattern mining for the generation of semantic annotations, semantic location privacy, and - more recently - the connection with the theories of complex networks and social analysis. The main results achieved so far are nicely summarized in the survey paper [4].

Somewhat surprisingly, one aspect that is largely ignored by the most recent literature regards the data management dimension of semantic trajectories. Put simply: how can we store and access semantic trajectories? How can we represent semantic trajectories through a rigorous data model? How can semantic trajectories interplay with raw trajectories and conventional data? These questions have been only marginally addressed. In fact no operational system enabling the management of semantic trajectories in real applications exists. We believe that this is a critical limitation especially in the light of the increasing availability of *big* raw trajectory data collected from mobile applications (e.g. LBS) that creates challenging opportunities for the application of this concept.

The research that we have undertaken in the context of the European initiative Cost Action MOVE[1] aims to fill this gap. Indeed the goal is not simply to take some existing definition of semantic trajectory and find the best way for implementing it on a DBMS, but rather to re-think of the notion of *semantically meaningful movement* while targeting the specification of a general, formal and operational framework. We imagine that in the long run this research could lead to the development of a novel class of software platforms for mobility data handling. The users of these systems will be able to organize and analyze mobility data in the same way that users now organize and analyze spatial data in a conventional GIS platform, e.g. Quantum GIS, or using one of the more recent platforms on cloud, e.g. GISCloud. While the idea in itself may sound not particularly innovative, just a restyling of GIS, we believe that these platforms, going beyond the notion of Moving Object DBMS, can greatly facilitate the development of novel and challenging applications. In what follows, the notion of semantic trajectory is presented; next the concept of symbolic trajectory is introduced along with the results achieved so far and major open issues.

2. SEMANTIC TRAJECTORIES

Early work on semantic trajectories was triggered by the experimental analysis of a set of raw trajectories about a group of birds [5]. By using the standard functionalities of a GIS, we found that the sequences of points, just pairs

[1]http://move-cost.info/

of timestamped coordinates, associated with birds identifiers were actually representing the migration routes from Central Europe to Africa and vice versa. Such discovery, that was somewhat unexpected, inspired the proposal of a novel model for the high level representation of movement. Since this first result, research developed along different directions, including the following:

- *Conceptual modeling.* The first conceptualization was centered on the notions of stop and move [5]. A stop represents a temporary suspension of the movement, while a move is the transfer from one stop to another stop. While this conceptualization is appropriate in many applications, there is increasing evidence that stop-and-move is just one of the possible mobility patterns. For example Yan et al. [7] present an approach to extract and represent the sequence of activities from raw trajectories. In the light of these experiences, a novel conceptual model has been recently proposed which enables the attachment of any kind of meaning (not just stop and move) to sequences of points [4].

- *Extraction of mobility patterns.* A major research direction regards the mining of mobility patterns to automatically annotate semantic trajectories. Early work by Alvares et al. [1] focuses on the identification of stops and moves. Numerous approaches can be found in literature, either explicitly related to the notion of stop-and-move or developed within different communities. A comprehensive survey can be found in [4].

- *The privacy of mobility patterns.* A different issue is to preserve the privacy of sensitive mobility patterns such as the presence in places, e.g. hospitals and religious buildings, that might reveal sensitive information about moving individuals. This problem is particularly challenging in on-line applications, e.g. LBS and geo-social networks, whereas the privacy mechanism has to rely on partial knowledge of the movement (past and current positions are known, but not future positions). The privacy of mobility patterns in an open issue [2]. An approach in this direction, focused on the protection of specific mobility pattern, i.e. sensitive places, is presented in [8].

3. SYMBOLIC TRAJECTORIES

Semantic trajectories are often considered the result of an analytical process conducted on raw trajectories. We believe that the notion of semantic trajectory is valuable on its own, independently of how these trajectories are generated. For example, annotations can be deliberately attached by individuals (e.g. user can specify the transportation means) or even the annotation can be automatically attached by the location tracking system (e.g. locations in indoor settings have natural semantics, such as room 1 and building A). Moreover, even in those cases in which semantic trajectories are obtained from an analytical process, the problem remains of how to encode them in a machine readable form. This is the focus of our current research that we briefly present in what follows.

3.1 The data model

We have defined a simple generic data model able to capture different types of semantics called symbolic trajectory [6]. In essence the idea is to represent semantic information in terms of names or labels. For example an activity (running, walking) and points of interest (Colosseum, Louvre) can be straightforwardly described by labels while sensor readings, e.g. temperature, need first to be turned into qualitative values, e.g. high, low temperature. Formally, a symbolic trajectory is an ordered sequence of pairs

$$(i_1\ l_1), ..(i_n\ l_n)$$

called *units* when each unit $u_j = (i_j\ l_j)$ consists of a time interval i_j and a label l_j. The label l_j describes the movement in the time interval i_j. Symbolic trajectories are provided as abstract data types and integrated into the ADT model defined in [3]. For example a symbolic trajectory describing places and the transportation means used to reach those places, can be as follows:

```
(2013-01-17-9:02:30 2013-01-17-9:05:51) "home")
(2013-01-17-9:05:51 2013-01-17-9:08:44) "bus")
(2013-01-17-9:08:44 2013-01-17-9:50:02) "train")
(2013-01-17-9:50:02 2013-01-17-17:50:02) "work")
....
```

The core technical contribution is a novel language for pattern matching and rewriting on symbolic trajectories. The pattern language enables the extraction of subsequences from symbolic trajectories. Patterns are defined as regular expressions that can be matched by single units or sequences of units. For example, the query: *Which are the trajectories in which the individuals take more than 1 hour to move from home to work?* can be solved specifying the following pattern:

```
*(_ home ) Z* (_ work)*// getDuration(Z.time)> 3600
```

where:

- Z is a variable denoting a sequence of units, the symbol $*$ denotes a sequence of zero or more units,
- $(\_\ home)Z * (\_\ work)$ is the pattern
- $getDuration(X.time) > 3600$ is the condition that must be met by the matching sequences, in this case the duration in seconds of the transfer from home to work.

An important feature of the language is that it is embedded into an existing Moving Object DBMS (i.e. Secondo). The pattern language at work is illustrated in a video[2].

4. CONCLUDING REMARKS

Capturing and representing the meaning of movement is a challenging issue that calls for novel solutions. We are working on the definition of the symbolic trajectory data model for the representation of time-varying textual descriptions. A number of issues are still open. For example, a major issue is integrating - whenever it is meaningful - the symbolic dimension with the geometric dimension of the movement. Another major issue regards the usability of the system that is fundamental for an effective deployment of symbolic trajectories in real applications.

[2]http://molle.fernuni-hagen.de/DfnA/SymbolicTrajectories.mp4

5. REFERENCES

[1] L. Alvares, V. Bogorny, B. Kuijpers, B. de Macedo, J.and Moelans, and A. Vaisman. A model for enriching trajectories with semantic geographical information. In *Proc. ACM GIS*, GIS '07, 2007.

[2] M.L. Damiani. *European Data Protection : Coming of Age ?*, chapter Privacy enhancing techniques for the protection of mobility patterns in LBS: research issues and trends. Springer, 2013.

[3] R. H. Güting, M. Böhlen, M. Erwig, C. Jensen, N. Lorentzos, M. Schneider, and M.Vazirgiannis. A foundation for representing and querying moving objects. *ACM Trans. Database Syst.*, 25(1):1–42, 2000.

[4] C. Parent, S. Spaccapietra, C. Renso, G. Andrienko, N. Andrienko, V. Bogorny, M.L. Damiani, A. Gkoulalas-Divanis, J. Macedo, N. Pelekis, Y. Theodoridis, and Z. Yan. Semantic trajectories modeling and analysis. *ACM Comput. Surv.*, 45(4):42:1–42:32, Aug. 2013.

[5] S. Spaccapietra, C. Parent, M.L. Damiani, J. de Macedo, F. Porto, and C. Vangenot. A conceptual view on trajectories. *Data Knowl. Eng.*, 65(1):126–146, 2008.

[6] F. Valdés, M.L. Damiani, and R. Güting. Symbolic trajectories in secondo: Pattern matching and rewriting. In *DASFAA*, 2013.

[7] Z. Yan, D. Chakraborty, C. Parent, S. Spaccapietra, and K. Aberer. Semitri: a framework for semantic annotation of heterogeneous trajectories. In *Proc. EDBT*, 2011.

[8] E. Yigitoglu, M.L. Damiani, O. Abul, and C. Silvestri. Privacy-preserving sharing of sensitive semantic locations under road-network constraints. In *IEEE MDM*, 2012.

Towards a Framework for Semantic Exploration of Frequent Patterns

Behrooz Omidvar Tehrani[‡], Sihem Amer-Yahia[†‡],
Alexandre Termier[‡], Aurélie Bertaux[††], Eric Gaussier[‡],
Marie-Christine Rousset[‡]
[‡]LIG; [†]CNRS; [††]INRIA
France

[‡]firstname.lastname@imag.fr
[††]aurelie.bertaux@inria.fr

ABSTRACT

Mining frequent patterns is an essential task in discovering hidden correlations in datasets. Although frequent patterns unveil valuable information, there are some challenges which limits their usability. First, the number of possible patterns is often very large which hinders their effective exploration. Second, patterns with many items are hard to read and the analyst may be unable to understand their meaning. In addition, the only available information about patterns is their support, a very coarse piece of information. In this paper, we are particularly interested in mining datasets that reflect usage patterns of users moving in space and time and for whom demographics attributes are available (age, occupation, etc). Such characteristics are typical of data collected from smart phones, whose analysis has critical business applications nowadays. We propose pattern exploration primitives, *abstraction* and *refinement*, that use hand-crafted taxonomies on time, space and user demographics. We show on two real datasets, Nokia and MovieLens, how the use of such taxonomies reduces the size of the pattern space and how demographics enable their semantic exploration. This work opens new perspectives in the semantic exploration of frequent patterns that reflect the behavior of different user communities.

1. INTRODUCTION

Nowadays, large amounts of user-generated content representing behavioral data are made available. This is particularly true for data generated by users carrying a mobile phone and moving in different geographic regions. The large size of such data hinders its effective exploration. Fortunately, user-generated data contains repetitive behavior that can be discovered using frequent pattern mining, a common method for discovering hidden patterns that capture some regularities or correlations in the data. Those patterns are used in decision support and can be displayed to analysts for further exploration. There has been extensive work on optimizing algorithms for mining patterns from large datasets. The problem is that there can be millions of automatically discovered patterns, hindering their analysis. Moreover, such patterns can be very long (*i.e.* containing many items), making them difficult to interpret.

Our goal is to develop a method that leverages the richness of underlying data to determine which patterns the analyst has to focus his attention on, and how to interpret those patterns. The two questions we ask ourselves are: *i*. How to explore frequent patterns that characterize subparts of a pattern space in a data-centric way, by giving legible and useful information to the analyst? The possibility to organize items forming a pattern along space and time taxonomies is a new opportunity to reduce the size of the pattern space and the length of patterns. *ii*. How to help an analyst better interpret a pattern by going beyond the notion of support? The availability of demographics information such as users' age and occupation enables a semantic exploration of the space of support users of a pattern.

The most simple way to explore patterns is skimming through the list of frequent patterns, *i.e.*, those for which there is enough evidence in the underlying dataset, or in other terms, those whose support (number of users who exhibit the behavior illustrated in the pattern), is above a given threshold. However, just like in Web search, a list that is more than few dozens of patterns long is infeasible to exploit effectively. Instead, we define two pattern exploration primitives each of which operates on a single pattern at a time. When applied to a pattern, *abstraction* reduces its size and as a side-effect, the size of the pattern space. *Refinement*, on the other hand, highlights different subsets of users forming the support of a pattern, making it more understandable to the analyst.

1.1 Abstraction and Refinement Examples

We propose to use two very different datasets, each of which contains a rich set of usage data and user demographics. Our focus is on Nokia, a small dataset (38 users) that contains application usage data on smartphones and that was made available to the research community as part of a challenge. Nokia contains one year of smartphone usage traces, from GPS position to applications used with a millisecond resolution. We also validate our approach on

MovieLens which is a well-known movie rating dataset (we use the 1M ratings set). All patterns presented in this paper are real ones screeched from our datasets.

Pattern abstraction exploits hand-crafted domain taxonomies. We illustrate it on examples. Our first example is in the context of Nokia. A pattern of the form {*Females between 39 to 50 years old use Email, Bluetooth and Contacts at noon*} could be abstracted into {*Females between 39 to 50 years old use Desktop Communication at noon*} if the collective usage of the applications in the original pattern covers that of a more general *Desktop Communication* class in the taxonomy. This abstraction makes use of a taxonomy on applications that dictates the semantics of abstraction.

In our second example, patterns represent correlations between movies rated by the same users in MovieLens. The support of a pattern is the number of users who rated all the movies in the pattern. Using abstraction on one pattern at a time, multiple patterns can be rewritten into the same abstracted form if ratings of items in the patterns cover most ratings for their parent node in a time taxonomy. For example, a pattern of the form {*Independence Day (ID4), Total Recall, Star Wars: Episode V are watched by users in IL, KS, NE and MO states of U.S.*} could be abstracted once to {*Independence Day (ID4), Total Recall, Star Wars: Episode V are watched by users in center of U.S.*} using location taxonomy if most ratings for those movies come from the states in the center of U.S. It could again be abstracted into {*Action movies are watched by users in center of U.S.*} if most ratings for *Action* movies are covered by those 3. This last abstraction relies on a movie genre taxonomy.

The intuition behind refinement is to enable the understanding of users that constitute the support of a pattern. The pattern {*The Fugitive, Terminator 2, Men in Black, The Matrix*} in MovieLens has a support of 1054 users. Refinement reveals more information on those users by providing a mechanism for exploring their demographics. To enable that, we need a mechanism that identifies which subsets of users in the support of a pattern we need to focus on. In our example, if for instance most support users of the pattern are [28-33] years old, this age bracket would qualify as a refinement for that pattern. To enable that, refinement is a primitive that relies on a notion of *saliency* for user demographics in order to examine the distribution of values of different demographics attributes in a pattern.

1.2 Contributions

We formalize pattern abstraction and refinement as two pattern exploration primitives. We then study experimentally the potential of our primitives in reducing the space of patterns, making them more compact and hence more readable, and in providing a better understanding of the support users of patterns. Our work lays the ground for exploring how time, geography and item taxonomies as well as demographics attributes enable pattern exploration using behavior semantics.

The paper is organized as follows. Section 2 describes our data model, patterns and primitives. Section 3 contains an evaluation of our primitives. Related work is reviewed in Section 4. Finally, Section 5 concludes with a discussion of ongoing and planned efforts.

2. FRAMEWORK

2.1 Data Model: Taxonomies

We are given a set of users U, items I, locations L and a database D of quadruples of the form $< u, i, l, t >$ where $u \in U$, $i \in I$, and $l \in L$ and t, a time-stamp, represent the location and time user u has used (opened, watched, rated, voted, etc.) item i.

Each user u is also described with attributes drawn from a set of attributes A representing demographics information such as *Gender* and *Age*. We refer to each attribute in A as a_i and to its values as v_j^i. The domain of values of attribute a_i is D_{a_i} with $D_A = \cup D_{a_i}$. For example, if we use a_1 to refer to *Gender*, it takes two values v_1^1 and v_2^1 representing *male* and *female* respectively.

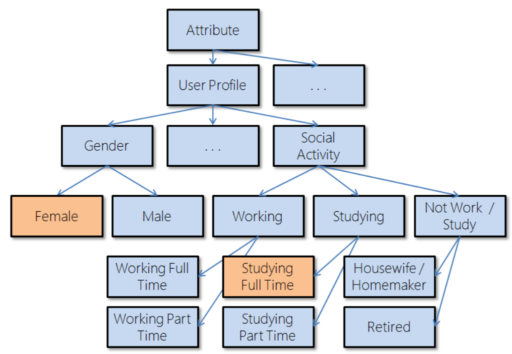

Figure 1: τ_A: User Attribute Taxonomy

User attributes, items, location and time, are organized in hand-crafted taxonomies. The values of each user attribute in A are organized in a taxonomy τ_A (Figure 1). Similarly, items in I (applications in Nokia and movies in MovieLens) and locations in L are organized into their respective taxonomies τ_I and τ_L . The set of all taxonomies is referred to as $\mathcal{T}$. We do not aim to show all the taxonomies we built for our datasets, rather we illustrate some examples that will be used later in the paper. In particular, the time taxonomy is omitted. Figure 2 shows a subset of the taxonomy we built for Nokia applications. Figure 3 shows a subset of the location taxonomy for MovieLens. Finally, Figure 4 shows the taxonomy for MovieLens movies.

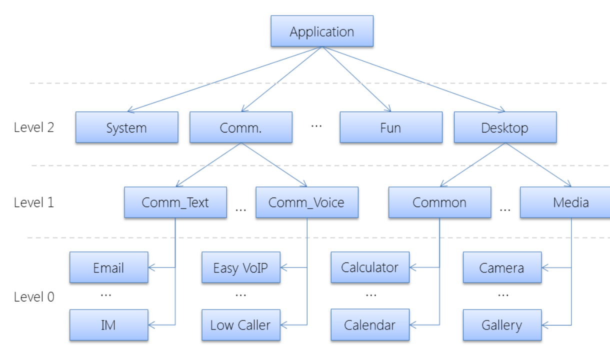

Figure 2: τ_I: Application Taxonomy

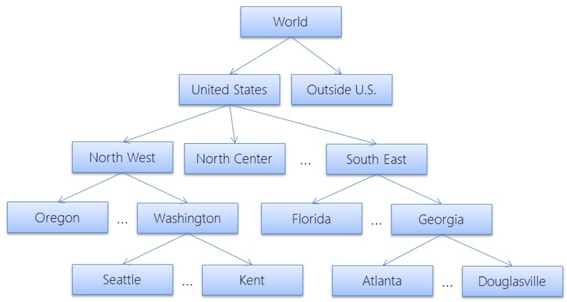

Figure 3: τ_L: Location Taxonomy

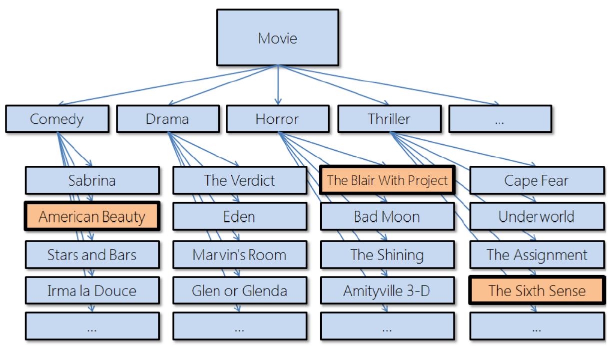

Figure 4: MovieLens Movie Taxonomy

2.2 Patterns

Given a database D, we are interested in finding patterns of the form $p = < v_1, \ldots, v_k, i_1, \ldots, i_m >$ where $\{v_1, \ldots, v_k\} \subseteq D_A$ and $\{i_1, \ldots, i_m\} \subseteq I$. In our context, patterns reflect user behavior and a pattern p can be read as users with attribute values $v_1, \ldots, v_k$ used items $i_1, \ldots, i_m$. The number of users in D satisfying p is referred to as *the support of* p and denoted $support(p) = |users(p)|$. Typically, only patterns satisfying a minimum support value σ are retained.

Table 1 contains example patterns and their support retrieved from Nokia and MovieLens, using a closed frequent itemset mining algorithm [1].

Dataset	Pattern	Sup. (%)
Nokia	*Female, Age 39-45, Calculator, Calendar, Bluetooth, Clock, Messaging*	13
MovieLens	*The Fugitive, Terminator 2, Men in Black, The Matrix*	17

Table 1: Example Patterns

2.3 Abstraction Primitive

A pattern mining algorithm explores a very large space (exponential in the number of items) and can return long patterns (the lower the support threshold, the longer the patterns). In order to enhance pattern readability, we propose to use semantic information provided in taxonomies to abstract items in a pattern into their parent item in the taxonomy. The intuition behind abstraction is simple yet powerful. Our abstraction method is not merely syntactic and relies on a taxonomy-based usage measure and reflects a way of *approximating the interest of users*. This approximation could be applied to items, time of day or to location.

2.3.1 Definitions

We define the *usage* of an item i for a set of users $V \subseteq U$, $usage(V, i) = |\{< u, i > \in D \mid u \in V\}|$ as the number of times users in the set V used item i. The usage of an item i wrt a pattern p, $usage(users(p), i)$ is the number of times the item i has been used by users who satisfy p.

Definition 1. **Taxonomy-Based Usage**. *Given a set of sibling items $i_1, \ldots, i_n$ and their parent item $\hat{i}$ in the taxonomy τ_I, their taxonomy-based usage in a pattern p, denoted $Pusage(p, i_1, \ldots, i_n) = \frac{\Sigma_{i_i} usage(users(p), i_i)}{usage(users(p), \hat{i})}$, is the proportion of usage between sibling items and their parent in τ_I.*

The intuition of taxonomy-based usage is that if most of the usage of a given item is that of some of its children in the taxonomy τ_I, those children could be replaced by their parent in all patterns they appear in thereby reducing the size of those patterns and making them more readable.

Definition 2. **Valid Pattern Abstraction**. *Given an abstraction threshold ρ and a pattern p containing sibling items $i_1, \ldots, i_n$ whose parent in τ_I is $\hat{i}$, we say that a pattern p_a is a* valid abstraction *of a pattern p iff $Pusage(p, i_1, \ldots, i_n) \geq \rho$ and $\forall i_i, i_i \notin p_a$ and $\hat{i} \in p_a$.*

Definition 3. **Maximal Pattern Abstraction**. *Given an abstraction threshold ρ, we say that a pattern p_a is a* maximal abstraction *of a pattern p iff p_a is a valid abstraction of p and $\nexists i_1, \ldots, i_n \in p_a$ s.t. $Pusage(p, i_1, \ldots, i_n) \geq \rho$ is satisfied.*

Let us now illustrate the definitions above on our datasets.

2.3.2 Pattern Abstraction in Nokia and MovieLens

In this section, we show some examples of abstraction using taxonomies. In Nokia, pattern $p = \{$*Studying Full-time*, *Female*, *FG Thread*, *WLAN Wizard*, *Calculator*, *Calendar*, *Bluetooth*, *Contacts*, *Log*, *Web*, *Text message*, *Messaging*$\}$ has a support equal to 4. Given an abstraction threshold of 50%, we obtain *a maximal abstraction* of p using application taxonomy into $p_a = \{$*Studying Full-time*, *Female*, **System**, *WLAN Wizard*, *Calculator*, *Calendar*, **Desktop Communication**, **Web App**$\}$ where the highlighted applications are parent classes in the taxonomy.

The pie charts A, B and C in Figure 5 show usages that enable a recursive abstraction of pattern p_1 into p_{a1}. In pie chart A of Figure 5 we can see that 87.68% of usage for item *Desktop Communication* is for its children items *Bluetooth*, *Contacts*, *Log*, *Text Message* and *Messaging*. Pie charts B and C of Figure 5 contain two other usages, one showing that 50% usage of *System* items is for *FG Thread* and another showing that 53.61% of usages of *Web App* items, is for *Web*. Finally, in pie charts D and E of Figure 5, we show two examples of non-valid abstractions given a threshold equal to 50%. We see that 9% usage of *Configuration* items, is for *WLAN Wizard* and that 22% of usage of *Desktop Common* items, is for *Calculator* and *Calendar*. Thereby, none of those could be abstracted in pattern p.

As another example, pattern p=$\{$*Engineer*, *Age 18-45*, *Batman*, *Jurassic Park* in *NY, MA, MN, MI, OH, TN states of U.S.*$\}$ becomes abstracted to p_a=$\{$*Engineer*, *Age 18-45*,

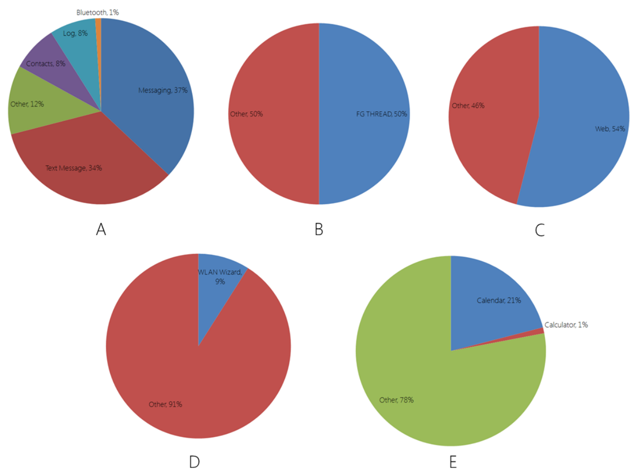

Figure 5: Item Usages for Abstraction

Batman, *Jurassic Park* in **North of U.S.**} using location taxonomy. The abstraction is possible in location because most ratings for movies in the above pattern come from users in the location specified in the abstracted pattern.

2.4 Refinement Primitive

The goal of refinement is to provide an informed exploration of users in the support of a pattern. User attributes constitute an opportunity for doing so. Given the set of users in the support of a pattern, we propose to identify demographics attributes that can be used for further exploration. Just like abstraction, refinement is a simple yet powerful primitive as it constitutes a way to characterize pattern users by exploiting the richness of their demographics. Our refinement method is not merely syntactic and relies on *computing a saliency measure* in order to best determine which subset of users in the support of a pattern is most interesting to further explore.

2.4.1 Definitions

Many pattern interestingness measures have been suggested in the literature as summarized in [2]. However, they were designed to pick representative patterns in a large pattern space. None of them was designed to explore users in the support of a pattern.

DEFINITION 4. **Attribute Saliency**. *The saliency of an attribute* a_i *wrt a pattern* p*,* $sal(a_i, p)$*, is a measure of interestingness of the attribute* a_i *for users in* $users(p)$*.*

We intentionally keep the definition of saliency general in order to explore different ones. Alternative measures are variance, entropy, or a measure that computes the ratio between distribution of values of pattern users for an attribute (say age) with that same distribution for all users in U. Such a measure aims at selecting user attributes for which pattern users differ from all users.

$sal(a, p)$ can be calculated using *standard deviation* or *entropy* measures. Standard deviation is the one we use in this paper in our examples and experiments and it measures the amount of variation or dispersion from the average, in a list of values.

Having a low standard deviation score for an attribute means its values tend to be very close to the average. Also having a high score for an attribute means its values are spread out over a large range of values.

DEFINITION 5. **Valid Pattern Refinement**. *Given a saliency threshold* μ *and a pattern* p*, we say that a pattern* p_r *is a* valid refinement *of a pattern* p *iff* $sal(a, p) \geq \mu$ *and* $a \in p_r$ *is a user attribute value.*

We propose to calculate all patterns that constitute valid refinements of a given pattern p and associate the k best refinements (that is, patterns) to p. Those refinements constitute alternative explorations of the support users of p.

2.4.2 Pattern Refinement in Nokia and MovieLens

Consider again pattern p = { *Studying Full Time*, *Female*, *FG Thread*, *WLAN Wizard*, *Calculator*, *Calendar*, *Bluetooth*, *Contacts*, *Log*, *Web*, *Text Message*, *Messaging*} in NOKIA, with a support equal to 4 users. We show how standard deviation can be used to explore users of this pattern.

We report in Figure 6, two attributes having an especially high saliency. On the left, *Age* distribution of the pattern users is shown. We can see that 3 out of 4 users are in age category [28-33]. This non-homogeneous distribution leads to high saliency for attribute *Age* for this pattern. That is why standard deviation is a good measure for *Age* in pattern p. This information indicates the existence of a super-pattern {p, *Age 28-33* } with support 3 that may be of interest for further exploration.

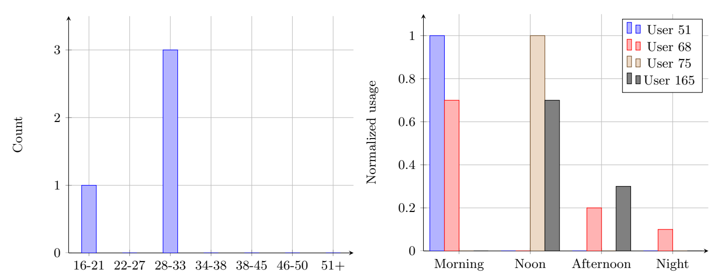

Figure 6: Saliency Input for *Age* (left) and *Calculator* (right)

On the right part of Figure 6 the time of day usage for the *Calculator* application is illustrated by all 4 users in pattern p. Once again, this usage is not homogeneous, leading to a high saliency. Users 51 and 68 mostly use *Calculator* in the morning, while users 75 and 165 mostly use it at noon. This indicates two sub-populations of 2 users, and the existence of two refinements {p, *Calculator Morning* } and {p, *Calculator Noon* }, each supported by 2 users. Mining again the dataset with a minimum support threshold of 2 would discover these super-patterns. They may have additional attributes giving more information on the demographics and specific application usage of these sub-populations.

3. EVALUATION

The goal of this section is to evaluate *abstraction* and *refinement* primitives on NOKIA and MOVIELENS datasets. We propose quantitative and qualitative evaluations. We discuss some interesting results for each evaluation.

3.1 Datasets

Nokia consists of data from smartphones of some participants in the course of more than one year. For each user, all records of phone events and sensors like application usage, calendar, contacts, and call-logs are logged with a time-stamp. Personal information is anonymized in the data. In our study, we focused on application usage: the opening of applications by users indicating what they use their smartphones for, at any time of day. This dataset also includes responses to a questionnaire by some users in the experiment. Demographic attributes like gender, age group, and profession come from that questionnaire. Application usage records consist of an application ID and a time-stamp of when it was used. After removing some core system applications, we ended up with 170 applications.

MovieLens is the dataset published by the GroupLens research group[1]. We used the 1M ratings version that contains 1,000,209 anonymous ratings of 3,952 movies by 6,040 users. Rating records consist of a user ID (between 1 and 6040), movie ID (between 1 and 3952), a rating (based on a 5-star scale) and time. Each user provided at least 20 ratings. When user X has rated movie Y, it means X has watched Y. This is how we define the usage or consumption of a movie by a user. For each user, gender, age group, occupation and zip-code are provided. All demographic information is provided voluntarily by users.

We pre-processed the datasets and ran the LCM closed frequent itemset mining algorithm [1] with a minimum support threshold of 7% that resulted in 74723 patterns for Nokia and 50,299,230 patterns for MovieLens.

3.2 Abstraction Evaluation

In order to evaluate the benefit of abstraction, we propose to explore abstraction volume and pattern space reduction as described below.

3.2.1 Abstraction Volume

As seen in Definition 2, the abstraction primitive only abstracts group of items of a pattern if a condition is met. We want to evaluate how often this condition is met, depending on the abstraction threshold ρ chosen. We thus define an *abstraction volume* measure, which evaluates for each pattern the ratio between the number of abstractions performed (given $\rho > 0$) and the maximal number of abstractions possible (case of $\rho = 0$).

Given N the number of occurred abstractions in the pattern p and M the total number of classes of the taxonomy that have at least one of their child items in p, the abstraction volume of p denoted by θ is equal to (N / M * 100). We perform *abstraction volume* experiment on patterns from Nokia mined with minimum support threshold of 7%. Patterns may include demographic information and applications. We applied the *abstraction* method using different *abstraction* thresholds ρ varying from 0 % to 100 %. Figure 7 shows the result of this experiment. The evolution of *abstraction volume* can be categorized into three different periods by two cutting points $M_1 = 15\%$ and $M_2 = 60\%$.

Before M_1 (where the abstraction threshold ρ is between zero and 15 %), we observe a very mild slope in the diagram and the *abstraction volume* decreases only 10 %. It shows that in Nokia, low values of ρ lead to many abstractions.

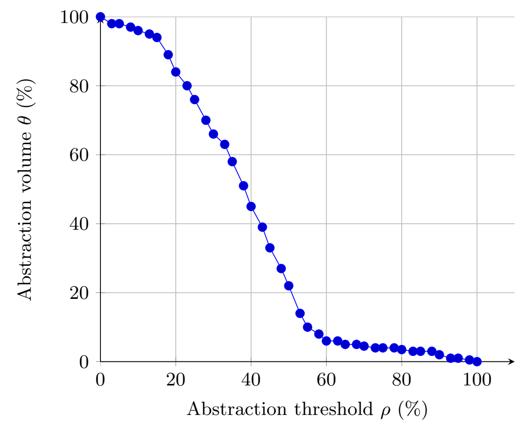

Figure 7: Abstraction Volume

Choosing ρ in this range causes to have many abstractions with less attention to the usage of attributes. It will abstract nearly syntactically except for some extreme cases where usage is very low. Therefore, it is useful to select an abstraction threshold in this range only when data does not provide many usage information.

After M_2 (where ρ is higher than 60 %), we observe that the *abstraction volume* decreases drastically and it remains very close to zero. It means that in this range, the number of abstractions done is very low. Thus, choosing the abstraction threshold in this range is useless for simplifying the analysis.

Between M_1 and M_2 (where ρ is between 15 to 60 %), we observe that the plot has a derivative close to -1. Thus changing ρ in these values gives a predictable reduction in the number of abstractions performed. We thus consider that this range of ρ threshold values is the most interesting, and we will focus on it for most of the following experiments.

3.2.2 Pattern Space Reduction

Applying abstraction on distinct patterns will sometimes result in the same abstracted pattern. Hence, in addition to reducing pattern size, a beneficial side effect of the abstraction primitive is to reduce the size of the pattern space. We want to evaluate the scale of this pattern space reduction experimentally.

Given a support threshold σ and an abstraction threshold ρ, the pattern space reduction is equal to $1 - \frac{|P_a|}{|P|}$ where P_a is the set of all abstracted patterns and P is the set of initial (not abstracted) patterns. For Nokia, we generated the set of maximal abstracted patterns using 4 different values for σ i.e. 10, 25, 50 and 75% by varying ρ from 0% (i.e. syntactic abstraction) to 100% (*i.e.* no abstraction). The result is shown in Figure 8.

As an example, using abstraction with $\sigma = 25\%$ and $\rho = 20\%$, the pattern space reduces to half of its initial size. The three periods mentioned in Figure 7 with the cutting points $M_1 = 15\%$ and $M_2 = 60\%$ are also visible in Figure 8, with a pattern space reduction between 20 and 30% in the most interesting range $[M_1, M_2]$. When fixing the abstraction threshold ρ, the lower the support threshold, the higher the reduction of the pattern space. However, for low support values the gain in reduction is from lowering the support threshold. This can be explained by the

[1]http://www.grouplens.org/

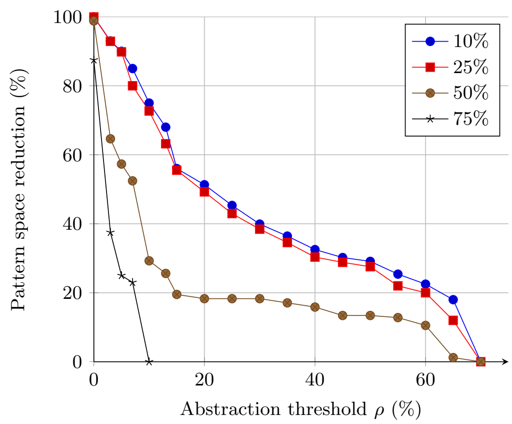

Figure 8: Pattern Space Reduction

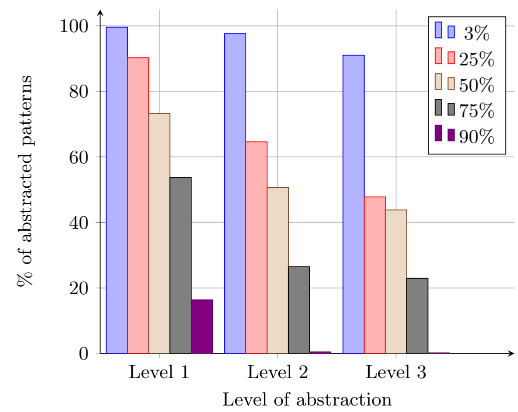

Figure 9: Abstraction per Level for Nokia

fact that with lower support thresholds, longer patterns are produced, many of them having mostly the same items (can differ by one item or two only). Thus abstraction is more likely to abstract those patterns to the same pattern. However the lower the support value, the smaller the number of users supporting the patterns, which reduces the differences in usage values (recall that Nokia has only 38 users).

These results show that the space reduction given by abstraction is not negligible, and can help reduce the burden on the analyst.

To reach maximal abstraction, in worst case an item may be abstracted at most 3 times, coming from the depth of the Nokia application taxonomy (as shown in Figure 2). It is interesting to see the influence of successive iterations of the abstraction primitive, and the distribution of abstracted items in the different levels of the application taxonomy. This result is presented in Figure 9. For each application of the abstraction primitive, and thus each level of the taxonomy as shown by Figure 2, the percentage of patterns that got abstracted to a class of the taxonomy of that level is shown. The bars correspond to different abstraction thresholds ρ, the support threshold is fixed to $\sigma = 25\%$.

One can note that for too low abstraction thresholds ($\rho = 3\%$), 90% of patterns are abstracted to the top level of the taxonomy, which is the least informative: it confirms the poor interest of such low abstraction thresholds. Conversely excessively high abstraction thresholds ($\rho = 90\%$) lead to less that 20 % of patterns abstracted on the lower level of the taxonomy, and near no higher level abstraction: this is not enough to help the analyst. On the other hand, abstraction thresholds between the bounds M_1 and M_2 defined before lead to a reasonable percentage of patterns abstracted per level, with a decrease of more than 20% of patterns abstracted from level 1 to level 3. This indicates that the analyst will be presented with patterns containing a mixture of classes from the taxonomy, which is what is expected to help in the analysis.

3.3 Refinement Evaluation

The goal of refinement is to restrict the number of choices at each step of the exploration so that the analyst is not confronted with thousands of choices. Saliency allows to do it in a principled way, and only present potentially "interesting" choices to the analyst. We run two experiments to evaluate refinement. In the first one, we evaluate quantitatively the practical reduction in the number of choices. In the second one, we evaluate the quality of obtained refinements.

3.3.1 *Exploration Choice Reduction*

For each pattern p, we count the number of exploration choices (i.e. number of valid refinements of p), and compute the average result for all patterns. The saliency threshold is fixed to $\mu = 50\%$, and we vary the support threshold σ. The results are shown in Figure 10 for MovieLens.

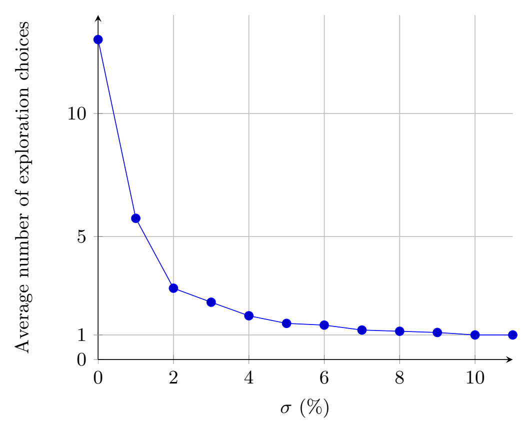

Figure 10: Average number of Exploration Choices for Patterns of MovieLens

As shown in Figure 10, with the lowest possible support threshold $\sigma = 0\%$, there exist in average 13 exploration choices for patterns of MovieLens, which is an extreme case. For more realistic support thresholds such as $3 \leq \sigma \leq 6\%$, there exist between 1 and 3 choices in average: the analyst is not overwhelmed with choices, but is often offered more than one choice, which suffices to allow different directions of exploration. One can note that in this experiment, with $\sigma \geq 10$ there is only 1 choice on average: saliency becomes too strict. For such higher values of support the saliency threshold should be relaxed to get more choices, but then it would be less adapted to lower support values.

3.3.2 *Refinement Qualitative Evaluation*

Beyond the quantitative aspect of refinement, we evaluate the quality of *refinement* primitive in an extensive user study. Through the survey, we evaluate the usefulness of refinement primitive by measuring if users find the provided histograms (plotting values of an attribute for support users) with the *refinement* primitive informative and prefer to observe the histograms that the primitive detects as more salient. We have used 2 patterns from Nokia and 2 from MovieLens.

In a comparative study, we present two histograms for each pattern (a salient and a random non-salient one) and ask the participant which histogram is more useful to be attached to the pattern.

In the second part of our study, we seek user feedback in order to independently evaluate the usefulness and meaningfulness of refinement. We present a pattern with the most salient histogram detected by our measure (standard deviation). We ask the participant if she considers the histogram informative. Table 2 shows overall results.

	Positive	**Negative**
Comparative (%)	69	31
Independent (%)	83	17

Table 2: Qualitative Evaluation

The *positive* option for the comparative study means preference for the salient histogram, and the *negative* option, otherwise. Also the *positive* option for the independent study shows the percentage of participants that have found the associated histograms informative, and the *negative* option, otherwise.

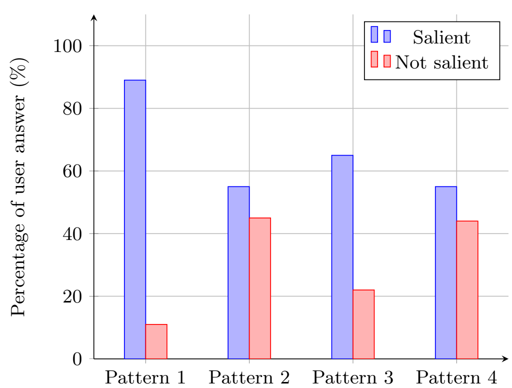

Figure 11: Comparative Refinement Evaluation

Figure 11 shows the percentage of responses for each pattern in the comparative evaluation. In all 4 patterns, people have preferred the salient histograms to random ones. We observe a high rate of positive answers (close to 90 %) for the first pattern. The value of saliency for the attribute of the histogram shown for this pattern is equal to 0.35 which counts as a high saliency. But we observe just a slight superiority of positive answers (around 10 %) for the second and fourth patterns, because the values of saliency for the attribute of the histogram associated to these patterns were not as high as the first pattern. It shows that the more salient an attribute is, the more people prefer its histogram.

The result above suggests to consider different saliency measures for different demographic attributes. We plan to explore that in future work.

Figure 12 shows the assessment of how informative histograms are for each pattern. Users found the histograms informative for all 4 patterns. For MovieLens patterns (3 and 4) the percentage of *not informative* responses is close to zero.

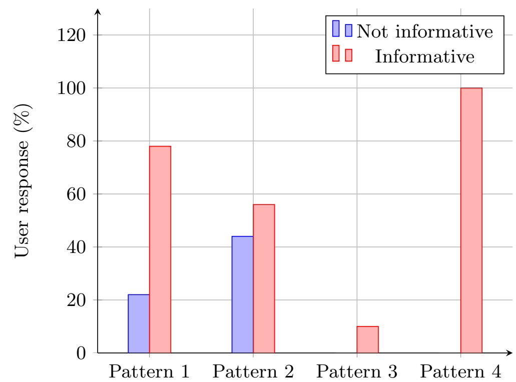

Figure 12: Independent Refinement Evaluation

4. RELATED WORK

There has been quite a bit of research activity to help analysts explore the space of patterns. We can however safely claim that none of them is data-driven and a lot of work is left to the analyst. Moreover, existing pattern interestingness measures have been used for selecting representative patterns but none of them was used to explore the support users of a pattern. Finally, in both our primitives, we exploit *usage* data that is not present in the transaction dataset given as input to frequent pattern mining algorithms, but that appears in the original data. This allows us to go beyond what just traditional methods do based on the transaction datasets.

B. Goethals et al. propose an interactive pattern exploration framework called MIME in [4] as an iterative process, so the analyst would be able to explore and refine the discovered patterns on the fly. In MIME, the analyst becomes an essential part of the mining algorithm as she has to select the items to include in the pattern for further exploration. Also, there is no data-driven navigation as in our case. The analyst is left alone to make an educated choice.

A *constraint-based mining* approach [5, 6] can also be seen as a pattern exploration mechanism where the analyst can iterativley tune the constraints to generate additional patterns to explore. Designing constraints is not an easy task and requires an appriori knowledge of the dataset.

In [7, 8], an approach is proposed to learn the model of prior knowledge of the analyst based on her exploration actions. In [7] the analyst has to order her exploration choice preference which puts more burden on the analyst. These methods can be complementary to ours to make a refinement biased towards previous choices by the analyst.

We exploit taxonomies for abstraction that helps reduce the space of patterns. Our *taxonomy-based usage* is an interestingness measure and it has the same principle as defined in previous work [9, 2]. The difference is that we calculate our measure for items in a pattern and not necessarily for a whole pattern. The reason is that we are not interested in pruning a pattern, but in abstracting parts of it. The method used in [10, 11] is a top-down approach and is the most similar work to our *abstraction* method.

The idea of *refinement* for semantic exploration can be categorized as an interestingness mining approach. In our work, salient demographics attributes can be visualized as histograms associated with a pattern. In [12, 13, 14], the idea is to mine a small set of interesting patterns using novel interestingness measures. Those measures are computed for whole patterns and are used to select representative ones as opposed to explore users of a given pattern. A complete list of measures used in the literature can be found here [2].

5. DISCUSSION

In this paper, we addressed the problem that arises when analyzing the behavior of a large number of users with frequent pattern mining, namely the discovery of a large number of patterns and the length of each pattern. We proposed abstraction and refinement primitives to navigate in the space of frequent patterns and better understand their users. Our evaluation on two real datasets showed that abstraction reduces the size of the pattern space and produces more readable patterns. It also shows the usefulness of refinement in guiding the analyst in the exploration of the behavior of well-defined sets of users as dictated by the data.

Our primitives are key for the implementation of an interactive exploration framework for frequent usage patterns. To do so we need to devise an algorithm that combines our primitives in such a way that provides to the analyst a number of alternative navigations in the space of patterns. Instead of relying solely on the lattice induced by pattern mining (pattern generalization and specialization), our algorithm could guide the analyst in exploring different users communities formed by patterns in the lattice. We hence envision an interactive framework within which users alternate between exploring patterns and exploring user communities induced by them and described by a combination of demographics attributes. For example, an analyst could from the pattern *Users of Communication and Web Search applications who live in Lausanne* and see two alternative subsets of those users, ones who are students and live on EPFL Avenue and use those Messaging applications, and ones who are stay-at-home users and use Google in the afternoon. For each subset, the analyst could ask to see other patterns. This process is iterative and requires the ability to compute patterns and communities on-demand. We are currently exploring the use of scalable indexing techniques to enable such flexibility.

Refinement is a principled way of identifying user communities of interest for which activity is known. We consider this a starting point for exploring subsets of users and plan to use the algorithms developed in [15] and in [16] to do so. There is an opportunity to specialize the exploration depending on the type of action users perform in the underlying dataset. For example, in the case of collaborative rating such as MovieLens, rating exploration may require to search for subsets of users whose ratings are uniform or polarized wrt to the movies contained in a pattern as in [15] whereas in the case of Nokia, the duration of usage of the set of applications embedded in a pattern is more appropriate. In our immediate future work, we plan to investigate the applicability of different action-aware exploration algorithms to complement pattern refinement.

6. REFERENCES

[1] T. Uno, M. Kiyomi, and H. Arimura, "LCM ver. 2: Efficient Mining Algorithms for Frequent/Closed/Maximal Itemsets," in *Workshop on Frequent Itemset Mining Implementations*, 2004.

[2] L. Geng and H. J. Hamilton, "Interestingness measures for data mining: A survey," *ACM Computing Surveys*, vol. 38, 2006.

[3] T. Uno, T. Asai, Y. Uchida, and H. Arimura, "Lcm: An efficient algorithm for enumerating frequent closed item sets," in *In Proceedings of Workshop on Frequent itemset Mining Implementations (FIMI03*, 2003.

[4] B. Goethals, S. Moens, and J. Vreeken, "Mime: a framework for interactive visual pattern mining," in *KDD*, pp. 757–760, 2011.

[5] F. Bonchi, F. Giannotti, A. Mazzanti, and D. Pedreschi, "Exante: Anticipated data reduction in constrained pattern mining," in *PKDD*, pp. 59–70, 2003.

[6] C. Bucila, J. Gehrke, D. Kifer, and W. M. White, "DualMiner: a dual-pruning algorithm for itemsets with constraints," in *Knowledge Discovery and Data Mining*, pp. 42–51, 2002.

[7] D. Xin, X. Shen, Q. Mei, and J. Han, "Discovering interesting patterns through user's interactive feedback," in *Knowledge Discovery and Data Mining*, pp. 773–778, 2006.

[8] T. D. Bie, K.-N. Kontonasios, and E. Spyropoulou, "A framework for mining interesting pattern sets," *Sigkdd Explorations*, vol. 12, pp. 92–100, 2011.

[9] C. J. Matheus, G. Piatetsky-shapiro, and D. Mcneill, "Selecting and reporting what is interesting: The kefir application to healthcare data."

[10] R. Srikant and R. Agrawal, "Mining Generalized Association Rules," in *Very Large Data Bases*, pp. 407–419, 1995.

[11] C. Marinica, F. Guillet, and H. Briand, "Post-processing of discovered association rules using ontologies," *CoRR*, vol. abs/0910.0349, 2009.

[12] M. Mampaey, N. Tatti, and J. Vreeken, "Tell me what i need to know: succinctly summarizing data with itemsets," in *KDD*, pp. 573–581, 2011.

[13] T. D. Bie, "Maximum entropy models and subjective interestingness: an application to tiles in binary databases," *Data Mining and Knowledge Discovery*, vol. abs/1008.3, pp. 1–40, 2010.

[14] P.-N. Tan, V. Kumar, and J. Srivastava, "Selecting the right interestingness measure for association patterns," in *Knowledge Discovery and Data Mining*, pp. 32–41, 2002.

[15] M. Das, S. Amer-Yahia, G. Das, and C. Yu, "Mri: Meaningful interpretations of collaborative ratings," *PVLDB*, vol. 4, no. 11, pp. 1063–1074, 2011.

[16] M. Das, S. Thirumuruganathan, S. Amer-Yahia, G. Das, and C. Yu, "Who tags what? an analysis framework," *PVLDB*, vol. 5, no. 11, pp. 1567–1578, 2012.

A Method for Activity Recognition Partially Resilient on Mobile Device Orientation

Nikola Jajac, Bratislav Predic, Dragan Stojanovic
University of Nis, Faculty of Electronic Engineering
18000 Nis, Serbia
{nikola.jajac, bratislav.predic, dragan.stojanovic}@elfak.ni.ac.rs

ABSTRACT

This paper demonstrates a method for activity recognition partially resilient on mobile device orientation, by using data from a mobile phone embedded accelerometer. This method is partially resilient on mobile device orientation, in such a way that a mobile device can be rotated around only one axis for an arbitrary angle. The classifier for activity recognition is built using data from one default orientation. This method introduces a calibration phase in which the phone's orientation is determined. After that, accelerometer data is transformed into the default coordinate system and further processed. The solution is compared with the method that built a classifier using data from multiple orientations. Three classifiers were tested and a high accuracy of around 90% was achieved for all of them.

1. INTRODUCTION

The advance of low-cost and low-power sensors has led to their massive integration into modern mobile devices that have become powerful sensing platforms. By using data from sensors like accelerometers, gyroscopes, digital compasses, light sensors etc. it is possible to describe the context of a mobile device user in much more detail than by using only location data. A richer description of a user context provides better adaptation of mobile content, services and resources, enabling the user to stay focused on the task at hand.

An activity of the user represents an important aspect of a context, because it directly impacts user's ability to interact with the mobile device and applications. By employing the information about the current activity, the mobile device can adapt its interfaces, filter the content it provides, or perform a specified action, to support the user in the best possible way.

The greatest possibilities for application of activity recognition systems lay in the healthcare domain. For example, such systems can be used for elderly care support or for long-term health/fitness monitoring [1]. Current methods for tracking activities, like paying a trained observer or relying on self-reporting are time and resource consuming tasks, and are error prone. An automatic system for recognizing activities could help reduce errors that arise from previously mentioned methods. Also, such system enables users to go about their daily routines, while the data collection and processing are done in the background, and do not interfere with their current activities.

In recent years a lot of work has been done on activity recognition from accelerometer data. Since an accelerometer is a standard part of modern mobile devices (like mobile phones and tablets), they can be used in activity recognition. An advantage in using these devices is that they are already commonly used by a lot of people that would not have to wear an additional device to perform activity recognition, which greatly increases the acceptance of such a system.

This paper focuses on recognizing activities from accelerometer data processed at a mobile phone. Activity recognition is formulated as a classification problem. This paper considers the healthcare domain for activity recognition system application. For this reason it is important to recognize physical activities, such as: walking, running, walking up/down stairs etc. Examples of possible users in this domain are the elderly and persons with certain disabilities. It can be assumed that such users keep their mobile phone relatively (but not completely) fixed. In the paper it is assumed that the phone position is fixed, but that the phone orientation is only partially fixed. The orientation is partially fixed in such a way that the axis that is perpendicular to the phone's screen is parallel to the ground and the phone can be rotated freely around that axis. For experiments in this paper, the mobile phone was worn in the right front pants pocket (as one of the places where people usually carry mobile phones) and the screen of the phone was facing the user. The orientation where the bottom side of the phone is facing the ground is considered the default orientation. First, three classifiers were built using data from multiple orientations. These classifiers were tested using data from the same orientations. Then, classifiers were built again, but using data from the default orientation only. These new classifiers were tested using data from multiple orientations, transformed into the default coordinate system prior to testing, according to the method proposed in this paper. After that, classifier performances from these two tests were compared.

Using results from previously mentioned tests, we created an application for activity recognition that runs on a mobile device in real time and tested the impact that data transformation has on the performance of such an application, specifically in terms of processor load. The processor load is important because if such application is to be accepted by users, it must not significantly decrease battery life, or the performance of the device in everyday tasks.

The proposed method represents only an intermediate step in the development of a method for activity recognition with no restrictions in mobile device orientation.

The rest of the paper is structured as follows: section 2 provides an overview of related work on activity recognition. Section 3 describes the process of accelerometer data collection. In sections 4 and 5, mixed orientation data and data reorientation processing approach, for activity recognition resilient on mobile device orientation, are presented and evaluated. Section 6 presents an evaluation of how much the activity recognition application and data transformation phase participate in the processor load. Section 7 gives the conclusions about the paper and outlines plans for future work.

2. RELATED WORK

In recent years there has been a lot of research related to recognizing activities from accelerometer data. In [2] authors used data from 5 biaxial accelerometers worn simultaneously on different parts of the body. Used accelerometers could detect acceleration up to ±10G. Accelerometers were mounted onto hoarder boards and firmly attached to different body parts. Data was collected from 20 subjects performing various everyday tasks without researcher supervision. The following features were computed on sliding windows of accelerometer data: mean, energy, frequency-domain entropy and correlation. A number of classifiers were trained and tested with the calculated data, where decision trees showed the best result, recognizing activities with an accuracy of 84%.

Ravi et al. in [3] attempted to perform activity recognition using a single triaxial accelerometer worn near the pelvic region. Data was collected by 2 subjects performing 8 different activities. Similarly to [2] the features were computed using the sliding windows technique. Four features were extracted: mean, standard deviation, energy and correlation. Extracted features were used to train and test 5 base-level classifiers, and in addition to that, 5 meta-level classifiers. Authors concluded that meta-level classifiers in general outperform base-level classifiers and that plurality voting, which combines multiple base-level classifiers, shows the best results. The authors also showed that out of the used features, energy is the least significant one, and that there is no significant change in accuracy when this feature is avoided.

Kwapisz et al. in [4] tried to recognize activities by using data from a single acceleration sensor, but they used data from an acceleration sensor embedded into a standard mobile phone. These accelerometers typically detect acceleration up to ±2G along three axes. Their research methodology follows the one in [2, 3]. The authors collected data from 29 subjects, extracted 6 basic features and tested 3 classifiers, where multilayer perceptrons showed the best result, recognizing activities with an accuracy of 91.7%. The authors showed that activity recognition can be performed successfully by using acceleration data from a mobile phone.

The unifying fact for papers [2 - 4], no matter if one or more accelerometers are used, is that the position and the orientation of the accelerometer is fixed while performing all of the examined activities. This fact can be probably expected in case of specialized devices as in [2, 3]. In case of using a standard mobile phone as in [4], the method puts strains on how someone carries a mobile device, which could decrease acceptability of such a system.

Sun et al. in [5] tried to recognize activities by using acceleration data from a mobile phone, in a setting where the position and the orientation of the phone vary. They restrict their hypothesis space to 6 possible positions (6 pockets) and 4 orientations of the mobile phone. The data from all position and orientation combinations were collected. The authors added acceleration magnitude at each sample, as an additional sensor reading dimension. By using collected data, several features were calculated: mean, variance, correlation, FFT energy and frequency-domain entropy. Calculated features were used to train and test SVM (Support Vector Machine) models. Generated SVM model recognizes activities with an accuracy of 93.1% throughout all tested positions and orientations.

Thiemjarus in [6] applied a different approach. Accelerometer was mounted on a belt-clip device which was worn by test subjects in a fixed position on a body, but which could be mounted in 4 different orientations. Data was collected by 13 subjects that performed a routine comprised of 6 activities. The first step in data analysis was device orientation detection. The orientation detection was also formulated as a classification problem. The features used for orientation detection were mean along Y and Z axis. Orientation detection was performed for an activity routine performance, which contains approximately 5 seconds of data for each tested activity, while activity recognition was window-based. The second step was signal transformation using the appropriate transformation matrix, and the third step was activity recognition itself. Author achieved a subject-independent classification accuracy of 90.9%.

A possible problem with the approach in [6] is that orientation detection is done on a data set that includes information about all tested activities. The open issue is how this approach can be applied in a real world scenario, when the user does not perform a specified activity routine. A unifying fact for [5, 6] is that the hypothesis space is limited to a number of orientations. Again, the open issue is how such a system would perform when given data from an unknown orientation. In this paper we propose a method that only partially limits the device orientation, in a way that the device can be rotated only around one axis, but for an arbitrary angle. This arbitrary rotation practically creates an infinite number of possible orientations, in contrast of previous approaches, which are all limited to a certain number of orientations. To achieve this, a calibration phase which precedes activity recognition is introduced.

3. ACCELEROMETER DATA COLLECTION

As a test device, a smartphone Samsung I9001 Galaxy S Plus which runs on Android operating system version 2.3.5 is used. The accelerometer embedded in this phone detects acceleration up to ±2G. Data from the accelerometer has three attributes: acceleration along X, Y and Z axis, represented by floating point values. Sampling rate for the accelerometer was set to *SENSOR_DELAY_FASTEST* to achieve the highest possible accuracy.

An application for recording data from the accelerometer has been developed. Data was collected by several test users. The recording process is as follows: while standing still, test user selects the activity he is going to perform and starts the recording. After that the user has ten seconds to place the phone in the pocket in the desired orientation. After ten seconds a beep sound is played and for two seconds the gravity vector is extracted from accelerometer data. After two seconds another beep sound is played and an average value for the gravity vector is saved to a file. To extract the gravity vector from accelerometer data a simple low-pass filter is used. The user can then start to perform the specified activity. Another 2 seconds after the second beep, the application starts to record acceleration data to another file. After finishing with the activity the user stops the recording.

Data was collected while performing 6 different activities:

- Standing
- Walking
- Running
- Walking up stairs
- Walking down stairs
- Sitting.

For each activity data was collected for the default orientation and for 3-4 other orientations, depending on the activity. Some of the non-default orientations matched between activities and some did not. To minimize mislabeling a portion of data was removed from the beginning and the end of each recording.

4. A MIXED ORIENTATION DATA APPROACH

The first approach to the free orientation problem we test is building a classifier from data collected from all orientations, very similar to [5]. In this approach, the next step after data collection is feature extraction. The features were extracted from accelerometer data using a window size of 512 samples with 256 samples overlapping between consecutive windows. Three features were extracted from each of the three axes, giving a total of nine attributes for building a classifier. The features extracted were:

- Mean
- Standard deviation
- Correlation.

We specified these features, that are calculated using data in time domain, because we apply the activity recognition system in real-time locally on a device. For this reason, the features should be relatively simple to compute, to reduce power consumption and processor load. The selected features do not require signal representation in the frequency domain, and thus can be computed relatively fast. Also, the mean is used in the standard deviation calculation, and the mean and the standard deviation are used in correlation calculation, which further increases computation speed.

Extracted features were used to train and test 3 classifiers available in the WEKA Data Mining Toolkit [7], which are commonly used in activity recognition [2 – 4, 6]:

- C4.5 decision tree
- Naïve Bayes
- K-nearest neighbors.

We are mainly interested in the performance of the decision tree classifier, since it requires the least amount of computation in the classification phase, which is important when the system is applied in a real-time locally on a device. For the testing we used 10-fold cross validation, and the results are shown in Table 1. All of the tested classifiers showed excellent results in recognizing activities, which is consistent with previous work [5]. The results are slightly better than in [5] which can be probably contributed to a specific data set and the fact that the data was collected by a single test user.

Table 1. Classifier accuracy – Mixed orientation data approach

Classifier	Accuracy (in %)
C4.5 Decision Tree	98.8
Naïve Bayes	99.5
K-nearest neighbors	99.8

5. A DATA REORIENTATION PREPROCESSING APPROACH

The second approach to the free orientation problem is based on building a classifier from data collected in the default orientation. In the classifying phase, transformation of data collected in various orientations into the default coordinate system is performed, prior to the feature extraction and classification.

The classifiers for testing are built in the same way as in the mixed orientation data approach, but now only data collected from the default orientation is used. To get the most precise results, it is important to test the classifier with data from all available orientations, which includes also data from the default orientation. For this reason, a portion of data from the default orientation was omitted in classifier building, and was used later in classifier testing. In this way we avoid overfitting, which can happen when the same data is used for training and testing. No data transformation was done on data used for building a classifier.

As previously mentioned, in this paper we assume that the phone can be rotated only around the Z axis, and consequently, we assume that there is no change in the acceleration along the Z axis when performing some activity in the default and non-default position. This means that we do not need to transform the Z coordinate, just X and Y coordinates. To achieve that, we use a rotation matrix for rotation around the Z axis for an angle θ. The rotation matrix is given in (1). To calculate angle θ we use the information about gravity vectors.

$$R = \begin{bmatrix} \cos(\theta) & -\sin(\theta) & 0 \\ \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (1)$$

By using the gravity vectors from all of the recordings in the default orientation, we computed an average gravity vector for the default orientation (gravity vector is defined as data from the accelerometer, by three attributes: X, Y and Z). Since the phone's screen is facing the user, the Z axis is practically parallel to the ground, so we take into account only the X and Y components of the vector.

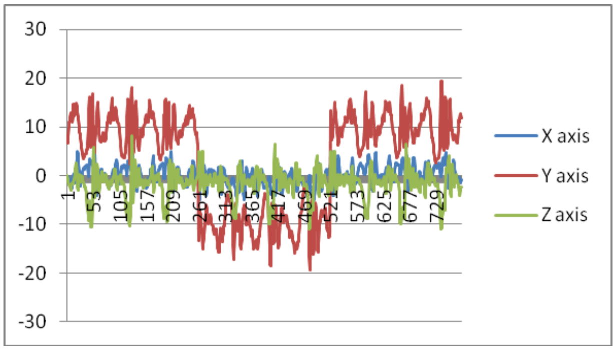

Figure 1. The accelerometer data collected while walking

In the classifying phase, we firstly calculate the difference between the angle of the average gravity vector in the default orientation and the angle of the gravity vector for the current orientation, and this difference is angle θ we need to transform accelerometer data into the default coordinate system. In the next step, each sample from the accelerometer is transformed into the default coordinate system using the rotation matrix given in (1). Figure 1 shows data from the accelerometer while walking. The vertical axis represents the acceleration in m/s^2. Samples from the acceleration sensor are represented along the horizontal axis. Samples 1-250 represent data when the phone is in the default orientation, samples 251-500 represent data when the phone is in a non-default orientation and samples 501-750 represent the same data as samples 251-500, but transformed into the default coordinate system. Transformed data is then used to extract features in the same way as in the mixed orientation data approach. For testing we used data from non-default orientations, and data from default orientation. Data from default orientation was treated the same way as data from non-default orientations, and was transformed accordingly to its gravity vector. We tested the same 3 classifiers as in the mixed orientation data approach, and the results are shown in Table 2.

Table 2. Classifier accuracy – Data reorientation preprocessing approach

Classifier	Accuracy (in %)
C4.5 Decision Tree	86.5
Naïve Bayes	91.5
K-nearest neighbors	95.0

The results obtained are lower than in the mixed orientation data approach, but are still above 90% threshold, except for the decision tree classifier. For this reason we analyze the decision tree classifier further. The confusion matrix for the decision tree classifier is shown in Figure 2.

```
  a   b   c   d   e   f   <-- classified as
 89   0   0   0   0  14 |   a = RUNNING
  6  31   2   2  32   1 |   b = SITTING
  0   0  36   0   0  10 |   c = STAIRS_DOWN
  0   0   6  39   0   5 |   d = STAIRS_UP
  0   0   0   0  55   0 |   e = STANDING
  0   0   0   0   0 251 |   f = WALKING
```

Figure 2. Confusion matrix for the decision tree classifier

It can be seen that a lot of instances that represent sitting are classified as standing, which is not intuitive, because these two activities should be easy to distinguish. This is a consequence of how WEKA generates the decision tree. When we look at the generated decision tree shown in Figure 3, we can see that sitting and standing are distinguished by mean value along the Z axis (MeanZ). WEKA makes the split on value -9.804189 which is the maximum value for MeanZ for sitting. It can be expected that the maximum value for MeanZ will vary for different recordings, since it can't be expected from someone to sit in exactly the same way every time. Since the minimum value for MeanZ for standing is -1.07289 the split should be done on the value -5.4385, which is halfway between the maximum for sitting and the minimum for standing. This would generate a much more robust tree with higher accuracy. Such decision tree would correctly classify all of the sitting instances, previously misclassified as standing. Consequently, decision tree accuracy increases to 92%.

```
StandardDeviationY <= 0.574771
|   MeanZ <= -9.804189: SITTING (37.0)
|   MeanZ > -9.804189: STANDING (34.0)
StandardDeviationY > 0.574771
|   CorrelationYZ <= -0.215697: RUNNING (55.0)
|   CorrelationYZ > -0.215697
|   |   MeanZ <= -2.416916
|   |   |   CorrelationYZ <= 0.206422: STAIRS_DOWN (18.0)
|   |   |   CorrelationYZ > 0.206422: STAIRS_UP (18.0)
|   |   MeanZ > -2.416916: WALKING (115.0)
```

Figure 3. Generated decision tree

To demonstrate the benefits of this approach, the same classifier was tested with the same test data, but this time no data transformation was performed prior to classification. Based on the results shown in Table 3, it can be concluded that a classifier built using data from only one orientation, cannot classify instances from other orientations with a high success rate. With the reorientation preprocessing included, the classification accuracy results (shown in Table 2) increase significantly, which demonstrates the advantage of this approach, compared to the one assuming a fixed orientation at all times.

Table 3. Classifier accuracy – Single orientation classifiers without reorientation

Classifier	Accuracy (in %)
C4.5 Decision Tree	54.4
Naïve Bayes	48.5
K-nearest neighbors	58.8

To compare the data reorientation preprocessing approach against the mixed orientation data approach, the classifier built with the mixed orientation data approach was retested with a data set consisting of classifier training data, as well as data including orientations that were not used in classifier training. Accuracy of the mixed orientation data approach, when handling data including unknown orientations, can be analyzed in this manner. Evaluation results are shown in Table 4 and are comparable to the results achieved with the data reorientation preprocessing approach. It can be concluded that the mixed orientation data approach can also handle data from previously unknown orientations, but with a decrease in accuracy compared to the data reorientation preprocessing approach.

Table 4. Classifier accuracy - Mixed orientation data approach with unknown orientations

Classifier	Accuracy (in %)
C4.5 Decision Tree	85.0
Naïve Bayes	83.8
K-nearest neighbors	88.5

6. EVALUATION OF MOBILE CPU LOAD

Using results from the previous test we built an application for activity recognition in real-time locally on the device. The application implements the data reorientation preprocessing approach and uses a prepared decision tree as a classifier. To test how much the application and data transformation phase participate in the processor load we used DDMS (Dalvik Debug Monitor Service).

We ran DDMS for thirty seconds while the application was active on the device. The results are shown in Figure 4. The first line of the figure represents the main thread of the application in which all of the application processing is done. Different methods are represented with different colors on the timeline. This figure focuses on the period between two feature extractions. The boxes marked with the number 1 represent feature calculation. We notice three color groups, where each one, looking from the left to the right, represents calculation of mean, standard deviation and correlation. The box marked with the number 2 represents data transformation. We can see how data transformation is performed whenever a new sample is read from the accelerometer and that feature extraction is performed only when the window shifts for the specified number of samples.

The Figure 5 shows the timeline again but this time in more details. Similar to Figure 4, the box marked with the number 1 represents mean calculation (green color on the timeline), and the box marked with the number 2 represents data transformation. It can be seen that data transformation requires a small portion of processor time, compared to processor time required to calculate just one feature, but is performed more often.

The entry point of the application is the *onSensorChanged* function, which is called whenever a new sample from the accelerometer is read. This function encapsulates all of the application processing and participates in the processor load with 22.9%. Also, the function *transformData*, which encapsulates all of the data transformation, participates in the processor load with 1.3%. It can be concluded that although data transformation is performed more often than feature extraction, it doesn't increase processor load significantly.

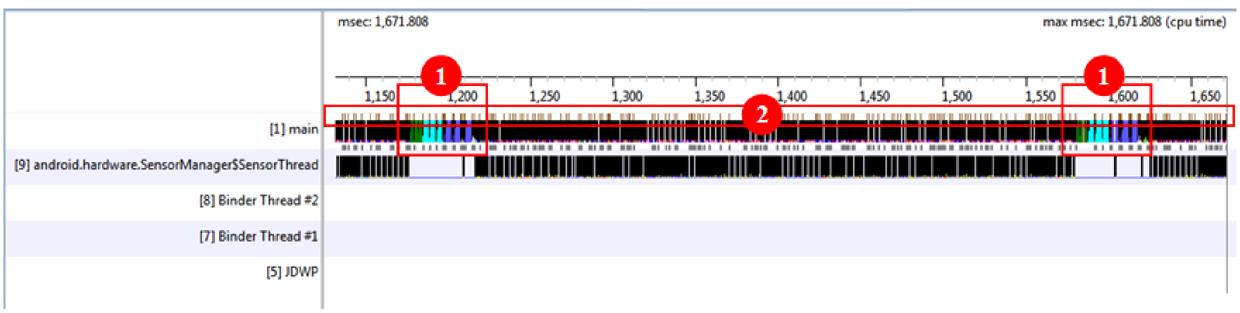

Figure 4. Result of the DDMS tool applied on the activity recognition application

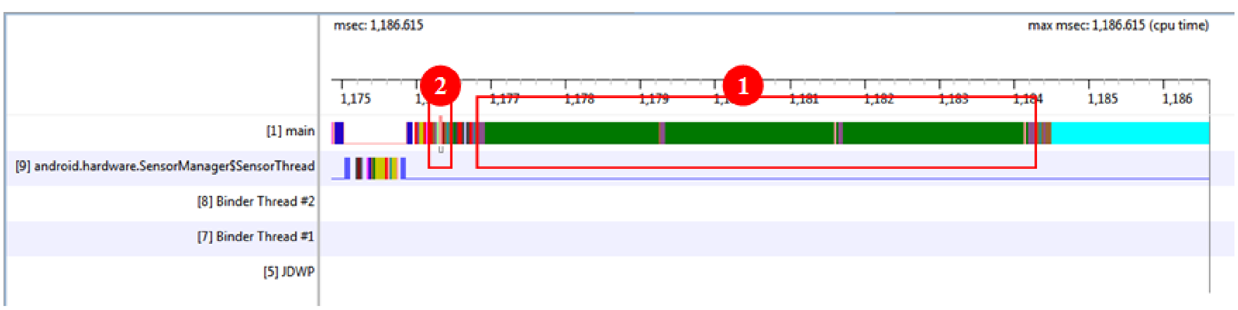

Figure 5. Result of the DDMS tool in more detail

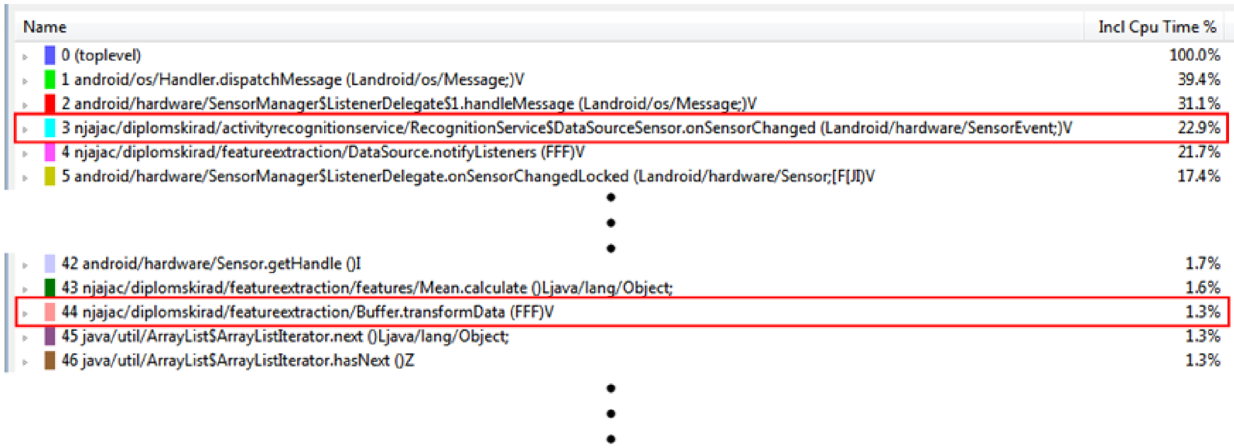

Figure 6. Participation of individual functions in processor load

7. CONCLUSION AND FUTURE WORK

In this paper a method for orientation independent activity recognition from accelerometer data is described. An accelerometer embedded in a mobile phone was used. Orientation of the phone was partially fixed in such a way, that the phone could be rotated only around one axis, but the angle of orientation was arbitrary. To determine the angle of rotation a calibration phase was introduced, in which the user has to stand still for a couple of seconds with the phone placed in the desired orientation. In this period the gravity vector is extracted and the difference between the angle of that gravity vector and the angle of the average gravity vector in the default orientation is calculated. This difference is the angle of rotation of the phone. After that the user can start to perform activities freely. The data from the accelerometer is transformed into the default coordinate system, the features are extracted and the activity recognition is performed.

This method showed slightly reduced accuracy compared to the method when a classifier is built from data collected from various orientations, when data from the predefined orientations only is considered, but the results are still above the threshold (accuracy above 90%). When data from not predefined orientations is considered as well, the proposed method demonstrates increased accuracy. The most significant advantage of this method is that it requires data collection from only one orientation, so less data is required for training. Also it makes no assumption on the orientation in the classifying phase; there are no predefined orientations, so the system will work with data from any orientation. The drawback is the existence of the calibration phase, so the process in not fully transparent. This probably makes this method unusable in some areas of application, like in elderly care for example, but in others, like in fitness monitoring, we believe that is acceptable because the calibration phase is very short and requires very little effort from the user. In return the user can place the phone in any orientation.

In this paper the phone orientation is assumed to be partially fixed, which still limits the user to a certain degree. Demonstrated method represents only an intermediate step in development of a method which would acquire a full three-dimensional orientation in the calibration phase and impose no limits in the phone's orientation.

8. ACKNOWLEDGMENTS

This paper was realized as a part of the project "Studying climate change and its influence on the environment: impacts, adaptation and mitigation" (43007) financed by the Ministry of Education, Science and Technological Development of the Republic of Serbia within the framework of integrated and interdisciplinary research for the period 2011-2014.

9. REFERENCES

[1] J. Lester, T. Choudhury, G. Booriello. A Practical Approach to Recognizing Physical Activities. *Proceedings of the 4th International Conference on Pervasive Computing*, 2006, pp. 1-16.

[2] L. Bao, S. S. Intille. Activity recognition from user-annotated acceleration data. *Proceedings of the 2nd International Conference on Pervasive Computing*, 2004, pp. 1-17.

[3] N. Ravi, N. Dandekar, P. Mysore, M. L. Littman. Activity recognition from accelerometer data. *Proceedings of the Seventeenth Conference on Innovative Applications of Artificial Intelligence (IAAI-05)*, 2005, pp. 1541-1546.

[4] J. R. Kwapisz, G. M. Weiss, S. A. Moore. Activity Recognition using Cell Phone Accelerometers. *ACM SIGKDD Explorations Newsletter*, 2010, vol. 12, issue 2, pp. 74-82.

[5] L. Sun, D. Zhang, B. Li, B. Guo, S. Li. Activity Recognition on an Accelerometer Embedded Mobile Phone with Varying Positions and Orientations. *Proceedings of the 7th International Conference on Ubiquitous Intelligence and Computing*, 2010, pp. 548-562.

[6] S. Thiemjarus. A Device-Orientation Independent Method for Activity Recognition. *International Conference on Body Sensor Networks*, 2010, pp. 19 - 23.

[7] I. H. Witte, E. Frank. *Data Mining: Practical Machine Learning Tools and Techniques*, 3rd ed. Morgan Kaufmann, 2011

Extending Augmented Reality Mobile Application with Structured Knowledge from the LOD Cloud

Betül Aydin
Grenoble Informatics Lab.
681, rue de la Passerelle
38402 Saint Martin d'Hères
France
Betul.Aydin@imag.fr

Jérôme Gensel
Grenoble Informatics Lab.
681, rue de la Passerelle
38402 Saint Martin d'Hères
France
Jerome.Gensel@imag.fr

Philippe Genoud
Grenoble Informatics Lab.
681, rue de la Passerellem
38402 Saint Martin d'Hères
France
Philippe.Genoud@imag.fr

Sylvie Calabretto
INSA de Lyon
20, avenue Albert Einstein
69621 Villeurbanne Cedex
France
Sylvie.Calabretto@liris.cnrs.fr

Bruno Tellez
Claude Bernard Uni. Lyon 1
43 bvd du 11 Novembre 1918
69622 Villeurbanne Cedex
France
Bruno Tellez@liris.cnrs.fr

ABSTRACT

ARCAMA-3D (Augmented Reality for Context Aware Mobile Applications with 3D) is a mobile platform that allows us to overlay a 3D representation of the surroundings with augmented reality objects. In this paper, we show how these 3D objects, which are overlaid on the real view captured by the camera of the mobile device, are coupled with the Linked Open Data (LOD) cloud. With this approach, the data aggregation for a mobile augmented reality system is provided using the interconnected knowledge bases on the Web. This offers the possibility to enrich our 3D objects with structured information on the cloud. The objects that act as an augmented reality interface are used to provide an interactive access to these information. This approach provides an opportunity for people who publish information on the LOD cloud to interlink their data with 3D urban models. In order to achieve this, we propose an extensible data model that takes into account the temporal evolution of real world entities (such as buildings, monuments, *etc.*) and we publish our 3D models using this data model.

1. INTRODUCTION

Ubiquitous mobile applications rely on context and location aware approaches in which the system takes into account the changing context of the user to provide information dynamically. However, constraints related to the small displays of mobile devices and the methods used in order to look for information during the mobility require providing the most relevant information when the user expects a quick answer. For example, in order to discover the surroundings using a mobile application, the systematic and exhaustive presentation of all Points of Interest (PoIs) will not only hamper the readability of these collected information, but also risk to divert the user who observes her surroundings.

In this paper, we present a context-aware Augmented Reality (AR) approach which minimizes the effort from the user to access and interact with the information that may be of her interest. For this, we interconnect our application with the Linked Open Data [8] which, by structuring the knowledge available on the Web, enables a semantic approach in the information search and discovery.

Augmented reality consists in an interactive medium that overlays the real world with some objects modelled in 3D and keeps their alignment in real time [3]. By adding a virtual layer on the real world view, the goal is to improve and enrich the user's perception.

LOD uses the infrastructure of the World Wide Web to publish and link datasets that can then be explored and exploited by applications accessing to them. While achieving some of these objectives still remains a research interest, many datasets from very different domains (media, biology, chemistry, economics, energy, *etc.*) are already published and are constantly being enriched. A portion of these data contains geo-localized information that can be exploited by the applications, especially while the user is moving. However, few mobile augmented reality applications have explored this possibility.

ARCAMA-3D (Augmented Reality for Context Aware Mobile Applications with 3D) [1], [2], our AR system, superimposes transparent 3D representations of the real world objects, with which the user can interact, on the view captured by the mobile device, such as a *smartphone*. By enriching the 3D models with thematic and temporal metadata published on the LOD cloud, ARCAMA-3D proposes answers to several applicable ubiquitous scenarios in which the augmented reality interface effectively helps the user during her exploration of the real world based on her choices and preferences.

The paper is organized as follows. In Section 2, we present some work that integrates a location-based approach in an

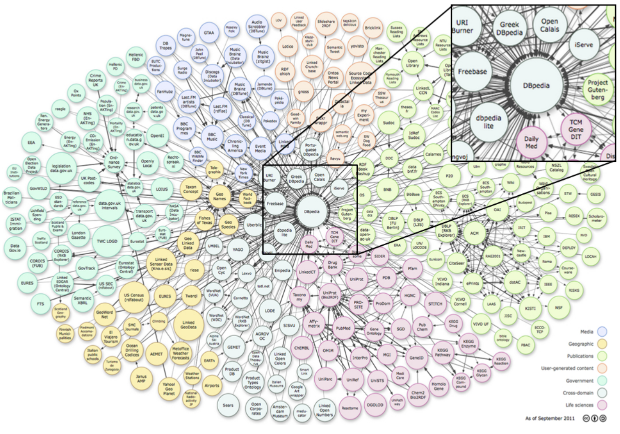

Figure 1: LOD Cloud

augmented reality system, as well as the aspects that differentiate our system from existing ones. Section 3 describes our previous work, the platform ARCAMA-3D and introduces the data model (ontology) that we defined in order to interconnect 3D objects with the LOD. We also show in this section how this interconnection, on the one hand, extends the use of ARCAMA-3D, and on the other, contributes to the expansion of LOD cloud with structured information currently unrepresented in this cloud. In Section 4, we draw a general view of the ARCAMA-3D system, and we conclude in Section 5.

2. BACKGROUND AND RELATED WORK

The most well-known mobile augmented reality applications are Layar[1] and Wikitude[2]. They provide the user with some local points of interests (POIs) by querying particular servers where such information is stored. The user's context in these applications is limited to the *location* and the *direction* in which the user is moving. Using these information, the *databases* and *servers* are queried. In Wikitude, the user discovers information about interesting places, famous landmarks and other POIs. In order to get such information, the application communicates with some servers such as Wikipedia, Twitter, Instagram, world heritage list, *etc.* Layar integrates information from social networking platforms such as Twitter and Qype, as well as with photo sharing web services such as Flickr. In both applications, the location-based data is overlaid on the real view as an augmented reality tool.

For a person who visits a place, finding and providing *all* the information by using her geolocation as the unique criterion, may not be considered as satisfactory. The context-aware mobile augmented reality platforms should also consider filtering the information to keep only the relevant information that fits to the expectations or the interests of the user. To this end, Semantic Web technologies can be used to provide a meaningful exploration of the information available on the Web with general and/or specialized knowledge expressed in a formal and structured manner.

Linked data[5] is a term used to describe a recommended practice to represent, share and interconnect information using Semantic Web technologies: URIs, RDF, SPARQL. Linked Open Data constitutes a community effort to publish datasets available under open licence respecting the principles of *linked data* [8]. Among them, DBpedia, which contains the structured information extracted from Wikipedia, plays a central role in interlinking many other datasets, such as Freebase, LinkedGeoData, GeoNames, Flickr, Revyu, or even FOAF profiles (see Figure 1).

Related to the problem of discovering the surroundings during the mobility of the user, mSpace mobile [15] is one of the early mobile applications that uses the Semantic Web to support the exploration of information resources. It is a London city guide with connections to different domains or sources of interests (movies, music, *etc.*). Using the location information of the user, the application queries different knowledge bases (IMDB, video archives of the BBC, *etc.*) and displays the results in an information box with textual descriptions. Although, this application does not provide any augmented reality experience and its user interface does not allow a fast and efficient exploration, it remains interesting for the use and exploration of multiple interconnected knowledge databases.

DBpedia Mobile, is a context-sensitive browser running on mobile devices [4]. Designed for the exploration of a touristic place, it displays on a map information about the nearby locations that can be found in the DBpedia dataset. The user can then explore information associated with these locations through links between DBpedia and other data sources (for instance, the associated Wikipedia pages). It also allows users to publish their location, photos and comments on the Linked Data. We note here that the interface does not use augmented reality, but a 2D map which necessitates constant attention of the user to read the information. Moreover, *all* the information within the vicinity of the user is provided to her without any preliminary selection.

In contrast, in [13] several possible use cases of linked data are discussed in the context of mobile applications based on augmented reality. [14] explores the use of datasets related to search and view information about cultural heritage. Using the geolocation and the direction in which the user is moving, the nearby PoIs are identified using the LOD cloud. The authors have focused here on merging and aligning multiple resources associated with the same PoI, as well as the semantic enrichment of the erroneous data (errors in georeferenced data, human annotations, *etc.*).

In conclusion, these different research projects show the potential of using the Linked Open Data in mobile applica-

Figure 2: A representation of the ARCAMA-3D application. (a) on the left: real view (b) on the right: a 3D model of the surroundings is overlaid on the real view with 'red' objects indicating some interactive objects providing access to information.

[1]http://layar.com/

[2]http://wikitude.com/

tions where the information access is based on the geolocation of the user. However, the risk of displaying *all* the georeferenced information, in particular in augmented reality applications, may expose the mobile user to some cognitive overload [10].

3. CONTRIBUTION

In our previous work, we have proposed a mobile augmented reality platform, ARCAMA-3D, which is based on the use of a lightweight 3D model [1]. The real-world objects are represented by a semi-transparent 3D geometric model superimposed on the real view captured by the camera of the mobile device (see Figure 2). Thus, the real world is represented by the 3D objects on the user interface. Interacting with these objects, the user can ask questions like "*what is it?*", "*Is there something interesting (for me) around me?*" *etc.*

The ARCAMA-3D application continuously acquires geolocation data as well as orientation data using the embedded sensors (GPS, accelerometer, gyroscope, *etc.*) of the mobile device (smartphone, tablet PC, *etc.*) on which it runs. We ensure the alignment of the 3D model on the real scene by refining the captured data using a Kalman filter and merging these data [9], [6]. The fusion of sensor data allows an accurate estimation of the position, hence, superimposes the 3D model on the real view, and maintains this superimposition all along the movement of the user [2].

When we observe users' demands from an AR mobile application, we notice that they expect a 'transparent' augmented reality that should be there only when needed and that should give only meaningful information, following the recommendations made in [10]. ARCAMA-3D displays the augmented reality objects as transparent, and indicates the objects in different colors if they contain interesting information for the user with respect to a given topic of her interest. By indicating the *existence* of information in this way, we invite the user to interact with these objects.

The 3D models exploited by ARCAMA-3D play a key role here. They are indeed the support of the augmented reality objects that provide access to information. We propose to represent these 3D models by means of an ontology that captures the characteristics of 3D models (levels of detail, time period that they cover, *etc.*). Thus, such 3D models can be linked to ontological descriptions of 3D objects by establishing a valuable link between the real world entities described on the cloud. To achieve this goal, we integrate the geometric descriptions (3D models) of a real-world entity corresponding to a nearby Object of Interest (OoI) [1] of ARCAMA-3D in an RDF graph. This graph uses a controlled vocabulary defined in an OWL ontology (`arcama-owl`). This ontology defines a generic model for ARCAMA-3D; and the RDF graph corresponding to a given entity can be seen as an instance of this generic model.

The representation of 3D models used by ARCAMA-3D brings two advantages:

- it allows to connect - and thus enrich - these 3D models with other structured datasets (especially Freebase, DBpedia, *etc.*)
- symmetrically, it offers the possibility to other datasets of the LOD cloud to bind their information with 3D models and re-use them.

3.1 3D Data Model

We have created a 3D data model where we store the geometric representations of real-world entities and the information associated with them. Our data model is extensible, and the representation of an OoI can be associated with temporal and thematic information. Also, in the design of a 3D data model, we think that it is necessary to incorporate spatial (what portion(s) of space the 3D model covers), temporal (which time period(s) the 3D model belongs to) and thematic (what is (was) the role (function) of that real world object) characteristics in the model. We adopt a similar approach to that described in [11] which proposes a spatio-temporal conceptual model focusing on the orthogonality of spatial, temporal and structural instances. In this work, the authors apply this approach to non-urban geographic features (rivers, lakes, townships, *etc.*) or 2D spatial data relating to urban objects (the surface of a castle, *etc.*). We integrate this approach in our 3D data model to allow the exploitation of spatial, temporal and thematic information related to the OoIs during the AR experience of the user.

To illustrate this model of 3D data using a concrete example, let us consider a real world entity: the *Hagia Sophia* edifice in Istanbul, Turkey. The choice of this entity is suitable to illustrate our approach, since the architecture and the role (the function) of the building have both evolved over time. Thus, on the Wikipedia page dedicated to this monument[3], we can find in the infobox (a table that summarizes the article on the right corner of each Wikipedia page), the information related to this evolution.

The DBpedia resource representing Hagia Sophia[4] assigns successive roles to it as *eastern orthodox cathedral*, *mosque* and *museum* (see Figure 3). However, no indication related to the relevant time periods of these functional changes nor its architectural changes is presented in this description. The purpose of our data model is to fill these gaps. Thus, the proposed model takes into account both the temporal evolution in space (geometry) and thematic information (role/function) that describe the OoIs. Therefore, for a given OoI, we need multiple representations of the object, each of them corresponding to a change in the *thematic* and/or the *geometric* dimension. Figure 3 shows the evolution for the Hagia Sophia, and the different representations (*i.e.* R_1, R_2, R_3, R_4 in the figure) that the model should be able to hold for this OoI. As expected, each of these four representations correspond to an architectural change (geometric) and/or a functional change (role) of the OoI.

3.2 Arcama-owl Ontology

In order to construct these multiple representations, we have defined an OWL ontology that describes a generic model for an OoI. The UML class diagram in Figure 4 shows the different concepts (classes) and relations between these concepts (properties) defined in this ontology.

For our system, this class model (and its instantiations) plays an important role: it allows the interconnection between the historical 3D models associated to OoI and the information related to these OoIs found on the LOD.

- The `Entity` class corresponds to an OoI (the real-world entity Hagia Sophia, for example). This model al-

[3] `http://en.wikipedia.org/wiki/Hagia_Sophia/`
[4] `http://dbpedia.org/resource/Hagia_Sophia/`

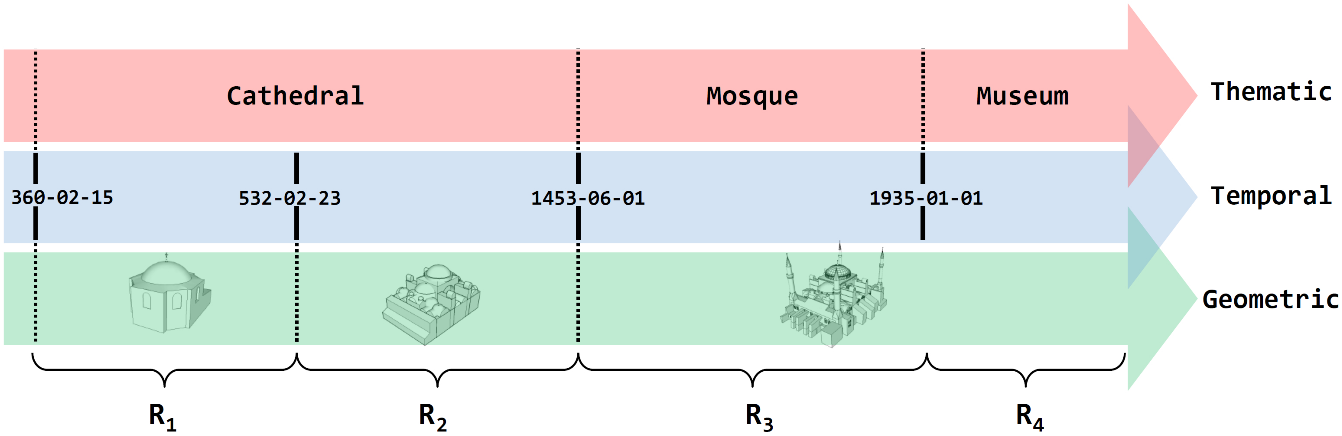

Figure 3: Representations of *Hagia Sophia* according thematic and geometric changes.

lows to combine several temporal representations of the same entity (see class `TemporalRepresentation`). The `represents` association allows to interconnect this model with other resources described in other datasets and published on the LOD cloud (for instance, the DBpedia resource[5] or the Freebase resource[6] for the *Hagia Sophia*).

- The `TemporalRepresentation` class is at the heart of our model. It corresponds to a given representation of the entity. This representation has a temporal validity interval defined by OWL-Time[7] with `hasTimePeriod` association. A temporal occurrence aggregates spatial and thematic attributes that describe the entity during this time interval. These attributes are described respectively by `Geometry` and `Role` classes. Any changes in the `Geometry` and/or `Role` information require creating another `TemporalRepresentation` associated with a new time interval.

- The `hasRole` association sets the role(s) of a TemporalRepresentation object through the `Role` class which holds the properties about the function of the object (museum, monument, hall, church, *etc.*). This information is used to classify the temporal representations according to the thematic criteria chosen by the user. Accordingly, it allows to filter temporal representations that will be displayed as the augmented reality interface and allows the access to additional information related to the that entity. All the roles presented here are defined by another ontology that we have created and currently has 125 classes corresponding to different types of architectural structures, for example, religious buildings (cathedral, mosque, temple, *etc.*) or historical sites (castle, monument, pyramid, *etc.*). To facilitate the interconnection with resources from other datasets defined with their own ontology, a mapping between roles and classes of some of these ontologies is defined (currently with DBpedia and Freebase datasets). Initially, we had planned to base it solely on the roles used in the DBpedia ontology, but it turned out, on the one hand that the proposed DBpedia structure was not fully satisfactory, and, secondly, that the roles of the resources could be of variable qualities (*i.e.* depending on the quality of the information provided by Wikipedia contributors). Using our own ontology and establishing a mediator mechanism allowed us to free ourselves from these drawbacks.

- In order to describe the geometry associated with a time instance, we use the composite design pattern [7] that combines geometry in a tree structure that represents a part-whole hierarchy. `Geometry` is an abstract class with two subclasses ElementaryGeometry and `CompoundGeometry`.

- The `CompoundGeometry` class allows to define a complex geometry composed of several different geometries that can be either complex geometries or elementary geometries.

- The `ElementaryGeometry` class describes a leaf node in a tree structure of geometries. As for a specific purpose, we want to be able to offer different 3D models at different levels of detail, `ElementaryGeometry` class aggregates one or more `GeometryFile` objects. Each `GeometryFile` object is described by a URI that provides access to a physical file containing a description of the 3D object in a usual format (KML, CityGML, VRML, COLLADA *etc.*).

- The `LevelOfDetail` class describes the level of detail of a `GeometryFile`. We have defined three levels of detail where *LOW* corresponds to a rectangular prism. It corresponds to the LOD1[8] of CityGML. *MEDIUM* corresponds to a much detailed 3D geometry with no texture. *HIGH* corresponds to texturized 3D objects (it can even represent a laser scanned object).

Figure 5 presents a part of the RDF graph corresponding to an instantiation of this model for the example of *Hagia Sophia*. In order to facilitate the interpretation of this RDF graph in relation to the generic model presented above, we

[5] `http://dbpedia.org/resource/Hagia_Sophia`
[6] `http://www.freebase.com/m/0br5q`
[7] `http://www.w3.org/TR/owl-time/`

[8] Please note that this acronym is not related to the Linked Open Data appreviation (LOD) that we often use in our article.

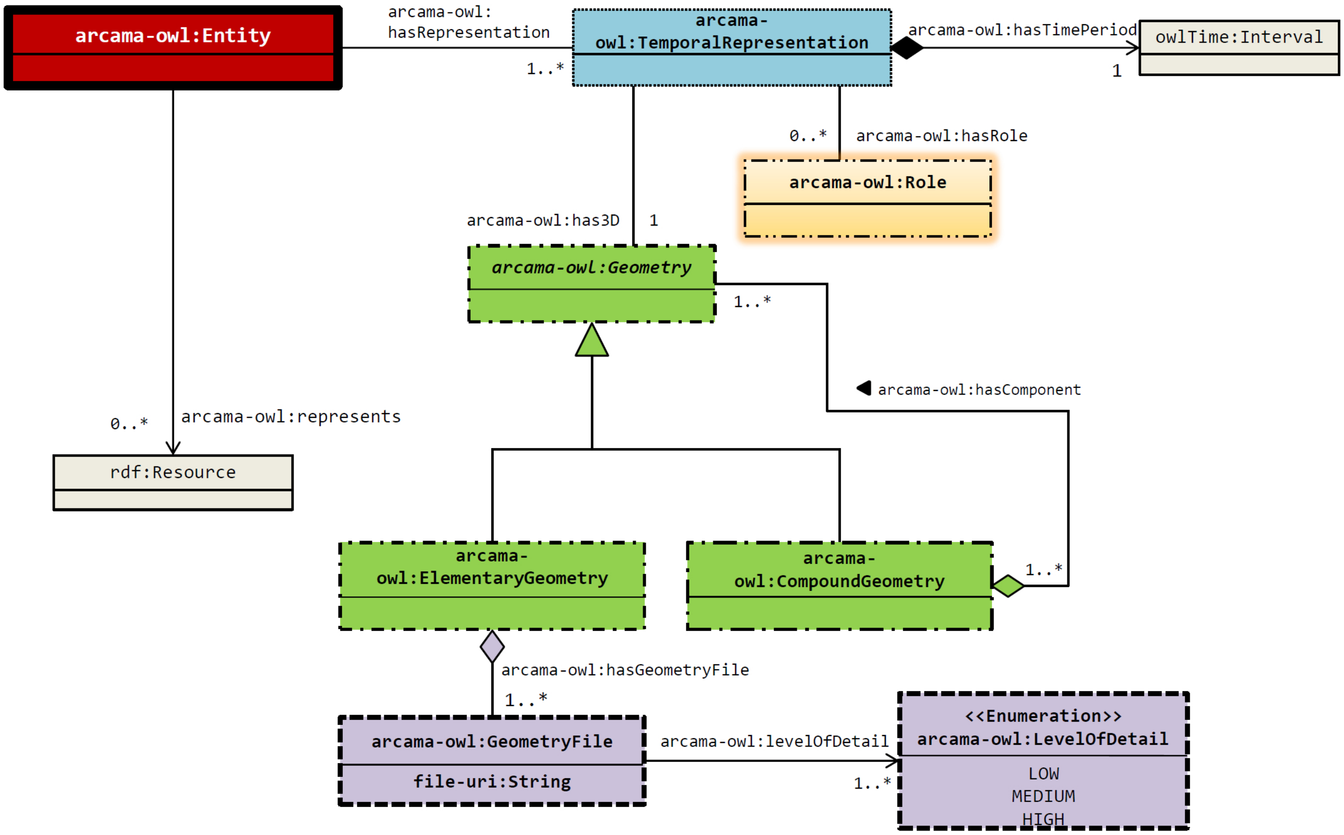

Figure 4: UML class diagram of the arcama-owl ontology.

choose to color and format the resource nodes that belong to the same `arcama-owl` class in Figure 4 with the same color and line format.

In this RDF graph (Figure 5), *Hagia Sophia* is represented by two temporal instances: `model:HagiaSophia#TI1` and `model:HagiaSophia#TI2`. The former one is the instantiation of the model in interval R_2 of Figure 3 and the latter one is of R_3. As it can be seen in the temporal evolution of the real world entity, first temporal instance represents a change in the geometry of the entity, whereas the second one represents the changes both in the role and the geometry.

4. ARCAMA-3D

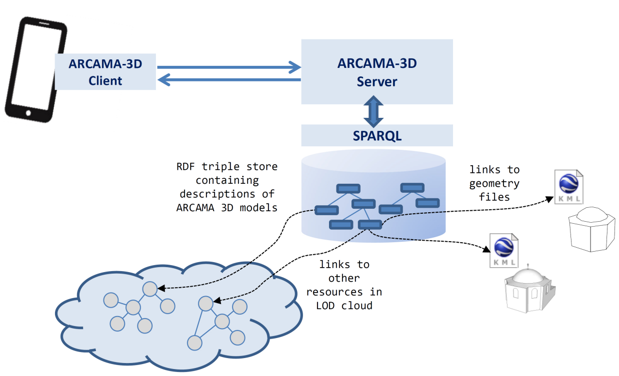

Figure 6: The architecture of ARCAMA-3D

The mobile device (ARCAMA-3D client) connects to the ARCAMA-3D server (see Figure 6). Using geolocation information and user thematic preferences (architecture, history, museum, *etc.*), the server queries the triplestore containing descriptions associated with 3D models using SPARQL. Then, the corresponding 3D geometries are superimposed on the actual view in a semi-transparent form, and in different colors if they have some results from the SPARQL query associated with them. The user can then interact with the 3D objects by selecting the object on the screen of the mobile device (3D picking technique). This interaction gives access to information related to this object in the LOD cloud (photos interlinked with that object on the LOD cloud, historical 3D models and roles of the buildings during different time periods, *etc.*). Then the user accesses these information.

5. CONCLUSION AND FUTURE WORK

New datasets are continuously being added contributing to the evolution of the LOD cloud which is a semantically interconnected knowledge base. We describe in this paper a system dedicated to ubiquitous mobile applications based on augmented reality which includes temporal, thematic and spatial representations to interlink 3D models with the information available on the LOD cloud. The interconnection with the LOD cloud not only helps to represent 3D datasets on the cloud, but also improves the integration process of the augmented reality system. Temporal aspects supported by the proposed model reflect the evolution of real-world objects in their form, function, location, *etc.*

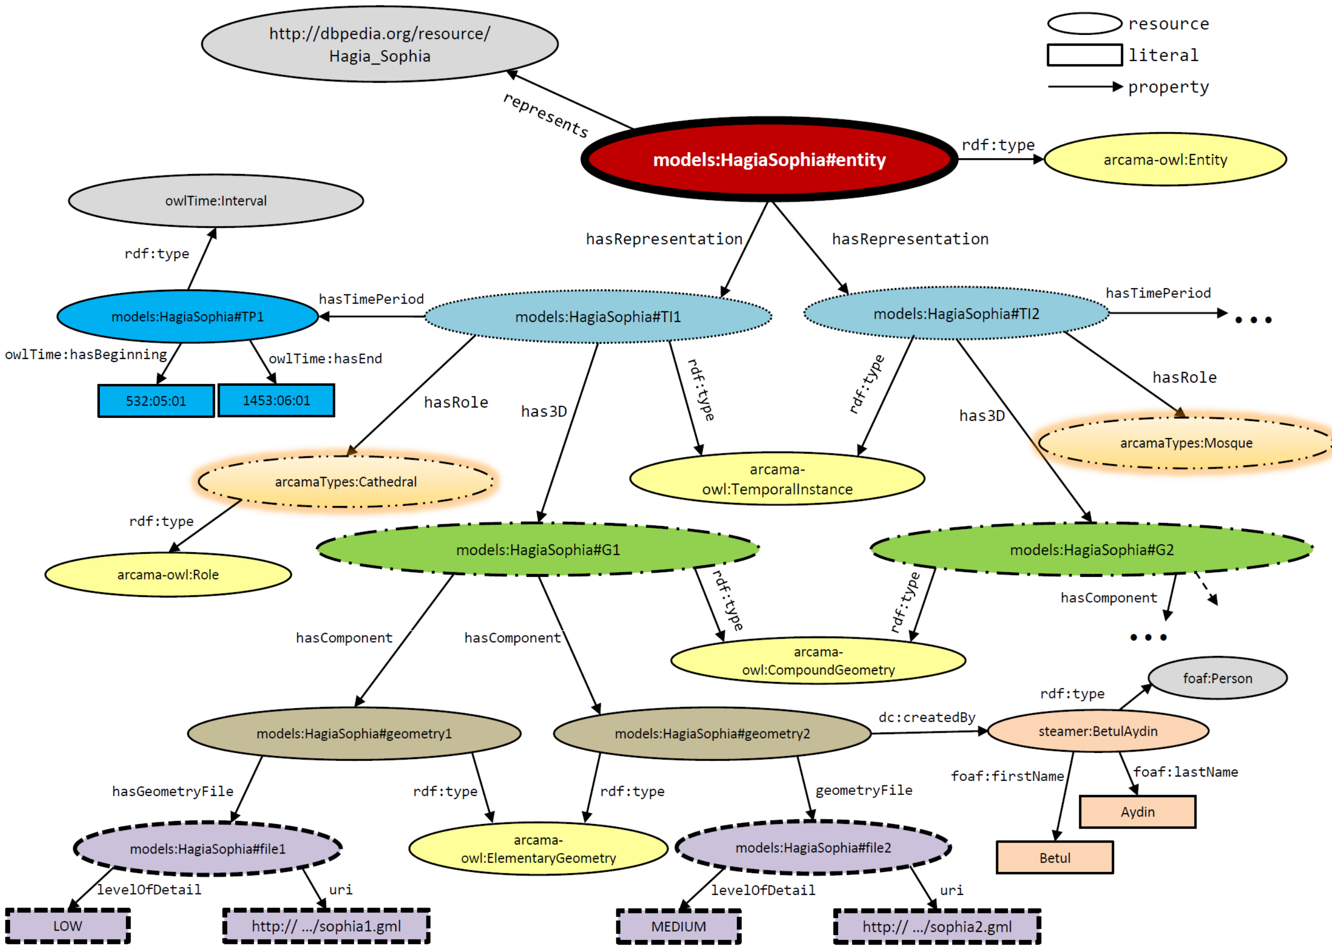

Figure 5: RDF graph representing *Hagia Sophia* with different temporal instances.

We now intend to improve the representation of the identity of objects in the real world and the characterization of their evolution over time, adapting some of our previous work [12].

6. ACKNOWLEDGMENTS

This research is granted by the Region Rhône-Alpes.

7. REFERENCES

[1] B. Aydin, J. Gensel, S. Calabretto, and B. Tellez. Arcama-3d - a context-aware augmented reality mobile platform for environmental discovery. In *W2GIS*, pages 17–26, 2012.

[2] B. Aydin, J. Gensel, S. Calabretto, and B. Tellez. Arcama-3d: Une application mêlant réalité augmentée et 3d pour la présentation d'informations contextuelles. In *8èmes journées francophones Mobilité et Ubiquité (UBIMOB), Anglet, France*, pages 34–41, 2012.

[3] R. Azuma. A survey of augmented reality. *Presence*, 6(4):355–385, 1997.

[4] C. Becker and C. Bizer. Exploring the geospatial semantic web with dbpedia mobile. *J. Web Sem.*, 7(4):278–286, 2009.

[5] T. Berners-Lee and M. Fischetti. *Weaving the web - the original design and ultimate destiny of the World Wide Web by its inventor*. HarperBusiness, 2000.

[6] R. G. An intelligent multi-sensor system for pedestrian navigation. *Journal of Global Positioning Systems*, 5(1–2):110–118, Sept. 2006.

[7] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design patterns: elements of reusable object-oriented software*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1995.

[8] T. Heath and C. Bizer. *Linked Data: Evolving the Web into a Global Data Space*. Morgan & Claypool, 1st edition, 2011.

[9] R. E. Kalman. A New Approach to Linear Filtering and Prediction Problems. *Transactions of the ASME - Journal of Basic Engineering*, (82 (Series D)):35–45, 1960.

[10] T. Olsson, T. Kärkkäinen, E. Lagerstam, and L. Ventä-Olkkonen. User evaluation of mobile augmented reality scenarios. *JAISE*, 4(1):29–47, 2012.

[11] C. Parent, S. Spaccapietra, E. Zimányi, P. Donini, C. Plazanet, C. Vangenot, N. Rognon, and P. A. Crausaz. MADS, modèle conceptuel spatio-temporel. *Revue Internationale de Géomatique*, 7(3–4):317–352, Dec. 1997.

[12] C. Plumejeaud. *Modèles et méthodes pour l'information spatio-temporelle évolutive*. These, Université de Grenoble, Sept. 2011.

[13] V. Reynolds, M. Hausenblas, A. Polleres, M. Hauswirth, and V. Hegde. Exploiting Linked Open Data for Mobile Augmented Reality. In *W3C Workshop: Augmented Reality on the Web*, June 2010.

[14] C. Van Aart, B. Wielinga, and W. R. Van Hage. Mobile cultural heritage guide: location-aware semantic search. In *Proceedings of the 17th international conference on Knowledge engineering and management by the masses*, EKAW'10, pages 257–271, Berlin, Heidelberg, 2010. Springer-Verlag.

[15] M. Wilson, A. Russell, D. A. Smith, A. Owens, and M. C. Schraefel. mspace mobile: A mobile application for the semantic web. In *User Semantic Web Workshop, ISWC2005*, 2005.

Vanet-X: A Videogame to Evaluate Information Management in Vehicular Networks

Sergio Ilarri
IIS Department
University of Zaragoza
Zaragoza, Spain
silarri@unizar.es

Eduardo Mena
IIS Department
University of Zaragoza
Zaragoza, Spain
emena@unizar.es

Víctor Rújula
IIS Department
University of Zaragoza
Zaragoza, Spain
han.vikktor@gmail.com

ABSTRACT

Vehicular Ad Hoc Networks (VANETs) are attracting considerable research attention, as they are expected to play a major role for Intelligent Transportation Systems (ITS). Thus, according to a recent survey by ABI Research[1], about 62% of new vehicles will be equipped with vehicle-to-vehicle (V2V) communications by 2027. Vehicular networks offer new opportunities for the development of interesting mobile applications for drivers, but at the same time they also bring challenges from the data management point of view. Thus, for example, techniques should be developed to estimate the relevance of the information exchanged among the vehicles and to propagate the relevant data in the network efficiently and effectively. As testing the proposals in a real large-scale scenario is impractical, simulators are often used.

In this paper we present *Vanet-X*, an online multiplayer driving videogame that we have developed to help in the difficult evaluation task of data management strategies for VANETs. The idea behind the proposal is to exploit the potential of players around the world driving vehicles in the videogame to effortlessly collect data that can be used to extract some conclusions and fine-tune the proposed data management strategies. So, for example, the videogame allows to evaluate if a certain data management strategy is able to provide useful information to the driver/player (i.e., if the presented information represents an advantage for him/her). We argue that this videogame can be a good complement for existing simulators. As a proof of concept, we have performed some preliminary tests that show the potential interest of the proposal.

1. INTRODUCTION

The widespread availability of mobile devices and the development of wireless communication technologies (such as Wi-Fi, WAVE, etc.) have encouraged the development of services for drivers within the context of *Intelligent Transportation Systems* (*ITS*). In particular, *Vehicular Ad Hoc Networks* (*VANETs*) have become an attractive research area [1, 14, 15, 20, 24, 26, 30]. In these vehicular networks, the vehicles can exchange information directly by using short-range wireless communication technologies. This decentralized architecture provides some advantages over other solutions such as the use of 3G communications: e.g., no need of an infrastructure, quicker transmission of safety-related data in the vicinity, localized communications without the need of a centralized server, and free of charges (which also encourages the participation of peers in the network). Numerous types of events can be relevant for drivers (e.g., accidents, traffic congestions, an ambulance asking the right of way, available parking spaces, etc.). These events can be exchanged in the vehicular network and stored locally by the vehicles. Then, a query processor can periodically evaluate the interest of those events and decide if they should be shown to the driver; there may be implicit queries (e.g., information about an accident in the direction of travel will be relevant for any driver) and explicit queries (e.g., a driver may indicate his/her interest in finding an available parking space or in receiving information about other specific types of events).

However, although VANETs offer interesting opportunities for the development of data services for drivers, they also bring new challenges. Thus, several difficulties arise from the point of view of data management [5]. As an example, estimating the relevance of events in order to disseminate them effectively and efficiently in the network is a challenge [2]. Similarly, disseminating information about a scarce resource (e.g., an available parking space) to many vehicles can lead to competition situations among them to try to reach the resource [7]. As a final example, the relevance of events must also be considered in order to decide if a specific event received by a vehicle should be shown to the driver or not [3].

A big challenge is how to evaluate the data management techniques proposed. Evaluating them in a real scenario with a significant number of vehicles is simply impractical and expensive. Therefore, simulations are frequently used in this field. However, even with simulations the evaluation task can be very time-consuming. For example, many proposals depend on a number of parameters that can be fine-tuned for a given scenario (e.g., see [2, 31]), and determining a good choice of parameters for general evaluation is quite challenging. On the other hand, crowdsourcing strategies where users play the role of drivers could help to

[1] `http://www.abiresearch.com/press/v2v-penetration-in-new-vehicles-to-reach-62-by-202`.

introduce human behavior and facilitate new tests initiated by the users themselves.

So, in this paper we propose a complementary approach that can be used in conjunction with the use of simulators. In particular, we argue that we can benefit from players having fun with a driving game to easily collect interesting data that can be used to extract some conclusions and fine-tune the proposed data management strategies. The videogame is inspired by the classic videogame *Rally-X* (`http://www.klov.net/game_detail.php?game_id=9259`, videogame released in 1980) but it is a new development, with different goals, game modes, and spirit. So, the basic idea is that the vehicles can receive information through the vehicular network and different data management techniques can be plugged in the videogame (e.g., different data dissemination strategies). Data received from other vehicles, if evaluated as interesting by the local query processor in the car, are shown on a radar and can provide a competitive advantage to the player. During the game, a variety of data are collected (e.g., number of messages received by the vehicles, network overhead, time required by the vehicles to complete their goals in the game, etc.), that can be analyzed later. So, while playing, players contribute to collect data for a variety of scenarios, and these data can be exploited to evaluate the effects of particular data management strategies.

The structure of the rest of this paper is as follows. In Section 2 we describe the high-level architecture of the videogame and its features. In Section 3 we summarize the main behaviors implemented for the computer-managed vehicles. In Section 4 we present some basic aspects about the way the data are collected for later analysis. In Section 5 we present the results of the first experiments that we have developed as a proof of concept. In Section 6 we present some related work. Finally, in Section 7 we present our conclusions and some lines of future work.

2. ARCHITECTURE AND FEATURES

Vanet-X is a car videogame that can be played by multiple players connected to the Internet (see Figure 1 for a snapshot showing parking spaces).

2.1 Main Features

We summarize some features of the game as follows:

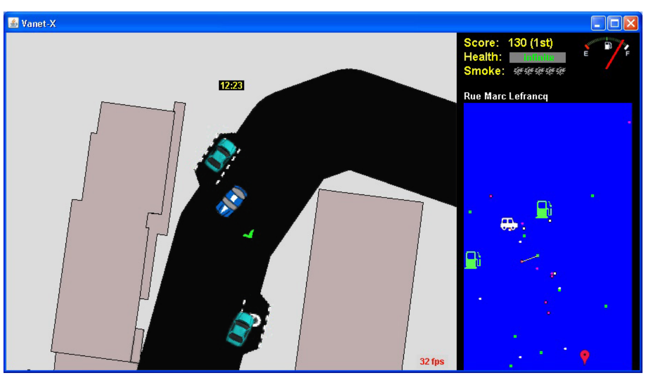

Figure 1: Cars and parking spaces in Vanet-X

- It is implemented in Java as a Java applet, so only a Java Virtual Machine and a browser is needed to play. A desktop application version is also available.

- Both real (human) and computer-controlled players can participate in the game. Human players can join a game through the Internet.

- The game can be configured to execute on a server and create new games when necessary. Alternatively, the computer of any user can play the role of a server and start a new game that other users can join.

- Any real map can be used in the game, by selecting and downloading the data of the desired area from *OpenStreetMap* (`http://www.openstreetmap.org/`).

- To increase the playability, real maps are combined with some extra elements, such as enemy cars, smoke emission devices to disturb enemies (see Figure 2), evolution of events in game time rather than in real-world time, higher maximum speeds for cars controlled by humans, when the driver has a task to go to a certain building he/she has to park nearby and then go by foot to the destination (he/she will be a vulnerable target for enemy cars, that will try to hit him/her, as shown in Figure 3), the car can get damaged and be repaired by paying a certain price (points accumulated during the game), there is infinite or limited fuel depending on the game mode (requiring refueling in a petrol station when running out of fuel in the second case), etc.

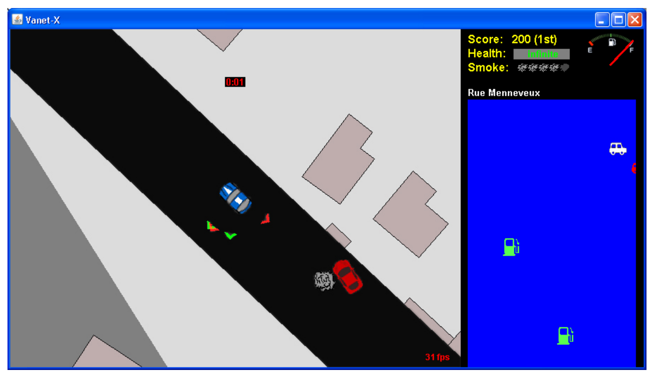

Figure 2: Trying to escape from an enemy vehicle

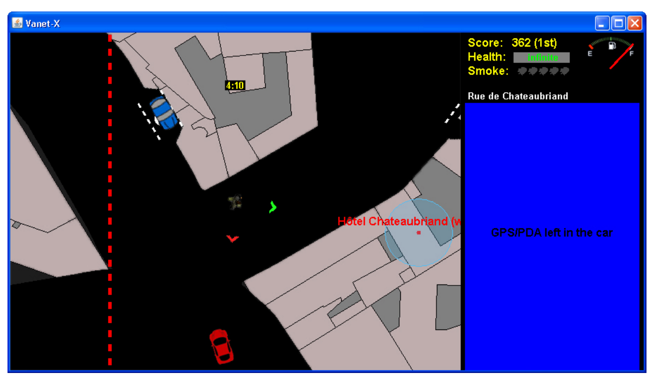

Figure 3: Driver going by foot

- A wide range of game modes is available (see Table 1). Thus, for example, we offer games where the goal is to collect some items along the roads, and task-based games where the players have to complete a series of goals in sequence (with tasks such as parking the vehicle, going to a certain building/business or address, etc.) as soon as possible to win the game. As shown in the table, some game modes can be cooperative, competitive, or both. For the tasks implying going to a certain location, the task may require reaching that location with the car, park and then get there by foot, or just park as near as possible (in this last case, the score for completing the task will be inversely proportional to the distance between the parking location and the final destination). Competitive games involves from 1 to 4 teams in the game, being the winner the team that obtains more points during the game.

Game mode	Multiplayer mode	Inmortal	Possibility to get out of the car and walk	Infinite fuel
Capture the flag (capture 5 flags)	cooperative competitive	no	no	configurable
Capture the enemy cars (1 or more)	cooperative competitive	yes	no	no
Solve tasks (1-3 tasks)	cooperative competitive	no	yes	configurable
Survival (1 or 2 tasks)	cooperative competitive	no	yes	configurable
Park (find one available parking spot)	competitive	yes	yes	no

Table 1: Summary of game modes

- Some default data management strategies, inspired by the work performed in the VESPA project [2, 3, 4, 6, 7], have been implemented. Different tuning parameters can be modified through the graphical user interface of the videogame (see Figures 4 and 5). Moreover, the design of the videogame allows an easy integration of other data management alternatives.

Figure 4: Data management: basic options

- There is a "radar" (e.g., on the right part of Figure 2 we show a basic radar, and on the right part of Figure 1 a radar in debug mode that shows some extra elements about the scenario) that can provide some information to the players. For example, a player can see the following on the radar: his/her location, the petrol stations, and the destination location (if any). Besides, if the option to use a data sharing strategy for that vehicle has been enabled, it will also show data about interesting events received from other vehicles, such as free parking spaces, enemy vehicles, items to pick up (e.g., flags in Figure 6), priority vehicles like ambulances, etc.

Figure 5: Data management: advanced options

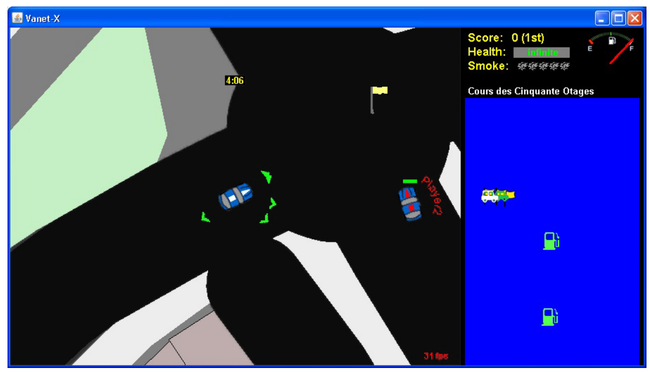

Figure 6: Picking up flags during the game

From a more technical point of view, we have used the Java programming language to develop the video game. Besides, some auxiliary libraries have been useful. For example, we use *Apache Xerces2 Java* (`http://xerces.apache.org/#xerces2-j`) to extract data from the XML files obtained from OpenStreetMap, *JLayer* (`http://www.javazoom.net/javalayer/javalayer.html`) to decode and reproduce MP3 files for the game music, *Guava-12.0* (`https://code.google.com/p/guava-libraries/`), etc.

2.2 Basic Architecture

The basic architecture of the videogame is presented in Figure 7 (the part concerning the collection of statistics about the game is not shown here, as it will be described in Section 4). At a high-level, we can briefly describe the main components as follows:

- A *client* application receives commands from the player, sends them to the server, and receives from the server information about the objects that should be rendered on the screen (see Figure 8).
- The *server* receives the input from the clients, updates the current status of the game (e.g., by considering the movements performed by the vehicles and the tasks that they complete), and generates new goals and events as needed (see Figure 9). The server is multi-threaded, with a thread per vehicle that performs a

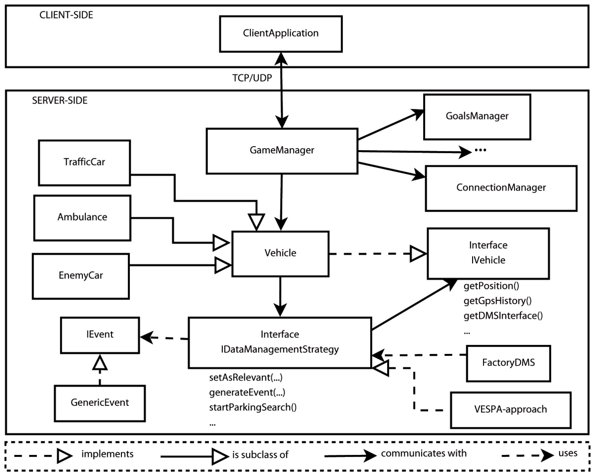

Figure 7: Basic architecture of Vanet-X

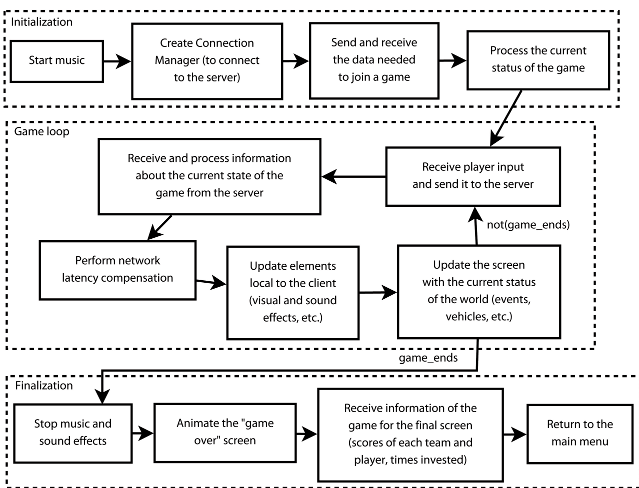

Figure 8: Basic functioning of a client

basic cycle of "while a vehicle is alive, perform actions and check for potential collisions".

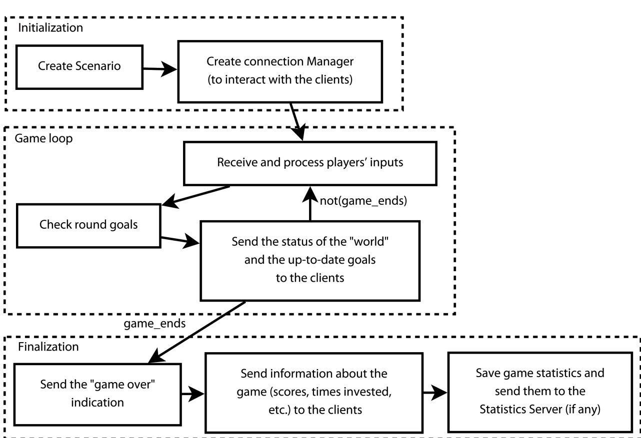

Figure 9: Basic functioning of the server

- An interface *IDataManagementStrategy* declares the methods that should be implemented by a *data management strategy* to allow its integration with the videogame (e.g., a method to define the types of events that are interesting for the driver, a method to generate an event, etc.).
- Another interface *IVehicle* is implemented by the *vehicles* to allow interacting with them (e.g., to obtain a reference to the data manager in the vehicle or to obtain information about the GPS location).

Any data management strategy can potentially be integrated in this framework, as long as it implements the interface *IDataManagementStrategy* and calls the appropriate methods to inform the vehicles (interface *IVehicle*). So, we can easily plug in different alternative data management techniques for testing.

3. BEHAVIORS OF THE VEHICLES

We have implemented several behaviors for the vehicles controlled by the computer, which adapt the steering behaviors proposed in [23]. In particular, we consider the following basic behaviors:

- *Seek* implies directing the vehicle towards a certain static target, by adjusting its direction and speed.
- *Flee* is the opposite behavior to *Seek*, as it implies getting as much further as possible from the target.
- *Pursuit* is similar to *Seek*, but in this case the target is a moving object. So, the expected movement of the target is estimated, to try to catch it.
- *Evasion* is the opposite behavior to *Pursuit* (i.e., based on *Flee* instead of *Seek*).
- *Arrival* implies the progressive reduction of speed as the vehicle approaches the target.
- *Obstacle avoidance* provides vehicles with the ability to dodge vehicles and other obstacles.
- *Wander* generates a random trajectory, to represent a vehicle traveling around with no clear objective. This is useful, for example, to represent a vehicle that is searching for an available parking space in the vicinity.
- *Path following* allows a vehicle to circulate within the boundaries of a certain path.
- *Unaligned collision avoidance* is a behavior that tries to avoid the collision of vehicles moving in different directions. Thanks to this behavior, vehicles can estimate a potential collision risk with other vehicles in the near future, to try to avoid it.

Of course, all the vehicles exhibit the whole set of behaviors at the same time, applying a priority ordering in case several behaviors could be applied at the same time and are in conflict to each other. Based on the previous basic behaviors, we have defined the schema of a normal behavior for different types of vehicles: enemy cars (that try to catch the players or flee from the players, depending on the game mode), ambulances (as representatives of emergency vehicles which may ask the right of way), and traffic cars (that represent neutral cars in the game). As an example, the basic behavior of traffic cars is shown in Figure 10.

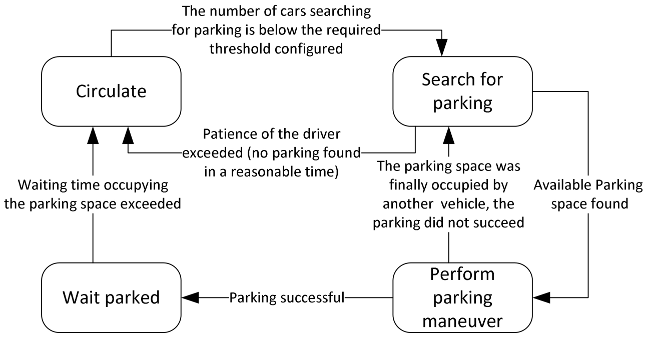

Figure 10: Basic behavior of a traffic car

4. DATA COLLECTION AND EXPLOITATION

In this section, we summarize the strategy applied for data collection during the game and the corresponding exploitation of results. If a certain configuration option that activates the collection of statistics during the game is enabled, several data are collected: data about the scores obtained by the players, the time needed by vehicles (the ones controlled by humans as well as those managed by the computer) to perform certain tasks (such as parking), and other measures about the performance of the data management strategy applied (e.g., events created, events that are considered relevant by each vehicle, etc.). When the game ends, all these data are stored in several files on the game server, along with a file that contains information about all the configuration parameters used in that game (e.g., game mode, configuration parameters used for the data management strategy considered, the wireless communication range simulated, etc.).

To centralize the data collected, it is possible to set up a *Statistics Server*, which is a process executing continuously on a certain computer. In this way, the clients playing the game automatically connect to the Statistics Server when a game ends, in order to communicate the statistics collected during the game. Besides, it is possible to connect to the *Master Server* by using a terminal client (called *Statistics Client*) that allows seeing and modifying the configuration parameters as well as retrieving the statistics files generated. Another option is to avoid the use of a Statistics Server and collect the statistics in the computer that plays the role of a server for a game. If we consider configuration settings where there is a predefined game server and all the clients connect to it to start a new game or join an existing game, this option also keeps the statistics in a single location. However, if there are several game servers then the statistics would have to be centralized manually.

Figure 11 provides an overview of the way the different components of the game, and particularly the Statistics Server, are distributed in a network. Notice that we actually distinguish between a *Master Server* and a *Game Server*. The Master Server is executing on the server machine and a client first connects to it (so, it is the entry point for clients); then the Master Server checks if a Game Server is available and if not it creates one; finally, it returns the port number of its Game Server to the client, as the client will interact with the Game Server during the game.

It should be noted that, as we collect information about the performance of human players, the skills of those players

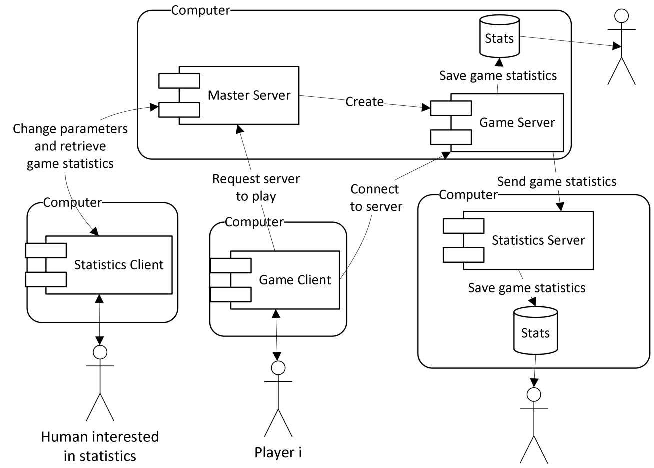

Figure 11: Deployment of components in a network

with the game will have an impact on the results and this has to be taken into account when exploiting the results. Indeed, directly comparing the achievements of several human players without considering their game skills could lead to wrong conclusions. For example, *player1* without a data sharing system could perform better than *player2* with a data sharing system, but we should not necessarily conclude that the use of such a data sharing system is harmful. In other words, we should always compare players with the same skills. For this reason, each human player is assigned a certain *skill level* (which may change along time, as the player improves his/her performance) and the statistics about players are tagged with the skill level corresponding to that player. Besides, players that have a skill level below a certain threshold are (by default) not allow to participate in games with collection of statistics enabled, as performance data about them are assumed to be unreliable and besides their clumsiness could interfere with the normal development of the game. The skill level of a player is computed based on his/her ability to complete missions in the game (tasks per time unit).

5. EXPERIMENTAL EVALUATION

We have performed a few preliminary experiments to evaluate the interest of our proposal. As a use case for testing, we focused on the case of available parking spaces, as these are events that represent scarce resources, which implies additional challenges for data management (i.e., the competition among vehicles should be minimized).

5.1 Data Management Strategies

As a data sharing strategy for the vehicles, we considered the following options.

5.1.1 VESPA-P*: VESPA With No Reservation*

First, we adapted the proposal in [2], developed in the context of the system *VESPA* (*Vehicular Event Sharing with a mobile P2P Architecture*) [4, 6], which is based on the computation of an *Encounter Probability* (*EP*).

The EP between a vehicle and an event estimates the likelihood that the vehicle will meet the event, based on geographic computations that estimate the spatio-temporal relevance of the event. For example, the relevance decreases with the distance between the event and the vehicle, with the time since the event was generated (e.g., consider the case of information about an available parking space, which can be unoccupied only for a limited amount of time), and the direction of the vehicle (e.g., if it is approaching the event or not). In particular, the directions of both the vehicle and the event are estimated and several *penalty coefficients* (α, β, γ, and ζ) are used to weigh the importance of four estimated parameters: the minimum distance to the event over time (Δd), the time until the closest position to the event (Δt), the age of the event at the closest position (Δg), and the angle between the vehicle and the event (c).

So, when a vehicle receives an event it computes its EP and disseminates the event again if the computed EP exceeds a certain *dissemination threshold* (DT). The intuition is that vehicles should disseminate data that are relevant for them (as those data are also probably relevant for the neighboring vehicles). Two other thresholds are managed: the *storage threshold* (ST) and the *relevance threshold* (RT). The ST determines the minimum value of the EP for an event to be stored locally in the vehicle, and the RT the minimum value needed to show the event to the driver.

Besides, the proposal in [2] proposes a contention-based approach for data dissemination in order to limit the network overhead in the dissemination of messages (basically, when there are several candidate vehicles to re-disseminate an event, the message will be disseminated only by the vehicle located further away from the vehicle that disseminated the message previously). Several parameters are used in the protocol, such as D (the maximum time to wait before redifffusing) and D' (time to wait for an acknowledgement that a message sent previously was received by some other vehicle).

5.1.2 VESPA+P*: VESPA With Reservation Protocol*

Communicating the availability of a single parking space to many vehicles could lead to an unfruitful competition among the vehicles to try to reach the same parking space, leading to dissatisfaction of the drivers and parking times that could even exceed those that would be obtained if no data sharing system were used. For this reason, the work presented in [7] proposed an enhancement to the previous approach *VESPA-P* for the case of scarce resources such as parking spaces. It provides an allocation protocol that coordinates a procedure that ensures that the information about an available parking space is communicated to a single interested vehicle.

5.1.3 Blind*: No Data Sharing*

Finally, we also considered an approach where no data sharing strategy is used. In this case, the vehicles receive no information and the only data available for the drivers are what they see with their own eyes. For vehicles trying to find available parking spaces, this will lead to a *blind search*.

5.2 Experimental Settings

The basic configuration of the videogame for the experimental evaluation is as follows. The communication range considered for the vehicles is 200 meters and a maximum of 50% of the vehicles are assumed to be equipped with a data sharing application. The penalty coefficients used to compute the EP for VESPA are: α=1/1500 ($\Delta d \leq 500$ meters), β=1/180 ($\Delta t \leq 60$ seconds), γ=1/360 ($\Delta g \leq 120$ seconds), and ζ=1/270 ($c \leq 90°$); these are parameters that can be considered for a "medium" (not small, not large) dissemination area, according to [2]. The RT and the DT are both set to 75%, and the ST is 60%. The query processor on each vehicle re-evaluates the relevance of the events received with a refreshment period of 2 seconds, showing on the radar the events that are considered relevant. For the dissemination protocol, D is set to 1 second and D' to 2 seconds.

5.3 Experimental Results

We have simulated a varying number of vehicles moving in an area of 1 squared kilometer around the street "Sophie Oury" in the city of Valenciennes (France). In this scenario, we measured the time needed by the vehicles to find free parking spaces near certain destinations. In Figure 12 we show the reduction on the average time needed by a human player to find an available parking space near the target. The experimental results show the interest of sharing data among the vehicles (with both *VESPA-P* and *VESPA+P*), as these data can later be shown on the radar to provide interesting information to the drivers. Besides, according to these results, using a reservation protocol to avoid the competition problem (*VESPA+P*) is particularly beneficial.

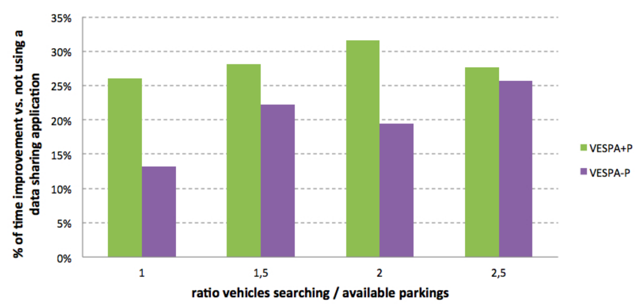

Figure 12: Time to park by a human

In Figure 13 we compare the performance of human players (vehicles controlled by humans) and computer players (vehicles controlled by the computer), by showing the reduction on the average time needed to find an available parking space near the target when using *VESPA+P*. According to these experimental results, we can see that the human players get more benefit from the use of the data sharing strategy. The difference may be due to the way the artificial intelligent behavior of the computer vehicles is implemented.

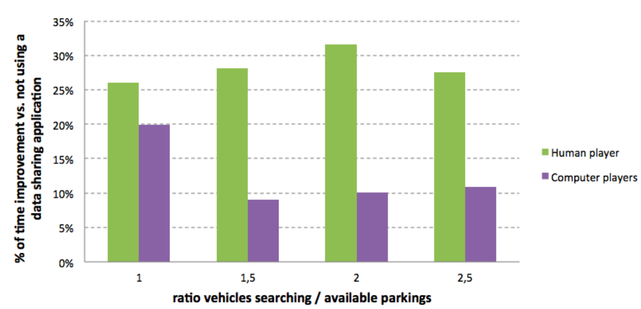

Figure 13: Time to park: human vs. computer

The experimental results obtained correspond to data collected during a total of 14 hours playing the videogame (about 400 parking actions by the human player during this game time). The results are consistent with our intuition and with other experimental results obtained previously by using a simulator. Nevertheless, more tests are needed to validate the results and evaluate other scenarios. For example, we started to obtain some first preliminary results with games played by more than one human player. It is also interesting to perform experiments with other types of events (e.g., accidents, obstacles on the roads, etc.); with information about them, drivers could try to avoid those hazards and so decrease the total travel time.

6. RELATED WORK

As far as we know, this is the first attempt to develop a videogame whose hidden purpose is to help with the evaluation of information management strategies for vehicular networks.

Nevertheless, the idea of trying to benefit from human actions to improve or evaluate a system is not new. Exploiting the power of people to perform large-scale tasks that are costly, time-expensive, or hard, is called *crowdsourcing* [33]. For example, *mCrowd* [32] benefits from sensors available on iPhone devices to perform collaborative tasks such as image tagging or road traffic monitoring. As another example, *reCAPTCHA* [29] exploits *CAPTCHAs* [22] (Completely Automated Public Turing test to tell Computers and Humans Apart), as a security measure to avoid web access to programs, in order to recognize words from scanned books that are challenging for OCR (Optical Character Recognition) systems. According to [10], "The practice of crowdsourcing is transforming the Web and giving rise to a new field".

Particularly relevant for our work with Vanet-X are those proposals that achieve the crowdsourcing results through the use of a videogame. A notable example is the *ESP game* [27], where players implicitly help to label images while playing the game. The use of videogames as learning tools is a clear example of the benefits of using educative videogames; as an example, *CodeSpells* [11] is a fantasy videogame where players have to write spells in Java. Other games with a hidden purpose exist, as commented in [28]. The multiplayer online game *Planet PI4* [16] intends to serve as a testbed environment for Peer-to-Peer (P2P) game architectures. It is also interesting to mention that the term *gamification* has appeared to denote a variety of software that is inspired somehow by videogames [8, 9].

There exist some driving videogames that, as Vanet-X, are based on the use of real road maps or city layouts, such as *Mini Maps* (`https://apps.facebook.com/minimaps/`) and *Push-Cars 2: On Europe Streets* (`http://www.push-cars.com`). However, unlike in Vanet-X, in these games the players do not contribute to any crowdsourcing task or data management strategy evaluation.

Finally, a good number of simulators of vehicular networks and mobility generators have been developed, such as *TraNS* [21], *SUMO* [19], *Veins* (Vehicles in Network Simulation) [25], *GrooveNet* [17], or *VanetMobiSim* [13]. Some interesting surveys can be found in [12, 18]. As commented along the paper, we argue that the videogame-based approach can be an interesting complement (but not a replacement) to the use of existing simulators to evaluate information management strategies for vehicular networks. Besides, mobility generators and vehicle simulators could potentially be used to generate neutral traffic for Vanet-X.

7. CONCLUSIONS AND FUTURE WORK

We have developed a videogame that can be used to evaluate data management strategies for vehicular networks, as a complement to existing simulators. Whereas the opportunity of crowdsourcing through a videogame is attractive, several challenges arise. Thus, the goal of developing a fun videogame required the introduction of several elements that would not appear in a real scenario (like enemy cars), which could have an impact on the results, but on the other hand this will attract people to play. Moreover, the results obtained can depend not only on the benefits offered by the data management strategy but also on the ability of the specific player. So, whereas the videogame can provide an ideal tool to collect many data for a variety of scenarios, the experimental results obtained have to be judged with caution (e.g., we label the collected data with the skill level of the player). Even with these limitations, we argue that the videogame helps to collect with less effort data that can be used to fine-tune a protocol and/or obtain some initial conclusions, prior to the evaluation in more realistic scenarios.

Additional information regarding the videogame is available at `http://sid.cps.unizar.es/Vanet-X/`, including a playable version of the videogame, some videos, and screenshots. This is a first step that shows the potential interest of exploiting videogames to evaluate data management strategies for vehicular networks. As future work, we would like to optimize and improve the videogame, as well as to develop a complete methodology and architecture to collect the data, evaluating the interest of the results obtained in other scenarios and in a larger scale.

8. ACKNOWLEDGMENTS

This research work is currently supported by the CICYT project TIN2010-21387-C02-02 and DGA-FSE. The data management strategy adapted and used as an example in the videogame has been proposed in the context of the VESPA project, and we would like to warmly acknowledge the collaboration with Dr. Thierry Delot in that project.

9. REFERENCES

[1] J. J. Blum, A. Eskandarian, and L. J. Hoffman. Challenges of intervehicle ad hoc networks. *IEEE Transactions on Intelligent Transportation Systems*, 5(4):347–351, 2004.

[2] N. Cenerario, T. Delot, and S. Ilarri. A content-based dissemination protocol for VANETs: Exploiting the encounter probability. *IEEE Transactions on Intelligent Transportation Systems*, 12(3):771–782, 2011.

[3] T. Delot, N. Cenerario, and S. Ilarri. Vehicular event sharing with a mobile peer-to-peer architecture. *Transportation Research Part C: Emerging Technologies*, 18(4):584–598, 2010.

[4] T. Delot and S. Ilarri. Data gathering in vehicular networks: The VESPA experience (invited paper). In *Fifth IEEE Workshop On User MObility and VEhicular Networks (LCN ON-MOVE 2011)*, pages 801–808. IEEE Computer Society, 2011.

[5] T. Delot and S. Ilarri. Introduction to the Special Issue on Data Management in Vehicular Networks. *Transportation Research Part C: Emerging Technologies*, 23:1–2, 2012.

[6] T. Delot and S. Ilarri. The VESPA Project: Driving advances in data management for vehicular networks. *ERCIM News*, (94):17–18, July 2013. Special Theme on “Intelligent Vehicles as an Integral Part of Intelligent Transport Systems”.

[7] T. Delot, S. Ilarri, S. Lecomte, and N. Cenerario. Sharing with caution: Managing parking spaces in vehicular networks. *Mobile Information Systems*, 9(1):69–98, 2013.

[8] S. Deterding, D. Dixon, R. Khaled, and L. Nacke. From game design elements to gamefulness: Defining “gamification”. In *15th International Academic MindTrek Conference: Envisioning Future Media Environments (MindTrek'11)*, pages 9–15. ACM, 2011.

[9] S. Deterding, M. Sicart, L. Nacke, K. O'Hara, and D. Dixon. Gamification: Using game-design elements in non-gaming contexts. In *2011 Annual Conference on Human factors in Computing Systems (CHI'11) – Extended Abstracts*, pages 2425–2428. ACM, 2011.

[10] A. Doan, R. Ramakrishnan, and A. Y. Halevy. Crowdsourcing systems on the World-Wide Web. *Communications of the ACM*, 54(4):86–96, 2011.

[11] S. Esper, S. R. Foster, and W. G. Griswold. On the nature of fires and how to spark them when you're not there. In *44th ACM Technical Symposium on Computer Science Education (SIGCSE'13)*, pages 305–310. ACM, 2013.

[12] J. Harri, F. Filali, and C. Bonnet. Mobility models for vehicular ad hoc networks: A survey and taxonomy. *IEEE Communications Surveys & Tutorials*, 11(4):19–41, 2009.

[13] J. Härri, F. Filali, C. Bonnet, and M. Fiore. VanetMobiSim: Generating realistic mobility patterns for VANETs. In *Third International Workshop on Vehicular Ad Hoc Networks (VANET'06)*, pages 96–97. ACM, 2006.

[14] H. Hartenstein and K. P. Laberteaux. A tutorial survey on vehicular ad hoc networks. *IEEE Communications Magazine*, 46(6):164–171, 2008.

[15] G. Karagiannis, O. Altintas, E. Ekici, G. J. Heijenk, B. Jarupan, K. Lin, and T. Weil. Vehicular networking: A survey and tutorial on requirements, architectures, challenges, standards and solutions. *IEEE Communications Surveys & Tutorials*, 13(4):584–616, 2011.

[16] M. Lehn, C. Leng, R. Rehner, T. Triebel, and A. Buchmann. An online gaming testbed for peer-to-peer architectures. *ACM SIGCOMM Computer Communication Review*, 41(4):474–475, 2011.

[17] R. Mangharam, D. S. Weller, and R. Rajkumar. GrooveNet: A hybrid simulator for vehicle-to-vehicle networks. In *Second International Workshop Vehicle-to-VehicleCommunications (V2VCOM'06)*, pages 1–8, 2006.

[18] F. J. Martinez, C. K. Toh, J.-C. Cano, C. T. Calafate, and P. Manzoni. A survey and comparative study of simulators for vehicular ad hoc networks (VANETs). *Wireless Communications & Mobile Computing*, 11(7):813–828, 2011.

[19] J. E. Michael Behrisch, Laura Bieker and D. Krajzewicz. SUMO – Simulation of Urban MObility: An overview. In *The Third International Conference on Advances in System Simulation (SIMUL'11)*, pages 63–68. IARIA, 2011.

[20] S. Olariu and M. C. Weigle, editors. *Vehicular Networks: From Theory to Practice*. Chapman & Hall/CRC, 2009.

[21] M. Piorkowski, M. Raya, A. L. Lugo, P. Papadimitratos, M. Grossglauser, and J.-P. Hubaux. TraNS: Realistic joint traffic and network simulator for VANETs. *SIGMOBILE Mobile Computing and Communications Review*, 12(1):31–33, 2008.

[22] C. Pope and K. Kaur. Is it human or computer? Defending e-commerce with Captchas. *IT Professional*, 7(2):43–49, 2005.

[23] C. W. Reynolds. Steering behaviors for autonomous characters. In *Game Developers Conference*, pages 763–782. Miller Freeman Game Group, 1999.

[24] M. L. Sichitiu and M. Kihl. Inter-vehicle communication systems: A survey. *IEEE Communications Surveys & Tutorials*, 10(1–4):88–105, 2008.

[25] C. Sommer, R. German, and F. Dressler. Bidirectionally coupled network and road traffic simulation for improved IVC analysis. *IEEE Transactions on Mobile Computing*, 10(1):3–15, 2011.

[26] Y. Toor, P. Mühlethaler, A. Laouiti, and A. de La Fortelle. Vehicle ad hoc networks: Applications and related technical issues. *IEEE Communications Surveys & Tutorials*, 10(1–4):74–88, 2008.

[27] L. von Ahn and L. Dabbish. Labeling images with a computer game. In *SIGCHI Conference on Human Factors in Computing Systems (CHI'04)*, pages 319–326. ACM, 2004.

[28] L. von Ahn and L. Dabbish. Designing games with a purpose. *Communications of the ACM*, 51(8):58–67, 2008.

[29] L. von Ahn, B. Maurer, C. McMillen, D. Abraham, and M. Blum. reCAPTCHA: Human-based character recognition via web security measures. *Science*, 321(5895):1465–1468, 2008.

[30] T. L. Willke, P. Tientrakool, and N. F. Maxemchuk. A survey of inter-vehicle communication protocols and their applications. *IEEE Communications Surveys & Tutorials*, 11(2):3–20, 2009.

[31] B. Xu, A. M. Ouksel, and O. Wolfson. Opportunistic resource exchange in inter-vehicle ad-hoc networks. In *Fifth IEEE International Conference on Mobile Data Management (MDM'04)*, pages 4–12. IEEE Computer Society, 2004.

[32] T. Yan, M. Marzilli, R. Holmes, D. Ganesan, and M. Corner. mCrowd: A platform for mobile crowdsourcing. In *Seventh ACM Conference on Embedded Networked Sensor Systems (SenSys'09)*, pages 347–348. ACM, 2009.

[33] M.-C. Yuen, I. King, and K.-S. Leung. A survey of crowdsourcing systems. In *Third International Conference on Privacy, Security, Risk and Trust (PASSAT 2011) and Third International Conference on Social Computing (SocialCom 2011)*, pages 766–773. IEEE, 2011.

HealthNet: A System for Mobile and Wearable Health Information Management

Christoph Quix, Johannes Barnickel, Sandra Geisler, Marwan Hassani,
Saim Kim, Xiang Li, Andreas Lorenz, Till Quadflieg,
Thomas Gries, Matthias Jarke, Steffen Leonhardt, Ulrike Meyer, Thomas Seidl
UMIC Research Cluster at RWTH Aachen University, Germany
http://www.umic.rwth-aachen.de/

ABSTRACT

Medical health care is undergoing a significant change of paradigm. Moving health care from health centers to home environments poses new challenges for acquisition, management and mobile exchange of information. The HealthNet project at RWTH Aachen University has developed a prototype which addresses these new challenges: a Body Sensor Network (BSN) collects information about the vital functions of a patient while she is in her home environment; the integration of smart textile sensors increases the acceptability of such technology; mobile communication and data management enables the exchange of health data between patients and doctors; data stream mining techniques tuned for mobile devices provide immediate feedback of the collected data to the user; and finally, advanced security and privacy features increase user acceptance and cope with legal requirements. This paper summarizes the challenges and achievements in the development of this prototype.

*C. Quix, S. Geisler, X. Li, M. Jarke, and A. Lorenz are with Informatik 5 (Information Systems) at RWTH Aachen University, Germany (geisler,jarke,lixiang,lorenz@dbis.rwth-aachen.de).

†C. Quix and M. Jarke are also with the Fraunhofer Institute for Applied Information Technology FIT, St. Augustin, Germany (christoph.quix,matthias.jarke@fit.fraunhofer.de).

‡J. Barnickel and U. Meyer are with the IT Security group at RWTH Aachen University, Germany (barnickel,meyer@itsec.rwth-aachen.de).

§M. Hassani and T. Seidl are with Informatik 9 (Data Management and Exploration) at RWTH Aachen University, Germany (hassani,seidl@cs.rwth-aachen.de).

¶S. Kim and S. Leonhardt are with the Philips Chair for Medical Information Technology (MedIT) at RWTH Aachen University, Germany (s.kim,leonhardt@hia.rwth-aachen.de).

‖T. Quadflieg and T. Gries are with the Institute for Textile Technology (ITA) at RWTH Aachen University, Germany (till.quadflieg,thomas.gries@ita.rwth-aachen.de).

Keywords

Health Information Management, Data Acquisition, Data Analysis, Body Sensor Network

1. INTRODUCTION

The demographic change with a growing population of elderly people and the associated increase of health care related expenses require new models for health care management. Moreover, there is a growing group of health-aware people that would like to take more personal responsibility for their own health, e.g., by monitoring their vital parameters during sport activities. New innovative technologies are necessary to fulfill these new requirements. Mobile and remote health monitoring has demonstrated positive influence on patients disease courses, especially for chronic diseases [20, 13], and promises high cost reductions [16]. While various systems have been proposed to measure the physiological state of mobile users, most of these systems are restricted to a certain set of sensors, or can monitor only a few vital parameters [6].

In this paper, we describe an extendable and flexible monitoring system for the case study of physiological state monitoring of runners. The system has been developed in the context of the HealthNet project [14][1], which addresses interdisciplinary challenges such as sensor network design, manufacturing of smart textiles, information exchange, data mining, security and privacy, and mobile communication. The HealthNet project is part of the UMIC Research Cluster at RWTH Aachen University which focuses at Ultra high-speed Mobile Information and Communication systems supporting the demands of future mobile applications and systems. In the prototype developed by the HealthNet project team, the vital functions of athletes (or patients) are monitored by a BSN (e.g., ECG, skin humidity / temperature, activity) which are partly integrated into textiles. These sensors produce data streams that are integrated, consolidated, and aggregated in a device which acts as a peer in a network. Other trusted participants in the network are, for example, other runners or trainers who want to observe the performance of a runner. In a medical scenario, other peers in the network might be doctors or nursing staff who monitor the state of a patient while (s)he is at home. Furthermore, data can be stored in a server system for long-term monitoring and analysis. An intensive monitoring of vital parameters of patients is especially important after they have

[1] http://dbis.rwth-aachen.de/cms/projects/UMIC/healthnet

been released from hospital. Changes in environment and medication often result in expensive re-hospitalizations of patients which could be avoided by more detailed observation of vital parameters [10]. Thus, both scenarios (sports and medicine) share a common basis; in addition, specific features like the identification of critical situations are relevant for both cases, although the definition of a critical situation is different. Nevertheless, the same techniques for data analysis can be applied. Furthermore, merging the acquired sensor data with additional information such as position, time, or weather conditions improves the expressiveness of pure health data and can lead to new insights.

Using a mobile communication infrastructure (e.g., UMTS, LTE, or Wi-Fi), mobile devices can communicate with each other such that peers can easily exchange health information. Especially, the mobility of patients is improved as detailed monitoring can now be performed at home: periodically or in the case of important events, the device sends the collected and pre-processed data to information systems maintaining patient health records (e.g., hospital information systems) which can be accessed by medical experts.

The main challenges in this project are

- the design and development of wireless medical sensors, which are able to monitor vital functions of a person,
- the integration of theses sensors into textiles and development of electronic units as textiles (e.g., conductive paths) for unobtrusive and comfortable usage,
- the integration of the data collected by various sensors in one data stream, and
- the analysis, mining, and aggregation of the sensor data to detect emergency events, to reduce communication costs, and to predict near future.

We addressed these challenges in the HealthNet project and report in this paper our experiences in developing an integrated prototype. Sectioin 2 first describes the requirements analysis which we have done with a group of athletes. An overview of the prototype system and its architecture is given in Section 3. The main components of the prototype are an intelligent T-Shirt with integrated conductive leads/electrodes (cf. Section 4) and a Body Sensor Network (IPANEMA, Integrated Posture and Activity NEtwork by Medit Aachen) which aggregates multiple data streams from a range of sensors (cf. Section 5). The data are transmitted wirelessly over a Bluetooth interface to a mobile device for visualization and a first lightweight analysis (cf. Section 6). A more detailed analysis of the data is done on a server which receives periodically or in case of peculiar events data from the mobile device. We also performed a case study in a running event of which we will briefly summarize the results in Section 8.

2. REQUIREMENTS AND USE CASES

For gathering of requirements, four active runners on semi-professional level were interviewed. All interviews were conducted by two interviewers with one interviewee. The interviews used a unique set of 14 questions regarding mobile health monitoring, and eight questions regarding a stationary counterpart. The questions especially targeted the usability of smart phones as supporting device in runs and the information needs of the runners and the willingness to share information during training and competitions. All interviews took about one hour, and were recorded for post-interview analysis.

2.1 Participants

The interviewed runners are male, three between 20 and 25 years old, and one between 30 and 35 years old. All runners participate in competitions on national level. The disciplines range from 3000 meters steeplechase to marathon distance, and triathlon. All interviewed persons do intensive training between 5 and 15 sessions per week.

2.2 Information Requirements

2.2.1 Personal Information Sources

All interviewed persons consider their self-assessment as the most important source of information, which is even more reliable than any physiological measure. They stated to ignore measured high peaks of their pulse if feeling good, and also low measures if feeling bad. That means, that their subjective rating of their state is more important for them than an objective measurement. They treat many technologies as fun, which could be more interesting to increase motivation in mass sports.

2.2.2 Medical Information Sources

In addition to the self-assessment, all interviewed runners were interested in heart rate (they all measure it in training). The data could be used to trigger a notification if exceeding upper or falling below lower personal limits. Furthermore, the sportspersons consider breathing rate and oxygen absorption relevant to detect exhaustion in advance. Sweat analysis could lead to an estimation of water balance of the body, useful for reminding the runner to drink or to determine the amount of liquid to drink for convalescence after a training session. Determining the blood sugar level could signal a low sugar level or hunger knock[2]. Last but not least, all interviewed runners had measured lactate in the past. It is probably the best indicator for the current personal fitness. Noticeably, none of the interviewed runners associated any value with blood pressure information, even if explicitly asked by the interviewers.

Most measurements listed above require settings incompatible with daily outside use. Some request tests in medical labs, some include in addition blood analysis (such as blood sugar level, lactate) which is not compatible with mobile use. All interviewed runners agreed that therefore it will be challenging or impossible to apply these measurements in their training sessions or competition.

2.2.3 Track, Time and other Information

Speed measurements have a strong influence on the running speed. All runners pointed out that speed measures from cars or bicycles are not usable because of the meaningless unit (mph, km/h) and low precision. They request measures of time needed for the last lap (on cycle tracks), the last 400 meters or the last 1000 meters (all preferable in a unit of minutes:seconds) to adjust their personal running speed accordingly.

[2]a completely run out of energy, also known as "bonk" or "hit the wall"

For uphill sections, the absolute distance and remaining distance of the uphill part are valuable for all interviewed runners. The gradient is less important because of the low absolute number.

Other information, like weather conditions, weather forecast or condition of the ground are important in preparation of training or competition; it is of no value while being on the move.

2.2.4 Personalization

All interviewed persons request methods for personalisation of the measurements and accompanied items, such as frequency, upper/lower border.

2.3 Mobile Monitoring

2.3.1 Use of Technology

All interviewed runners had applied technology for monitoring heart rate; all interviewed runners knew technology for gathering track data (i.e., GPS). None used other technology, like step counters or sensors in shoes. Only one person carries the mobile phone in training sessions, in a back pocket together with keys. They do their sports without listening to music.

All interviewed runners track heart rate in training, only two do the same in competitions. Two do not track the heart rate in competition mainly because of loosing comfort, i.e., chest belt slipping out of place and making the runner feeling confined. The interviewed runners do not agree to carry any additional device. In competition, none of them would be willing to carry a mobile device.

2.3.2 Carrying a Mobile Device

Carrying a mobile device while doing sports is considered burdensome. There must be a reasonable benefit from doing so. It must not require any attention by the runner, it must not swing (e.g., on a neck strap), it must not disturb the rhythm of arms, legs or breathing (the latter nearly excludes speech interfaces). The device must be lightweight, small, waterproof and shock resistant. The touch-sensitive surface, if any, must come in a sweat resistant cover.

The shape and feeling of a watch was considered most appropriate, as applied in current monitoring systems for heart rate. It can be worn at one arm and operated with the hand of the other arm. If more functions are to be integrated, the only sensible way of carrying a larger mobile device seems to be a pocket at the arm. It supports a similar way to operate it using the hand of the arm not carrying the device.

2.3.3 Operating a Mobile Device

Operation of a 1-button-watch was considered sufficient; nevertheless the operation of buttons of a mobile device were considered to require too much attention and too fine granular movements for hand and finger. A mobile device at the arm can be similarly operated by touch on the display.

The interviewed persons see the problem with touching the display that it might get dirty and smeared by the runner's sweat, making checking current values from the display impossible. Because of disruption of rhythmic breathing, speech-based operation is only considered feasible for a few short commands.

2.4 Sharing of Information

2.4.1 Live-sharing

Live-sharing information with others is considered a minor issue by the interviewed runners. Together with personal trainers, post-processing (for long distance runs) and frequent analysis after smaller sessions (e.g., in interval training) was seen to be more important than live data transmission. One interviewee had the idea that the trainer might interrupt over-pacing of a runner in a hopeless intermediate state of a competition, especially if it is one in a row of competitions. All interviewed persons declined to lively share personal or medical information with other external persons like friends, training mates, online communities, or event organisers or competitors. It was only acceptable for notification in case of emergency.

To receive information from others, trainers and supporters call out time information and intermediate state of the competition to the runner on track. The interviewed runners think that receiving more information, e.g., about personal state of competitors, is rather distracting. One of the interviewed persons stated a value of knowing intermediate state of competition within the same age group, in particular if persons nearby are of the same or another group like the runner. Getting the positions of the team mates was considered not interesting, neither in training nor in competition.

As an open question, the interviewees brainstormed about other ideas for valuable live-sharing of information. As a result, it could be valuable for optimisation of the handover in relays, especially in long distance relays. It would be of value to the successor to know the personal health state of the predecessor in order to adjust warming and preparation phase. If the predecessor is in good shape, the estimated arrival time is earlier than if the person is in bad shape, influencing the point of time to start preparation.

2.4.2 Post-event sharing

After a training session or competition, the runners were open to share track and time data with team mates and online communities, which is already implemented by portals like `http://www.gpsies.com`.

2.5 Persistent Storage

Post-processing of the collected data is very important to all interviewed runners. They asked to file all information to a computer system for persistent storage. They all use a kind of training diary, two use already computer applications for this purpose.

2.5.1 Connecting with PC

The interviewed runners asked for easy connection with the PC, and easy to handle download.

2.5.2 Post-Processing

All interviewed runners do intensive performance analysis combining tracking data, time data, health information, and comments on personal feeling. If applicable they compare current data with past datasets for recurring events, competitions, tracks, or distances. The main goal of the analysis is identification of flaws in performance (absolute speed, endurance, power to go uphill) requesting updates of the training method and plan.

The triathlete analyses shifts in performance of the single disciplines, e.g., intensively training one discipline has contradictory influence on the performance in the other two.

One runner mentioned to use the post-processing also to estimate lifespan of used hardware, e.g., professional running shoes that loose suspension after 3000 km of use, demanding for replacement to prevent damage from tendons and ligaments.

2.6 Use Cases

Based on the requirements analysis, several use cases were identified which are described in this section. The use cases are grouped in four categories: Sensor management, mobile monitoring, sharing, and archiving.

2.6.1 Managing Sensors

The sensor managing use cases describe the setup, configuration and maintenance of the set of sensors delivering information to the system. The actor usually is the user. In addition, other persons or organizations might perform the use cases as well, e.g., an emergency doctor who adds a sensor after the user had an accident, or a physician who adjusts the upper border of a physiological parameter to raise notification earlier. The actor employs a plug-in / plug-off mechanism to add or remove sensors to the network; this should be as automatic as possible. The added sensors perform registration and de-registration at the controlling component of the sensor network. Configuration of the sensors should be also possible, so that user can adjust the properties of the sensor (e.g., sampling rate, sensor identifier, measure unit, data transmission rate) to his/her personal needs.

2.6.2 Monitoring

The monitoring use cases describe the use of a mobile system to monitor the health status. The user employs the system for observing specific parameters, being informed about the current status and alarming himself or another entity during a personal activity. The user can also turn off all monitoring and notification functions by muting the device.

2.6.3 Sharing

The group of sharing use cases describes the information exchange between all parts of the system with external entities (e.g., server or other users). It applies to sharing information while being mobile as well as sharing information from the other parts of the system like the archive. The group contains:

2.6.4 Archiving

The archiving use cases describe the use of and retrieval from a persistent storage. The user employs a stationary device (such as a laptop or desktop PC) to search for information of a specific type, date and time, activity, or value. The archiving use-cases are:

3. SYSTEM ARCHITECTURE

The HealthNet prototype is based on a BSN integrated into a textile platform (i.e., T-shirt) measuring the physiological state of a person. An overview of the system is illustrated in Fig. 1. A registry server manages the communication between different peers in the network. The sensor data is received by a smartphone via Bluetooth which sends the data to other peers in the network. Other peers in the network are an advanced data mining & analysis service or other trusted parties such as trainers and doctors.

In the current prototype, the BSN consists of an ECG sensor, a combined temperature/humidity sensor, two 3D acceleration sensors, and a master node. The master node collects the data from the individual sensors and sends it to the smartphone. Conductive yarn acts as electrodes as well as leads. The signals are received by the ECG sensor attached to the shirt. The sensor processes the ECG and infers the current heart rate from it.

On the smartphone, a mobile application integrates the health data with data measured by the phone, such as the current GPS position. The mobile application also visualizes, stores, and analyzes the data. If enabled by the user the integrated data is sent via UMTS or Wi-Fi (IEEE 802.11) to a registry server which distributes it to registered third parties, such as a trainer, a doctor, or a server analyzing the data. The architecture also allows sending feedback and results of the analysis of the data to the users smartphone.

4. TEXTILE PLATFORM

The state-of-the-art electrodes used for most medical applications are, for example, disposable electrodes glued onto the skin. These electrodes are coated with electrolyte-gel to improve the conductivity. The advantages of those electrodes are low contact impedance and a fixed position. However, they are not suitable for a continuous long-term measurement because the electrolyte-gel can dry and may also cause allergic reactions. Moreover, the wires between electrodes and the sensor exacerbate the handling for untrained users. To achieve the aim of a continuous and mobile monitoring system, another solution has to be found.

Textile electrodes could be a good alternative for the standard ones. They can be used for long-term measurements because they are not coated with electrolyte-gel. The yarn for the textile electrodes must possess high conductivity, good elastic behavior to assure a good skin conformance and it should be biocompatible due to the constant skin contact. Another advantage of textile electrodes is that these electrodes can be integrated into garments which lead to a very high mobility of the whole system and intuitive handling. Mobility can be further increased by using textile integrated conductive paths instead of cables. A reversible interface is necessary to remove the sensor node before washing. However, textile electrodes also have disadvantages: the contact impedance is higher and movement causes motion artifacts.

Suitable yarns matching all requirements mentioned above have been researched and tested. The best one was a silver-coated polyamide yarn. A circular foam padded textile electrode with a radius of 2.5 cm was used. In addition to the ECG electrodes, the same material was also used to manufacture the textile conductors (see Fig. 3). The textile conductors were applied to the outside of the T-shirt with metal push buttons to connect both electrodes and the sensor. Preliminary results with this T-Shirt show the suitability of textile electrodes for the application as ECG electrodes.

5. BODY SENSOR NETWORK

Body Sensor Networks (BSN) usually consist of a varying number and diverse types of sensors. They are wirelessly connected either to each other, called mesh network, or to a central master node, called star network. The acquired data is then transferred over wide area networks (WAN) to

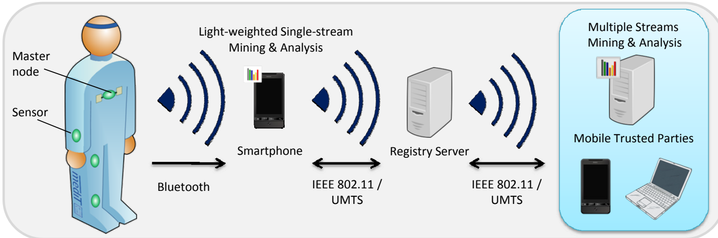

Figure 1: Overview of the system architecture

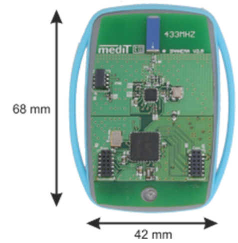

Figure 2: BSN node

Figure 3: Sensor shirt

central data and health service providers for further processing. This section focuses on the challenges in developing the medical sensors, connecting them in a BSN, and processing the measured signals.

Bringing health status monitoring to personal health care environments presents a new set of challenges: devices have to be small, unobtrusive and easy to handle. Preferably, they need no or only minor interaction and are connected via wireless technology to the supervising medical professional or health care center.

The IPANEMA BSN is designed to be easily modified for different application scenarios, e.g., cardio-vascular monitoring or hydration status monitoring [11]. It is small (68 x 42 mm, see Fig. 2), light (30 g) and wireless enabled. A sensor node consists of a base board which includes a low power microprocessor (MCU, MSP430F1611, Texas Instruments), power management circuitry, and a low power radio transceiver (CC1101, Texas Instruments). Modularity is ensured by using a pair of connectors to attach different sensor extensions. Two connectors (Samtech Inc.) enable the use of digital (SPI, UART, I2C) sensors, five analog-to-digital converter inputs and three interrupt capable inputs. The MCU is running at 8 MHz with an additional precision 32.768 kHz crystal for the real time clock. It is powered by a lithium polymer battery which can be recharged over an on-board MicroUSB connector.

The sensors of the current prototype produce a raw data stream of about 14 kbit/s which is transmitted over a 433 MHz ISM band transceiver with a proprietary protocol.

The network is structured in a star topology. The leaves are formed by a flexible number of modules which can be equipped with different types of sensors. The sensor data is send over-the-air to a central master module. The main tasks of the master node include network management, data transfer to a mobile device and creating time synchronization beacons for the sensor nodes.

6. MOBILE APPLICATION

The goal in the design of the architecture of the mobile application was to have a very flexible and extensible system. As explained before, the HealthNet project is not limited to a particular application domain, our solution should be applicable in a healthcare domain as well as in a sports domain. To allow easy customization and adaptation to new domains, we identified four main components for the mobile application on the smartphone (cf. Fig. 4).

The *HealthNet Controller* is the central unit for managing the set of active sensors, and notifying dependants if measures changed value or the composition of the network changed. The *Data Cache* stores recent sensor data in a

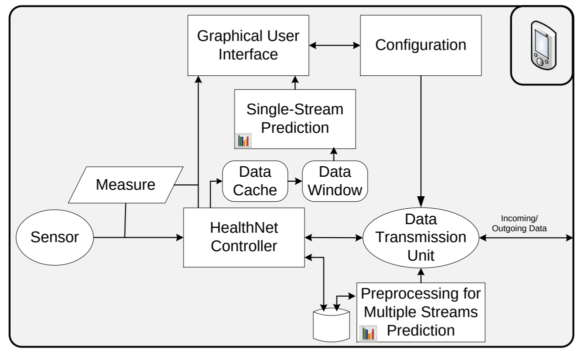

Figure 4: Architecture of the mobile application

Data Window such that a single-stream prediction over a short timeframe is possible. The windows are implemented as a circular data structure - if a window is full, the latest incoming data will flush the oldest data. Furthermore, the cache stores also all data (if desired by the user) such that the data can be uploaded to a server for detailed data mining and analysis later on.

The *Data Transmission Unit* (DTU) takes care of the information exchange among different stakeholders. Four methods of sending data to authorized entities have been implemented. Any external entity must prove eligibility to receive any data from the mobile application. The DTU supports three communication modes:

1. *Request-response:* an external entity requests information from the mobile application. The DTU retrieves the requested data from the cache and transmits the response. This is for example done when a trainer wants to see detailed data about a runner.

2. *Time-based submission:* A fixed interval after which a selected data set is sent, e.g., data is sent from the runner to a trainer only every 10 seconds to reduce required bandwidth and communication costs.

3. *Direct transmission:* The relevant data is transmitted directly to the receiver. This mode is used for audio feedback from a trainer to a runner.

Due to the modular design, peer mobile applications use roughly the same architecture, with the only difference that these applications receive data via the DTU and not from sensors.

On the user interface level, the data which is received from the sensors or other peers is managed according to the use cases as described in section 2.

6.1 Data Analysis

To get the maximum benefit of the HealthNet application the measured data has to be analyzed to detect critical situation or events, and to make a short term predictions. Data mining techniques in this context are restricted by two important constraints: (i) the data is a continuous stream and has to be analyzed in real-time; persistent storage and long time series of data are not available as in classical data mining tasks, (ii) the resources (CPU power, battery life, memory) of the mobile device are very limited.

To cope with the problem of limited resources, we developed: (i) an adaptive technique for anytime classification, which is capable of both, classifying under varying time and resource constraints, and incrementally learning from data streams to adapt to possible evolutions of the underlying data stream [15], (ii) and a novel in-network distributed sensor data clustering technique that efficiently aggregates similar sensor readings using coordinators [8].

Context prediction is an emerging topic in the field of data mining, e.g., predicting the location of mobile objects was a frequently tackled subtask of mobile context prediction in recent researches. For scenarios of managing health information of mobile persons, the prediction of the near future health status of persons is at least equally important to predicting their location. A first method for predicting the next health context of mobile persons equipped with body sensors and a mobile device has been developed and implemented [7]. The proposed PrefixSpan-based method searches for sequential patterns within multiple streaming inputs from the body sensors as well as other contextual streams that influence the health context.

Our main observation is: frequent sequential patterns appearing in rules containing multiple streams, are completely built using frequent patterns existing in each single stream. Thus, predicted values were directly presented to the user in the mobile application using a light-weighted resource-aware algorithm that was implemented locally on the user's mobile phone. More accurate predicted values were sent to the user from a multiple stream prediction algorithm which was implemented on a server using the preprocessed frequent patterns on each stream (cf. Fig. 4).

6.2 Security and Privacy

A rigorous evaluation of security and privacy risks was done, requirements were derived from it, and the implementation was developed accordingly [3]. The measured data is kept confidential at all times: during collection, in storage, and during transmission within and between all components of the system. To reduce the risk of data extortion from stolen devices, secure authentication methods are used both for wireless links as well as user interfaces on the devices themselves. Generally, data may only be read by persons authorized by the user. Finally, no more data than required for a given monitoring application shall be stored.

Confidentiality during data collection is achieved by using ZigBee AES-128 encryption between the sensor nodes and the master node, and Bluetooth encryption E0 is used between the master node and the smartphone. Confidentiality during data storage is achieved using AES-128 encryption on the devices, so that no data can be recovered wrongfully by someone who has physical access to a device. During communication between trusted devices, we do not rely on the security mechanisms of the technologies used (e.g., UMTS, LTE, WLAN) because the data must not be revealed to the network operators, and wireless technologies such as UMTS, LTE, and WLAN typically only encrypt the air interface. Instead, all data transfers apply AES-128 encryption and message authentication codes on the application layer.

In wireless connections to trusted parties, all parties are identified using certificates with shared keys. The implementation of the encryption is transparent to the application as standard interfaces of the Android SDK are used to implement secure storage and communication. In addition, we found that the authentication and encryption mechanisms had no significant influence on battery life or performance of the handheld device.

7. RELATED WORK

The interest in mobile healthcare applications started with systems like [19, 12] supporting professionals (like physician, nurse, therapist, or midwives) to enter, receive and exchange information about their patients. Systems for professional users in hospitals like [2] considered specific design aspects to support local mobility in the hospital by interconnecting PDA, laptop and desktop computers. Examples of systems for non-professional users are the self-monitoring application for overweight people [22], alcohol consumption monitor [4], or dietary advisor [9]. The results of these studies point to a high degree of monitoring by those using a mobile monitoring device compared to other monitors. In difference to the aim of the HealthNet project, these systems are not equipped for continuously monitoring vital parameter in silent mode.

7.1 Textile Sensor Platforms

A reasonable idea to integrate real-time monitoring into daily life activities are the application of wearable or textile sensor platforms. This section therefore reviews the integration of sensors into garments, such as sport shirts or similar. In [5] two types of textile sensor platforms are distinguished: while *textile sensors* are realized by special yarns, *non-textile* or *textile-integrable* sensors are singular units which are applied to the garment, e.g., printed onto the textile. The advantage of textile sensors is that these can be produced in one manufacturing process [17]. A disadvantage is that current technologies for textile sensors have to be moistened to deliver acceptable results [5].

To integrate multiple vital parameters into one textile platform, several sensors are combined to form a sensor network. Often, a master component controls the network and centralizes data acquisition, short-term storage and transmission. These can be realized either wired or wireless.

7.2 Textile-based Monitoring Applications

The MyHeart-project[3] led by Philips was dedicated to the prevention, diagnosis and therapy of cardiovascular diseases The monitoring is based on sensors integrated into daily life textiles, such as undergarment. A sensor shirt has been

[3]`http://www.hitech-projects.com/euprojects/myheart/`

developed using conductive and piezoresistive yarn for monitoring of heart (ECG) and respiratory activity (impedance pneumography), core and skin temperature with non-textile sensors and an accelerometer [1]. The shirt has been used for monitoring during outdoor activities and at home. A proprietary user device or PDA is used for interaction [21].

The respiratory sensing technology was also used in the Wealthy project[4] [18, 17]. For the data processing and transmission a relatively heavy and big Portable Patient Unit (250g) was connected with the sensors by wires. The data is transmitted from the PPU via GPRS to a central system analysing and visualizing the data.

A project that supports medical treatment and behaviour of elderly people suffering from cardio-vascular disease is described in [23]. The system comprises a front worn array of body sensors, a user interaction system for a PDA for displaying information and entering simple answers and a back-end system for professionals analysing data and providing feedback.

7.3 Products for Sports Monitoring

Commercial products are available on the market in particular to support ambitious sports(wo)men. The products do not aim on sophisticated measuring medical data. Usually, it is considered sufficient to provide heart rate and calories burned, and location and time related information. The often use wrist or chest bands.

A large set of wrist-mounted computers is available for example from Polar, ranging from low-end technology for beginners to high-end systems for professionals like the RS800[5]. They receive body signals from chest straps, display and store the information on the watch, and allow for downloading and post-processing with the personal computer. A similar system is the Garmin Forerunner[6]. It monitors time, distance, pace, heart rate and calories burned. As Garmin's unique selling proposition, it additionally tracks the position of the sportsperson by the use of a high-sensitive GPS receiver built into the wrist watch. The GPS antenna is partially integrated into the watchstrap. The heart rate is measured by the use of a chest strap. The system supports different profiles, e.g. for swimming, cycling, and running of triathlons.

A more sophisticated system is the adidas miCoach[7]. It is an integrated system to plan, work-out, and analyse personal training. As the main part of the system, it combines three components to support the work-out: an auditive display (miCoach Pacer) for heart rate measures, speed and distance which reacts to the speed; a bundle of a chest belt measuring heart rate and a wristwatch (miCoach Zone); an application running on the user's mobile phone for coaching (miCoach Mobile).

As a main advantage, the textile strap for monitoring the heart rate can be replaced by two different bra's (adidas supernova glide/sequence bra) or a shirt (adidas supernova cardio shirt). Nike+[8] is a training system similar to miCoach developed by Nike and Apple. The main differences to miCoach is that Nike+ combines the features from Nike's running shoes with an integrated step counter (with the drawback of getting depended on the Nike's brand), and that it uses the iPod instead of a mobile phone (with the same drawback of dependency).

8. CONCLUSIONS AND LESSONS LEARNED

We implemented an end-to-end prototype for a runner scenario (training and competition mode) with one or more runners and a trainer. Case studies with the implemented prototype have been conducted during the Lousberglauf 2011 & 2012 (a local running event in Aachen with about 2000 participants). A team of five runners has been equipped with the sensors and smartphones. In addition, a trainer monitored the performance of the runners using also a smartphone. Data communication and management did not cause any problems; the trainer could always see the position and vital parameters of the runners. Due to excessive motion artifacts during running, we used standard electrodes for the run. In the meanwhile, we did some additional measurements with a new version of the textile electrodes in a lab environment on a treadmill which gave better results. We also improved the algorithm for inferring the heart rate from the raw ECG data, such that it is less sensitive to movement artifacts. This improved the data quality in the second case study in 2012, but the data quality is still too low for deriving health-related advices.

We have shown in this project that health monitoring using mobile wearable sensor networks is feasible. Data management and analysis can be done in real-time although the data is coming at a high frequency. Security and privacy issues have been addressed by implementing suitable encryption and authentication mechanisms into the application. In another related project (Nanoelectronics for Mobile AAL-Systems[9]), a similar approach for data management has been developed in the context of Ambient Assisted Living (AAL). Some results (e.g., the architecture of the mobile application in Fig. 4) have been applied also in this project.

However, we have seen that with the current technology, problems like data management, analysis, security, and privacy can be solved as mobile devices are powerful enough in terms of CPU and communication bandwidth. The real challenges are at the two ends of the data processing flow: firstly, the sensor data must have *very high* quality to be useful in any kind of application (for sportspeople or patients), false alarms will be annoying, missed alarms might be fatal; secondly, the potential users have to be convinced about the usefulness of such technology. In our interviews, the sportspeople were sceptic about the benefit of such an application. The same applies also to elderly people who might be even more reluctant in wearing any device that monitors them.

Acknowledgements

This work was supported by the DFG Research Cluster of Excellence on Ultra High-Speed Mobile Information and Communication UMIC (`http://www.umic.rwth-aachen.de`) at RWTH Aachen University.

9. REFERENCES

[1] O. Amft and J. Habetha. Smart medical textiles for monitoring patients with heart conditions. In *Smart*

[4] `http://www.wealthy-ist.com`
[5] `http://www.polar.fi/en/products/maximize_performance/running_multisport/RS800CX`
[6] `http://www8.garmin.com/marathon/forerunner/`
[7] `http://www.micoach.samsungmobile.com/`
[8] `http://nikerunning.nike.com/`
[9] `http://www.mas-aal.eu/`

textiles for medicine and healthcare: Materials, systems and applications. Woodhead Publishing Ltd., 2007.

[2] J. Bardram, T. A. K. Kjær, and C. Nielsen. Supporting local mobility in healthcare by application roaming among heterogeneous devices. In L. Chittaro, editor, *Proc. ACM MobileHCI'03*, volume 2795 of *Lecture Notes in Computer Science*, pages 161–176. Springer, 2003.

[3] J. Barnickel, H. Karahan, and U. Meyer. Security and privacy for mobile electronic health monitoring and recording systems. In *IEEE International WoWMoM Workshop on Interdisciplinary Research on E-Health Services and Systems*, pages 1–6, 2010.

[4] R. L. Collins, T. B. Kashdan, and G. Gollnisch. The feasibility of using cellular phones to collect ecological momentary assessment data: Application to alcohol consumption,. *Experimental and Clinical Psychopharmacology*, 11(1):73 – 78, 2003.

[5] A. Ehrmann, E. Flacke, A. Brücken, and A. Büsgen. Intelligent umgarnt. Textile Sensorik. *Physik in unserer Zeit*, 41(6):282–286, 2010.

[6] Y. Hao and R. Foster. Wireless body sensor networks for health-monitoring applications. *Physiological Measurement*, 29:R27, 2008.

[7] M. Hassani and T. Seidl. Towards a mobile health context prediction: Sequential pattern mining in multiple streams. In *Proc. Intl. Workshop on Managing Health Information in Mobile Applications (HIMoA)*, pages 55–57, 2011.

[8] M. Hassani and T. Seidl. Resource-aware distributed clustering of drifting sensor data streams. In *Proc. 4th Intl. Conf. on Networked Digital Technologies (NDT)*, 2012.

[9] S. Intille. Just-in-time technology to encourage incremental, dietary behavior change. In *Proc. AMIA'03*, page 874. Wiley, 2003.

[10] V. Jones, V. Gay, and P. Leijdekkers. Body sensor networks for mobile health monitoring: Experience in europe and australia. In *Proc. Intl. Conf. Digital Society (ICDS)*, pages 204 –209, 2010.

[11] S. Kim, L. Beckmann, M. Pistor, L. Cousin, M. Walter, and S. Leonhardt. A versatile body sensor network for health care applications. In *Proc. Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP)*, pages 175 –180, 2009.

[12] J. Kjeldskov and M. B. Skov. Supporting work activities in healthcare by mobile electronic patient records. In M. Masoodian, S. Jones, and B. Rogers, editors, *Computer Human Interaction*, volume 3101 of *Lecture Notes in Computer Science*, pages 191–200. Springer Berlin / Heidelberg, 2004.

[13] C. Klersy, A. D. Silvestri, G. Gabutti, F. Regoli, and A. Auricchio. A meta-analysis of remote monitoring of heart failure patients. *J. American College of Cardiology*, 54(18), 2009.

[14] P. Kranen, D. Kensche, S. Kim, N. Zimmermann, E. Müller, C. Quix, X. Li, T. Gries, T. Seidl, M. Jarke, and S. Leonhardt. Mobile mining and information management in healthnet scenarios. In *Proc. Conf. on Mobile Data Management (MDM)*, pages 215–216. IEEE, 2008.

[15] P. Kranen, E. Mller, I. Assent, R. Krieger, and T. Seidl. Incremental learning of medical data for multi-step patient health classification. In C. Plant and C. Bhm, editors, *Database Technology for Life Sciences and Medicine*, pages 321–344. World Scientific Publishing, 2010.

[16] R. Litan. Vital signs via broadband: Remote health monitoring transmits savings, enhances lives. Technical report, Better Health Care Together, 2008.

[17] M. Pacelli, G. Loriga, N. Taccini, and R. Paradiso. Sensing fabrics for monitoring physiological and biomechanical variables: E-textile solutions. In *Proc. of 3rd IEEE/EMBS Intl. Summer School on Medical Devices and Biosensors*, pages 1–4. IEEE, 2007.

[18] R. Paradiso and D. De Rossi. Advances in textile technologies for unobtrusive monitoring of vital parameters and movements. In *Proc. EMBC'06*, pages 392–395. IEEE, 2006.

[19] J. Pollard, S. Rohman, and M. Fry. A web-based mobile medical monitoring system. In *Proc. Intl. Workshop on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications*, pages 32 –35, 2001.

[20] J. Riley and M. Cowie. A clinical perspective on remote monitoring of chronic disease. *Eurohealth*, 5(1), 2009.

[21] R. Schmidt, J. Habetha, and M. Harris. Myheart - fighting cardio-vascular diseases by prevention and early diagnosis. In L. Azevedo and A. R. Londral, editors, *HEALTHINF (1)*, pages 296–300, 2008.

[22] C. C. Tsai, G. Lee, F. Raab, G. J. Norman, T. Sohn, W. G. Griswold, and K. Patrick. Usability and feasibility of PmEB: a mobile phone application for monitoring real time caloric balance. *Mob. Netw. Appl.*, 12:173–184, March 2007.

[23] E. Villalba, M. Arredondo, A. Moreno, D. Salvi, and S. Guillen. User interaction design and development of a heart failure management system based on wearable and information technologies. In *Proc. EMBC'06*, pages 400 –403, 2006.

A clinical quality feedback loop supported by mobile point-of-care (POC) data collection

Christopher A. Bain
Director of Health Informatics, Alfred Health
Melbourne, Australia
+61 3 9076 3079
christopher.bain-informatics@alfred.org.au

Tracey Bucknall
Professor of Nursing, Deakin University; Clinical Chair in Nursing, Alfred Health
Melbourne, Australia
+61 3 9076 5391
t.bucknall@alfred.org.au

Janet Weir-Phyland
Executive Director – Nursing, Alfred Health
Melbourne, Australia
+61 3 9076 2039
j.weir-phyland@alfed.org.au

ABSTRACT

A POC clinical information system was designed and deployed on a Motion tablet to allow data collection at the bedside by specialist auditors. Collected data is then stored in a Microsoft SQL Server Environment in real time, allowing immediate feedback through corporate level reporting. The reporting framework consists of self-service non-interactive reports and scorecards, and self-service interactive reports (dashboards) where users can seek detailed information from an organizational level down to a patient level. Information is presented both numerically and graphically, including in a traffic light paradigm, to highlight clinical risks and their relative urgency. Increased information provision to end users via the intranet scorecard and dashboard systems was able to efficiently and effectively close the audit and feedback loop within the organization. Previous variability in information provision was systematically improved to push information to the appropriate decision makers in a timely and efficient manner. Since this health service is one of the first across the country to undergo accreditation under new national standards, the system as described is a novel approach to meeting this compliance challenge. This paper will outline the design features of the system, implementation and training challenges, and proposed future directions for the application and the collected data.

1. INTRODUCTION

The aim of this work is to improve the provision of quality of care information to healthcare executives, managers and clinicians, in order to target organizational strategies to achieve clinical improvements and increase patient safety. This initiative has been undertaken in the context of new nationally driven quality and safety standards for healthcare in Australia [1]. This aim has been achieved through the use of a new mobile data collection platform linked to tailored business intelligence tools. Business intelligence tools are known to have a place in assisting in the management of facilities and systems in range of industries [2,3] and healthcare is no different [4,5].

Recently the Australian Government, through its healthcare quality agency, the Australian Commission for Quality and Safety in Healthcare, released a series of clinically focused standards ("the standards") that health services need to meet in order to be accredited. Failure to meet the standards can have adverse effects on the financial state and reputation of health services.

These 10 standards [1] cover the areas of:

- Governance
- Partnering with consumers
- Infection control
- Medication safety
- Patient identity checking
- Clinical handover
- Blood product usage
- Pressure injury prevention
- Managing deterioration in patients condition and
- Falls prevention.

Typically, in order to supplement the often limited range of data available to measure compliance with standards like these, additional ad-hoc audit work has been performed to attempt to measure the necessary dimensions of care. One of the practical and information management challenges imposed by this approach is the need to randomly audit numerous patients for each different kind of audit - pressure care, falls management, drug interactions and so forth. As a result, such audits often need to be conducted in a rolling fashion through multiple clinical areas so as not to create an implementation burden for frontline staff. However this in turn delays feedback to these staff and reduces the frequency of measurement - effectively creating a series of "photographs" in relation to clinical quality and safety, rather than a constantly running "movie".

2. RELATED WORK

As mentioned above, a variety of clinical audits have been a long standing part of clinical practice improvement in healthcare [6,7,8,9]. Historically, many clinical audit [10,11] processes-designed to collect information about an organizations or individuals compliance with clinical quality standards (eg - care measures to prevent pressure ulcers)- have been centered around primary data collection on paper based structured or unstructured assessment tools. An example is the Surgical Tool for Auditing records (STAR), used for auditing clinical records in the UK [12]. Historically, data has then typically been secondarily entered into electronic systems for subsequent analysis, reporting and benchmarking. Obviously such processes are error prone and far from ideal from an information management perspective. They also introduce a potential delay between primary data collection and subsequent feedback to key staff that can influence practice and improve outcomes for patients.

There has certainly been much work done around the world on how to better support the measurement of clinical care - be that the measurement of the outcomes of care or the processes of care. Examples include the work of Wong and Giallonardo [13], Roodpeyma et al [14] and that of Nseir et al [15]. In addition there has been work done on describing systems for, or approaches to, storing healthcare data for subsequent analysis and reuse [16]. This could be at a hospital level [17], or at a broader community level [18].

In addition, as previously mentioned, there had been work done in the business intelligence space in healthcare [19,20] with further work acknowledging the value added by investments in such systems. Examples are the paper by Rufer [21] on the value added by such systems in the context of radiology services, as well as the work by Moore et al [22]. Karami et al [23] have also published in the area of business intelligence in support of radiology services

In the space of standards for healthcare, there has been a lot of work published. Ryan et al [24] examined the role of standards in credentialing specialist documentation practitioners in healthcare, whilst Drew and Funk [25] published around practice standards for electrocardiograph (ECG) monitoring in hospital settings. There are also numerous publications around treatment standards of different kinds – for example the work by Davies et al [26] providing guidelines for treatment of patients with spleen related problems.

Finally, the other relevant context here is the use of mobile devices in healthcare. An extensive literature review published in 2013 showed that there are a wide number of contexts in which mobile devices and mobile apps are being used in healthcare [27]. In their investigation of many thousands of apps, the authors discovered the usage of commercial apps in a range of disease conditions including: (in descending order of number of apps) diabetes, depression, migraine, asthma, low vision, hearing loss, OA (osteoarthritis), and anemia. This is despite some of the concerns over the security and privacy implications of mobile devices in the healthcare context [28].

In this work, the worlds of healthcare standards and care measurement, mobile device usage and business intelligence in healthcare are considered together through the lens of a specific case study In this case study, the nursing division of a large tertiary - quaternary Australian health service conceived of a modular, electronic audit tool to allow data collection at the bedside. Data collection at the bedside using a range of devices is a topic of interest amongst nursing staff given the nature of their work [29]. In this case, the senior nursing staff had a vision of a system that would dramatically improve the efficiency and validity of the audit process, as well as of resultant intelligence about care that could improve outcomes for patients in a range of dimensions. Technical staff at the health service shared the vision and played a key role in bringing the system to life. The commencement of the standards program acted as a further catalyst for the work. In this paper we will describe this work, and also elements of the storage and reporting systems that allow the collected data to be used to improve clinical and managerial practice.

3. APPROACH

A POC clinical information system was designed and deployed on a Motion tablet (Figure 1) to allow data collection at the bedside by specialist auditors. Collected data was then warehoused in a Microsoft SQL Server Environment in real time, allowing immediate feedback through corporate level reporting. The reporting framework consists of self-service non-interactive reports and scorecards (Microsoft Reporting Services), and self-service interactive reports (Qlikview) where users can seek detailed information from an organizational level down to a patient level. Information is presented both numerically and graphically, including in a traffic light paradigm, to highlight clinical risks and areas for action.

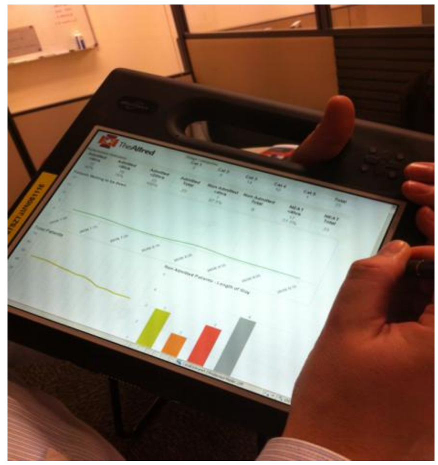

Figure 1 – A Motion Tablet Device

The application was developed using an iterative prototyping approach with a phased implementation, commencing with a pilot amongst the core nursing staff involved in the development process. The modules contained within the application are as follows:

- Module 1 – Patient Demographics and Identity Checking
- Module 2 - Bed Area
- Module 3 - Infection Prevention
- Module 4 - Medication Safety
- Module 5 - Patient Identification
- Module 6 - Clinical Handover
- Module 7 - Blood & Blood Products
- Module 8 - Pressure Injury Prevention
- Module 9 - Deteriorating Patients
- Module 10 - Falls Prevention
- Module 11 – Nutrition

As can be seen from this list, there is an alignment, but not an absolute correlation, between the modules and the list of government standards. This is in part because the desired scope of the POC application was broader than just those topics mandated by the standards.

The context in which the POC application was developed was characterized by several constraints. Foremost of these was a limited time frame- approximately 6 months from the need being raised to an outcome being expected. In addition, there was little if any, additional funding available to support the work. However, the importance of establishing an electronic tool to support clinical audit data collection and reporting for the purposes of demonstrating compliance with new national healthcare standards cannot be overstated

3.1 System Features

The POC system (Figure 2) is a Microsoft Excel-VBA application and hence can be run on any platform capable of running Excel with macros enabled. One of the drivers for the use of this platform was the fact that the available programming resource had already demonstrated proof of concept in this space with another, smaller audit application, and the responsible managers were impressed with the results. In addition, it could be seen that development on this platform would enable completion of the project under the immense time pressures surrounding the project, and within existing resource allocation. At project initiation it was always accepted that the application may need to be ported to a more stable long term platform when time allowed.

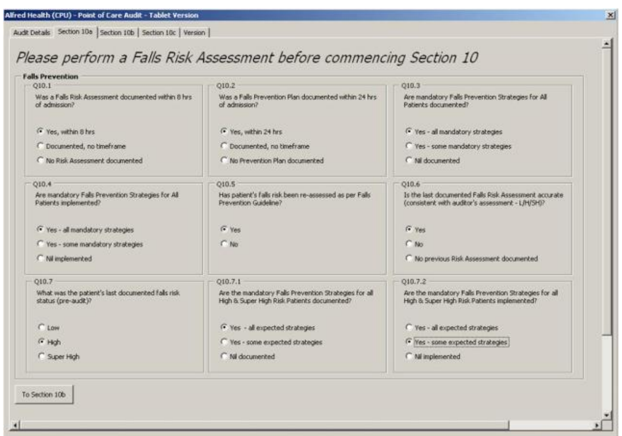

Figure 2 – The POC application – Module 10- Falls Prevention

One of the reasons for this is the limitation of the approach to security implicit in this product. Security is not achieved at an application level, but rather at a folder access level. So the Motion devices are only accessible to the small range of trained and approved users (the auditors) and the application file is in a folder that can only be accessed by those trained users. Whilst far from ideal, this was felt to be adequate for the intended use of the application in the short to medium term, particularly given that the application does not allow viewing of, or writing to, the official medical record of the patient.

3.2 Graphical User Interface (GUI)

From a graphical user interface (GUI) perspective, the POC application was specifically designed to be deployed on a Motion tablet, but can be run on a PC, laptop, computer on wheels (COW) or other devices, thru minor GUI adaptations that do not affect the core flow and function of the application.

Staff undertaking clinical audits may be time poor, and at times overawed by the amount of data to be collected in the audit process. With these issues in mind, the GUI has several specific design features to enhance the usability of the application. These design features include a modular design - at this stage users can select a "full audit" - all 10 modules; or a single audit module from a pick list. Another relevant GUI design feature in this regard is the progressive exposure of questions to users as they proceed down the relevant logic chain. So, for example, within a given module, if the larger number of questions a user may have to answer is 20, they will not be presented with these all at once. Each question will only be displayed on the screen as the last one is answered. This reduces the cognitive load on the user and allows them to focus on accurately answering the question at hand, rather than mentally skipping ahead to answer the next question or the one after that.

As is the case with most good surveys, a consistent feature across all POC modules is a highly structured question format with codification of answers, and with minimal need for, or opportunity to enter, free text answers. Data entry is supported by stylus with a rapid system response time.

3.3 Information Management Issues

Patient identity validation is undertaken automatically by the application, by checking the entered patient identification number against the organizational data warehouse. The patient's gender and date of birth are then retrieved so the user can ensure they are auditing the correct patient. As well as ensuring that the collected data will be linkable with other data about the patient, this reduces the risk of incorrect patient details being collected as can be the case with traditional primary data collection on paper.

Very significantly, the POC application was designed from the outset to be connected to what is known as the organizational "information grid". This grid- analogous to the electricity grid, where data is equivalent to electricity- is designed to allow the reliable and predictable flow of data and information to users irrespective of how or where it is generated. The aforementioned data warehouse is one of the key components that has been constructed inside "the grid".

The services available to the business from "the grid" include data extraction and distribution, non-interactive (static) reporting through Microsoft Reporting Services, and interactive (dynamic) reporting through Qlikview. It is these latter 2 services that the POC application leverages off to provide feedback to stakeholders regarding collected data.

3.4 Mobile Deployment Issues

There is a growing awareness of the utility of mobile devices and applications deployed on them across all of healthcare [30]. There are already a range of clinical areas in which mobile device usage has established a stronghold such as in Nutrition [31], Radiology [32] and Emergency care [33].

It is also important to note however, that there are a number of potential issues in relation to deploying and using healthcare applications on mobile devices, be they smart phones, handheld tablets or larger and heavier tablets like the Motion tablet. These include the need for a medical grade tablet device choice at the bedside versus more popular choices such as the iPad. In a bedside setting it is critical that end user devices are physically robust but that they can also withstand trauma and are able to be wiped down to prevent the spread of infection. Some pieces of computing equipment are certainly known to harbor infectious organisms [34]. The Motion tablet was the device of choice in our facilities based on criteria such as these. In addition, this device was also suitable as a platform on which to deploy other core applications including the Cerner Millennium clinical system which the health service uses.

4. RESULTS

The POC system has had a strong history of use since its release in April 2013.

4.1 Usage Statistics

Since the commencement of the POC audit program, in excess of 635 unique patients have been randomly audited in a 4 month period across 3 campuses of an academic health service. This number will grow substantially over time. This has been across 21, 12 and 4 wards respectively at the 3 main campuses of the health service-these being the main acute campus, the aged and sub-acute care campus, and the community based hospital campus.

4.2 Reporting

The interactive report (Figure 3) to which the data contributes has had 675 views in 4 months – equating to about 42 views per week.

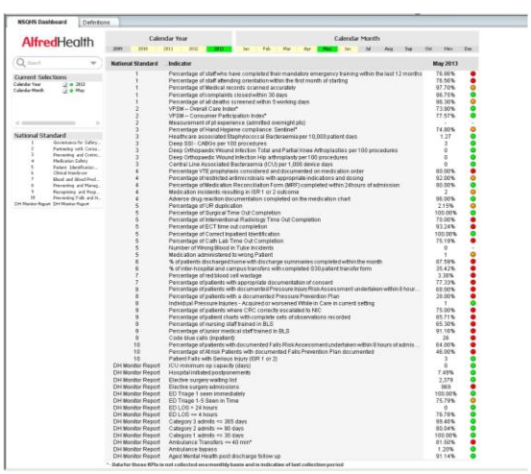

Figure 3 – NSQHS Interactive report (or Dashboard) – DevelopmentVersion (Qlikview)

The longest standing non–interactive report (Figure 4) to which the data contributes has had 1382 views in 4 months – equating to about 80 views per week. Below is an example of the non-interactive report showing the summary traffic lights across the entire organization for all measures captured in the POC application.

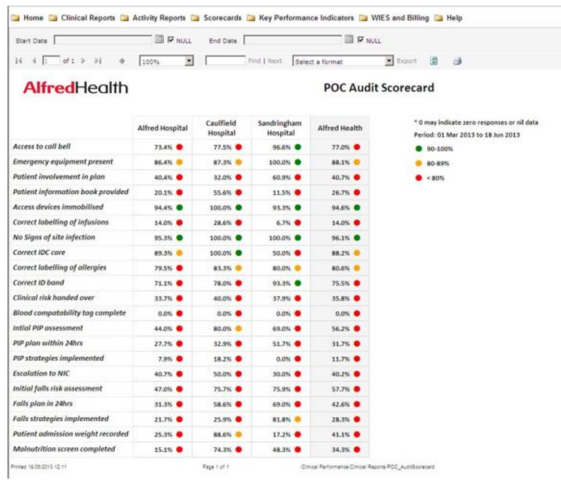

Home | Clinical Reports | Activity Reports | Scorecards | Key Performance Indicators | WIES and Billing | Help

AlfredHealth — POC Audit Scorecard

	Alfred Hospital	Caulfield Hospital	Sandringham Hospital	Alfred Health
Access to call bell	73.4%	77.5%	96.6%	77.0%
Emergency equipment present	86.4%	87.3%	100.0%	88.1%
Patient involvement in plan	40.4%	32.0%	60.9%	40.7%
Patient information book provided	20.1%	55.6%	11.5%	26.7%
Access devices immobilised	94.4%	100.0%	93.3%	94.6%
Correct labelling of infusions	14.0%	28.6%	6.7%	14.0%
No Signs of site infection	95.3%	100.0%	100.0%	96.1%
Correct IDC care	89.3%	100.0%	50.0%	88.2%
Correct labelling of allergies	79.5%	83.3%	80.0%	80.6%
Correct ID band	71.1%	78.0%	93.3%	75.5%
Clinical risk handed over	33.7%	40.0%	37.9%	35.8%
Blood compatability tag complete	0.0%	0.0%	0.0%	0.0%
Intial PIP assessment	44.0%	80.0%	69.0%	56.2%
PIP plan within 24hrs	27.7%	32.9%	51.7%	31.7%
PIP strategies implemented	7.9%	18.2%	0.0%	11.7%
Escalation to NIC	40.7%	50.0%	30.0%	40.2%
Initial falls risk assessment	47.0%	75.7%	75.9%	57.7%
Falls plan in 24hrs	31.3%	58.6%	69.0%	42.6%
Falls strategies implemented	21.7%	25.9%	81.8%	28.3%
Patient admission weight recorded	25.3%	88.6%	17.2%	41.1%
Malnutrition screen completed	15.1%	74.3%	48.3%	34.3%

* 0 may indicate zero responses or nil data
Period: 01 Mar 2013 to 18 Jun 2013
90-100%
80-89%
< 80%

Figure 4 – The POC Scorecard – Development Version (MS Reporting Services report)

Both kinds of reports can be viewed on a variety of devices, as they are intranet based, including on the Motion tablets themselves. In the case of the non- interactive reports, these are updated as soon as the data is submitted from the POC application.

4.3 Stakeholder Acceptance

Stakeholder acceptance of the solution - the POC application, the Motion tablets and the reporting options - has on the whole been excellent. One of the reasons for this, on the management front at least, is that the solution replaces a mixture of manually executed processes, with a large administrative burden, that resulted in incomplete and insufficient detailed data stored in disparate Excel spreadsheets.

Instead, the management staff responsible for running the clinical audit program, and those charged with driving towards improved outcomes off the back of it, now have more complete and robust data that can be fed back to relevant staff in reports almost instantaneously. In addition that data is stored in a centrally supported, robust technical environment and can be kept for as long as needed and re-used, along with other relevant data from "the grid" for a range of research and evaluation activities in addition to its core and immediate purpose.

Some of the direct feedback obtained from stakeholders has been very encouraging. For example one of the nurse managers stated “It is great that we can see the data instantaneously and can use it to inform our practice”. The auditors, who were nurse educators, made positive statements such as “(the) ward appreciate receiving the feedback as we complete the audit and can address issues at the time at a local level". Another auditor saw the solution and its use as "A great educational opportunity. We are getting to know what the wards learning needs are by having the ability to look at the data as it is entered".

In relation to the Motion tablets themselves, some comments included that they were "user friendly" and "easy to use". Others felt that it was "great not to have to double handle data!" when compared to the old method of primary data recording on paper then transcribing into an electronic source. Certainly the broader literature has many examples of tablet devices being well received amongst users. For example, work from the US has examined the positive benefits of distracting children with iPads as an alternative to sedation or anesthesia for medical procedures [35] and another study from Singapore [36] describes positive experiences, on the whole, of radiologists viewing images on iPads versus on traditional workstations.

5. DISCUSSION

Increased information provision to end users via the intranet scorecard and dashboard systems was able to efficiently and effectively close the audit and feedback loop within the organization. Previous variability in information provision was systematically improved to push information to the appropriate decision makers in a timely and efficient manner. Since this health service is one of the first across the country to undergo accreditation under the new national standards, the system as described is a novel approach to meeting this compliance challenge.

There were a number of specific lessons learned from the development and implementation activities in this project.

5.1 Requirements Elicitation

One of the key lessons in relation to the project and its relative success was the nature of the requirements elicitation process. Even though the software development approach could be best described as iterative prototyping, the many dozens of questions across all audit modules, and their complex conditional logic, were captured and kept up to date in a semi- formal requirements document. This document became critical in maintaining good communication and clarity between the development team and nursing subject matter experts (SME's).

The other critical aspect of this part of the project was the very direct and regular access that the developer had to the SMEs both on terms of explanations about requirements and feedback on released development versions of the software.

These 2 things combined to ensure a stable, well tested application at release that needed little post release work.

5.2 Software Bugs

Despite the measures outlined above, released software can still have bugs, even when developed with an abundance of resources applied to the process. The key software issue discovered in the case of the POC application was the identity checking function. In short the function appeared to work well when checking against test databases and in small scale production use. However when greater production use began there were cases where known inpatients could not have their details checked and hence could not have their care audited as the application was designed with this check as a "gate keeping" step.

An investigation identified the problem which was remedied by having the check look at more than one underlying database for auditable patients. Once this was done the bug was resolved. Ironically this was not due to a fault with the application itself but with an inherited architectural decision in the underlying databases.

5.3 Wireless Coverage

Another unforeseen issue was that of wireless coverage in ward areas. This was a problem especially in the main campus of the health service. In short, multiple staff using the application in multiple locations highlighted the fact that wireless coverage was somewhat patchy. Users soon found ways around this, but the problem expedited an investigation by the IT department of wireless coverage across the main building.

6. FUTURE WORK

6.1 The POC Application

Some of the issues for consideration in relation to the future of the POC application and its use include porting the application to a more robust software platform, and potentially expanding the range of devices on which the application is deployed (eg - Windows based phones or smaller, handheld tablets), and increasing the range of usage of the collected data. It is also highly likely that the business will drive changes in the core application - for instance an 11th module relating to audit of timely patient access to care is already under development.

The health service has an IT architecture team and that team has begun investigating a more robust software platform on which the same data collection functionality could be delivered. It is highly likely, that in line with the standard approach to patient centric data collection used in the institution, the Cerner Millennium Clinical system, will serve as that platform. A key challenge will be the extent to which the Cerner platform can support some of the GUI features of the current POC application, For example, the progressive "unhiding" of questions as users navigate through the application. There is no question however, that deployment on this platform would allow great levels of security around the use of the application and, critically, direct integration with the remainder of an individual's clinical record.

6.2 Mobile Device Usage

There are also opportunities off the back of this work to further examine the use of mobile devices across our institution, and how that usage compares to best practice from around the world. This is of interest from both a heath and IT research perspective, and an operational perspective. Such opportunities are also important in shaping future "bring your own device" (BYOD) strategies for the organization, particularly given the part time (senior staff) and or transient (junior staff) nature of much of its medical workforce.

6.3 Information Management and Usage

What will be critical, in terms of assessing the long term impact of the POC application, is to see the effect of the tool and the feedback loop to clinicians and managers, in terms of reducing bad practice and maximizing good practice. So for example in performing falls risk assessments, or in reducing harder end points like hospital acquired pressure ulcers. An example of such an evaluation is the previous quoted work by Tuffaha et al [12]. Let us consider a specific example of how the collected data has been, and will continue to be used to improve patient care. In Figure 5 below, the percentage of audited patients having had a falls risk assessment upon admission is tracked from pre-POC implementation (February) thorough to the time of writing (July). In that time that percentage has risen from 42% to 76%.

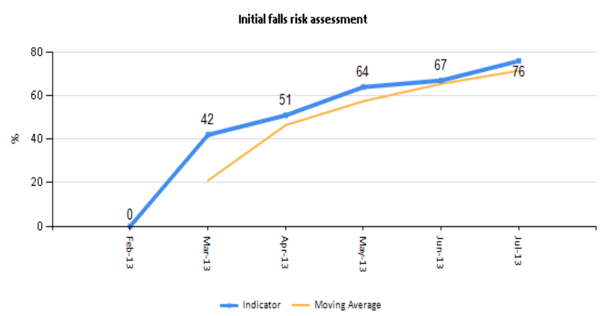

Figure 5 – Percentage rate of Initial Falls Risk Assessment completion

Now lets us consider Figure 6 below. In this graph we can see that the percentage rate of audited patients having had a falls plan documented within 24 hours of admission has risen from 42% to 63% over the same time period. These trends may partially be explained by an increasing sample base over time, but also will to some extent reflect efforts, based on the early results of the audit, to raise the rate of both of these key markers of good nursing care. It is critical to note that such trend data, which can also be examined by sub-group (ward, campus) was previously not readily available, and certainly not in a timely fashion with automated reporting. Hence clinicians, managers and executives have been "flying blind" to some extent in relation to how well their clinical workforce has been performing in regards to these 2 key pieces of documentation. Given the known problems relating to falls in hospitals all over the world [37,38,39,40], this is a non-trivial issue.

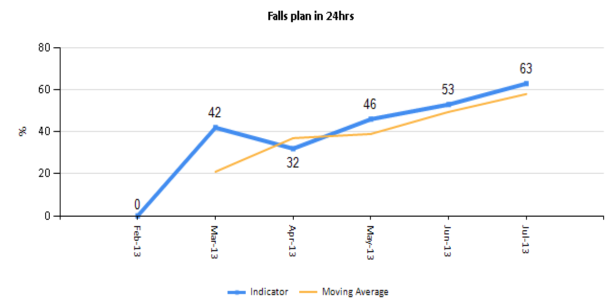

Figure 6 – Percentage rate of completion of Falls Plan within 24 hrs of admission

There are also great opportunities to leverage off, and cross reference the data collected by the POC application with, data in "the grid". The best example, although there are several, is in the area of pressure ulcers and pressure ulcer prevention. Specifically, data about the completeness of implementation of strategies regarding pressure ulcer prevention, and data about the incidence of pressure ulcers – collected from the POC application; will be able to be compared with data from International Classification of Disease (ICD) codes of admitted patients, and data on adverse events in patients (including pressure ulcers), each from their own separate sources. Such comparisons will not only enable a complete picture of pressure care at the hospital to be built, but also will have the effect of driving up overall data quality in this area of measurement right across our health service.

Furthermore, use of grid data will enable interesting correlations between, for example, shift by shift staffing, or hospital occupancy; and the impact on the processes and outcomes of care as they relate to individuals on various hospital wards. Work like this would be of a particularly novel nature.

7. CONCLUSIONS

This paper outlines how a novel POC software application deployed on a mobile device has been able to successfully supplant much paper based auditing practice resulting in strong uptake and positive feedback from users and other key stakeholders. The coupling of this application – albeit that it is likely to be a temporary deployment vehicle - with a robust corporate approach to information management, has paved the way for greater amounts of data collection, and greater data quality in this space, and in turn improved outcomes for patients. Another essential benefit has been an increase in the ability of the health service to demonstrate its compliance with new national healthcare standards.

8. ACKNOWLEDGEMENT

The 3 main authors would like to acknowledge the efforts of our other key contributors – Suzanne Metcalf and Pam Ingram from Nursing Management, and Lucy Nie from the Health Informatics department, for their key roles in bringing the system to life.

In addition, all the authors would like to acknowledge the work of Dr. Roger Hawkins – Manager, Analytics in the Clinical Performance Unit of the health service for his original work on the POC software application. Finally, we would like to acknowledge Mr Zaf Alam and Mr Hien Le from the Health Informatics and IT departments respectively for their work in establishing the reporting around the POC data.

9. REFERENCES

[1] "NSQHS Website". http://www.safetyandquality.gov.au/our-work/accreditation/ Accessed 18/6/2013.

[2] F. Mohammadi and N. Hajiheydari. "How Business Intelligence Capabilties Contributed Managerial Decision Making Styles". *International Journal of e-Education, e-Business, e-Management and e-Learning.* 2012 Vol 2, No 1; February pp 28-33.

[3] J.K. Ho. "Strategic Innovation through Business Intelligence: Linking Competitive Forces to Profitabilty" *International Journal of e-Education, e-Business, e-Management and e-Learning.* 2012 Vol 2, No 4; August pp 280-285.

[4] I. Cattinelli, E. Bolzoni, M. Chermisi, F. Bellocchio, C. Barbieri, F. Mari, C. Amato, M. Menzer and E. Gatti. "Computational Intelligence for the Balanced Scorecard; Studying performance trends of Haemodialysis Clinics" *Artif Intell Med.* 2013 Jun 11. pii: S0933-3657(13)

[5] J. Van De Graaff and A.Cameron."Quest for Business Intelligence in Healthcare". *Healthc Financ Manage.* 2013 Feb;67(2):44-6, 48.

[6] R.A. Ralph and M. Delaney. "Dual purpose cataract audit: improvement of care, improvement of documentation." *QRB Qual Rev Bull.* 1979 May;5(5):29-31.

[7] J.D. Vail. "Retrospective anesthesia nursing audit-process or outcome criteria?" *Mil Med*. 1979 Aug;144(8):537-9.

[8] D.K. Dhariwal and A.J. Gibbons. "The CRABEL Score – Setting standards in Maxillofacial Medical Note Keeping". *Br J Oral Maxillofac Surg.* 2004 Jun;42(3):200-2.

[9] M. Chopra, S. Stirling, D. Wilkinson, C. Connolly and D.McCoy. "Paediatric admissions to a rural South African hospital: Value of hospital data in helping to define intervention priorities and allocate district resources". *S Afr Med J.* 1998;88(6 Suppl):785-8.

[10] S. Gilam and A.N. Siriwardena. "Frameworks for Improvement: Clinical Audit, the Plan-Do-Study-Act Cycle and Significant Event Audit". *Qual Prim Care*. 2013 May;21(2):123-30.

[11] J. Barlow and G. Krassas. "Improving management of Type 2 Diabetes-Findings of the Type2 Care Clinical Audit". *Aust Fam Physician.* 2013 Jan-Feb;42(1-2):57-60.

[12] H. Tuffaha, T. Amer, P. Jayia, C. Bicknell, N. Rajaretnam and P. Ziprin."The STAR score: a method for auditing clinical records". *Ann R Coll Surg Engl.* 2012 May;94(4):235-9.

[13] C. Wong and L. Giallonardo. "Authentic leadership and nurse-assessed adverse patient outcomes".*J Nurs Manag*. 2013 Jul;21(5):740-52. doi: 10.1111/jonm.12075.

[14] S. Roodpeyma, M. Hekmat, M. Dordkhar, S. Rafieyian and A. Hashemi. "A prospective observational study of paediatric cardiac surgery outcomes in a postoperative intensive care unit in Iran". *J Pak Med Assoc*. 2013 Jan;63(1):55-9.

[15] W. Nseir, S. Haj, B. Beshara, J. Mograbi and O. Cohen. "Seeking out high risk population: the prevalence characteristics and outcome of diabetic patients of arab ethnicity hospitalized in internal medical and acute coronary units in Israel". *Int J Endocrinol.* 2013;2013:371608. doi: 10.1155/2013/371608. Epub 2013 Jun 18.

[16] P. Embi, S. Kaufman and P. Payne. "Enabling Knowledge-Driven Health Care". *Circulation.* 2009; 120: 2393-2399

[17] A. Post, T. Kurc, S. Cholleti, J. Gao, X. Lin, W. Bornstein, D. Cantrell, D. Levine, S. Hohmann and J. Saltz. "The Analytic Information Warehouse (AIW): a platform for analytics using electronic health record data". *J Biomed Inform*. 2013 Jun;46(3):410-24. doi: 10.1016/j.jbi.2013.01.005. Epub 2013 Feb 9.

[18] D. Berndt, A. Hevner and J.Studnicki, "The Catch data warehouse: support for community health care decision-making". *Decision Support Systems*, Volume 35, Issue 3, June 2003, Pages 367-384, ISSN 0167-9236.

[19] P. Bartsch, T. Lux, A. Wagner and R. Gabriel. "Business Intelligence and Information Systems in Hospitals - Distribution and Usage of BI and HIS in German Hospitals". *Stud Health Technol Inform*. 2013;190:191-3.

[20] J. Van De Graaff and A. Cameron. "Quest for business intelligence in health care". *Healthc Financ Manage*. 2013 Feb;67(2):44-6, 48.

[21] D. Rufer. "Advancing RIS features while addressing MU. New business intelligence tools enable improved care and profitability". *Health Manag Technol*. 2013 May;34(5):20. No abstract available.

[22] K. Moore, K. Eyestone and D. Coddington. "How business intelligence can improve value". *Healthc Financ Manage*. 2012 Oct;66(10):112-4.

[23] M. Karami, R. Safdari and A. Rahimi. "Effective radiology dashboards: key research findings". *Radiol Manage.* 2013 Mar-Apr;35(2):42-5.

[24] J. Ryan, K. Patena, W. Judd and M. Niederpruem. "Validating competence: a new credential for clinical documentation improvement practitioners". *Perspect Health Inf Manag.* 2013 Apr 1;10:1g. Print 2013.

[25] B. Drew and M. Funk." Practice standards for ECG monitoring in hospital settings: executive summary and guide for implementation". *Crit Care Nurs Clin North Am.* 2006 Jun;18(2):157-68, ix.

[26] J. Davies, R. Barnes, D. Milligan and British Committee for Standards in Haematology. Working Party of the Haematology/Oncology Task Force. "Update of guidelines for the

prevention and treatment of infection in patients with an absent or dysfunctional spleen". *Clin Med.* 2002 Sep-Oct;2(5):440-3.

[27] B. Martínez-Pérez, I. de la Torre-Díez and M. López-Coronado. "Mobile health applications of the most prevalent conditions by the world health organization: review and analysis". *J Med Internet Res.* 2013 Jun 14;15(6):e120.

[28] B.M. Silva, J.J Rodrigues, F. Canelo, I.C. Lopes and L. Zhou. "A data encryption solution for mobile health apps in cooperation environments". *J Med Internet Res.* 2013 Apr 25;15(4):e66.

[29] E. Carlson, C. Catrambone, K. Oder, S. Nauseda, L. Fogg, B.Garcia, F.M. Brown Jr, M.E. Johnson, T.J. Johnson and J. Llewellyn. "Point-of-Care technology supports bedside documentation". *J Nurs Adm.* 2010 Sep;40(9):360-5.

[30] R. Watson. "Report highlights potential benefits of mobile communication devices for health". *BMJ.* 2013 Jun 6;346:f3623.

[31] S. Holubar and L Harvey-Banchik. "A review of the use of handheld computers in medical nutrition". *Nutr Clin Pract.* 2007. Aug;22(4):428-35.

[32] A.E. Flanders, R.H. Wiggins 3rd and M.E. Gozum. "Handheld computers in Radiology". *Radiographics.* 2003 Jul-Aug;23(4):1035-47.

[33] A.N. Khan, J.Frank, R.Geria and S.Davidson.. "Utilization of Personal Digital Assistants (PDAs) by Paediatric and Emergency Medicine Residents". *J Emerg Med.* 2007 May;32(4):423-8. Epub 2007 Apr 16.

[34] W.A. Rutala, M.S. White, M.F. Gergen and D.J. Weber. "Bacterial contamination of Keyboards: Efficacy and Functional Impact of Disinfectants". *Infect Control Hosp Epidemiol.* 2006 Apr;27(4):372-7. Epub 2006, Mar 29.

[35] A. McQueen, C. Cress and A. Tothy. "Using a tablet computer during pediatric procedures: a case series and review of the 'apps'". *Pediatr Emerg Care.* 2012 Jul;28(7):712-4.

[36] S. John, A.C. Poh, T.C. Lim, E.H. Chan and R. Chong. "The iPad tablet computer for mobile on-call radiology diagnosis? Auditing discrepancy in CT and MRI reporting". *J Digit Imaging.* 2012 Oct;25(5):628-34.

[37] N. Warshawsky, M. Rayens, K. Stefaniak and R. Rahman. "The effect of nurse manager turnover on patient fall and pressure ulcer rates". *J Nurs Manag.* 2013 Jul;21(5):725-32. doi: 10.1111/jonm.12101.

[38] K. Swartzell, J. Fulton and B. Friesth. "Relationship between occurrence of falls and fall-risk scores in an acute care setting using the Hendrich II fall risk model". *Medsurg Nurs.* 2013 May-Jun;22(3):180-7.

[39] R. Schwendimann, H. Bühler, S. De Geest and K. Milisen. "Characteristics of hospital inpatient falls across clinical departments". *Gerontology.* 2008;54(6):342-8. doi: 10.1159/000129954. Epub 2008 May 6.

[40] M. Caldevilla, M. Costa, P. Teles and P. Ferreira. "Evaluation and cross-cultural adaptation of the Hendrich II Fall Risk Model to Portuguese". *Scand J Caring Sci*. 2013 Jun;27(2):468-74. doi: 10.1111/j.1471-6712.2012.01031.x. Epub 2012 Jun 27.

Mobile objects and sensors within a video surveillance system: Spatio-temporal model and queries

Dana Codreanu, Ana-Maria Manzat, Florence Sedes
Université de Toulouse – IRIT – UMR 5505
118 Route de Narbonne, 31062 Toulouse Cedex 9, France
{codreanu, manzat, sedes}@irit.fr

ABSTRACT

The videos recorded by video surveillance systems represent a key element in a police inquiry. Based on a spatio-temporal query specified by a victim, (e.g., the trajectory of the victim before and after the aggression) the human operators select the cameras that could contain relevant information and analyse the corresponding video contents. This task becomes cumbersome because of the huge volume of video contents and the cameras' mobility. This paper presents an approach, which assists the operator in his task and reduces the research space. We propose to model the cameras' network (fixed and mobile cameras) on top of the city's transportation network. We consider the video surveillance system as a multilayer geographic information system, where the cameras are situated into a distinct layer, which is added on top of the other layers (e.g., roads, transport) and is related to them by the location. The model is implemented in a spatio-temporal database. Our final goal is that based on a spatio-temporal query to automatically extract the list of cameras (fixed and mobile) concerned by the query. We propose to include this automatically computed relative position of the cameras as an extension of the standard ISO 22311.

1. INTRODUCTION

The number of video surveillance cameras increases in public and private areas (e.g., in train and metro stations, on-board of buses and trains, inside commercial areas, inside enterprises buildings). For example, some estimations show that there are more than 400000 cameras in London and that only the RATP also known as Régie Autonome des Transports Parisiens (English: Autonomous Operator of Parisian Transports) surveillance system comprises around 9000 cameras in Paris. In these conditions, any person that lives and walks in those two big European capitals is likely to be captured many times during a day (up to 300 times in London) by several video surveillance systems (e.g., the traffic surveillance cameras, the cameras in the subway, and the cameras of a commercial centre). The only markers available for all these videos are the id of the camera (eventually GPS coordinates) and a local date/timestamp that are not homogenous throughout the different systems.

A great majority of the existing video surveillance systems are manual or semi-automatic (they employ some form of video processing but with significant human intervention) [11]. Taking into account the huge amount of video contents that need to be handled, the purely manual approach (agents watching the videos and detecting events) becomes insufficient. The main objective in the video surveillance domain is to provide users with tools that could assist them in their research by reducing the research space and therefore the response time. These tools depend on the research context and complexity (e.g., real time surveillance of big events, police inquiry) [22].

Our work is situated in the context of the police inquiry which involves an a posteriori processing of the data in order to help the investigator to highlight (isolate) the relevant elements (e.g., persons, events). To do that, the investigators dispose of the set of recorded videos from different video surveillance systems (e.g., public, private, RATP). In order to assist the investigators in their tasks, it is important that the different outputs of the systems are interoperable, which is not currently the case. The interoperability between any video surveillance systems from the simple ones with only few cameras to the large scale systems is the main goal of the standard ISO 22311[1]. It specifies a format for the data which can be exchanged between the video surveillance systems in the inquiry context.

This standard does not consider the video surveillance cameras' mobility or their fields' of view modification. In fact, at the beginnings of video surveillance systems the cameras were placed in fixed locations in order to monitor indoor and outdoor places. With the improvements in the hardware and software technologies, on-board cameras are more and more employed in mobile vehicles (e.g., buses, police cars). This cameras' mobility makes the task of security agents even more difficult in the context of an inquiry, when they have to analyse a huge amount of video contents and to have supplementary knowledge on the system's characteristics (e.g., the bus timetables, the city transport plan) in order to select the most appropriate video contents.

In this context, our goal is to provide users with tools that could assist them in their research and reduce the research space. In order to achieve this objective, in this article, we propose an extension of the ISO 22311 standard in order to take into account

[1] http://www.iso.org/iso/fr/catalogue_detail.htm?csnumber=5346

the cameras' mobility. We consider the video surveillance system as a multilayer geographic information system, where the cameras are situated on a distinct layer, which is added on top of the other layers (e.g., roads, transport) through the location. We implemented our solution using a spatial database in order to select the cameras that might have acquired video contents corresponding to a user's spatio-temporal query.

The remainder of this paper is organized as follows. After a review of related work concerning the three aspects addressed in this paper, video surveillance systems, standard ISO 22311 and mobile objects modelling in the Section 2, Section 3 presents our multilayer modelling approach. This model is implemented using a spatio-temporal database. Some queries that can be answered based on this database are presented in Section 4. Finally, Section 5 concludes and discusses possible future research.

2. STATE OF THE ART

2.1 Video Surveillance Systems

The generic schema of a video surveillance system is illustrated in Figure 1. The content is captured and stored in a distributed manner and analysed in a control centre by human operators that watch a certain number of screens displayed in a matrix (the Video Wall in Figure 1).

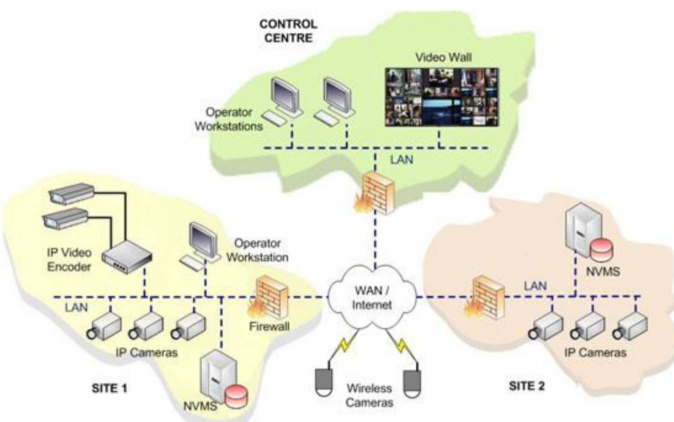

Figure 1: Video surveillance system's schema

There is a big diversity of cameras and sensors that constitute the acquisition part of surveillance systems and a heterogeneity of their installation contexts (e.g., on the halls or platforms of railway or metro stations, on-board of trains and buses, on the streets, in commercial centres or office buildings). Therefore, we have fixed and mobile cameras having different technical characteristics (most of the time dynamic) (see Figure 2 for an example of such cameras) [14]:

- Camera type: optical, thermal, infrared
- Sensor type and dimension: CMOS, CCD
- Transmission type: analogous/ IP
- Angle of view (horizontal and vertical), focal distance, pan-tilt-zoom, field of view orientation, visible distance etc.

Figure 2: Examples of video surveillance cameras having the same position but different fields of view

We started by analysing the way a query is processed in a video surveillance system today. When a person (victim of an aggression for example) files a complaint, he is asked to fill a form describing the elements that could help the investigators to find the relevant video segment (the Figure 3 illustrates an example of such form). Based on the spatial and the temporal aspects of the query, the surveillance operator uses his own knowledge concerning the spatial disposal of the cameras' network in order to select the most relevant video contents. Then he analyses these contents by playing them on the different screens that he has in front of him. The monitors themselves show no spatial relationship of any kind, only the numbering of the cameras is in a somewhat logical order.

Location: Paris
Date and time: 14 July 2013 between 10h and 12h
Event: stolen bag
Victims trajectory: Rivoli Street, from Louvre to
Metro Chatelet: entrance, hallway, elevator, train
Clothes color, distinguishing signs

Figure 3: Example of a form filled by a victim

Therefore, the operators' tasks become cumbersome taking into consideration the huge volume of video contents to be analysed, the mobility and the different characteristics of cameras. Moreover, in the current systems, most of the stored contents is not exploitable because of the recording's low quality. This lack of quality is often caused by inappropriate installation of cameras, bad shooting, bad illumination conditions etc. The operator has no a priori knowledge on the quality of the video contents and thus he loses time by visualizing the low quality contents also.

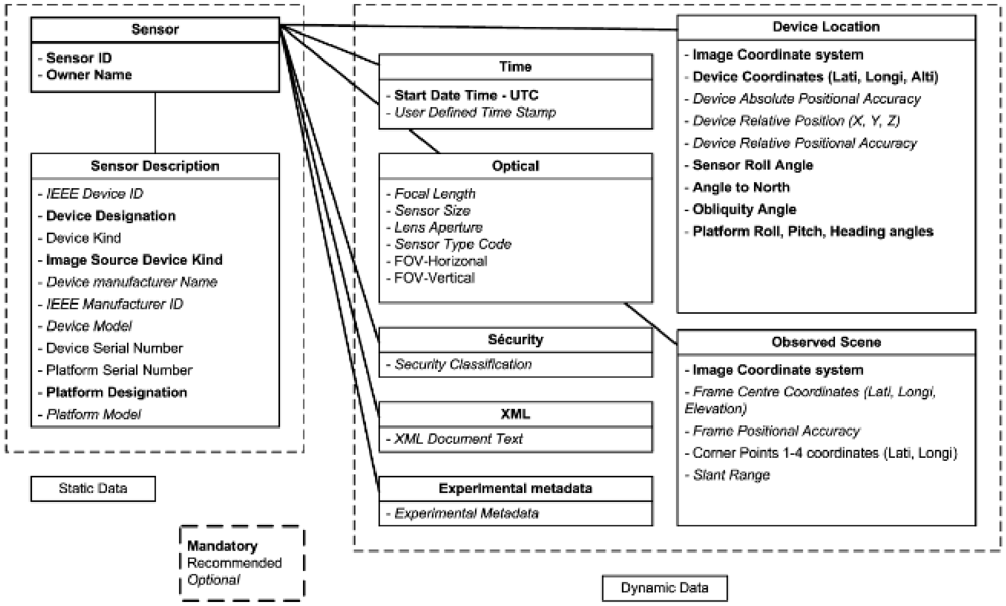

Figure 4: ISO 22311 sensor description

The video surveillance domain has seen a big number of commercial systems developed [8]. In the research area, many projects were developed as well: CROMATICA [5], CARETAKER[2] [3], VANAHEIM[3] for the indoor static video surveillance, and SURTRAIN [20], BOSS[4] [13], PROTECTRAIL[5] projects for the on-board mobile surveillance. All these heterogeneous projects concentrate on the development of the system's physical architecture and of better detection algorithms in order to obtain a fully automatic system [12], [24].

We can summarize by saying that there is a growing concern in the research and industrial environments for developing algorithms for video content analysis (VCA) in order to automatically index content and detect objects (e.g., abandoned packets or luggage) and events (e.g., intrusions, people or vehicles going the wrong way) [16] or to draw operators' attention to events of interest (e.g., alarms). However, solutions for assistance to a posteriori investigation are at a lesser stage of maturity, and to date most of the data remain unexploited.

In this article, we are going to address also the lack of interoperability between different surveillance systems. In the context of an inquiry, the police might need to analyse data from different sources (systems), so it is important that the different outputs of the systems to be interoperable. As a consequence, the big actors of the domain started to unify efforts in order to standardize the structure of folders and of metadata files generated by video surveillance systems. A result of these efforts is represented by the ISO 22311 standard that proposes a structure for the data issued from video surveillance systems and the metadata needed to exploit that data.

In the following, we are going to present the ISO 22311 standard, especially the part concerning the description of the cameras characteristics and mobility. We are going to highlight the interesting elements which relate to our research.

2.2 Standard ISO 22311

The Standard ISO 22311 defines an interoperability format for the data generated by video surveillance systems and for the metadata needed to exploit these huge volumes of data.

The audio visual packages (containing audio, video or metadata files) have to be structured hierarchically (in files, folders and groups of folders) according to time intervals in Coordinated Universal Time (UTC). For each group of folders it is mandatory for the system to provide a XML description of the source(s) (e.g., cameras, GPS, video analysis tools), codec(s), file formats and a temporal index enabling an easy access to the content.

The current technologies and processing power enable the analysis of video content and the extraction of metadata describing objects, events, scenes etc. This analysis depends on the acquisition context (e.g., the position of the camera, the image quality, the type of sensors). Therefore, the standard distinguishes between the systems, those that can generate such metadata (i.e., level 2 systems) and provides a general structure and dictionary for describing sensors and events (i.e., metadata).

As in this paper we are going to address the problem of cameras' geo-localization we present the schema for the sensors description in Figure 4.

[2] http://cordis.europa.eu/ist/kct/caretaker_synopsis.htm

[3] http://www.vanaheim-project.eu/

[4] http://celtic-boss.mik.bme.hu/

[5] http://www.protectrail.eu/

Each camera has an absolute location (GPS coordinates) as more and more of the installed cameras have an embedded GPS transmitter. But, there are many cases when the GPS is not enough because: (1) we need to model the position of the camera with regards to the video surveillance system and not to the world; (2) in some situations, for example in indoor environments, the GPS positions do not provide a good precision.

In the context of a video surveillance system:

- The mobile cameras are embedded in buses, train, police cars;
- The movement of these vehicles is constrained by a road network and a transportation network.

By analysing the standard, we can notice that it defines a relative position for a camera that is today a simple link to an image (the plan of the network of cameras or of a building). This kind of location is not easily exploitable. Furthermore, the standard does not consider the video surveillance cameras' mobility. In order to overcome these issues, we propose to extend this standard through a multilayer modelling approach, where the network of cameras is put on top of a transportation network.

In the following, we present a state of the art of the mobile objects modelling as the cameras' mobility management represents the main focus of this paper.

2.3 Mobile Objects Modelling

With the technology's evolution, the mobility became very important in the context of video surveillance systems. Not only the objects (e.g., persons, cars) are moving in the monitored scene, but also the surveillance cameras are moving. The great majority of the research papers concerning the mobile objects in the video surveillance domain concentrate on the video content analysis in order to detect and track the objects, to interpret their behaviour and to understand the visual events of the monitored scene [10]. Thus, the mobility of the cameras is not exploited.

In the field of moving objects, a mobile object means the continuous evolution of any object over the time, in terms of position and dimension [21]. This movement of the mobile objects can be effectuated in an unconstrained environment [18] (e.g., for hurricanes, fires) or in a constrained environment [17] (e.g., cars move on road and transportation networks).

In the video surveillance domain, the objects are moving in a constrained environment, mainly by the road network. This environment is represented as a graph-based model [6], [15], [25], where the vertices are junctions and the edges are the roads between the two junctions. [9] considers also the connectivity at each junction in order to represent the road network. [19] extends the model proposed by [9] in order to consider the predefined trajectories that some objects could have (e.g., buses). [7] proposes a mobile object data model where they consider the road and rail networks. [2] takes into account the transport network in a city as a graph and they add to each graph vertex the transport modes available (i.e., pedestrian, auto, urban rail, metro, bus).

In the management of mobile objects, a major issue is the storage of the objects' spatio-temporal positions. Several strategies can be considered: using the spatio-temporal data types defined by [9] (e.g., moving points, moving lines, moving regions), or using the dynamic attributes [23] (e.g., motion vector) which enables to limit the size of the data that has to be stored and queried.

As far as we know, the video content's mobility is not taken into account in the video surveillance domain. In this article, we want to exploit the advances in the field of mobile objects and apply them in the video surveillance domain in order to consider the mobile aspect of surveillance cameras.

3. Extension of the Standard 22311 for the management of cameras mobility

As you could see in Section 2.2, the Standard 22311, defines a fix position of video surveillance camera, through the GPS coordinates and a link to an image containing the plan of the network. In order to overcome this issue, we propose to compute a relative position with regards to a map which will enable us to:

- Model the distances between the cameras and select the relevant cameras for a certain trajectory;
- Model the connections between the cameras (e.g., possible path between camera1 and camera2 but not between camera2 and camera3);
- Model trajectories for mobile cameras;
- Model the fields of view and the maximum detection distances of fixed and mobile cameras.

In order to achieve this goal we took our inspiration from the domain of GIS (Geographical Information Systems) [4] and mobile objects modelling.

By considering the video surveillance system as a GIS we benefit from the separation between the conceptual layers. Thus at any time, a new layer can be added without modifying the existing layers.

In our approach, we propose a four layer model: (1) Road network, (2) Transportation network, (3) Objects and (4) Cameras network. The Figure 5 illustrates the UML model for the first three layers.

The "*Road network*" layer, presented in blue in Figure 5, is based on the graph modelling approach well-known in the literature. The road network is considered as an undirected graph G= (V, E), with V a set of vertices and E a set of edges defined according to the granularity level that we want to consider (for a big boulevard of a European capital for example we can consider each segment of the road, each segment between two intersections or the entire boulevard). Each vertex has an identifier and a 2D position. Each edge is determined by two vertices.

The "*Transportation network*", presented in yellow in Figure 5, is also based on a graph model. At this level, the vertices of the transportation network are intersections between roads, and bus stations. Each transportation vertex has a position with regards to a road segment. Ordered sequences of transportation vertices constitute sections, which form lines (e.g., bus lines). The advantage of our approach with regards to the ones proposed in the state of the art [9] is that we have two independent graphs that are connected to each other through the positions of transportation vertices. That way if the buses stations are modified or new buses lines are introduced we do not have to recompute the underlying road graph.

The "*Objects*" layer, presented in red in Figure 5, models the positions of fixed and mobile objects with regards to the underlying layers.

The Fixed Object has a position on a road segment. Its position is defined as a distance from each end of the segment. For this kind of objects, we adopt the same localisation as the one proposed by [9].

In the case of Mobile Objects (e.g., buses, police cars, persons), the position changes in time. Each object will periodically transmit its position using different strategies (e.g., each Δt seconds, each time the object is changing the segment, when the object's position predicted by the motion vector deviates from the real position by more than a threshold [23]) that are out of the scope of this article. We suppose that we periodically receive updates containing time-stamped GPS points that we transform into a relative position with regards to the road network (i.e., the segments). We use this information to reconstitute object's trajectory.

We distinguish two types of mobile objects: objects that move freely within the road and transportation networks (e.g., car, person) and objects of which trajectories are constrained by a "line" (e.g., buses).

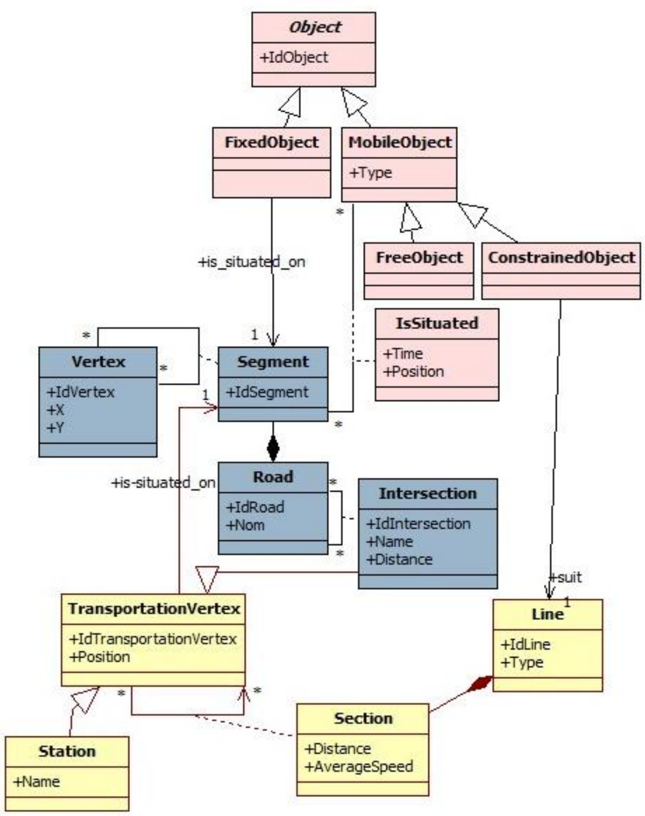

Figure 5: "Road network", "Transportation network" and "Objects" layers

On top of all these layers, we model a video surveillance cameras' network. A simplified schema of this model is illustrated in Figure 6.

The cameras' network is composed of fixed and mobile cameras. The fixed cameras have a 2D position that is given at installation time. The mobile cameras are associated with mobile objects (e.g., buses) and their trajectory is the same as the object's one.

The new generation of digital surveillance cameras has embedded GPS transmitters and even compasses. The technologies developed around these cameras make it possible to automatically extract information from the camera related to its orientation, pan, tilt, zoom, focal distance, compression parameters etc.

Based on all these elements it is possible to model the field of view for each camera and track its modifications in time. The field of view is computed based on four parameters [1]: the 2D position, the viewable angle, the orientation and the visible distance. A schema of a 2D field of view proposed by [1] is shown in Figure 7.

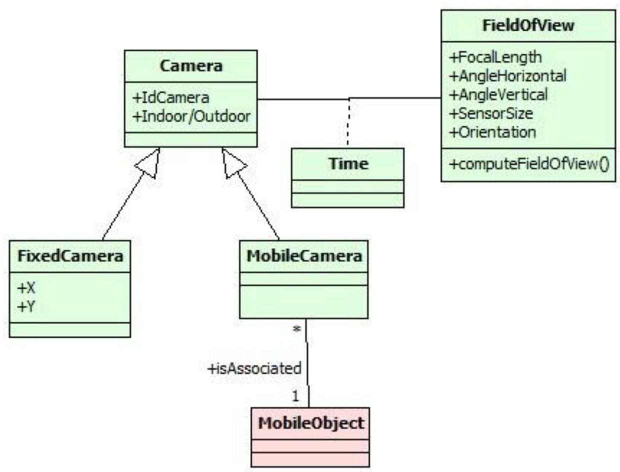

Figure 6: "Cameras network" layer

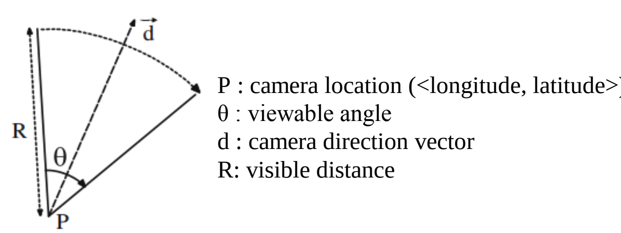

Figure 7: Illustration of the field of view model in 2D [1]

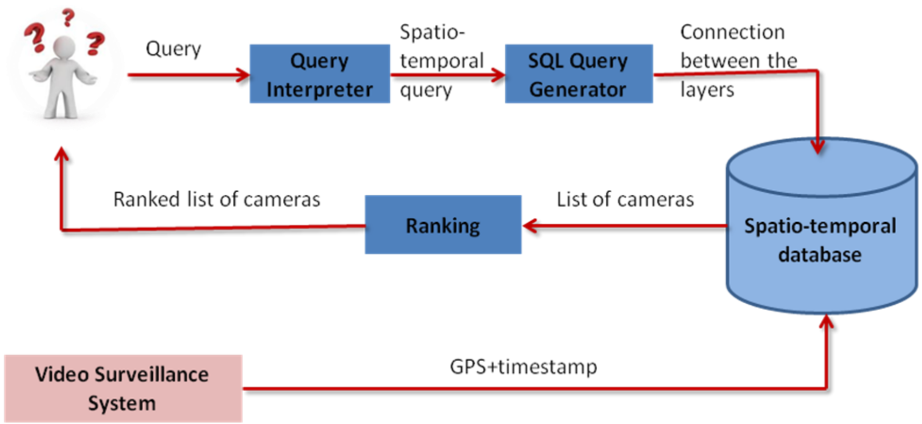

Figure 8: General architecture of the system

In order to select the most appropriate attributes to describe a video surveillance camera, we studied the sensor description proposed by the ISO 22311 standard, SensorGML[6], KML[7]. We separated the identified camera's properties in two categories: properties that could be modified over the time, and fixed characteristics.

Thus, the extension of the standard ISO 22311 is realised at three levels:

- Taking into account the road and transportation networks as a graph and not as an image;
- Taking into account the camera's relative position and its mobility on the networks;
- Taking into account the camera's characteristics change over the time.

Our model is implemented in a spatio-temporal database that can be queried by users in order to retrieve the relevant cameras for a given trajectory. The originality of our research work is given by:

- the fact that it combines different spatio-temporal information (e.g., road network, transportation network, objects' positions) and computation (e.g., trajectories, field of view) within the same database;
- the twofold mobility, of the target objects and of the cameras.

In the next section we present the general architecture of the tool that could assist the video surveillance operators in their research based on our spatio-temporal database and some examples of queries.

4. Spatiotemporal database and queries

Based on the presented model, our goal is to automatically select the cameras (fixed and mobile) that could contain relevant video content with regards to the user query (their field of view intersected the query trajectory).

More precisely, the idea is to compare a spatio-temporal query of the user (e.g., Rivoli Street from Louvre to Metro Chatelet the 14th of July between 10h and 14h) with the trajectories stored in our database and, for a better precision, with the cameras fields of view. The Figure 8 illustrates the generic architecture of a system based on our spatio-temporal database for assisting the video surveillance in their research.

From the Figure 8 it is easy to observe that there are two main questions when developing such system: How to query the system? and How to update the system?. As explained in the previous section our work addresses only the querying aspect that we are going to describe in the following.

First, a Query Interpreter module will transform the user query (e.g, Rivoli Street from Louvre to Metro Chatelet the 14th of July between 10h and 14h) in a spatio-temporal query. By spatio-temporal query we understand a sequence of road segments and a time interval that will be further transformed in a SQL query, by the SQL Query Generator module. The SQL query is executed on the database having as a result a list of cameras. Based on some image quality parameters a score per camera can be computed and the initial list can then be ranked according to this relevance score.

In the following we present two examples of spatio-temporal queries executed on our database implemented in Oracle Spatial [8]:

- *The first selects the fixed cameras of which geometry (field of view) intersects the geometry of the Rivoli street;*

```
SELECT IdCamera
FROM FixedCamera
WHERE SDO_RELATE(
                camera_geom,
                (SELECT street_geom
                        FROM Road
                        WHERE Name ='Rivoli‘ ),
               'mask=OVERLAPBDYDISJOINT  querytype=WINDOW'
)='TRUE';
```

[6] http://www.opengeospatial.org/standards/sensorml

[7] http://www.opengeospatial.org/standards/kml

[8] http://www.oracle.com/fr/products/database/options/spatial/index.html

- *The second selects the mobile cameras that are associated with the buses that crossed the street within the given time interval.*

```
LET TimePeriod = Timestamp(hour(2013,1,14,10), hour(2013,1,14, 12));
SELECT ObjetID
	FROM ConstrainedObject
	WHERE Type.MobileObject= “Bus” AND TimePeriod.ConstrainedObject (atperiods (Timestamp, TimePeriod));

SELECT DISTINCT IdMobileCamera
	FROM ConstrainedObject, FreeObject, MobileCamera
	WHERE Intersect (MobileCamera.geom, ConstrainedObject.geom) AND Intersect (MobileCamera.geom, FreeObject.geom);
```

5. CONCLUSION

In this paper, we presented a spatio-temporal modelling approach of fixed and mobile cameras within a common transportation network. Taking our inspiration from the multilayer representation of the geographical information systems, we model spatial information about the road and transportation infrastructures and mobile objects’ trajectories in four independent layers: (1) Road network, (2) Transportation network, (3) Objects and (4) Cameras network.

Based on this modelling approach we also proposed a generic architecture for a system that could assist the video surveillance operators in their research. Starting from a sequence of trajectory segments and a temporal interval, such system generates the list of cameras that could contain relevant information concerning the query (that “saw” the query’s trajectory).

The need of such assisting tools was identified within the French National Project METHODEO. Among the project’s partners, we mention the French National Police, Thales and the RATP also known as Régie Autonome des Transports Parisiens (English: Autonomous Operator of Parisian Transports). Our approach has been validated and will be evaluated within the project.

Obviously, many questions are still left with no answer giving way to a large number of perspectives. We will present several of them in the following.

For now, our model considers only outdoor transportation and surveillance networks. We plan to extend our model to indoor spaces also in order to model cameras inside train or subway stations for example.

Our work is situated in the context of the a posteriori research in the case of a police inquiry. We would like to extend this context in the future in order to be able to process real time queries or to predict trajectories based on some statistics realized based on the stored data (e.g., average speed on some road segments).

Another perspective of our work is the improvement of the resulted cameras list by re-ranking it based on cameras’ characteristics (e.g., image quality, visible distance).

6. ACKNOWLEDGMENTS

This work has been supported by the ANR CSOSG-National Security (French National Research Agency) project METHODEO.

7. REFERENCES

[1] Ay S. A., Zimmermann R., and Kim S.O.. Relevance Ranking in Georeferenced Video Search. In *Multimedia Systems Journal*, 16, 2 (March 2010), Springer, 105-125.

[2] Booth J., Sistla P., Wolfson O., and Cruz I. F. A data model for trip planning in multimodal transportation systems. In *Proceedings of the 12th International Conference on Extending Database Technology: Advances in Database Technology*, ACM, 2009, 994-1005.

[3] CARETAKER Consortium. Caretaker puts knowledge to good use. In *European Public Transport Magazine*. 2008.

[4] Chang K.-T. *Introduction to Geographic Information Systems*, McGraw-Hill Higher Education, 2006, 450 pages

[5] Deparis J.P., Velastin S.A., and Davies A.C. Cromatica project. In *Advanced Video-Based Surveillance Systems*, *The Springer International Series in Engineering and Computer Science*. volume 488, 1999, 203-212.

[6] Dingin Z., and Deng K. Collecting and managing network-matched trajectories of moving objects in databases. *In Proceedings of the 22nd international conference on Database and expert systems applications*, LNCS Volume 6860, 2011, pp 270-279 2011, 270-279.

[7] El bouziri A., Boulmakoul A., Laurini R. Mobile Object and Real Time Information System Modeling for Urban Environment, *In Proceedings of the 26th Urban and Regional Data Management Symposium*, 2007, 403-413.

[8] GDANSK, KU. Deliverable 2.1 – Review of existing smart video surveillance systems capable of being integrated with ADDPRIV. ADDPRIV consortium, 2011, www.addpriv.eu

[9] Güting R.H., Almeida, V.T., and Ding Z. Modeling and Querying Moving Objects in Networks. In *VLDB Journal*, 15, 2 (2006), 165-190.

[10] Joshi K. A., and Thakore D. G. A Survey on Moving Object Detection and Tracking in Video Surveillance System, in *International Journal of Soft Computing and Engineering*, 2, 3 (July 2012), 44-48

[11] Kim I.S., Choi H. S., Yi K. M., Choi J. Y., and Kong S. G. Intelligent Visual Surveillance - A Survey. In *International Journal of Control Automation and Systems*, 8, 5 (april 2010), 926-939

[12] Lakshmi Devasena C., Revathí R., and Hemalatha M., Video Surveillance Systems-A Survey, in *International Journal of Computer Science Issues*, 8, 4(July 2011), 635-642

[13] Lamy-Bergot C., Ambellouis S., Khoudour L., Sanz D., Malouch N.,. Hocquard A, Bruyelle J-L., Petit L., Cappa A., Barro A., Villalta E., Jeney G., and Egedy K. Transport system architecture for on board wireless secured a/v surveillance and sensing. In *Proceedings of the 9th international conference on Intelligent Transport Systems Telecommunications*, IEEE, 2009, 564-568.

[14] Le Barz, C. and Lamarque, T. Video Surveillance Cameras. In *Jean-Yves Dufour, editor, Intelligent Video Surveillance Systems, chapter 4*, Wiley, novembre 2012. 33-46

[15] Liu K., Li Y., He F., Xu J., and Ding Z. Effective Map-matching on the Most Simplified Road Network. In *Proceedings of the 20th International Conference on Advances in Geographic Information Systems* (SIGSPATIAL '12). ACM, 2012, 609-612.

[16] Marraud, D. Cepas, B. and Reithler, L., Semantic browsing of video surveillance databases through Online Generic Indexing, in *Proceeding of the 3rd ACM/IEEE International Conference on Distributed Smart Cameras*, 2009, .1-8.

[17] McKenney M, and Schneider M, Spatial Partition Graphs: A Graph Theoretic Model of Maps. In *Proceedingsof the 10th Int. Symp. on Spatial and Temporal Databases*, LNCS 4605, Springer, 2007, 167–184.

[18] Parent Ch., Spaccapietra S., and Zimányi E. *Conceptual Modeling for Traditional and Spatio-Temporal Applications: The MADS Approach*. Springer-Verlag New York, 2006

[19] Popa I.S., and Zeitouni K. Modeling and Querying Mobile Location Sensor Data, In *Proceedings of the 4th Int. Conference on Advanced Geographic Information Systems, Applications, and Services*, IARIA, 2012, 222-231

[20] Priam Q.-C., Lapeyronnie A., Baudry C., Lucat L., Sayd P., Ambellouis S., Sodoyer D., Flancquart A., Barcelo A.-C., Heer F., Ganansia F., and Delcourt V. Audio-video surveillance system for public transportation. In *Proceeding of the 2nd international conference on Image Processing Theory Tools and Applications*, IEEE, 2010, 47-53.

[21] Schneider M. Moving Objects in Databases and GIS: State-of-the-Art and Open Problems, in *Research Trends in Geographic Information Science*. Springer-Verlag, 2009, 169-188.

[22] Sedes F., Sulzer J.F., Marraud D., Mulat Ch., and Cepas B.. A Posteriori Analysis for Investigative Purposes. In *Jean-Yves Dufour, editor, Intelligent Video Surveillance Systems, chapter 3,* Wiley, novembre 2012, 33-46.

[23] Sistla A.P., Wolfson O., Chamberlain S., and Dao S. Querying the Uncertain Posi-tion of Moving Objects. In: *Temporal Databases: Research and Practice*, LNCS 1399, Springer, 1998, 310–337

[24] Soro S. and Heinzelman W. A Survey of Visual Sensor Networks, in *Advances in Multimedia*, vol. 2009, 2009, 1-22

[25] Xu J., Guo L., Ding Z., Sun X., and Liu C. Traffic aware route planning in dynamic road networks. In *Proceedings of the 17th international conference on Database Systems for Advanced Applications*, LNCS 7238, 2012, 576-591

MappingSets for Spatial Observation Data Warehouses

José R.R. Viqueira
COGRADE - CITIUS
Universidade de Santiago de Compostela, Spain
jrr.viqueira@usc.es

Sebastián Villarroya
COGRADE - CITIUS
Universidade de Santiago de Compostela, Spain
sebastian.villarroya@usc.es

David Martínez
COGRADE - CITIUS
Universidade de Santiago de Compostela, Spain
david.martinez.casas@usc.es

José A. Taboada
COGRADE - CITIUS
Universidade de Santiago de Compostela, Spain
joseangel.taboada@usc.es

ABSTRACT

The amount of time evolving spatial data that is currently being generated by automatic observation processes is huge. In general, observation data consists of both heterogeneous spatio-temporal data and relevant observation metadata. The former includes data of Spatial Entities (cities, roads, vehicles, etc.) and data of temporal evolution of both properties of Spatial Entities (population of a city, position of a vehicle, etc.) and properties of space (temperature, elevation, etc.). Real uniform integrated management of all these types of data is still not achieved by current models and systems. The present paper describes the design of a data modeling and management framework for observation data warehouses. A hybrid logical-functional data model based on the concept of MappingSet and relevant language enables the specification of spatio-temporal analytical processes. The framework in currently being implemented.

1. INTRODUCTION

According to [16], properties of entities (called Features of Interest - FOI) are either exact values assigned by some authority (names, prices, geometry of a municipality, etc.) or estimated by some observation process (height, classification, color, etc.). Observation processes may be classified in various different ways [15]. *Physical Processes* produce their data in some spatial context. They are usually hardware sensing devices that perform measurements either locally or remotely. Besides, they may be installed in either static or mobile platforms. *Non-Physical Processes* are computations that may be defined in some mathematical way. Any process may be either *Time-triggered* o *Event-triggered*. The former perform their results at some predefined time frequency. The latter are started by some external event at any moment in time.

Observation data has an inherent temporal nature. Besides, in many cases FOIs are also spatial. Therefore, systems devoted to observation data analysis should cope with spatial and spatio-temporal data analysis. In particular, they should support relevant functionality for the management of Spatial Entities and Spatial Coverages, and their evolution with respect to time [9, 20, 6]. *Spatial Entities* are entities of a given application domain that have geometric valued properties (rivers, municipalities, cities, etc.). *Spatial Coverages* are sets of functions with a common spatial domain that describe the continuous or discrete variation over space of some specific phenomenon (temperature, humidity, elevation above sea level, etc.).

The amount of data that is currently being obtained from automatic observation processes is huge and the estimated tendency is to have an exponential growth during the upcoming years. The analysis of all these data to support appropriate decision making is key challenge for future information systems. Many application domains exist that would benefit from innovative technologies in this area, including environmental observation and monitoring, natural disaster management, e-health, etc.

Based on the above, in the present paper a data modeling and management solution is proposed that enables spatio-temporal analysis in data warehouses of observation data. In particular, a proposed E-R extension enables the insertion of observation metadata in spatial models at a conceptual level. At a logical level, a new data model based on MappingSets enables the integrated management of any kind of spatial and temporal data. A MappingSet is a collection of Mappings, in the functional programming sense, defined on a common domain. Both Spatial Entities and Spatial Coverages and both Time-triggered and Event-triggered observation data are modeled uniformly with MappingSets.

The remainder of this paper is organized as follows. Section 2 describes other pieces of work related to the proposed solution. The MappingSet based spatio-temporal logical model is described in Section 3. The conceptual level E-R extension for observation data is described in Section 4, as it is also its translation to the MappingSet based logical model. Section 5 illustrates the spatio-temporal analysis capabilities of the model for the definition of Non-Physical spatio-temporal analytical processes. Finally, Section 6 concludes the paper and outlines lines of future work.

2. RELATED WORK

The OGC defines an abstract specification of a data model for Observations and Measurements [16] in a Geographic Information context. Various types of observations are supported, according to the data type of their values. Simple observations include: i) measurements that combine a value of a real type with a unit of measure, ii) categories whose results are items of enumerated types, iii) counts of integer types, iv) truth observations of boolean type, v) time observations and vi) geometric observations. Complex observations are record structures that combine various simple observation types. Metadata of each observation is also represented in the model. In particular, each observation references its observation Process, the observed Property and its related FOI, the time instant when the observation applies to the FOI observed property (phenomenon time) and the time instant when the Process obtained the result value (result time). Notice for example that if a sample of water is obtained from a river and next analyzed in a laboratory two different observation time instants are involved. Optionally, other metadata, parameters, data quality information and observation context may also be provided.

In [4] a conceptual model to represent observation data semantics is defined. Annotating the conventional data models of available heterogeneous datasets with observation and measurement conceptual constructs enables their integration at a semantic level. Integrated query of heterogeneous observation datasets becomes therefore possible after the annotation process. A similar approach is followed by the E-R extension proposed in the present paper.

Observation data has always a temporal nature. Besides, the spatial components of observation data and metadata is centric to many application domains, such as those related to environmental observation and monitoring. Spatial and temporal extensions of conceptual and logical data models have to be considered. Examples of spatio-temporal conceptual models are [18, 17]. Relational and object-relational spatio-temporal extensions are defined in the area of Spatial Databases [9, 20] to support spatial entity management. Field [6] and array algebras [3] are behind spatial coverage and array management systems [14, 5, 2]. Integrated management of spatial entities and coverages is also objective of some approaches [19, 12], that incorporate different structures for those data types. Integrated management of entities and coverages in a uniform manner is achieved by the MappingSet data model proposed in the present paper.

Various different data management approaches are possible to deal with spatio-temporal observation data automatically generated by sensing devices. If we consider the data generated by each sensor as a virtual temporal relation, then the simplest approach is to consider *Materialized Views* of such virtual relations. Automatic maintenance of such views on the arrival of new data from sensors has to be solved by the system [8]. Automatically updating these views through Extraction Transformation and Load (ETL) processes on sensor data is the approach followed by the present framework.

A more sophisticated solution is to consider sensor data streams and to enable the continuous execution of queries on those input streams. Continuous query languages [1, 11] enable the definition of those continuous queries on both data streams and recorded relations. Operations to create relations from streams and streams from relations are at the core of those languages. A similar approach is followed by some languages specifically designed to access sensor networks [13, 7]. It is important to notice that spatial data, including spatial entities and spatial coverages and spatial analysis is not explicitly supported in these solutions.

3. SPATIO-TEMPORAL MAPPINGSET BASED DATA MANAGEMENT

This section introduces the MappingSet based data model that is the basis of the proposed framework. Temporal and spatial data types are first defined. Based on them Mappings and MappingSets are next formalized. Data management will be based on the intensional definition of MappingSets using both logical and functional paradigms.

Conventional data types include **Boolean**, **CString** (variable size character strings), **Int16**, **Int32**, **Int64** (integers), **Float32**, **Float64** (reals with floating point representation). Fixed point parametric type **Numeric(P,D)** consists of real numbers with a maximum of P digits, D of them are in the fractional part. In order to define temporal and spatial data types, 1D and 2D samplings are first formalized. Let R and I denote the set of real and integer numbers, respectively, then 1D and 2D samplings are defined as follows.

Definition 1. A **1D-sampling** S with resolution $r \in R$ and phase $p \in R$ is defined as the infinite subset of R
$\{x | x = i \cdot r + p, \forall i \in I\}$

Definition 2. Let vr_1, vr_2, vp_1 and vp_2 be four vectors in R^2 defined by respective directions D_1, D_2, D_1, $D_2 \in (-\pi, \pi]$ and respective magnitudes r_1, r_2, p_1 and p_2. A **2D-sampling** S with directions D_1, D_2, resolutions r_1, r_2 and phases $p_1 \in [-r_1/2, r_1/2]$, $p_2 \in [-r_2/2, r_2/2]$ is defined as the infinite subset of R^2

$$\{(x, y) \in R_2 | \\ x = (i_1 r_1 + p_1)\cos(D_1) + (i_2 r_2 + p_2)\cos(D_2) \wedge \\ y = (i_1 r_1 + p_1)\sin(D_1) + (i_2 r_2 + p_2)\sin(D_2), \\ \forall i_1, i_2 \in I\}$$

An element s of a 1D-sampling (2D-sampling) S is called a 1D-sample (2D-sample). Integer i, i_1, i_2 are called the *sampling coordinates* of s. $s(i)$, $s(i_1, i_2)$ denote respectively the 1D-sample and 2D-sample with sampling coordinates i and $(i_1, \ i_2)$. Figure 1 illustrates the above definitions with a geometrical representation.

Definition 3. **TimeInstant(D)** is defined as a finite subset of elements $s(i)$ of a 1D-sampling S with resolution 10^{-D} and phase 0 such that
$-2^{63} < i < 2^{63} + 1$
where each s(i) is interpreted as the time instant 1/1/1970+ $s(i)$ seconds. Maximum allowed D is 6 (microsecond).

Definition 4. **TimeInstantSample(D, R)** is defined as a finite subset of elements $s(i)$ of a 1D-sampling S with resolution $R \cdot 10^{-D}$ and phase $(R \cdot 10^{-D})/2$ such that
$-2^{63} < i < 2^{63} + 1$
where each s(i) is interpreted as the time interval [1/1/1970+ $s(i)$ seconds, $1/1/1970 + (s(i) + R \cdot 10^{-D})$ seconds). Again, maximum allowed D is 6 (microsecond).

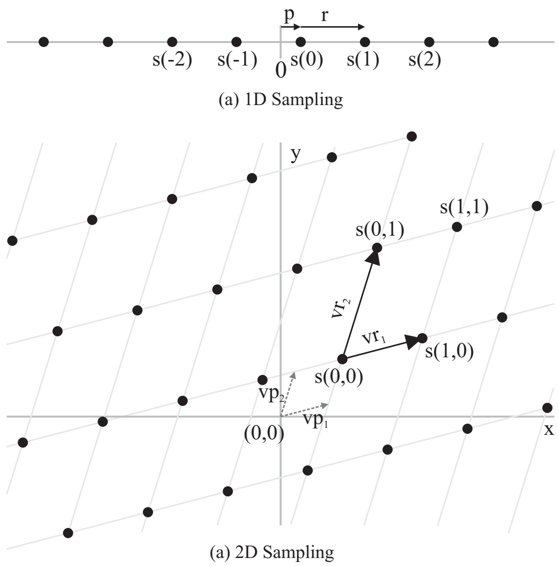

Figure 1: Illustration of 1D and 2D samplings.

Definition 5. **Date** is defined as a shorthand of TimeInstantSample(0, 86400).

Definition 6. **Point2D(P,D)** is defined as the finite subset of elements $s(i_1, i_2)$ of a 2D-sampling S with directions $D_1 = 0$ and $D_2 = \pi/2$, resolutions $R_1 = R_2 = 10^{-D}$ and phases $Ph_1 = Ph_2 = 0$ such that
$-10^P < i_1, i_2 < 10^P$

Definition 7. **Point2DSample(P,D,R)** is defined as the finite subset of elements $s(i_1, i_2)$ of a 2D-sampling S with implementation dependent directions D_1 and D_2, resolutions $r_1 = r_2 = K \cdot R \cdot 10^{-D}$ and phases $p_1 = p_2 = 0$ such that

1. $-10^P < i_1, i_2 < 10^P$
2. $K < max(\left|\cos(\frac{D_2-D_1}{2})\right|, \left|\sin(\frac{D_2-D_1}{2})\right|)$

TimeInstant and Point2D data types provide discrete representations for both time and space, where the user has control over the supported precision. Types TimeInstantSample and Point2DSample provide representations for temporal and spatial samplings at user defined resolution. It is noticed that each time instant is approximated by its closest lower TimeInstantSample, whereas each 2D point is approximated by its closest Point2DSample. It is out of the scope of this paper to demonstrate that K factor above ensures that any 2D point is approximated by a sample at a distance lower or equal to $R \cdot 10^{-D}$. Type castings are available for the above data types.

If T is either a numeric or temporal type, then data type **Interval(T)** is a new data type whose values are closed intervals over data type T. If t_1, t_2 are two elements of data type T, then $[t_1, t_2]$ is used to denote the relevant closed interval. Similarly, if S is spatial data type then the following geometric data types are also supported, based on the standard specification given by [10].

- **Geometry(S)**: Abstract type. Represents any vector geometry or set of geometries defined with elements of S.
- **LineString(S)**: Vector polylines defined by sequences of elements of S.
- **Polygon(S)**: Vector polygons, possibly with holes, whose borders are defined by sequences of elements of S.
- **GeometryCollection(S)**: Heterogeneous collections of Geometries.
- **MultiPoint(S)**: Homogeneous collection of elements of S.
- **MultiLineString(S)**: Homogeneous collection of elements of LineString(S).
- **MultiPolygon(S)**: Homogeneous collection of elements of Polygon(S).

Definition 8. If $ADT_1, ADT_2, \dots ADT_n$ are not necessarily distinct data types, $A_1, A_2, \dots, A_n$ are distinct names and RDT is a data type, then:

1. A **Mapping** with signature $M() : RDT$ is defined as a value of type RDT
2. A **Mapping** with signature
$M(A_1 : ADT_1, A_2 : ADT_2, \dots, A_n : ADT_n) : RDT$
is defined as a partial function
$M : ADT_1 \times ADT_2 \times ADT_n \rightarrow RDT$

Operations are syntactic sugar for Mappings. Implicit castings between compatible data types are applied during Mapping invocations, enabling transparent transformation between temporal and spatial elements of different resolutions by applying constant interpolation. Various primitive mappings and operations are provided by the model. However, formalizing a complete set of them is out of the scope of the paper. Informal descriptions of required primitive mappings will be given throughout the paper.

A MappingSet is nothing but a set of Mappings that share a common domain defined as a n-ary relation over data types. Formalism is given below.

Definition 9. Let $C_1, C_2, \dots, C_n$ be distinct names, ADT_1, ADT_2, ..., ADT_n be not necessarily distinct data types and RDT_1, RDT_2, ..., RDT_m be not necessarily distinct data types. Let also D be a n-ary relation with scheme $D(C_1 : ADT_1, C_2 : ADT_2, \dots, C_n : ADT_n)$ defined as a finite subset of $ADT_1 \times ADT_2 \times \dots \times ADT_n$. Then a **MappingSet** is defined in either of the three following forms:

1. A 1-tuple $MS = \langle D \rangle$.
2. A m-tuple $MS = \langle M_1, M_2, \dots, M_m \rangle$, where each M_i is a Mapping with signature $M_i() : RDT_i$ defined as a value of RDT_i.
3. A (m+1)-tuple $MS = \langle D, M_1, M_2, \dots, M_m \rangle$, where each M_i is a Mapping with signature $M_i(C_1 : ADT_1, C_2 : ADT_2, \dots, C_n : ADT_n) : RDT_i$ defined as a partial function $M_i : ADT_1 \times ADT_2 \times ADT_n \rightarrow RDT_i$.

The evolution with respect to time of spatial entities and spatial coverages may be modeled with appropriate MappingSets that contain both Domain and Mappings. n-ary relationships are also modeled with MappingSets, usually without Mappings. MappingSets without Domain are also useful to record short collections of key-value pairs that are common in the specification of configuration settings.

The Domains and Mappings of a MappingSet may be defined either extensionally or intensionally. If a extensional definition of the Domain is given, then both extensional and intensional definitions of Mappings are allowed. On the other hand, an intensional definition of the Domain may only be accompanied by intensional definitions of Mappings. Generally, an extensional definition is a sequence of all the elements of Domain and Mappings in some specific order. Both row-wise and column-wise orderings may be used. It is even possible to combine row and column-wise orders for different components and Mappings. If the data type of a Domain component is of some integer or sampling data type, then its extensional definition might be given in the form of a collection of sequence definitions. In general, a sequence definition has an start element, a size and a step. For example, for an integer data type, a sequence starting at 5, with size 4 and step 2 describes the following list $< 5, 7, 9, 11 >$. For a TimeInstantSample data types, a sequence starting at "$2013 - 05 - 0215 : 00 : 45.06$", with size 2 and step 30.42 describes the following sequence of type TimeInstant(2, 3042) ¡"$2013 - 05 - 0215 : 00 : 20.22$", "$2013 - 05 - 0215 : 00 : 50.64$"¿. [1]. For Point2DSample data types, starting element is fo type Point2D and step has to be given by two pairs (direction, resolution).

Spatio-temporal analysis is enabled through the intensional definition of Mappings and MappingSets. Mappings may be intensionally defined with functional, conditional and aggregate expressions.

Functional expression. A Mapping M with signature M(D): DT may be defined by a expression of the form

M(D) := e

where e is a functional expression of data type DT that may include variables referencing components of D, mappings, operations, constants and castings.

Conditional Expression. A Mapping M with signature M(D): DT may be defined by a expression of the form

$$\begin{array}{llll} \text{M(D)} := & \text{CASE} & b_1 & \text{THEN} & e_1 \\ & \text{CASE} & b_2 & \text{THEN} & e_2 \\ & \ldots & & & \\ & \text{CASE} & b_n & \text{THEN} & e_n \\ & [\text{OTHERWISE } e_{n+1}] & & & \end{array}$$

where each b_i is a *functional expression* that yields a value of Boolean type and each e_i is a *functional expression* that yields a value of type DT. The semantics are the obvious ones.

Aggregate Expression. A Mapping M with signature M(D): DT may be defined by a expression of the form

M(D):= agge
OVER {P}

where P is a domain relational calculus predicate and agge is an functional expression where variables bounded to MappingSet domains in P must be used as arguments of aggregate functions. Various aggregate functions are provided by the system including both statistical and rank functions. MappingSet domains may also be intensionally defined.

[1]Notice that the start instant of the sequence is automatically adapted to match the underlying time representation for type TimeInstant(2, 3042)

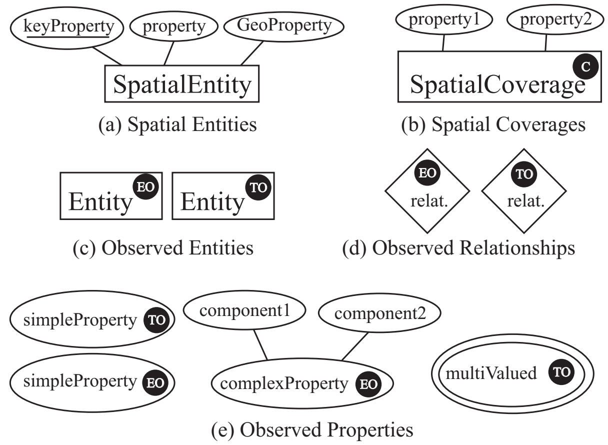

Figure 2: E-R Diagram Notation for Spatial and Observation Data.

Intensional Domain. Let e be a functional expression that yields a value s of either Interval(T) or Geometry(S) data type, whose base type T, S is either some integer type or some sample type. Then, SAMPLING(e) yields all the elements of type T or S contained in s. Based on this, the domain D of a MappingSet M may be defined by an expression of the form

$\{(e_1, e_2, \ldots, e_n)|P\}$

where P is a domain relational calculus predicate and each e_i is either a functional expression or an expression of the form $SAMPLING(e)$, where e is also a functional expression. Expressions e and e_i may include variable names bounded to MappingSet domain components in P. Given that nested structures are not allowed in the model, if an expression $SAMPLING(e)$ is used then the result relation has to be unnested.

4. MODELING OBSERVATION DATA WAREHOUSES

The data model described in this section captures observation data semantics and integrates them with spatial entities and coverages. An E-R extension is proposed in Subsection 4.1 to model observation metadata. The translation of such a conceptual model to the MappingSet based logical model is explained in Subsection 4.2.

4.1 Conceptual Data Model

Contrary to conventional metadata that is recorded at the level of entity and property types, some observation metadata has to be recorded at the level of entity and property instances, i.e., combined with the data itself. This is the case for example of observation time instants and observation processes.

An extension of the E-R model is next proposed to incorporate spatial and observation data semantics in conceptual models. *Spatial Entity types* are represented in diagrams as conventional entities (see Figure 2(a)). Spatial Coverage Types are represented as entities tagged with the symbol

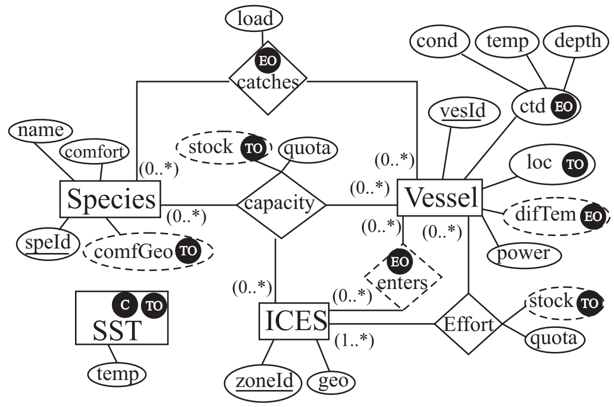

Figure 3: E-R Diagram of a Running Application Example.

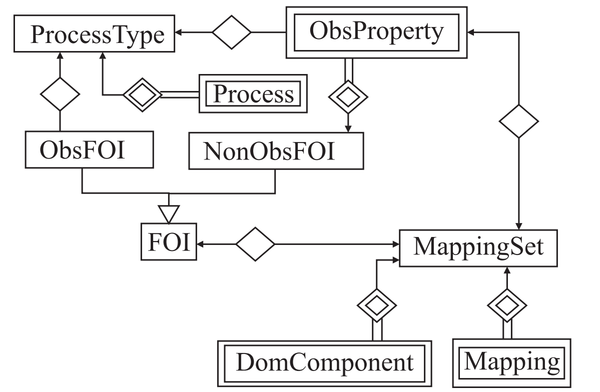

Figure 4: E-R Diagram of the Frameworks Catalog.

C (see Figure 2(b)). Entity Types, either spatial or not, and Coverages whose whole data is obtained through an observation Process are tagged with either symbol TO if it is a *Time-triggered Process* or symbol EO if it is an *Event-triggered Process* (see Figure 2(c)). Relationships resulting from observation processes are tagged in the same way (see Figure 2(d)). Finally, properties of either Entity or Coverage types that are obtained through observation processes are also tagged with the same TO and EO symbols, as it is shown in Figure 2(e) for simple, complex and multivalued properties.

To illustrate the use of the above notation the E-R diagram of a reduced running application example is given in Figure 3. Spatial Entity Type *ICES* records fishing zones defined by the International Council for the Exploration of the Sea (ICES). Spatial Coverage *SST* records Sea Surface Temperature at each location of the sea, daily produced by the Moderate Resolution Imaging Spectroradiometer (MODIS) sensor installed in the Terra and Aqua NASA satellites. Entity type *Vessel* records data of fishing vessels, including an identifier (*vesId*) and its engine *power*. Vessels incorporate CTD sensors that enable obtaining triples of water conductivity, water temperature and depth. Every time a ctd observation is performed a *Non-Physical Process* is executed that computes the difference with the value given by MODIS and provides it as a derived property *difTemp*. Vessels also incorporate GPS sensors from which locations are obtained every 30 seconds. Entity type *Species* records data of fishing species, including an identifier *specId*, species *name* and an interval of temperature values where the fish feels comfortable (property *comfort*). The derived property *comfGeo* records the geometry of the area of the sea where comfortable temperatures for the fish are located. This property is obtained by a *Non-Physical Process* from the SST data. Property *load* of relationship *catches* records the values measured by the vessel bascule for each species. The authorized fishing capacity of a vessel is given by two parameters. The Fishing Effort gives a measure of the number of days weighted by the vessel engine power that the vessel may stay in each zone. Relationship *Effort* records both the initial *quota* and the available one (property *stock*). Available Fishing Effort *stock* is obtained by a *Non-Physical Procress* using the quota and the vessels GPS information. The Fishing Capacity gives the kilograms of each species that the vessel may get from each zone. Again both *quota* is recorded and *stock* is computed by a *Non-Physical Process*.

The translation of the above model to the MappingSet logical model of the framework is explained in the following subsection.

4.2 MappingSet Based Logical Model

To support the implementation of the conceptual model of the previous section, observation metadata has to be added to the frameworks catalog. Thus, the catalog contains metadata of the defined Mappings and MappingSets and metadata related to the various observation processes, including observation properties and features of interest. The E-R diagram of such catalog structures is given in Figure 4.

Entity types *MappingSet*, *DomComponent* and *Mapping* record general metadata of the MappingSets. Entity type *FOI* records metadata of Features of Interest, and it references the MappingSet that records its data. FOIs that are fully generated by observation processes are registered in *ObsFOI*. The remainder FOIs, i.e., those that combine observed with non observed properties are represented by entity type *NonObsFOI*. Each observed property of such a FOI is represented by a weak entity of type *ObsProperty*, which references the MappingSet that records its data. Finally, *ProcessType* records metadata of the various types of observation processes registered in the framework. Metadata of each specific instance of each process type is recorded in weak entity type *Process*. Notice the difference between the process type "Vessel Bascule" that obtains values of *load* property of relationships *catches* and the specific bascule installed in each vessel that must be referenced from each observation.

The rules that enable the transformation of the conceptual model of the previous section to MappingSets are now given next. Each Entity Type, either Spatial or not, generates a relevant MappingSet, whose domain is defined by key properties and whose Mappings are defined by the remainder properties. See for example Entity Types Vessel, Species and ICES in Figure 3 and relevant MappingSets in Figure 5. Each Spatial Coverage generates a MappingSet, whose domain has just one component of some Point2DSample type and whose Mappings are generated from coverage properties. Each Relationship Type with cardinalities various to

```
MAPPINGSET Vessel
DOMAIN
 vesId: CString
MAPPINGS
 power(vesId:CString):Numeric(6,2)

MAPPINGSET Vessel_loc
DOMAIN
  obsTime: TimeIntantSample(0, 30),
  vesId: CString
MAPPINGS
  loc(phenTime: TimeIntantSample(0, 30),
      vesId:CString):Point2D
  process(obsTime: TimeIntantSample(0, 30),
          vesId:CString):CString

MAPPINGSET Vessel_ctd
DOMAIN
  obsTime: TimeIntant(0),
  vesId: CString
MAPPINGS
  cond(obsTime: TimeIntantSample(0, 30),
       vesId:CString):Numeric(4,1)
  condUOM(obsTime: TimeIntantSample(0, 30),
          vesId:CString):CString
  temp(obsTime: TimeIntantSample(0, 30),
       vesId:CString):Numeric(5,2)
  tempUOM(obsTime: TimeIntantSample(0, 30),
          vesId:CString):CString
  depth(obsTime: TimeIntantSample(0, 30),
        vesId:CString):Numeric(5,2)
  depthUOM(obsTime: TimeIntantSample(0, 30),
           vesId:CString):CString
  process(obsTime: TimeIntantSample(0, 30),
           vesId:CString):CString

MAPPINGSET Species
DOMAIN
 speId: CString
MAPPINGS
 name(speId:CString):CString
 comfort(speId:CString):Interval(Numeric(5,2))

MAPPINGSET ICES
DOMAIN
  zoneId: CString
MAPPINGS
  geo(zoneId:CString):Polygon(Point2D(9,2))
```

```
MAPPINGSET Capacity
DOMAIN
  vessel: CString
  ices: CString
  species: CString
MAPPINGS
  quota(vessel:CString, ices:CString
        species:CString):Numeric(7,3)
  quotaUOM(vessel:CString, ices:CString
           species:CString):CString

MAPPINGSET Effort
DOMAIN
 vessel: CString
 ices: CString
MAPPINGS
 quota(vessel:CString,
       ices:CString):Numeric(7,3)
 quotaUOM(vessel:CString,
          ices:CString):CString

MAPPINGSET Catches
DOMAIN
  species: CString
  vessel: CString
  obsTime: TimeInstant(0)
MAPPINGS
  load(species:CString, vessel:CString,
       obsTime:TimeInstant(0)): Numeric(7,3)
  loadUOM(species:CString, vessel:CString,
          obsTime:TimeInstant(0)): CString
  process(species:CString, vessel:CString,
          obsTime:TimeInstant(0)): CString

MAPPINGSET SST
DOMAIN
  loc:Point2DSample(9,2,100000)
  obsTime:Date
MAPPINGS
  temp(loc:Point2DSample(9,2,100000),
       obsTime:Date):Numeric(5,2)
  tempUOM(loc:Point2DSample(9,2,100000),
          obsTime:Date):CString
  process(loc:Point2DSample(9,2,100000),
          obsTime:Date):CString
```

Figure 5: MappingSets for a Running Application Example.

various generates a MappingSet whose domain is defined from the key properties of the participating Entity Types. Properties of those Relationship Types generate Mappings in such a MappingSet. See for an example Relationship Types capacity and effort in Figure 3 and MappingSets Capacity and Effort in Figure 5. If an Entity, Coverage or Relationship Type is tagged with the symbol TO, then a component named *obsTime* of some TimeInstantSample(D,R) data type is added to the MappingSet Domain to enable the recording of observation time.[2] Besides, a Mapping named *process* is also added to obtain the id of the process used to produce the observation. See for example Spatial Coverage Type SST in Figure 3 and relevant MappingSet SST in Figure 5. If symbol EO is used instead, then the data type of component obsTime is some TimeInstant(D). See for example Relationship Type catches in Figure 3 and relevant Catches MappingSet in Figure 5. In any of the above cases, an entity of type ObsFOI has to be added to the catalog with relevant relationships to its process type and MappingSet.

If a simple or complex property is tagged with symbol TO then such property is not added as a Mapping to the relevant MappingSet. Instead, a separate MappingSet is created for the property whose domain has components to reference the key of its Entity Type (FOI of the relevant observation) and has a component named *obsTime* of some TimeInstantSample(D,R) type to record observation time. The property itself is added as a Mapping to the MappingSet and an additional Mapping named *process* is added to record the id of the process that generates the observation. An example is loc property of Entity Type Vessel in Figure 3 and relevant Vessel_loc MappingSet in Figure 5. If symbol EO is used instead then the transformation is exactly the same except for the fact that Domain component *obsTime* is of some TimeInstant(D) type. For an example see ctd property of Vessel Entity Type in Figure 3 and relevant Vessel_ctd MappingSet in Figure 5. In any of the above cases an entity of type ObsProperty is added to the catalog, with appropriate references to its MappingSet, ProcessType and NonObsFOI.

Once the MappingSets are created and the required metadata are added to the catalog, the insertion of observation data may be started. ETL tasks are continuously executed to maintain the data warehouse updated with latest observation data, using extensional MappingSet definitions. Each observation is appended to the appropriate MappingSet with its observation time and reference to its process and FOI.

5. DEFINITION OF SPATIO-TEMPORAL ANALYTICAL PROCESSES

The capabilities provided by the framework for the intensional definition of MappingSets enable the specification of spatio-temporal analytical processes. These capabilities are now illustrated with some examples.

Example 1. Define a *Non-Physical Process* that obtains a derived observed property that computes the difference between the temperature measured by the CTD and the sea surface temperature produced for the same location by MODIS (see *difTemp* derived property of Vessel in Figure 3).

```
MAPPINGSET Vessel_difTem
DOMAIN
  {(obsTime, vesId) | Vessel_ctd(obsTime, vesId)}
MAPPINGS
  difTem(obsTime, vesId):=
    SST.temp(Vessel_loc.loc(obsTime, vesId), obsTime) −
    Vessel_ctd.temp(obsTime, vesId)
  difTemUOM(obsTime, vesID):=
    Vessel_ctd.tempUOM(obsTime, vesId)
  process(obsTime, vesID):= “difTemProcess”
```

In the expression above it is noticed that automatic castings of spatial and temporal types are performed during the evaluations of Mappings Vessel_loc.loc and SST.temp.

Example 2. Define a *Non-Physical Process* that detects when a vessel leaves an ICES zone to enter a new one (see *enters* derived relationship in Figure 3).

```
ICESFromLoc(loc):=
  singleton(zone)
  OVER {ICES(zone) ∧ within(loc, ICES.geo(zone))}

MAPPINGSET enters
DOMAIN
  {(vesId, ICESFromLoc(Vessel_loc.loc(obsTime, vesId)),
    obsTime) |
    Vessel_loc(obsTime, vesId) ∧
    ICESFromLoc(Vessel_loc.loc(obsTime, vesId)) <>
    ICESFromLoc(Vessel_loc.loc(predecessor(obsTime), vesId))}
MAPPINGS
  process(vesId, zoneId, obsTime):= “entersProcess”
```

In the above expression, Mapping *within*(g_1, g_2) yields true if geometry g_1 is within geometry g_2. Aggregate function *singleton*(S) yields the element contained in the unitary set S. Finally, Mapping *predecessor*(ts) yields the time sample that precedes time sample ts in its data type.

Example 3. Define a *Non-Physical Process* that produces a measure of the remaining fishing effort for each vessel and ICES zone for each of the preceding 60 days. Consumed fishing effort is obtained from the temporal evolution of vessel location data and ICES zone geometries (see derived property stock of relationship type Effort in Figure 3).

```
ICESFromLoc(loc):=
  singleton(zone)
  OVER {ICES(zone) ∧ within(loc, ICES.geo(zone))}

consumed_effort(vesId, zoneId, obsTime) :=
  ((count(obsTime2)*30)/86400)*Vessel.power(vesId)
  OVER {Vessel_loc(obsTime2, vesId2) ∧
    obsTime2 < obsTime ∧ vesId2 = vesId ∧
    ICESFromLoc(Vessel_loc.loc(obsTime2, vesId2)) = zoneId
    }

MAPPINGSET Effort_stock
DOMAIN
  {(vesId, zoneId,
    SAMPLING([cast(difTime(now(), 60 Days) AS Date),
      cast(now() AS Date)])) | Effort(vesId, zoneId)}
MAPPINGS
  stock(vesId, zoneId, obsTime):=
    Effort.quota(vesId, zoneId) −
    consumed_effort(vesId, zoneId, obsTime)
  stockUOM (vesId, zoneId, obsTime):=
    Effort.quotaUOM(vesId, zoneId)
  process(vesId, zoneId, obsTime):= “EffortStockProcess”
```

In the above expression, Mapping *now*() yields the current system time instant. Mapping *difTime*(t, i) subtracts time interval i from time instant t.

[2]Currently we restrict to phenomenon time semantics, however, it can be extended with result time and other required metadata.

Example 4. Define a *Non-Physical Process* that obtains the evolution with respect to to time during the last 7 days of the geometry of the comfort zone for each species. Comfort zone is obtained from the temperature interval defined for each species and the sea surface temperature generated by MODIS (see derived property comfGeo of entity type species in Figure 3).

MAPPINGSET Species_comfGeo
DOMAIN
{(speId,
SAMPLING([cast(difTime(now(), 7 Days) AS Date),
cast(now() AS Date)])) | Species(speId)}
MAPPINGS
comfGeo(speId, obsTime):=
vectorize(loc)
OVER { SST(loc, obsTime2) ∧ obsTime = obsTime2 ∧
within(SST.temp(loc, obsTime2), Species.comfort(speId))}
process(vesId, zoneId, obsTime):= "ComfortZoneProcess"

In the above expression, aggregate function *vectorize*(loc) obtains the vector geometry that surrounds the set of sample locations loc. Mapping within(e, i) yields true if element e is within interval i.

6. CONCLUSIONS AND FURTHER WORK

A data model and data management framework has been proposed spatio-temporal analysis of data in data warehouses of spatial observation data. The approach consists of an E-R extension for observation data to be used at a conceptual level and a new logical level model that combines logical and functional paradigms. The advantages of the approach can be summarized as follows:

- General purpose observation data and metadata coexist with application specific Spatial Entities and Coverages, enabling efficient analysis over the whole set.
- Few primitive Mappings combined with general purpose logical and functional expressions enable the integrated management of any kind of spatial and spatio-temporal data. Besides, both data and analytical processing is unified under the well known mathematical concept of function.
- Parametric temporal and spatial types enable the user to have control over the precision and resolution of underlying time and space representations.
- Specific constructs for the specification of sampled and non-sampled domain components together with the absence of nested structures simplifies efficient implementation.

Further work is mainly related to efficient implementation structures and algorithms and the extension of the framework to deal with continuous queries on sensor data.

7. ACKNOWLEDGMENTS

This work has been partially supported by the Spanish Ministry of Science and Innovation (TIN2010-21246-C02-02).

8. REFERENCES

[1] A. Arasu, S. Babu, and J. Widom. The cql continuous query language: semantic foundations and query execution. *The VLDB Journal*, 15(2):121–142, June 2006.

[2] P. Baumann, A. Dehmel, P. Furtado, R. Ritsch, and N. Widmann. The multidimensional database system rasdaman. In *Proceedings of the 1998 ACM SIGMOD international conference on Management of data*, SIGMOD '98, pages 575–577, New York, NY, USA, 1998. ACM.

[3] P. Baumann and S. Holsten. A comparative analysis of array models for databases. In T.-h. Kim, H. Adeli, A. Cuzzocrea, T. Arslan, Y. Zhang, J. Ma, K.-i. Chung, S. Mariyam, and X. Song, editors, *Database Theory and Application, Bio-Science and Bio-Technology*, volume 258 of *Communications in Computer and Information Science*, pages 80–89. Springer Berlin Heidelberg, 2011.

[4] S. Bowers, J. Madin, and M. Schildhauer. A conceptual modeling framework for expressing observational data semantics. In Q. Li, S. Spaccapietra, E. Yu, and A. Oliv, editors, *Conceptual Modeling - ER 2008*, volume 5231 of *Lecture Notes in Computer Science*, pages 41–54. Springer Berlin Heidelberg, 2008.

[5] P. G. Brown. Overview of scidb: large scale array storage, processing and analysis. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, SIGMOD '10, pages 963–968, New York, NY, USA, 2010. ACM.

[6] J. a. P. Cerveira Cordeiro, G. Câmara, U. Moura De Freitas, and F. Almeida. Yet another map algebra. *Geoinformatica*, 13(2):183–202, June 2009.

[7] I. Galpin, C. Brenninkmeijer, A. Gray, F. Jabeen, A. Fernandes, and N. Paton. Snee: a query processor for wireless sensor networks. *Distributed and Parallel Databases*, 29(1-2):31–85, 2011.

[8] A. Gupta and I. S. Mumick. Materialized views. chapter Maintenance of materialized views: problems, techniques, and applications, pages 145–157. MIT Press, Cambridge, MA, USA, 1999.

[9] R. H. Güting, M. H. Böhlen, M. Erwig, C. S. Jensen, N. A. Lorentzos, M. Schneider, and M. Vazirgiannis. A foundation for representing and querying moving objects. *ACM Trans. Database Syst.*, 25(1):1–42, Mar. 2000.

[10] International Organization for Standardization (ISO). *Information technology – Database languages – SQL multimedia and application packages – Part 3: Spatial. ISO/IEC 13249-3:2011*, 2011.

[11] N. Jain, S. Mishra, A. Srinivasan, J. Gehrke, J. Widom, H. Balakrishnan, U. Çetintemel, M. Cherniack, R. Tibbetts, and S. Zdonik. Towards a streaming sql standard. *Proc. VLDB Endow.*, 1(2):1379–1390, Aug. 2008.

[12] M. Kersten, Y. Zhang, M. Ivanova, and N. Nes. Sciql, a query language for science applications. In *Proceedings of the EDBT/ICDT 2011 Workshop on Array Databases*, AD '11, pages 1–12, New York, NY, USA, 2011. ACM.

[13] S. R. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong. Tinydb: an acquisitional query processing system for sensor networks. *ACM Trans. Database Syst.*, 30(1):122–173, Mar. 2005.

[14] M. Neteler and H. Mitasova. *Open Source GIS: A GRASS GIS Approach. Third edition.* Springer, New York, USA, 2008.

[15] Open Geospatial Consortium (OGC). *OpenGIS Sensor Model Language (SensorML) Implementation Specification*, 2007. http://www.opengeospatial.org/standards/sensorml.

[16] Open Geospatial Consortium (OGC). *Geographic Information: Observations and Measurements. OGC Abstract Specification Topic 20*, 2010. http://www.opengeospatial.org/standards/om.

[17] C. Parent, S. Spaccapietra, and E. Zimányi. Spatio-temporal conceptual models: data structures + space + time. In *Proceedings of the 7th ACM international symposium on Advances in geographic information systems*, GIS '99, pages 26–33, New York, NY, USA, 1999. ACM.

[18] N. Tryfona, R. Price, and C. Jensen. Chapter 3: Conceptual models for spatio-temporal applications. In T. Sellis, M. Koubarakis, A. Frank, S. Grumbach, R. Güting, C. Jensen, N. Lorentzos, Y. Manolopoulos, E. Nardelli, B. Pernici, B. Theodoulidis, N. Tryfona, H.-J. Schek, and M. Scholl, editors, *Spatio-Temporal Databases*, volume 2520 of *Lecture Notes in Computer Science*, pages 79–116. Springer Berlin Heidelberg, 2003.

[19] A. Vaisman and E. Zimányi. A multidimensional model representing continuous fields in spatial data warehouses. In *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, GIS '09, pages 168–177, New York, NY, USA, 2009. ACM.

[20] J. Viqueira and N. Lorentzos. Sql extension for spatio-temporal data. *The VLDB Journal*, 16(2):179–200, 2007.

To trust, or not to trust: Highlighting the need for data provenance in mobile apps for smart cities*

Mikel Emaldi
Deusto Institute of Technology - DeustoTech
m.emaldi@deusto.es

Oscar Peña
Deusto Institute of Technology - DeustoTech
oscar.pena@deusto.es

Jon Lázaro
Deusto Institute of Technology - DeustoTech
jlazaro@deusto.es

Diego López-de-Ipiña
Deusto Institute of Technology - DeustoTech
dipina@deusto.es

Sacha Vanhecke
Ghent University - iMinds - Multimedia Lab
sacha.vanhecke@ugent.be

Erik Mannens
Ghent University - iMinds - Multimedia Lab
erik.mannens@ugent.be

ABSTRACT

The popularity of smartphones makes them the most suitable devices to ensure access to services provided by smart cities; furthermore, as one of the main features of the smart cities is the participation of the citizens in their governance, it is not unusual that these citizens generate and share their own data through their smartphones. But, how can we know if these data are reliable? How can identify if a given user and, consequently, the data generated by him/her, can be trusted? On this paper, we present how the IES Cities' platform integrates the PROV Data Model and the related PROV-O ontology, allowing the exchange of provenance information about user-generated data in the context of smart cities.

1. INTRODUCTION

According to the "Apps for Smart Cities Manifesto"[1], smart city applications could be sensible, connectable, accessible, ubiquitous, sociable, sharable and visible/augmented. It is not a coincidence that all of these features can be found in a standard smartphone: the popularity of these devices makes them the most suitable to ensure access to the services provided by smart cities. As one of the main features of the smart cities is the participation of the citizens in their governance, it is not unusual that these citizens generate and share their own data through their smartphones. Reviewing the literature, some examples of apps that deal with user generated data can be found, like Urbanopoly [4], Urbanmatch [5] or popular mobile apps related to the 311 service in cities like Calgary, Minneapolis, Baltimore or San Diego, all of them available in Google Play. The IES Cities project goes one step beyond, providing an entire architecture to foster the development of urban apps based on Linked Open Data[2] provided by government, through user-friendly JSON APIs. All of these works that manage user-generated data have the same worry about these data: are they reliable? How can we know if can a given user and, consequently, the data generated by him/her can be trusted? Recently, the W3C has created the PROV Data Model [14], for provenance interchange on the Web. This PROV Data Model describes the entities, activities and people involved in the creation of a piece of data, allowing the consumer to evaluate the reliability of the data based on the their provenance information. Furthermore, PROV was deliberately kept extensible, allowing various extended concepts and custom attributes to be used. For example, the Uncertainty Provenance (UP) [8] set of attributes can be used to model the uncertainty of data, aggregated from heterogeneously divided trusted and untrusted sources, or with varying confidence. On this paper, we present how IES Cities' platform integrates PROV Data Model and the related PROV-O ontology [13], allowing the exchange of provenance information about user-generated data in the context of smart cities. The final aim is to enrich the knowledge gathered about a city not only with government-provided or networked sensors' provided data, but also with high quality and trustable data coming from the citizens themselves.

The remaining of the paper is organized as follows: in Section 2 the current state of the art on apps that deal with user data in the context of smart cities is presented. Section 3 outlines the main concepts about IES Cities project. Sections 4 and 5 describe the semantic representation of the provenance through a use case and the metrics to calculate the reliability of the data, respectively. Finally, in Section 6 the conclusions and the future work are presented.

2. RELATED WORK

The following works can be highlighted regarding smart cities' mobile applications. Urbanopoly [4] presents an app

*This research is funded by project CIP-ICT-PSP-2012-6 "IES Cities: Internet Enabled Services for the Cities accross Europe", under "The Information and Communication Technologies Policy Support Programme". More info at http://ec.europa.eu/information_society/apps/projects/factsheet/index.cfm?project_ref=325097

[1] http://www.appsforsmartcities.com/?q=manifesto

[2] http://linkeddata.org/

for smartphones which combines Human Computation, *gamification* and Linked Open Data to verify, correct and gather data about tourism venues. To achieve this, Urbanopoly offers different games to the users, like quizzes, photo taking contests, etc. Similar to Urbanopoly, Urbanmatch [5] can be found, a game in which the user takes photos about some tourism venues, in order to be published as Linked Open Data by the system. Another work that uses Human Computation for movie-related data curation is Linked Movie Quiz[3]. In [3], the authors present *csxPOI*, an application that allows its users to *collaboratively create, share, and modify semantically annotated POIs*. These *semantic POIs* are modelled through a set of ontologies developed to fulfill this specific task; and published following the Linked Open Data principles. csxPOI allows users to create custom ontology classes, modelling new POI categories, and to establish subclass, superclass or equality relationships among them. In addition to create new classes, users can link these categories to concepts extracted from DBpedia[4]. In order to detect duplicate POIs, csxPOI clusters the available POIs with the aim of finding similarities among them.

As can be seen, the authors that work with user-generated Linked Open Data have to deal with duplication, missclasification, mismatching and data enrichment issues; and, as previously described, the end-user has arisen as the most important agent in smart cities' environments. In the next sections we explain how the IES Cities project uses the Provenance Data Model to represent provenance information about user-generated data.

3. IES CITIES

'IES Cities'[5], is the last iteration in a chain of inter-related projects promoting user-centric and user-provided mobile services that exploit both Open Data and user-supplied data in order to develop innovative services.

The project encourages the re-use of already deployed sensor networks in European cities and the existing Open Government related datasets. It envisages smartphones as both a sensors-full device and a browser with increasing computational capabilities which is carried by almost every citizen.

IES Cities' main contribution is to design and implement an open technological platform to encourage the development of Linked Open Data based services, which will be later consumed by mobile applications. This platform will be deployed in 4 different European cities: Zaragoza and Majadahonda (Spain), Bristol (United Kingdom), and Rovereto (Italy), providing citizens the opportunity to get the most out of their city's data.

Remarkably, IES Cities wants to analyse the impact that citizens may have on improving, extending and enriching the data these services will be based upon, as they will become leading actors of the new open data environment within the city. Nonetheless, the quality of the provided data may significantly vary from one citizen to another, not to mention the possibility of someone's interest in populating the system with fake data.

Thus, the need for evaluating the value and trust of the user contributed data requires the inclusion of a validation module [12]. In other words, we should be able to express special meta-information about the data submitted by IES Cities' users. The idea that a single way of representing and collecting provenance could be internally adopted by all systems does not seem to be realistic today, so the actual approaches modelling their provenance information into a core data model, and applications that need to make sense of provenance information can then import it, process it, and reason over it [6].

In addition, when considering user-provided data measures for data consolidation have to be considered. Contributions from one user have to be cross-validated with contributions from other users in order to avoid information duplication and foster validation of others' data. Thus, data contributions from different users presenting spatial, linguistic and semantic similarity should be clustered [2]. Before a user contributes with new data, other user's contributions at nearby locations should be shown to avoid recreating already existing data and encourage additions and enhancements to be applied to the existing data. After contributing with new data, the data providing user should be presented with earlier submitted similar contributions both in terms of contents and location in order to confirm whether their new contribution is actually a new contribution or it is amending an earlier existing one. In essence, aids before and after editing new entries have to be provided and a two phase commit process for user provided data should be put in place to ensure that contents of the highest quality are always added. Future work in IES Cities will tackle these issues by providing REST interfaces to invoke services for clustering data entries and to retrieving related entries associated to a given one.

4. SEMANTIC REPRESENTATION OF PROVENANCE

To illustrate the semantic representation of trust and provenance data through the Provenance Ontology, a use case is presented: 311 Bilbao. This app uses Linked Open Data to get an overview of reports addressing faults in public infrastructures. From the data owner's point of view, the enrichment of datasets carried out by third parties (such as users of the 311 Bilbao app), revealed two problems: 1) the fact that data does not need to be approved before being published and that there is no mechanism to control the amount of data a citizen can add and 2) there is still the need for a way to differentiate the default trustworthiness of the different authors such as citizens and city council's staff. The following code represents the provenance of a user-generated report[6]:

```
@prefix foaf: <http://xmlns.com/foaf/0.1/> .
@prefix prov: <http://www.w3.org/ns/prov#> .
@prefix iesc: <http://studwww.ugent.be/~satvheck/IES/
schemas/iescities.owl> .
@prefix up:   <http://users.ugent.be/~tdenies/up/> .
@prefix :     <http://bilbao.iescities.org#> .

entity(:report_23456, [ prov:value="The paper bin is
broken" ])
wasGeneratedBy(:report_23456, :reportActivity_23456)
wasAttributedTo(:report_23456, :jdoe)
wasInvalidatedBy(:report_23456, :invActivity_639,
2013-07-22T03:05:03)

```

[3]http://lamboratory.com/hacks/ldmq/

[4]http://dbpedia.org

[5]http://iescities.eu

[6]The provenance data is represented using Provenance Notation (PROV-N). More information at http://www.w3.org/TR/prov-n/

```
15 activity(:reportActivity_23456, 2013-07-22T01:01:01,
16 2013-07-22T01:05:03)
17 wasAssociatedWith(:reportActivity_23456, :jdoe)
18
19 agent(:jdoe, [ prov:type='prov:Person', foaf:name=
20 "John Doe", foaf:mbox='<mailto:jdoe@example.org>' ])
21
22 entity(:report_23457, [ prov:value="It is incorrect,
23 another paper bin has replaced the old one, but 2
24 meters beyond" ])
25 wasAttributedTo(:report_23457, :jane)
26 wasDerivedFrom(:report_23457, :report_23456,
27 :invActivity_639, -, -, [ prov:type='prov:Revision' ])
28
29 activity(:invActivity_639, 2013-07-22T02:58:01,
30 2013-07-22T03:04:47)
31 wasAssociatedWith(:invActivity_639, :jane)
32
33 agent(:jane, [ prov:type='prov:Person', foaf:name=
34 "Jane", foaf:mbox='<mailto:jane@bilbao.iescities.org>'
35  ])
36 actedOnBehalfOf(:jane, :bilbao_city_council)
37
38 agent(:bilbao_city_concil, [ prov:type=
39 'prov:Organization', foaf:name="Bilbao City Council"
40 ])
```

On this piece of semantic information the `:report_23456` resource represents the report made by the user. This report is identified by its own and unique URL and provides information about the user that has made it and which activity that has generated this report (lines 8-13). The `:reportActivity_23456` shows details about the activity that generated the report, like when the user started reporting the issue and when it ended. At line 19 the information about "John Doe", the user that reported the fault, can be seen. In the example given, another user, Jane (lines 33-36), has revised the report made by John (lines 22-31). As the `actedOnBehalfOf` asserts, Jane is some kind of municipal worker of Bilbao City Council (line 38). As Jane's report has more authority agains John's report, John's report is invalidated as `wasInvalidatedBy` asserts. Allowing the semantic descriptions of the provenance of the reports made at 311 Bilbao app, the data generated by a concrete user can be reached through SPARQL [15] language queries.

5. PROVENANCE BASED RELIABILITY

There exist some approaches on how to calculate trust in semantic web using provenance information. IWTrust [16] uses provenance in the trust component of an answering engine, in which a trust value for answers is measured based on the trust in sources and in users. In [10] provenance data is used to evaluate the reliability of users based on trust relationships within a social network. [11] presents an assesment method for evaluating the quality of data on the Web using provenance graphs, and provides a way to calculate trust values based on timeliness. In [7] the authors propose generic procedures for computing reputation and trust assessments based on provenance information.

In [9] the authors identify 19 parameters that affect how users determine trust in content provided by web information sources, such as the authority of the creator of the information or the popularity and recency of that information, among others. Based on these factors, we have built a generic model for the measurement of a trust value in the context of IES Cities, in which the trust according to each factor is calculated independently:

$$trust(report) = \frac{\sum_{p=[auth,agree...]}^{n} \alpha_p * trust_p(report)}{n} \quad (1)$$

where p is the measured property and n is the total number of measured properties. α is a value between 0 and 1 to denote the relevance of this property, making the measure based on a certain property more or less relevant. $trust_p$ is a function that returns a value between 0 and 1 determining the trust of a given report according to a certain property.

Both the α values and the $trust_p$ functions can be defined by the developers using IES Cities platform, because both of them are dependant on the context and the need of the application domain.

To clarify, we are using this model in the 311 Bilbao use case. To that end, we have selected the most relevant trust-properties concerning our use case:

Authority: It refers to the fact that if a resource is created by an authority in a given context, this information is more reliable. For our use case a basic function like the following can be used:

$$trust_{authority} = \begin{cases} 0 & if user \neq authority \\ 1 & if user = authority \end{cases} \quad (2)$$

in which being authority can be checked with a SPARQL ASK query:

```
PREFIX prov: <http://www.w3.org/ns/prov#>
ASK { :jane prov:actedOnBehalfOf :bilbao_city_concil }
```

Popularity: The number of references and uses of a piece of information is a key aspect to determine its trust. In the case of 311 Bilbao we measure the popularity of a report based on the number of visits that the report receives, with the following formula:

$$trust_{popularity} = \frac{visits_{report}}{visits_{open\ reports}} \quad (3)$$

in which the number of visits of the report is normalized with the number of overall visits of opened reports at the moment.

Recommendation: Recommendation refers to importance that the ratings that other users gives to a given resource has in its trust. The function to measure the relevance of user ratings can be as sophisticated as the developer wants, but for our case we have selected a very naive and simple one, in which other users can vote the reports with +1 / -1 buttons and the trust value is calculed with this formula:

$$trust_{recommendation} = \frac{positive\ votes_{report}}{total\ votes_{report}} \quad (4)$$

Provenance / Reputation: In this case, *provenance* refers to the trust that the entities responsible for generating a piece of information may transfer information itself. A key aspect to measure the trust in a publisher is the reputation. There exist many approaches to measure the reputation of a user; some of them measure the reputation based on trust relationships between users [10], while some others like [7] are based the historical evidence of each user. For the our use case, we propose using the three-step procedure presented in [7]. In the 'evidence selection' step every report made by a given user are retrieved, in the 'evidence weighting' step the *recommendation* trust function is executed for every report, and in the last step all these trust values are aggregated through subjective logic to get the trustworthiness of a given user.

Recency / Timeliness: Timeliness can be defined as the the up-to-date degree of a data item in relation with the task at hand. We propose and adaption of [11] formula to measure timeliness, based on the work described in [1]:

$$trust_{authority} = (max(1 - \frac{currency}{volatility}), 0)^{sensitivity} \quad (5)$$

where *currency* is the difference between the time data is presented to the user and the time it was reported to the system. *Volatility* refers to the maximum amount of time a given report time should be active (for example, if a broken street lamp is reported, it should be repaired within a month at most), and *sensitivity* may change its value by observing the updates made over the status of the report: it would adopt a high value for data being constantly updated, and a low value for data that does not change often.

Other trust factors: Apart from the aspects identified in [9], the model is flexible enough to include other factors affecting the trust. In the case of 311 Bilbao mobile app, the geographical distance could be a key aspect of the truth, as reports talking about events happening near to where the user sends the report would be more reliable.

$$trust_{distance} = \frac{1}{geodistance(loc_{report}, loc_{reportedplace})} \quad (6)$$

The function for the calculus of the geographical distance has as input the geographic coordinates of the report, retrieved from the smartphone GPS sensor, and the geographic coordinates of reported place, obtained with geolocation services like Nominatim[7].

After applying our model we will get a trust value between 0 and 1, that could be inserted in the provenance graph with a triple, assuming the confidence level was '0.6', like *:report_23456 up:contentConfidence '0.6'* [8].

6. CONCLUSIONS AND FUTURE WORK

The proposed approach in this article will allow to evaluate the provenance of user-submitted data in IES Cities' platform. The metrics proposed will measure data trustworthiness level, providing an extra confidence layer in the project's framework. City council staff and platform administrators will be able to query data quality through SPARQL queries, retrieving only those results with a confidence level above a parameterised threshold.

The evaluation and validation of the proposed metrics against other implementations following the PROV-O ontology will be left for a future iteration on IES Cities, aggregating other significant metrics should they improve the provenance of the generated data.

7. REFERENCES

[1] D. Ballou, R. Wang, H. Pazer, and G. K. Tayi. Modeling information manufacturing systems to determine information product quality. pages 462–484, 1998.

[2] M. Braun, A. Scherp, and S. Staab. Collaborative semantic points of interests. In *The Semantic Web: Research and Applications*, page 365–369. Springer, 2010.

[3] M. Braun, A. Scherp, S. Staab, et al. Collaborative creation of semantic points of interest as linked data on the mobile phone. 2007.

[4] I. Celino, D. Cerizza, S. Contessa, M. Corubolo, D. DellAglio, E. D. Valle, and S. Fumeo. Urbanopoly – a social and location-based game with a purpose to crowdsource your urban data. In *Privacy, Security, Risk and Trust*, page 910–913, Amsterdam, 2012.

[5] I. Celino, S. Contessa, M. Corubolo, D. Dell'Aglio, E. D. Valle, S. Fumeo, and T. Krüger. UrbanMatch - linking and improving smart cities data. In C. Bizer, T. Heath, T. Berners-Lee, and M. Hausenblas, editors, *Linked Data on the Web*, volume 937 of *CEUR Workshop Proceedings*. CEUR-WS, 2012.

[6] D. Ceolin, P. T. Groth, W. R. van Hage, A. Nottamkandath, and W. Fokkink. Trust evaluation through user reputation and provenance analysis. In *Uncertainty Reasoning for the Semantic Web*, volume 900, page 15–26. CEUR-WS, 2012.

[7] D. Ceolin, P. T. Groth, W. R. van Hage, A. Nottamkandath, and W. Fokkink. Trust evaluation through user reputation and provenance analysis. In F. Bobillo, R. N. Carvalho, P. C. G. da Costa, C. d'Amato, N. Fanizzi, K. B. Laskey, K. J. Laskey, T. Lukasiewicz, T. Martin, M. Nickles, and M. Pool, editors, *URSW*, volume 900 of *CEUR Workshop Proceedings*, pages 15–26. CEUR-WS.org, 2012.

[8] T. De Nies, S. Coppens, E. Mannens, and R. Van de Walle. Modeling uncertain provenance and provenance of uncertainty in W3C PROV. In *International World Wide Web Conference*, page 167–168, Rio de Janeiro, Brazil, 2013.

[9] Y. Gil and D. Artz. Towards content trust of web resources. *Web Semantics: Science, Services and Agents on the World Wide Web*, 5(4):227–239, 2007.

[10] J. Golbeck. Combining provenance with trust in social networks for semantic web content filtering. In *Provenance and Annotation of Data*, pages 101–108. Springer, 2006.

[11] O. Hartig and J. Zhao. Using web data provenance for quality assessment. In *In: Proc. of the Workshop on Semantic Web and Provenance Management at ISWC*, 2009.

[12] O. Hartig and J. Zhao. Publishing and consuming provenance metadata on the web of linked data. In D. L. McGuinness, J. R. Michaelis, and L. Moreau, editors, *Provenance and Annotation of Data and Processes*, number 6378 in Lecture Notes in Computer Science, pages 78–90. Springer Berlin Heidelberg, Jan. 2010.

[13] T. Lebo, S. Sahoo, D. McGuinness, K. Belhajjame, J. Cheney, D. Corsar, D. Garijo, S. Soiland-Reyes, S. Zednik, and J. Zhao. Prov-o: The prov ontology. *W3C Recommendation, http://www. w3. org/TR/prov-o/(accessed 30 Apr 2013)*, 2013.

[14] L. Moreau, P. Missier, K. Belhajjame, R. B'Far, J. Cheney, S. Coppens, S. Cresswell, Y. Gil, P. Groth, G. Klyne, et al. Prov-dm: The prov data model. *Candidate Recommendation*, 2012.

[15] E. Prud'hommeaux and A. Seaborne. SPARQL query language for RDF, 2008.

[16] I. Zaihrayeu, P. P. Da Silva, and D. L. McGuinness. Iwtrust: Improving user trust in answers from the web. In *Trust Management*, pages 384–392. Springer, 2005.

[7] `http://wiki.openstreetmap.org/wiki/Nominatim`