

A Benchmark for Context Data Management in Mobile Context-Aware Applications*

Nadine Fröhlich Thorsten Möller Steven Rose Heiko Schuld
Databases and Information Systems Research Group
University of Basel, Switzerland
firstname.lastname@unibas.ch

ABSTRACT

Over the last few years, computational power, storage capacity, and sensing capabilities of mobile devices have significantly improved. As a consequence, they have undergone a rapid development from pure telecommunication devices to small and ubiquitous computing platforms. Most importantly, these devices are able to host context-aware applications, i.e., applications that are automatically adjusted to the current context of their user. This, in turn, requires sensing support on the device, the possibility to store context information, and to efficiently access this context information for the automated adaptation of applications. In this paper, we introduce a benchmark for context management in mobile context-aware applications. We present in detail the design and setup of the benchmark, based on an eHealth use case. The benchmark evaluation considers context queries on Android Nexus One cell phones and compares the performance of different settings including relational and object-oriented databases on the mobile device, and an RDF triple store on a stationary computer. The results show significant differences in the settings that have been evaluated and are thus valuable indicators for database selection and system design for mobile context-aware applications.

1. INTRODUCTION

Over the last few years, mobile devices have undergone a rapid metamorphosis from pure telecommunication devices to small and ubiquitous computing platforms. This is the result of a multitude of technical developments: i.) new types of powerful and energy-efficient processors; ii.) the significant increase in local storage capacity due to inexpensive, low latency flash memory cards; iii.) sophisticated sensing capabilities already embedded into off-the shelf devices (e.g., GPS sensors or acceleration meters). As a result of these developments, mobile devices are more and more in the focus of novel kinds of applications that aim at improving the way

*This work has been partly funded by the Hasler Foundation

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, to post on servers or to redistribute to lists, requires a fee and/or special permission from the publisher, ACM.

PersDB '10, September 13, 2010, Singapore
Copyright 2010 VLDB Endowment, ACM 000-0-00000-000-0/00/00.

users access data and information. A particular emphasis has been put in recent years on *context-aware applications*. These applications aim at automatically adapting the way information is accessed, processed, and/or presented to the current needs of their user, based on their context (e.g., preferences, location, devices at hand). As users are mobile, their context may rapidly change over time. Hence, in contrast to traditional earmarked applications running on stationary devices with users whose context is rather static, mobile applications are characterized by the intrinsic high frequency in which the users' context changes. For the management of sensed context and the exploitation of context for automated adaptation, e.g., by a rule engine, efficient and effective context data management is needed.

Motivation: eHealth Use Case. Consider, for instance, information access in a hospital. On a ward round, the doctor's time should mainly be spent for the interaction with their patients, rather than for searching in the clinical information system. For example, retrieving electronic patient records should be unobtrusive and efficient. When a physician enters a sick room, her mobile phone should automatically display the medical history of the patient(s) in this room. Annotations to the health record added via the device will be synchronized with the underlying clinical information system. In case she wants to share medical images with her patient for which the mobile device's display is too small, the images should be automatically transferred to the patient's bed-mounted multimedia device.

Challenges and Contribution. For the type of dynamic adaptations described before, the current context of a user needs to be sensed and stored for immediate and/or later exploitation. The amount of context data to store can vary significantly, depending on the frequency in which the user's context changes and the frequency in which it is sensed. Although today's mobile devices are much more powerful compared to devices as of some years ago, resources are still limited. Therefore, databases for storing and retrieving context data need to be as efficient as possible. We have designed a benchmark tailored to context data management for mobile applications (aligned to the eHealth use case presented above) in order to identify which data store is best suited for this task. This benchmark has been implemented using three different kinds of open source databases: i.) relational databases locally on the device and on a remote stationary server ii.) an RDF triple store on a remote server, and iii.) an object-oriented database locally on the device and on a

remote stationary server. We could only run the triple store remotely, with updates and queries submitted from a mobile device, as there is, to the best of our knowledge, currently no stable open source RDF-based implementation directly running on mobile devices. However, this configuration is very relevant, since context information will be subject to reasoning. To factor out communication costs when interacting with a remote server, we also ran the benchmark on the object-database and two relational databases remotely.

For the evaluation we used the context data model designed for the LoCa project [7], implemented for each of the data stores. Furthermore, we implemented a data generator to populate all databases with the same context data. The benchmark queries have been tailored to the different data models and schemas while keeping the original semantics as defined in the benchmark set-up.

The paper is organized as follows: Section 2 reviews related work. The benchmark setup is discussed in Section 3. The implementation of the benchmark is described in Section 4 and Section 5 presents and discusses the evaluation results. Section 6 concludes.

2. RELATED WORK

Existing OLTP and OLAP-style database benchmarks usually address data center-scales rather than mobile environments. The TPC-C and TPC-E Benchmarks [21] give preference to a relational data model, SQL as the query language, and aim at enterprise-scale database systems. Likewise, the Berlin SPARQL Benchmark [3] focusses on RDF triple stores, and the LUBM Benchmark [9] evaluates OWL knowledge base systems, both bound to SPARQL as the query language. The Pole Position [15] benchmark compares relational and object-oriented database engines and object-relational (O/R) mapping implementations, but not in a mobile setting.

In general, benchmarks for mobile environments are rare. There exists an open source benchmark [11] designed for the Android platform which compares the object-oriented database Perst [14] with SQLite. However, the evaluations are too limited to assess the overall performance characteristics of these systems. All queries access only one relation without joins, aggregations, subselects, etc. Therefore, they cannot be directly applied to the kind of schema and queries that have to be considered for context-awareness.

Our benchmark compares the performance of different approaches for storing context data. In [20], context models are compared w.r.t. simplicity and flexibility, while we focus in our work explicitly on performance aspects. [2] identifies requirements on context stores, especially in terms of historical context data, which are considered in our benchmark.

3. BENCHMARK DESIGN

In what follows, we introduce the data model designed for the benchmark and the query mix to be considered (more details on the benchmark specification and evaluation can be found in [8]).

Benchmark Data Model. For the benchmark, we exploit the LoCa data model for context data which is defined in accordance with the most commonly used definition for context (Dey et al. [6]). The LoCa context model is introduced in detail in [7] and summarized in Figure 1 and Table 1.

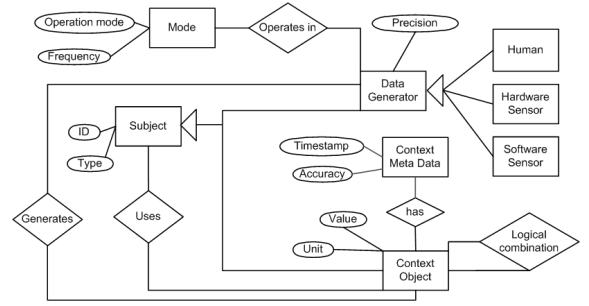


Figure 1: LoCa Context Model

Subject: specifies for whom the context is determined, e.g., the owner of a mobile phone
ContextObject: The context of the subject, e.g., the location of a physician. Context objects have a value, e.g., current GPS coordinates, which have a unit, e.g., degree (for longitude and latitude).
DataGenerator: the sensor that senses context data. Besides hardware (e.g., GPS) and software sensors (e.g., diary) we include manual human input.
Mode: includes metadata on DataGenerators (e.g., in which frequency a generator produces elements of a data stream.
ContextMetaData: metadata on ContextObjects (e.g., timestamp of sensed data, or accuracy of context data).
LogicalCombination: ContextObjects can recursively consist of ContextObjects.

Table 1: Description of the LoCA Context Model

Benchmark Test Load. The objective for the specification of the benchmark test load was to be as realistic as possible (in terms of the different types of queries, their mix, and the volumes of data), without giving preference to a very specific application. For this, we have analyzed typical eHealth environments (information access of physicians in the course of a day, based on statistical information of a medium-sized hospital) and have generalized these findings for the specification of our benchmark. Based on these numbers we estimated the amount of context data to be collected in the course of one year (see Table 2 column 3). On a mobile device, only context data of one user are stored so we scaled the data of column 3 down to one person (see Table 2 column 2). In order to make the benchmark as flexible as possible, we consider three settings with context data of i.) a year, ii.) a quarter, and iii.) a week (in the latter two cases, the numbers from Table 2 are scaled down accordingly).

3.1 Benchmark Queries

The benchmark queries we have defined reflect the way mobile users and context-aware applications interact with a context database. The query mix considers mainly read accesses to context data, as well as insert operations which create new context objects as the user’s context evolves.

Q1: Return a subject by a given ID. Such queries are often used, for instance for the identification of a physician.

Q2: Return the last recorded context object of a given type for a subject, e.g., the last blood pressure value.

*Cardinality changes over time.

Entity	data sets per person & year	data sets medium-size hospital/year
DataGenerator	851	15.500
Human	1	5.000
HardwareSensor	833	25.000
SoftwareSensor	17	500
Mode	851	5.000
ContextObject	185.000*	925.000.000*
ContextMetaData	185.000*	925.000.000*
LogicalCombination	185*	925.000*
Subject	1.800	27.000 (patients)

Table 2: Cardinalities of Entities in Benchmark

Query	Percentage	Query	Percentage
Query 1	10 %	Query 7	7 %
Query 2	7 %	Query 8	6 %
Query 3	7 %	Query 9	4 %
Query 4	6 %	Query 10	36 %
Query 5	3 %	Query 11	0.5 %
Query 6	13 %	Query 12	0.5 %

Table 3: Benchmark Query Mix

- Q3: Return the context object of a given type for a subject in a given time interval (day, week, month).
- Q4: Return the last recorded context object of a given type for a subject, generated by a given generator. It considers details on the sensor, e.g., to find out whether a blood pressure meter has the desired precision.
- Q5: Return the available data generators including type and precision that generate context objects of a given type for a subject (e.g., for the selection of a generator for a particular application).
- Q6: Return a subject of a given type with the same context object as a given subject, for instance when searching for patients belonging to a sick room or for devices available in the sick room in which the physician is currently situated.
- Q7: Return all available types of context objects to a given subject in alphabetical order.
- Q8: Return all logical combinations of context data for a given subject. The query result shows which context data are part of other context data.
- Q9: Return the number of data generators that generate context data per subject. By this query one is able to control the generators belonging to one subject.
- Q10: Insert a context object. This operation reflects all context change of a user, as sensed by the device.
- Q11: Update all context objects belonging to one logical combination.
- Q12: Delete a context object, e.g., for correcting mistakes of human data generators.

The order of the queries in the mix will be randomly chosen, but their occurrence over a longer interval is determined by the percentages given in Table 3. It should be noted that all queries, except for no. 10 (creation of context object) run sequentially, while the insertion is done automatically in the background, in parallel to the dynamic adaptation.

3.2 Performance Metrics

For the benchmark evaluation, we consider the following metrics, which are aligned with the Berlin benchmark [3].

Metrics for Single Queries

- Average Query Execution Time (aQET_x): Average time for executing an individual query of type *x* ten times with different parameters against the system under test (SUT).
- Minimum/maximum Query Execution Time (minQET_x, maxQET_x): A lower and upper bound execution time for queries of type *x*.
- Queries per Second (QpS_x): Average amount of queries of type *x* that were executed per second. This value is computed from the aQET_x values.

Metrics for Query Mixes

Overall Runtime (oaRT): Overall time it took the test driver to execute a certain amount of queries following the distribution in the mix against the SUT. Thereby, inserts are running in a parallel thread on the device. We decided to process three runs each with 300 queries each.

4. BENCHMARK SETUP

In this section, we describe in detail the setup of our benchmark evaluation.

4.1 Data Stores

We have selected three different kinds of data stores for the benchmark. All are well tested and stable open source systems that are widely adopted in practice.

- *Relational databases*
 - H2 (v. 1.2.136) in embedded mode and remotely
 - MySQL (v. 5.0.51a) running remotely
 - SQLite (version 3.5.9) [19] in embedded mode
- *RDF/OWL triple store* Sesame (v. 2.3.1) [16] with Storage and Inference Layer SwiftOWLIM (3.0 beta 12) [13] remotely (servlet, deployed into apache-tomcat-6.0.24).
- *Object-oriented database* db4o (v. 7.12) [4] in embedded mode and running remotely

Initially, we have chosen only SQLite, Sesame, and db4o for the benchmark. However, Sesame/SwiftOWLIM currently does not run on smart phones and SQLite cannot be used in a client/server (c/s) mode. Therefore, we added H2 to our benchmark as it supports both embedded and c/s mode. This makes results for remote operation comparable among the triple store and the relational databases. Finally, we also evaluated MySQL in c/s mode.

The relational database SQLite demands zero configuration, has a small system footprint and no external dependencies to other libraries which makes it highly appropriate for mobile devices. Furthermore, it is currently used in several applications (e.g., Firefox, Google tools). H2 [10] is an open source Java database also having a small system footprint (about 1 MB) and low memory requirements. It supports disk-based and in-memory databases in embedded and server mode. MySQL [12] is a highly popular and very widely adopted open source database, featuring a rich set of functionality.

As the LoCa approach considers the semantics of services when deciding on dynamic adaptations, we have also chosen a triple store for the benchmark. Out of the available open source RDF triple stores we have evaluated, Sesame has been the most promising tool as it is already widely used [16] and has proven to perform well enough for our purpose [3]. Sesame allows for easy extension via the so-called SAIL-API. As the back-end store implementations shipped with Sesame do not provide the OWL based reasoning we would like to use in our model, we used an alternative third party implementation called SwiftOWLIM, an in-memory triple store which implements rule-based forward-chaining strategy for inferencing and supports a subset of OWL DL.

The object-oriented database db4o has been chosen to avoid the impedance mismatch as the LoCa system is completely implemented in Java.

All tests were performed on Nexus One (N1) mobile phones using Google Android version 2.1 [1] platform. The devices come with 512MB RAM, 512MB ROM, and a Qualcomm QSD8250 CPU with 1GHz. When working in remote mode, instances of Db4o, MySQL, and Sesame/SwiftOWLIM ran on a standard server (Intel Core 2 - 6600 2.4GHz, 4GB RAM, 250GB SATA Seagate ST3250820AS, Ubuntu 9.04 x86_64). Mobile phones were connected to a 802.11b/g 54MBit wireless access point (DLINK DIR615), to which the server was connected via 100MBit Ethernet.

4.2 Data Generator

To supply the different types of databases with equivalent datasets, we implemented a parameterizable data generator that incorporates output modules for every target paradigm. In the first step, the raw data is generated in-memory using an integrated object-oriented data model according to the context data model presented in Figure 1. The actual data values are generated by exchangeable value generators, all relations are picked at random (evenly distributed). In the second step, this data is then transformed into the target data formats by special output modules. An intermediate raw format is stored to be able to create additional outputs later, that are equivalent to the formerly generated ones. Additionally, every output module creates equal sets of query parameter values that are used throughout the benchmark run. This ensures a maximum of comparability between the different platforms.

4.3 Implementation of the Data Model

The schemas for the different data stores are created according to the LoCa context data model. For the relational databases, a standard transformation has been applied. The inheritance is implemented using a horizontal partitioning to avoid joins and decrease execution time.

For the transformation to RDF, an OWL DL ontology has been defined¹. Entities are directly mapped to concepts (classes), relations among entities are mapped to either one or two object roles (properties) depending on whether both directions can be navigated, and attributes to data roles. Furthermore, we exploited additional modeling expressivity available in OWL: cardinality restriction constructors, disjointness of concepts, and functional, transitive, and irreflexive properties whenever appropriate.

For db4o, we transformed the context model to a Java class structure by mapping entities to classes and by us-

¹ Available via <http://on.cs.unibas.ch/owl/1.0/Context.owl>

ing collections for the relationships. To work in an object oriented manner, we applied the composite pattern to implement the inheritance hierarchy.

4.4 Implementation of Benchmark Queries

All benchmark queries have been formulated in the query languages supported by the respective data stores. For relational databases we used SQL, for db4o SODA, and for Sesame SPARQL.

In case of H2 and MySQL, we used suitable JDBC drivers to access the databases. SQLite comes as built-in embedded database with its own API that offers different ways to access the database, raw query (nearly plain SQL), and a structured interface for users with little SQL knowledge. For the benchmark, we used the raw query interface.

SPARQL has evolved as the de facto standard for querying RDF graph data. At the moment, there exist two major versions of SPARQL that are specified by the W3C. SPARQL 1.0 [17] has matured to a W3C recommendation and contains basic graph query constructs. At the time of writing SPARQL 1.1 [18] is still a working draft. It adds more advanced features like aggregations and data manipulation to the language. As our chosen triple-store only supports SPARQL 1.0, queries that change the dataset were difficult to handle. While query 12 can be easily implemented using a Sesame specific type of request ('transaction'), query 11 would have required to programmatically perform needed updates on the datasets. Thus, we chose to drop this query in the triple-store implementation of our benchmark. Also aggregating queries had to be implemented partly on the client side. To access the RESTful SPARQL interface that SESAME provides, we used the HttpClient implementation integrated into the Android SDK. All result data was encoded as JSON and parsed using androids own parser implementation.

Db4o supports at its APIs query by example (QbE), Native Queries (NQ) and SODA. We decided to use SODA because it is, according to the db4o documentation [5], up to two times faster than optimized NQ and five times faster than unoptimized NQ — the reason for this is that SODA is the underlying internal query API all the other query APIs are mapped to.

5. BENCHMARK RESULTS

This section summarizes the evaluation results for our benchmark.




Figure 2: Overview of Query Mixes

System	Week	Quarter	Year
H2 emb	26 (0.9)	1300 (23.8)	23251 (267.3)
H2 c/s	17 (0.9)	38 (23.8)	251 (267.3)
MySQL c/s	4 (0.8)	29 (18.8)	418 (186.0)
SQLite emb	386 (0.8)	– (17.5)	– (172.9)
SQLite emb (r)	384 (0.8)	– (17.5)	– (172.9)
OWLIM c/s	18 (1.6)	1025 (38.1)	– (384.1)
db4o emb	511 (1.5)	21100 (22.3)	– (126.5)
db4o c/s	– (1.5)	– (22.3)	– (126.5)

Table 4: Execution Times of Query Mixes in Seconds (in parentheses: Database Sizes in MB)

Query Mixes. Figure 2 and Table 4 show the results of the query mix evaluation in the different settings. The execution times of the relational databases vary considerably but except for SQLite, all query mixes could be processed. Only for H2 emb, we had to increase the JVM heap size to 80 MB which is however a rather unrealistic setting for the device. SQLite was comparably slow. H2 c/s and MySQL c/s manage to execute the query mix for one year in nearly the same time as SQLite needs for the query mix for one week (SQLite week: 386s, H2 year: 251s, MySQL year: 417s). This is most probably due to the lack of a sophisticated query optimizer for SQLite. Furthermore, H2 c/s proves to better perform on the largest data set (year), but MySQL c/s performs much better on the smaller data sets (week/quarter). As expected, the c/s architectures perform better than the embedded databases on our resource limited mobile devices. For SQLite, we have compared a setting with referential integrity with a setting without. The benchmarks results show only marginal differences. OWLIM c/s shows only an average performance, only slightly better than H2 emb – however, it is not able to process a query mix on the large data set (year). For db4o we could not manage to process a complete query mix. With db4o c/s, query 4 could not be executed although the query ran fine in the db4o embedded version – this seems to be the result of a marshaling/unmarshaling problem in db4o. The db4o emb version is slow, and needs a lot of space and JVM heap.

Single Queries. To find out why the query mixes showed significant performance variations on different systems, we have also analyzed the average execution time of frequently occurring queries. Query 6 has a fraction of 13% in the query mix. For this query, the triple store performs much better than most relational databases and even db4o c/s is for the large dataset (year) in the range of MySQL (see Figure 3).

Query 1 has also a major influence on the query mix (10%). SQLite performed best for this query while db4o emb performed comparably bad. db4o c/s performed better but could not reach the times of the relational databases; its performance is comparable to OWLIM (see Figure 4).

Query 9 is one of the most controversial queries in the benchmark since some systems have not been able to properly execute it in reasonable time (see Figure 5). This query runs fine on H2 c/s and MySQL but lasted extremely long on SQLite. One reason for this is might be the less elaborated optimizer of SQLite. H2 could perform this query but we needed to increase the JVM heap to 64 MB. OWLIM performs fair on small data sets but gets very slow with increased data sets.




Figure 3: Single Query Evaluation for Query 6




Figure 4: Single Query Evaluation for Query 1

Summary Results. The analysis of the performance for individual queries shows that the chosen systems have different strengths. A system that is very slow for one query can be fast for another query. For deciding what system is the best we have to care about the overall performance shown in the query mixes where the fraction of queries is considered. In the mixes, read-only queries (without Query 10) are running in parallel to inserts (Query 10). The execution of inserts usually lasts a few milliseconds (2 to 35 ms), except for OWLIM where the execution of inserts lasts more than 10000 milliseconds. The performance in c/s mode was mostly considerably better than in embedded mode. H2 needed for the query mix with the context data accumulated in one year 23251s while H2 c/s needed 251s. For db4o we could not run a complete query mix. But when comparing the single queries it is obvious that db4o c/s is mostly faster




Figure 5: Single Query Evaluation for Query 9

than db4o embedded. Query 9 seems to be an outlier as it performs better when executed in embedded mode than in c/s mode (see Figure 5). Although additional communication cost incur in c/s mode, powerful server-side hardware can partly or completely compensate this. Only for db4o, communication costs are in some cases higher than the benefit of more powerful servers. Finally, the benchmark shows that it is not feasible to store context data for a longer period in an embedded database on a mobile device.

Lessons Learned. *SQLite* is small and easy to use, but its optimizer is not as developed as that of more mature database systems. We had to manually optimize some of the queries to improve the query performance.

H2 supports referential integrity by default and works well in embedded and client/server setting. Most notable, *H2* caused least configuration problems.

MySQL, as the data store with the richest set of features, is also the most demanding system regarding configuration for remote access.

OWLIM, Sesame's built-in data store, can be easily created and maintained using the web based workbench. However, manual configuration was needed to create repositories working with an *OWLIM* store after the binaries were integrated into the Servlet distribution. As we chose the RESTful HTTP interface to access the stores, some conveniences of common database access layers had to be implemented on top of the standard http client, e.g., to circumvent memory restrictions when large result sets were received. This was the case for Query 8, where a lacking aggregation feature of SPARQL 1.0 had to be implemented on the client side.

db4o was comparatively slow, especially for Query 9 and it demanded much memory and JVM heap space (up to 64M/80M). For speeding up the tests, we used the SODA query language that is much faster than the recommended query interface (NQ). Actually, the performance of *db4o* also strongly depends on the design of the data model as inheritance and collections with a lot of associated objects slow down the query execution dramatically. We optimized our model, but in order to reach an appropriate speed we would have to completely redesign our model, especially by abandoning the object-oriented design (inheritance). In c/s mode, one query (no. 3), did not run at all, although we did not experience any problems in embedded mode. Although *db4o* performs well on a standard PC in c/s mode, we were faced with major problems regarding performance and stability in embedded mode on the mobile device. These problems are well taken by the *db4o* developers as they are currently aiming at decreasing the needed stack size to better support Android-based platforms.

6. CONCLUSIONS AND FUTURE WORK

Managing context data on mobile devices is an essential prerequisite for supporting dynamic, context-aware adaptations of applications. In the paper, we have presented a benchmark designed for the management of context data and we have reported in detail on the benchmark evaluation which considers relational and object-oriented databases and a triple store in embedded and/or client/server mode. The set-up considers different context data sets accumulated by mobile users on their devices in a realistic setting in the course of a week, a month, and a year. Relational data stores in c/s mode performed best in the benchmark. For

smaller data sets, the performance of embedded relational databases and the triple store (c/s) is sufficient.

In our future work, we will focus on the LoCa rule engine to provide support for the automatic adaptation of applications (workflows) and user interfaces, based on the local context store evaluated in our benchmark.

7. REFERENCES

- [1] Android Open Source Project. <http://source.android.com/>.
- [2] M. Baldauf and S. Dustdar. A survey on context-aware systems. *Int. Journal of Ad Hoc and Ubiquitous Computing*, 2(4):263–277, June 2007.
- [3] Berlin SPARQL Benchmark (BSBM), 2009. <http://www4.wiwi.fu-berlin.de/bizer/BerlinSPARQLBenchmark>.
- [4] db4o. <http://www.db4o.com/>.
- [5] db4o Reference. <http://developer.db4o.com/Documentation/Reference/db4o-7.12/java/reference/>.
- [6] A. K. Dey. Understanding and using context. *Personal and Ubiquitous Computing*, 5(1):4–7, February 2001.
- [7] N. Fröhlich, A. Meier, T. Möller, M. Savini, H. Schuldt, and J. Vogt. LoCa – Towards a Context-aware Infrastructure for eHealth Applications. In *Proc. of the 15th Int'l Conference on Distributed Multimedia Systems (DMS'09)*, 2009.
- [8] N. Fröhlich, T. Möller, S. Rose, and H. Schuldt. A Benchmark for Context Data Management in Mobile Applications. Technical Report CS-2010-002, Univ. Basel, 2010. http://informatik.unibas.ch/research/publications_tec_report.html.
- [9] Y. Guo, Z. Pan, and J. Heflin. LUBM: A Benchmark for OWL Knowledge Base Systems. *Web Semantics: Science, Services and Agents on the World Wide Web*, 3(2-3):158–182, October 2005.
- [10] H2 Database Engine. <http://www.h2database.com/html/main.html>.
- [11] McObject Benchmarks Embedded Databases on Android Smartphone. <http://www.mcobject.com/march9/2009>.
- [12] MySQL. <http://mysql.com/>.
- [13] OWLIM. <http://www.ontotext.com/owlim/>.
- [14] Perst – An Open Source, Object-oriented Embedded Database. <http://www.mcobject.com/perst/>.
- [15] PolePosition – Open Source Database Benchmark. <http://polepos.sourceforge.net/>.
- [16] Sesame – Open Source Framework for Storage, Inferencing and Querying of RDF Data. <http://www.openrdf.org/>.
- [17] SPARQL Query Language for RDF. <http://www.w3.org/TR/rdf-sparql-query/>.
- [18] SPARQL 1.1 Query Language. <http://www.w3.org/TR/2010/WD-sparql11-query-20100601/>.
- [19] SQLite. <http://www.sqlite.org/>.
- [20] T. Strang and C. Linnhoff-Popien. A context modeling survey. In *In: Workshop on Advanced Context Modelling, Reasoning and Management, UbiComp 2004, Nottingham/England*, 2004.
- [21] TPC - Transaction Processing Performance Council. <http://www.tpc.org/>.