Case studies on context-aware mobile multimedia services

Timo Ojala

MediaTeam Oulu research group
Department of Electrical and Information Engineering
University of Oulu
FI-90014 University of Oulu
timo.ojala@ee.oulu.fi

Abstract: This paper explores the design, implementation and evaluation of context-aware mobile multimedia services by presenting six case studies on different application domains. The case studies highlight the opportunities that mobile devices equipped with wireless data transmission provide for context-aware and ubiquitous computing. Special emphasis is placed on empirical evaluation of the services in the field in form of a user evaluation with genuine end users in real environment of use.

1 Introduction

The research community has introduced a large volume of context-aware systems [BDR07], which have adopted different definitions of the slippery notion of “context“ [Do04]. The systems typically involves mechanisms for context acquisition, preprocessing, representation, management and utilization in an application. The raw context data is acquired from different types of sources or “sensors“, where the term “sensor“ has to be understood very broadly. The raw context data may be subjected to different forms of preprocessing, for example to deal with missing data or to elicit higher level abstractions such as logical relationships or activities. Different models have been developed for representing context data in formats understood by computers and users, ranging from simple key-value pairs to extensive ontologies [Bo07]. Different types of centralized and distributed context management architectures have been proposed for storing and distributing context data to applications. They then utilize context for various purposes such as to adapting the user interface or to retrieving contextually relevant information.

When evaluating context-aware mobile systems, we have to make a distinction whether we evaluate the system architecture [OSW07] or the user interface [KS03]. Evaluating the user interface or usability of mobile systems is difficult due to extra interaction and evaluation challenges in the mobile domain. The interaction challenges include the mobile context of use, rich device functionality, small device size, lack of direct manipulation, and lack of standardization in handset and software design.

The evaluation challenges relate to data collection (technical, social and legal constraints) and uncontrollable variables (e.g. weather, ambient noise and attention destructors). Important decisions to be made regarding evaluation include what is evaluated (interface metaphors, mental models or UI elements), how (empirical vs analytical) and where (lab vs field). While a lab experiment does not provide a proper simulation of the true mobile context and the real-word factors affecting behavior and performance are missing, a field experiment in turn is time consuming, expensive, and suspect to data collection difficulties and uncontrollable external variables.

The six case studies presented in this paper demonstrate the utilization of context data in different types of context-aware mobile multimedia services. Each service employs a simple key-value context model and centralized context management architecture for storing or retrieving contextually relevant information. First five case studies originate from the Rotuaari – Context-Aware Mobile Multimedia Services research project [Ro03]. Each of them involved empirical evaluation in the field in form of a user evaluation with real users in true environment of use. It is a fundamental usability assessment method providing direct information about how the system is used and what are the exact problems [Ni94]. The sixth case study, panOULU Luotsi, is a joint effort of the Wireless Cities project [Wi05] and the panOULU network [pa09]. Some of the services are now in “production“ use, either as a public service or as a commercial product.

2 SmartRotuaari

SmartRotuaari was an early demonstration of the new mobile multimedia services that emerging wireless broadband Internet would eventually facilitate. SmartRotuaari comprised of a wireless multi-access network, SmartWare architcture for deploying context-aware mobile multimedia services, a web-based CPI (Content Provider Interface) for content management and a collection of functional prototype services [Oj03]. The design and implementation of the system started in 2002, leading to empirical evaluation in form of a large-scale field trial executed at downtown Oulu in fall 2003. SmartRotuaari was motivated by the needs of both companies (i.e. mobile service providers, technology providers, retailers etc.) and consumers as end users of mobile services. While companies have the most comprehensive knowledge about their business (“market pull”), some of them, especially smaller retailers, are not necessarily aware of the possibilities offered by the new technology (“technology push”). These aspects were studied in form of surveys, joint workshops and one-to-one discussions with local businesses.

The multi-access network comprised of a GPRS network and a WLAN (IEEE 802.11b). The WLAN emulated the emerging wireless broadband Internet access with much higher data transfer rate and lower latencies than GPRS. We built the WLAN ourselves, managing to install 11 WLAN access points around downtown Oulu by the start of the first field trial in late August 2003. Later the WLAN network was expanded, contributing to the founding of the panOULU network in October 2003 [pa09], which eventually proved to be the most valuable outcome of the project.

The SmartWare architecture included server components for service access, user positioning, instant messaging, real-time presence management and database access. The architecture facilitated using different context attributes represented as key-value pairs in service provisioning: *time* (e.g. time range when a service was active), *location* (absolute and relative locations of the user and/or a service provider, for example a service could be triggered if the user was within a specific distance of a service provider, location of the user could be provided by GPS, WLAN positioning or manual entry), *weather* (real-time weather observation, temperature and wind speed divided to non-overlapping ranges), *user profile* (e.g. age, marital status, education, occupation, income, personal interests entered upon registration) and *user presence* (availability and mood set to any of the seven predefined values).

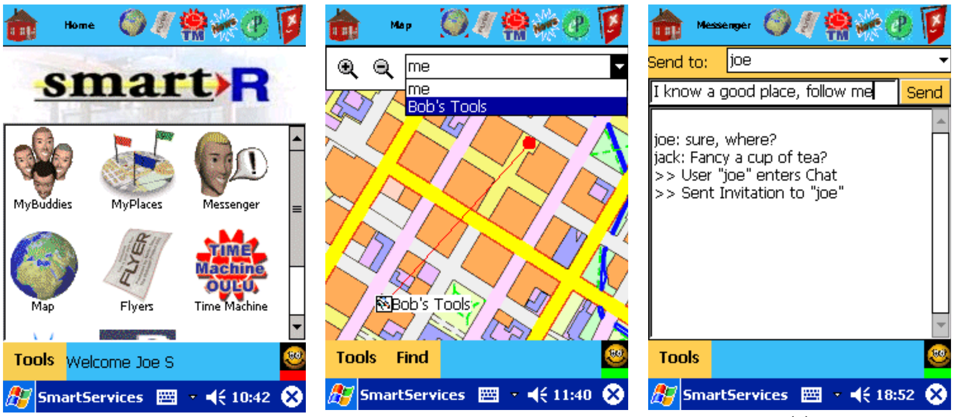

(a) (b) (c)

Figure 1. Example UI screenshots of the SmartRotuaari services: (a) "desktop"; (b) map-based guidance; (c) personal communication.

The prototype services were implemented with a monolithic Java client for PDA's. They included a service directory, map-based guidance, personal communications, Time Machine Oulu (see Section 3), mobile ads (see Section 5), personalized news and mobile payment. The "desktop" shown in Fig. 1(a) provided access to individual services. Map-based guidance (Fig. 1(b)) provided visualization of the location of a place relative to the user's current location shown as red dot. A place could refer to any entry in the service directory, a 'buddy', or a location of personal interest specified earlier by the user. Personal communications was supported in form of peer-to-peer and group chat, which were implemented as simple text-based chat (Fig. 1(c)).

The user could maintain a list of ‘buddies’ and invite them to a chat. Further, the user could set his presence status (‘mood’) to any of the predefined alternatives. The user could choose whether his location and/or presence status were shown to his ‘buddies’ and other users. Personalized news feed was provided so that the user could designate whether (s)he found a particular article interesting or not. A personalization engine then updated the user’s profile accordingly and provided first those of the incoming news that matched the user’s interest profile. Mobile payment was implemented with an external micropayment server, which facilitated payments for on-line content with real money.

A subset of the prototype services, service directory, map-based guidance, mobile ads and TimeMachine Oulu, were included in the field trial running from late August 2003 till the end of September 2003. The field trial was coordinated from an office established in a small hut placed at the very heart of downtown Oulu (Fig. 2). The office was staffed with researchers, who persuaded passers-by to sign up as test users, helped test users in creating a user profile and in using the iPAQ and services, and collected feedback via a questionnaire and occasional interviews. Each test user was awarded with a voucher to a nearby café after the test session.

(a) (b)

Figure 2. (a) Field trial office at the Rotuaari pedestrian street; (b) a test user is signing up.

Recruiting voluntary test users from the general public proved rather difficult. Eventually we had few hundred test users, of which 193 respondents filled in an extensive questionnaire. The rather small spatial coverage of the WLAN network was the most serious technical limitation. The test users had difficulties in understanding that they could lose connectivity indoors or when leaving the downtown, which led to various usability problems. Nevertheless, the field trial provided a very valuable lesson in organizing a large field experiment in a public laboratory.

3 TimeMachine Oulu

TimeMachine Oulu provided a dynamic, interactive and context-aware 3D virtual model nhof historical Oulu, to be used with the web browser of a WLAN-equipped PDA [POO03]. The development of the model was motivated and facilitated by the availability of high quality historical data of the buildings in Oulu between two devastating fires in 1822 and in 1882. The 3D virtual model was generated dynamically from the building database using VRML. The model was simplified (Fig. 3), not photorealistic, adapted to the limited rendering power, network bandwidth and storage of the mobile device, and the simple lighting model of the VRML. The user could "time travel" in the model by increasing/decreasing the "current year". The model was context-aware, placed at the current location of the user determined by WLAN positioning of the mobile device. The model was interactive in the sense that you could move around in the model and access the objects, for example to ask who lived in a particular building in the "current" year.

Figure 3. Two views of a particular location at downtown Oulu in 1827 (left) and 2003 (right). The old church from 1826 is still standing in the background. The building in front of it is only "recently" constructed.

TimeMachine was tested with a small task-based user-based evaluation at downtown Oulu. Ten test users were asked to complete four tasks using TimeMachine: T1: Go to year 1826; T2: Move around the church; T3: Find any red, yellow or green building and name its owner and insurance number; and T4: Find out the width and length of mayor Appelgren's house. After completing the tasks the test users completed a questionnaire addressing various aspects of the perceived usability and user experience. All ten users were able to complete tasks T1, T2 and T3 successfully, while seven were able to complete task T4. Seven users thought that TimeMachine was slow and six users found the visual quality of TimeMachine sufficient. To conclude, TimeMachine was a prime example of an interesting application made possible by high quality historical content.

4 SmartLibrary

SmartLibrary was a location-aware mobile library service, assisting library patrons in finding books and other objects. Back in 2003, SmartLibrary was the first OPAC (Online Public Access Catalogues) search interface tailored for mobile devices atop of Voyager, a widespread library management system. The traditional solution in libraries is to classify the books into holdings and shelf classes. This solution works well if the user is familiar with the shelf classification. For larger libraries, however, there can be tens of holdings, hundreds of classes and thousands of shelves. This results in especially novice library users consulting the library personnel for personal guidance, which consumes the library's resources.

SmartLibrary version 1 [ARO03] provided map-based guidance to the target bookshelf on a PDA equipped with WLAN (IEEE 802.11b) connectivity. The user's location could be estimated with WLAN positioning, which enabled dynamic guidance of the user towards the desired book. The service was a completely software-based solution, which could be provisioned atop a WLAN installed for wireless Internet access, without any additional hardware. SmartLibrary version 1 was deployed in the main library of University of Oulu on top of OULA, an online catalog based on Voyager. PDA's web browser was used to browse the OULA-pda, a web interface tailored for small devices such as PDA's. The query results provided by the OULA-pda were augmented with a link, which started a separate Java-based guidance application. Fig. 4(a) illustrates the definition of a query for a book authored by Tolkien. Fig. 4(b) shows the presentation of the entries matching the query. By clicking the "Locate"- link of the book of interest the user got access to the map-based guidance visualized in Fig. 4(c). In the map view the red dot denotes the location of the user, while the rectangular icon shows the shelf where the requested book resides.

(a) (b) (c)

Figure 4. (a) The query definition UI of OULA-pda. (b) The results of the query. (c) Visualization of the map-based guidance to the shelf containing the book searched for.

We conducted a task-based user evaluation of the SmartLibrary service with 32 randomly selected library patrons. Each user was given two tasks: first, to find a certain book by Tolkien, and second, any issue of a particular economic science periodical. Half of the users completed the first task using public desktop library terminals providing shelf classification, and the other half using SmartLibrary providing map-based guidance. The terminal was changed between the tasks. After the tasks were completed, the users were asked to fill in a questionnaire containing multiple-choice questions and space for feedback. The users were asked which of the two methods they would prefer for finding books in the library and why. All males and 64 % of females chose map-based guidance, which they also found less laborious to use than shelf classification. As expected, the assessment of the laboriousness of shelf classification correlated heavily with the users' experience in using it. The main usability problems were related to the guidance application, which was slow and worked only in particular PDA models. Switching between the web-based OULA-pda and the separate guidance application was considered awkward. The users also had difficulties in orienting themselves on maps, due to poor graphics of the maps.

SmartLibrary version 2 [Ai04] was re-designed to address the problems of the first version. The user interface was provided for the (X)HTML browsers in desktops, PDAs and high-end mobile phones, hence integration with web-based OPACs would be seamless. The graphics of the floor plan maps were designed to be simple and clear. Different symbols on the map were color-coded: walls and other fixed structures were drawn with black, book shelves with blue, and tables with yellow. Target areas and their names were superimposed on the map. The users could also position themselves relative to pre-defined landmarks, e.g. a circulation desk and stairs, shown on the map upon request. A separate web-based CPI (Content Provider Interface) was provided for the purpose of maintaining information of shelf-classes and landmarks.

SmartLibrary version 2 was tested with a task-based user evaluation, where library patrons and staff were asked to find three books with different means: with the user's own way, SmartLibrary with a public desktop terminal, SmartLibrary with a PDA over WLAN connectivity, and SmartLibrary with a mobile phone over GPRS data connection. Most novice users preferred SmartLibrary over the shelf classification, whereas more experienced patrons and library staff preferred the shelf classification. The users considered SmartLibrary most useful on public desktop terminals, i.e. just map-based guidance without the WLAN positioning of the user.

SmartLibrary is still in 'production' use at the libraries of the University of Oulu in two different forms. Map-based guidance is provided as a web service for locating shelf-classes, study rooms, equipment and other resources, and it has become quite popular among library patrons. Also, an online search web interface tailored for mobile devices is provided. However, the WLAN positioning of the mobile device has been discarded, due to the limited range of supported WLAN devices and high maintenance cost.

5 Mobile advertising

Three different context-aware systems for permission-based mobile advertising were developed in the Rotuaari project. The first advertising system was part of the SmartRotuaari system described in Section 2. A mobile ad was a SMIL 2.0 compliant multimedia message delivered to a PDA over WLAN connection. The ad was authored and activated by the advertiser using a web-based CPI. The advertiser defined the ad profile, i.e. the context attributes (time, relative location, weather, user profile and/or user mood) that were required for the ad to "trigger". Similarly, the user could define whether (s)he wanted to receive ads and from what relative range from her/his current location. If there was a match, the ad was delivered to the user's device. 18 local companies were recruited to produce free ads promoting their products and services, of which 12 eventually provided ads (Fig. 5) during the field trial. The test users' feedback supported our own observation of the limited added value provided by the bulk of the mobile ads, which were content to just provide the name and address of a store, for example. [Oj03]

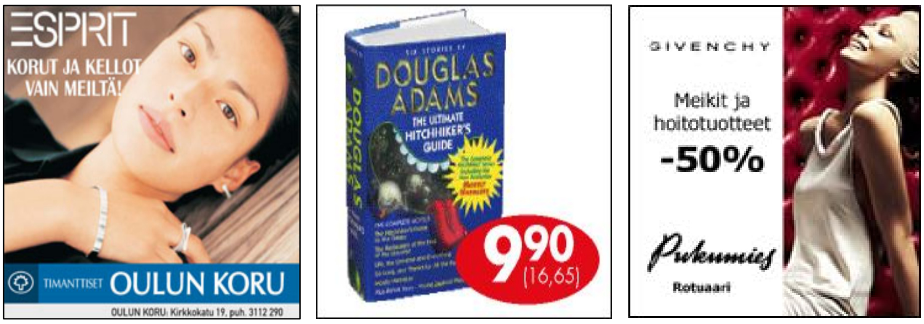

(a) (b) (c)

Figure 5. Example mobile ads: (a) jewelry store advertising a jewelry brand; (b) discount offer from a bookstore; (c) cosmetics discount ad.

In the second advertising system ad delivery was based on a telco grade MMSC and the user device was a mobile phone equipped with GPRS data connection. The mobile ads were again created with a web-based CPI, including tailoring of the content for different user devices and the specification of the ad profile. If there was a match between the ad profile and the user profiles, the ad was delivered either as MMS messages or as XHTML Mobile Profile pages with WAP Push. The advertising system was employed in two large field trials in 2004 and in 2005, and in a third smaller field trial in 2006. For example, in the 2004 trial 44 businesses were recruited as advertisers, producing 81 different ads that were delivered 11370 times in total to 610 test users. The advertising system was hampered by technical difficulties such as variations in the implementation of MMS players in different phones. [Ko07a]

The third mobile advertising system called B-MAD was based on Bluetooth and WAP Push [Aa04]. A Bluetooth sensor was configured to periodically scan for the globally unique Bluetooth device addresses of passing by Bluetooth user devices. The sensor sent the BT addresses of detected devices over a WAP connection to an ad server, together with a location identifier. The ad server mapped the BT addresses to the user phone numbers and checked from the database if there were any ads associated with the location that had not been delivered to the user. The undelivered ads were delivered as XHTML Mobile Profile pages using WAP Push. User's location was the only context attribute used for triggering ad delivery. The B-MAD was evaluated with a small-scale field trial where Bluetooth sensors were placed in eight stores providing 11 ads in total. Two ads from a particular store were temporally related so that if the user stayed nearby the store for a certain time after receiving the first ad, (s)he would receive a second ad including a gift certificate. 35 test users were asked to walk a designated route bypassing the stores and the delivery of ads was logged in the server. Due to the long Bluetooth scanning delay the average positioning latency was 25 seconds and the sensor was not guaranteed to detect every Bluetooth device passing by at walking speed. Further, ad delivery with WAP Push took almost 12 seconds on average. The combined total latency of 37 seconds meant that at times the user received the ad fairly far away from the sending store.

The three mobile advertising systems and their evaluation in five field trials provided us and advertisers valuable lessons in mobile advertising. The advertisers and the few agencies they used for authoring mobile ads on their behalf had great difficulties in understanding how personal channel a mobile device is. They were mostly content to just reproduce the ads they used in printed mass media. They were also reluctant to personalize the ads and consequently the customers found the ads rather useless. At the beginning we also provided too many context attributes to choose from, which led to sparse context spaces with few hits per ad. We learned that there was no room for technical glitches, which were partly our own doing and partly due to problems in commercial products. Further, we learned that being so much ahead of time was not necessarily a bonus, as other stakeholders had difficulties in buying our vision. Nevertheless, some advertisers reported that they managed to create new business by participating in our trials.

6 Mobile Fair Diary

Mobile Fair Diary (MFD) was designed to allow a housing fair visitor to make a personalized digital recording of his/her visit to the fairground for later use [Ko07b]. In a typical housing fair a large number of different types of buildings and their décor are on display for a given period of time. Typically, a visitor spends a day at the crowded fairground under a hectic schedule, visiting possibly dozens of houses and exhibition booths. Most of the offered information is made available as paper brochures, resulting in a pile of paper being carried home for later use. The housing fair also typically comes with an accompanying website containing online information about the houses and other things. Further, many visitors are equipped with notebooks and a personal digital camera for taking photos of interesting objects.

We can identify many problems in the conventional setup of a housing fair. Firstly, it does not support the user in collecting, from a range of different sources (e.g., brochures, the web and personal digital cameras), the specific detailed information of an object that is of personal interest to the user in a manner which would incorporate automatic hyperlinking between different sources. This is a very acute problem to be solved, as the total amount of information available at a housing fair is immense, but only a fraction of it may be relevant to a particular user. Secondly, the big pile of brochures and miscellaneous photos does not support efficient retrieval of relevant information at a later date, when that information would be needed in designing a kitchen renovation, for example. Thirdly, the current setup does not support efficient sharing of information with other users, for example with friends who might also be planning to renovate their kitchens. Fourthly, the many month long marketing period before the fair requires that all printed brochures are produced well ahead of the houses and their interior/exterior being finalized. This means that computer generated models are used in brochures instead of actual photos, or that illustrative photos are omitted altogether.

The MFD was designed to allow a housing fair visitor to capture the relevant information and to support more timely dissemination of information.The MFD combined a mobile phone and a desktop PC into a novel hybrid interface for collecting, storing and sharing information. The phone application was used for taking context-aware notes such as visual codes, photos, dictations and text. The notes were uploaded onto a website where they were presented in a contextual order for browsing, organizing and sharing with friends to be used on a later occasion. The website also automatically linked user-made notes with ready-made content packages of houses and related online resources.

(a) (b) (c) (d)

Figure 6. Screenshots of the UI of the MFD phone application: (a) main view; (b) adding a note; (c) a visual code is read; (d) a 2nd level visual code has been read and location is updated.

When starting the phone application for the first time, the user was required to type in a phone number or an email address or both so that the user could be delivered the user ID and password needed for accessing the website. After typing in the contact information, the user was presented with the main view showing the number of entries made at the current location and the total number of entries (Fig. 6(a)). The phone application had dedicated user interfaces for reading visual codes, taking photos, making dictations and writing text memos (Fig. 6(b)). The phone application also facilitated browsing and deleting the notes.

When a note was taken, it was associated with related contextual metadata comprising of the current location and a timestamp. The current location was indicated by reading a 2-D visual code of known locations with the phone's camera (Fig. 6(c)). Upon reading a visual code the user location was automatically updated (Fig. 6(d)). The MFD employed a nested two-level hierarchy of visual codes so that a 1st level code designated a house or an outdoor object such as a piece of art. The 2nd level code associated with a 1st level code of a house designated objects inside the house, for example a room or a piece of furniture. The 1st level codes were placed at the entrances, which were supposed to be read upon entering the houses. If the user forgot to read the 1st level code at the entrance, but then read a 2nd level code indoor, the phone application automatically recognized the corresponding 1st level code associated with the 2nd level code and updated the location accordingly. Furthermore, the user could also change the location by selecting it manually from the list of 1st level codes. The phone application was configurable so that an XML file specifying all visual codes and their UI labels was automatically downloaded when the application was started for the first time. The notes and their metadata were uploaded to a server at the end of the visit to the housing fair.

Figure 7. The MFD website presented notes in contextual order with hyperlinks to related information.

The MFD website facilitated fluent browsing, organizing and sharing of the notes taken with the mobile phone. The notes were presented in contextual order so that the objects corresponding to the 1st level codes read by the mobile phone (i.e., the houses and outdoor objects visited) were visualized in the left-hand frame with the title and thumbnail image in chronological order (Fig. 7). One of them was designated as the current object of interest, for which the right-hand frame then showed all notes taken in chronological order. The website also automatically augmented the notes of a particular object with hyperlinks to the online content package provided by us for that object and to other related online resources such as the website provided by the fair organizers. The user could rearrange the notes into any order that the user saw fit and deleting any needless notes. Furthermore, the user could compose personal collections of notes augmented with textual annotations. The user could also download the diary from the website to a local file as a ZIP-package.

The MFD was empirically evaluated in a large-scale field trial at the national Finnish Housing Fair in July-August 2005. The annual fair is the most prominent activity undertaken by The Finnish Housing Fair Co-operative Organization, a non-profit association, whose mission is to improve the quality of construction in Finland. Different types of buildings and their décor were on display for a period of 30 days, after which the actual occupants of the houses moved in. The fairground consisted of over fifty houses and apartment buildings constructed on an area of 13 hectares (Fig. 8(a)). We were granted permission to place a 1st level visual code at the entrance of each house and apartment building (Fig. 8(b)) and on outdoor art pieces. Furthermore, a limited number of 2nd level visual codes were allowed (3-8 per house, 43 in total) inside the eight houses circled in Fig. 8(a).

(a) (b)

Figure 8. (a) A map of the fairground, where the circled houses were equipped with 2nd level visual codes indoors; (b) A 1st level visual code placed at an entrance of an apartment building.

The field trial was coordinated from an office located near the main entrance to the fairground. If people had a compatible phone, they could acquire the phone application themselves by sending an SMS to a given phone number, downloading the application from the web, or visiting the field office. If their phone was not compatible they could loan one, without charge, for the duration of their visit at the fairground. Each test user was rewarded with ice cream upon returning to the desk and uploading their diary to the website. Qualitative data was collected through an online questionnaire available at the website and by interviewing ten randomly selected users. The test users were motivated to fill in the questionnaire with a raffle. Quantitative data comprised of the statistics of the notes and their metadata, together with the log of the users' actions on the website.

During the one-month trial 349 test users uploaded their diary to the website and 169 (48.4%) of them filled in the online questionnaire. Our small sample had similar profile to that of the whole Housing Fair clientale of 121110. The 349 diaries contained 27258 notes. 302 (86.5%) of the 349 diaries contained more than 10 notes, which can be regarded as actual usage of the MFD. An individual user took about 78 notes on average, over half of the users took more than 50 notes and 25 (7.2%) heavy users took more than 200 notes. The 27258 notes were of different types as follows: 24044 photos, 1869 indoor objects (2nd level visual codes), 831 dictations and 514 text notes. The large number of photos is simply explained by the fact that a photo has clearly the highest information value with respect to the amount of work needed. The low number of text notes is explained by the difficulty of typing notes with the 12-key keypad of a mobile phone. The low amount of dictations is at least partially explained by the challenging social setting, for it is understandably difficult to dictate given the crowd and ambient noise at the housing fair.

Here I present just some key observations from the data while the reader should see [Ko07b] for in-depth analysis. When we conducted a statistical test between the numbers of different notes and selected respondent attributes (gender, age, prior experience in using smart phones, prior experience in using the Internet), the only statistically significant finding was that females took more text notes. This means that, for example, the lack of any prior experience in using smart phones had no effect on the usage of the phone application, which testifies to a successful design in terms of learnability and efficiency. When we tested for any statistically significant differences in the questionnaire data with respect to respondent demographics, the following findings were made. Female users thought that the service was easier to learn (average score of 4.43 on 5-point Likert scale) and more self-evident to use (4.38) than male users (4.11 / 4.08). Female users were also more confident that the service was useful (4.60) than male users (4.35). Female users also felt that they gained more advantage from using the service (4.32 vs. 3.97). Quite surprisingly, the users with no prior experience of smart phones felt that the MFD was easy to learn (4.61) than those with prior experience (4.36). They also felt that the service was more self-evident to use (4.36 vs 4.10). As a whole, the data demonstrated excellent user satisfaction and identified unconventional enthusiastic user groups for a mobile service such as middle-aged women.

The MFD concept was commercialized by the startup company founded by the researchers of the project and is now available as a commercial product Entre Exhibitor.

7 panOULU Luotsi

panOULU Luotsi (~ Pilot in English, [pa09b]) is a location-based information mash-up provided for the users of a large municipal wireless network in the City of Oulu in northern Finland [Ku08]. The panOULU (public access network Oulu) network is provided by a consortium of five public organizations and four ISPs [Oj08][pa09]. As of now, the network totals about 1050 WLAN (IEEE 802.11a/b/g) access points, which provide indoor and outdoor coverage in locations deemed relevant for public access. The city center and its immediate surroundings are blanketed with a large outdoor WLAN mesh network, but otherwise the coverage is provided in hotspot manner. In its coverage area panOULU provides open (no login/authentication/registration) and free (no payment) wireless internet access to the general public, as long as you have a WLAN-equipped device. In September 2008, 15127 devices used the network, totaling 370000 sessions and 13.9 million online minutes. Up to 40% of the users in a given month are visitors, as determined from the usage patterns and by detecting devices that had not been seen in the network before. The proportion of multi-mode mobile handsets equipped with WLAN radios has been growing steadily so that they make up ~25% of the devices today. The usage of the network is still very much nomadic, not mobile, as only ~5% of sessions can be considered mobile.

The motivation behind Luotsi was to provide one obvious access point to useful information, which was before fragmented across several web sites and presented in different formats. Thus, the user had to implicitly know what information s/he needed and where to get it from. In many cases the user had to resort to multiple services before acquiring the necessary information. Some providers had made efforts to gather more comprehensive data in their site, but the data was still rather limited and outdated at times due to the lack of administration. Further, while one site might offer a map on the location of the desired target, another might not, and the user would have to manually enter the address in a map service, thus leaving the service s/he was originally using. The motivation was equally supported by the increasing coverage and usage of the panOULU network. The over 15000 users represent a considerable population in a city of 130000 people. Further, large proportion of them are visitors, who need up-to-date information of the services, places and events in a foreign city. In addition to providing a single access point to topical and relevant local information, other design goals included re-using existing information feeds instead of reinventing them, no client application should be needed besides a web browser, and the information should always be up-to-date. The design goals could be achieved with a web mash-up, which builds on the aggregation of various information feeds and real-time positioning of the user.

The high-level software architecture of panOULU Luotsi is shown in Figure 9. To access the vast amount of distributed information and to bypass the problem of outdated data on different sites, all content into Luotsi is acquired in form of XML-based feeds such as, but not limited to RSS or ATOM. The varying presentation formats of the feeds are dealt by the XML aggregator illustrated as the DataMapper in Figure 9, which maps different heterogeneous information feeds into the Luotsi database without any changes to the application source code.

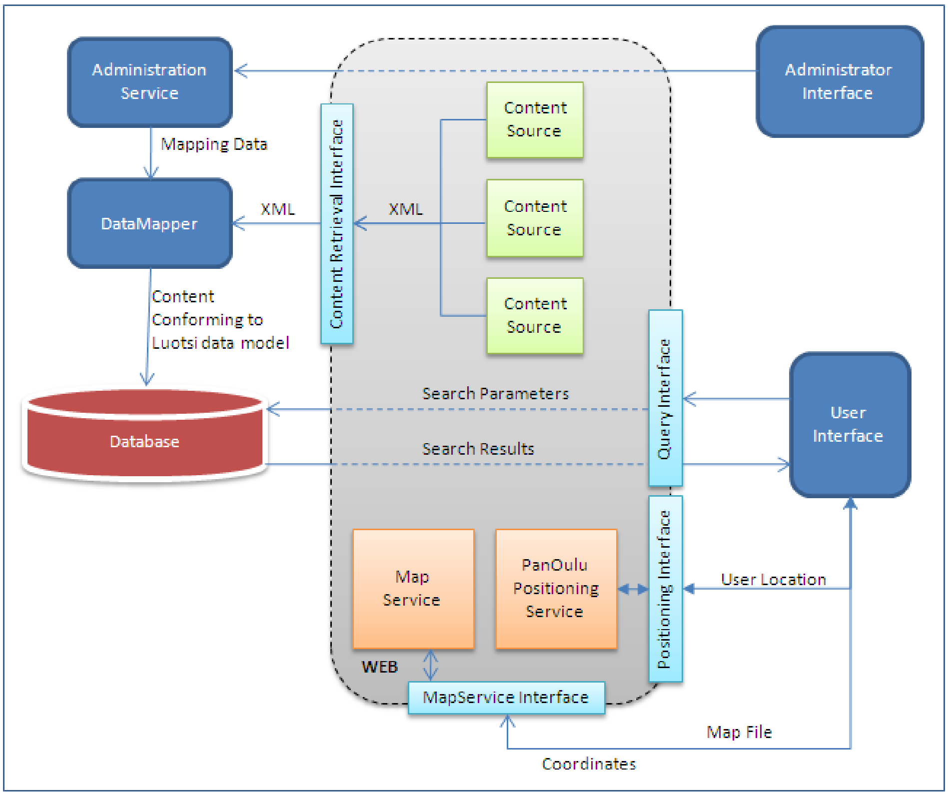

Figure 9. The high-level software architecture of panOULU Luotsi.

Luotsi has an administrator interface for registering any XML-based feed to be registered in the system. When registering a feed, the administrator receives a parsed presentation of how the feed is formulated (the XML tags). The administrator manually maps the attributes to corresponding fields in the Luotsi data model. For example, two feeds may describe a place with different XML structures such as <name></name> and <placeName></placeName>. Both are mapped to 'place_name' field in Luotsi data model. The mapping needs to be done only once for a feed unless, of course, the structure of the feed changes. The mapping rules are saved in a configuration file. The registered feeds are then automatically updated periodically by their provider, thus ensuring that all data is always up-to-date.

The location of the user is estimated as the location of the WLAN access point the user device is currently connected to. Typically, the user's actual location is within a 50 meter radius of location of the access point. This network-based AP ID positioning has several benefits over other methods such as GPS positioning: the user device needs no additional software or hardware, the first location estimate is obtained very quickly, and it works also indoors provided network coverage is available.

Currently, registered feeds include movie listings from Finnkino (the main movie provider in Finland), City of Oulu's event calendar and news feed, fast food restaurants from kaenkky.com (a local website listing and reviewing all fast food restaurants in Oulu region), service directory and local sights listings from ouluthisweek.net (local service provider producing both online and printed information services), and the latest news from Kaleva (the local main newspaper). Aggregated, these feeds provide a very comprehensive directory of services and events that both local people and visitors might need when moving around Oulu. Luotsi also provides a location aware weather report based on a cluster of micro weather stations installed around the city. The weather data provided by the stations is available as a public web service. Luotsi automatically determines the weather station nearest to the user's current location and displays the data in easy-to-read manner, thus adding another element to the mash-up.

The feeds are retrieved periodically by the Content Retriever, mapped by the DataMapper into Luotsi data model based on the rules in the configuration file and the mapped data is stored in the Database. The Database serves as a proxy (cache), thus making later queries much faster than if the data was always fetched from the individual feeds and mapped on the fly. The Database is a mySQL database with a table structure conforming to Luotsi data model. The data model is flexible to support the needs of varying feed structures. Fields not found from a feed are left blank, and cross-referencing is supported. This is useful if, for instance, an event-item does not have explicit location information (e.g. coordinates), but contains a logical name or street address, which can be mapped to the coordinates using the City of Oulu's geodetic service. Similarly, relevant tags of the emerging geoRSS format (e.g. <geo:lon> and <geo:lat>) could be mapped to a location in our data model, although the current content feeds do not use them.

The panOULU Luotsi web user interface is illustrated in Figure 10. It is implemented with the jQuery AJAX library using the model-view-controller design paradigm. The presentation layer is separated from the application logic stored in separate PHP files that are called from the UI. The layout is divided into two columns so that the left column shows the location of the user and objects of interest. The user's current location is presented by placing a 'guess circle' around the icon representing the user on the map, indicating that the user's actual location is within the guess circle. Object categories are presented with their specific distinct icons, e.g. a plate indicates a restaurant. The map file is retrieved by the MapServiceInterface component from the City of Oulu's online map service, which offers an open web inteface for retrieving maps. The map is initially centered at the user's location, but it can be freely moved in any direction after that. The maps outside the current view are fetched asynchronously in the background, so that when the user moves off the current map view, map loading times are minimized. The map also supports a 4-level zoom, allowing the user to zoom in or out at any point. Luotsi is primarily intended to be used via the panOULU network, which facilitates the AP ID positioning. However, Luotsi is a public web service, and thus can be used via any other access network, as well. If the positioning fails because the user is not connected to panOULU network, the user is asked to enter his/her current or desired location manually, e.g. as a street address or pointing a location on the map.

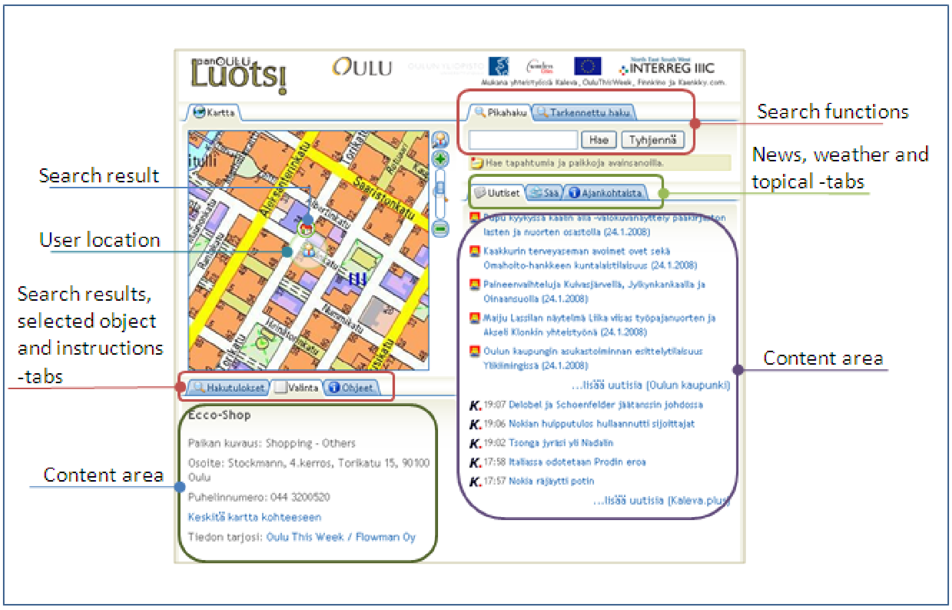

Figure 10. The browsing interface of panOULU Luotsi.

Luotsi also employs automatic proximity-based search functionality. Once the user location has been established, the system automatically displays five nearest objects on the map. The objects found with the proximity-based search are hidden once the user makes a search for something else. Below the map is a window that can display different information based on the selected tab. Initially, the content area displays the ‘instructions’ tab with general information on the use of the service, but once the user searches for something or selects one of the automatically shown objects, the content changes to either a list of search results, or details on the selected object, respectively.

We have been logging the usage of Luotsi since its launch in October 2007. By September 2008, a total of 6658 sessions was recorded, averaging 555 sessions a month. About one third of the sessions were from devices connected to the panOULU network, two thirds from elsewhere in the Internet.

8 Conclusion

I conclude with a brief discussion on some important lessons that I have learned from the case studies and related research efforts. The lessons may also help to understand why commercially relevant context-aware applications remain sparse, most notably GPS navigators and few other location-based applications, despite the long-term enthusiasm for context-aware computing in the research community.

Design for maximum usefulness. With usefulness I refer to the combination of utility and usability [Ni94]. A system with high utility satisfies some concrete need(s) of the end user. Their identification calls for understanding the “situation of concern”, where things are not quite right and which could be resolved with the proposed system. We can then quantify our success in system design and implementation by measuring or predicting the usability of the system with various usability factors. The Mobile Fair Diary was a prime example of a service satisfying a concrete need of capturing the relevant information from a housing fair. Consequently, the test users expressed high subjective satisfaction despite the assorted design flaws.

Content is king. The design and implementation of useful services requires high quality content, which can be expensive and difficult to obtain. The same applies to contextual data, as well. If you have to produce the content yourself, then you have to allocate sufficient resources for that purpose. The TimeMachine Oulu is a good example of an application made possible by high quality historical data. Similarly, the Google Maps with its extensive map data and functional APIs have enabled thousands of web mash-ups.

Understand existing business practices and value networks. Our difficulties with mobile advertising were partly due to the concept being simply too far ahead of the state-of-the-art in business-to-customer advertising. At the same time we failed to motivate and educate the advertisers to use the mobile advertising channel in the “correct” manner. We had similar difficulties with the Mobile Fair Diary, which we offered to the housing fair as plain “technology push” without any consideration for a potential business model. Consequently, the many commercial actors of the existing value network involved in the deployment were not motivated, which led to delays and problems in the design and implementation of the service and the execution of the field trial. Finally, the research community seems often happily forgetting that a long-term production deployment of a service calls for a viable business model. A fundamental shortcoming in all presented case studies was that the user did not have to pay anything for using the services, though most of them can be envisioned to be provided as free services.

Do not forget maintenance. The responsibilities do not end with the first deployment of a service. On the contrary, the subsequent maintenance can require lots of manpower and other resources, depending on the scalability, fault tolerance and maintainability of the service infrastructure. Maintenance of context data can be difficult, as well. For example, even if you manage to somehow acquire high quality user profile data, it is well known that users are generally not motivated to keep the data up to date. Based on our own experience research organizations and projects are traditionally not well equipped to do maintenance on long-term basis. For example, we simply ran out of resources to keep the WLAN positioning system based on signal strength fingerprints functional. Similarly, maintaining dedicated server resources and components for long-term “production” deployments has proven challenging, as the resources allocated per project basis disappear when the project runs out. Thus, take maintenance into account in designing your deployment and resourcing.

Share your infrastructure. The research community values novel contributions, not high quality engineering which is a prerequisite for successful experimental field deployments. Consequently, there is lots of needless “reinventing the wheel”, where a particular application problem is tackled the umpteenth time with a slightly incremental “novel” approach justifying a scientific publication, but not much else. This could be alleviated by researchers sharing their implementations, assuming they would be of sufficiently high quality for that purpose and no IPR issues would prevent sharing. The panOULU network is a prime example of how sharing your infrastructure can benefit the whole community. [RS05]

Empirical evaluation in the field pays off. I am a big fan of empirical user evaluation of prototype implementations in true environment of use by genuine end users. Despite its methodological pitfalls and practical challenges it remains the only way to involve all factors affecting overall user acceptance, most importantly the actual usage context. Field trials are expensive and time consuming for many reasons. In terms of engineering cost there can be a big gap between a one-time steering group demo and months-long deployment exposed to the diverse general public. Setting up a field trial in public space may require dealing with lots of bureaucracy and logistic challenges. A field trial inevitably requires lots of manpower for different tasks such as help desk, content production, technical support and maintenance. A field trial may expose you to public scrutiny, which is not always a pleasant experience. A large-scale field trial assumes using off-the-shelf mass market technology, thus approving its limitations in comparison to more advanced research equipment. Further, involving the general public as test users can be very difficult, as people are simply too busy with their real lives and reluctant to install your research software in their mobile phones. However, if you despite all these warnings wish to engage in a field trial, the following recommendations may prove useful. First establish your theoretical framework: what is your situation of concern, what are your objectives, what are your hypotheses, what is being evaluated and how, what data is needed for the evaluation, what are your evaluation criteria and their theoretical foundations? Think carefully about your service offering: it is typically much more difficult to sell a large service portfolio than an individual well-focused service. When designing the field trial, plan well ahead and prepare for the unexpected. Allocate sufficient time for the integration testing of your deployment before the start of the field trial. Remember to budget ample resources for PR, help desk, technical support and maintenance during the trial. Have a proper PR strategy to recruit people as test users, to manage the expectations of the general public, media and other stakeholders, and harden your skin for public scrutiny. Make sure that you obtain the data needed for your theoretical framework and store it in a safe place. Allocate ample time for analysing the data after the field trial, for otherwise your hard work can go waste, which would be a real shame. Finally, try to decide before the field trial whether your service(s) are going to be available after the conclusion of the field trial or not. If some remain, then you have to allocate resources for their maintenance. The real proof of a successful trial deployment is when the users continue to work with a system after the trial is officially over [Wa06].

References

[Aa04] Aalto, L; Göthlin, N; Korhonen, J; Ojala, T: Bluetooth and WAP Push based location-aware mobile advertising system. Proc. Second International Conference on Mobile Systems, Applications and Services, Boston, MA, 49-58, 2004.

[ARO03] Aittola, M; Ryhänen, T; Ojala, T: SmartLibrary - Location-aware mobile library service. Proc. Fifth International Symposium on Human Computer Interaction with Mobile Devices and Services, Udine, Italy, 411-416, 2003.

[Ai04] Aittola, M; Parhi, P; Vieruaho, M; Ojala, T: Comparison of mobile and fixed use of SmartLibrary. Proc. 6th International Conference on Human Computer Interaction with Mobile Devices and Services, Glasgow, Scotland, 383-387, 2004.

[BDR07] Baldauff, M; Dustdar, S; Rosenberg, F: A survey on context-aware systems. International Journal of Ad Hoc and Ubiquitous Computing 2(4):263-277, 2007.

[Bo07] Bolchini, C; Curino, C; Quintarelli, E; Schreiber, F; Tanca, L: A data-oriented survey of context models. SIGMOD Record 36(4):19-26, 2007.

[Do04] Dourish, P: What we talk about when we talk about context. Personal and Ubiquitous Computing 8(1):19-30, 2004.

[KS03] Kjeldskov, J; Stage, J: New techniques for usability evaluation of mobile systems. International Journal of Human-Computer Studies 60(5-6): 599-620, 2003.

[Ko07a] Komulainen, H; Mainela, T; Tähtinen, J; Ulkuniemi, P: Retailers' different value perceptions of mobile advertising service. International Journal of Service Industry Management 18(4):368-393, 2007.

[Ko07b] Korhonen, J; Ojala, T; Ristola, A; Kesti, M; Kilpelänaho, V; Koskinen, M; Viippola, E: Mobile Fair Diary - Hybrid interface for taking, browsing and sharing context-aware notes. Personal and Ubiquitous Computing 11(7):577-589, 2007.

[Ku08] Kukka, H; Ojala, T; Tiensyrjä, J; Mikkonen, T: panOULU Luotsi: A location based information mash-up with XML aggregator and WiFi positioning. Proc. 7th International ACM Conference on Mobile and Ubiquitous Multimedia, Umeå, Sweden, 2008.

[Ni94] Nielsen, J: Usability Engineering. Morgan Kaufmann, San Francisco, 1994.

[OSW07] Oh, Y; Schmidt, A; Woo, W: Designing, developing, and evaluating context-aware systems. Proc. 2007 International Conference on Multimedia and Ubiquitous Engineering, Seoul, Korea, 1158-1163, 2007.

[Oj03] Ojala, T; Korhonen, J; Aittola, M; Ollila, M; Koivumäki, T; Tähtinen, J; Karjaluoto, H: SmartRotuaari - Context-aware mobile multimedia services. Proc. 2nd International Conference on Mobile and Ubiquitous Multimedia, Norrköping, Sweden, 9-18, 2003.

[Oj08] Ojala, T; Hakanen, T; Salmi, O; Kenttälä, M; Tiensyrjä, J: Supporting session and AP mobility in a large multi-provider multi-vendor municipal WiFi network. Proc. Third International Conference on Access Networks, Las Vegas, NV, 2008.

[pa09a] panOULU network. http://www.panoulu.net.

[pa09b] panOULU Luotsi, http://luotsi.panoulu.net.

[POO03] Peltonen, J; Ollila, M; Ojala, T: TimeMachine Oulu - Dynamic creation of cultural-spatio-temporal models as a mobile service. Proc. Fifth International Symposium on Human Computer Interaction with Mobile Devices and Services, Udine, Italy, 342-346, 2003.

[RS05] Sharp, R; Rehman, K: The 2005 UbiApp Workshop: What makes good application-led research? IEEE Pervasive Computing 4(3):80-82, 2005.

[Ro03] Rotuaari – Context-Aware Mobile Multimedia Services research project, 2003-2006, http://www.rotuaari.net.

[Wa06] Want, R: Build what you use. IEEE Pervasive Computing 5(3):2-3, 2006.

[Wi05] Wireless Cities project, 2005-2007, http://www.wirelesscities.org/.