PosQ: Unsupervised Fingerprinting and Visualization of GPS Positioning Quality

Mikkel Baun Kjærgaard[1] and Kay Weckemann[2]

[1] Department of Computer Science
Aarhus University, Denmark
mikkelbk@cs.au.dk

[2] Institute for Informatics
Ludwig-Maximilian University Munich, Germany
kay.weckemann@ifi.lmu.de

Abstract. GPS positioning does not provide pervasive coverage and the accuracy depends on the local environment. When deploying and managing position-based applications it is important to know when to depend on GPS and when to deploy supplementary means of positioning, such as local or inertial positioning. This paper proposes PosQ, a system for unsupervised fingerprinting and visualization of GPS positioning quality. PosQ provides quality maps to position-based applications and visual overlays to users and managers to reveal the positioning quality in a local environment. The system reveals the quality both as it changes over time, in 2D and 3D, and for each type of GPS receiver. Our evaluation provides evidence that the collected quality maps are accurate, that they remain informative over time, that they capture the differences among GPS receivers, and that they can be efficiently collected by participating devices.

Keywords: global positioning system, positioning quality, fingerprinting, efficiency.

1 Introduction

Position-based applications use context information to provide position-dependent functionality. They are developed for a variety of domains e.g. location-based games [4,7], indoor and outdoor services for museums and amusement parks and professional applications for different working domains [10]. The positioning technologies used to enable such applications have to balance coverage, positioning accuracy as well as deployment and operation costs.

GPS positioning is today using state-of-the-art receivers available in most outdoor areas with high accuracy and in some indoor areas with lower accuracy as reported by Kjærgaard et al. [8]. This study also concludes that indoor GPS positioning coverage and accuracy mainly depend on the local environment, e.g. the building materials and number of walls, with some fluctuations due to the movement of people, satellites and atmospheric variations. However, in practice, it is difficult for people to judge for their local environment where and when

M. Griss and G. Yang (Eds.): MOBICASE 2010, LNICST 76, pp. 176–194, 2012.

they can rely on GPS, e.g. indoor availability can depend on room sizes, number of windows, types of walls and roofs and the existence of inner court yards or atriums. In the negative case they have to augment the GPS positioning with, e.g., WiFi [3,11], Powerline [15], GSM [18] or inertial [5] positioning. Secondly, it is difficult to judge what happens if the local environment changes later on, e.g. changes to the roof or new windows, movement of walls or changes in the number of people present. As manual evaluation of GPS availability is both time consuming and has to be continuously repeated it is relevant to automate this process. Thirdly, people using an application and entering an area are interested in judging the positioning quality and maybe decide to switch to inertial positioning [5]. Fourthly, position-based applications can use the information to adjust their functionality to the accuracy in an area [16] and the differences among GPS receivers.

This paper proposes PosQ, a system for unsupervised fingerprinting and visualization of GPS positioning quality. The system provides quality maps to position-based applications and visual overlays to users and managers to reveal the positioning quality. PosQ reveals the quality in three dimensions: As it changes over time, in 2D and 3D, and for each type of GPS receiver. The system enables unsupervised fingerprinting by running in the background of a GPS enabled device without any need of attention by the user. It enables people and applications to judge the GPS positioning quality in a local environment with low effort, e.g., in a mixed indoor-outdoor environment in the scope of a city. The PosQ system consists of four key elements:

- GPS receivers that provide estimates of the positioning quality
- PosQ collectors that fingerprint the positioning quality unsupervised as estimated on participating GPS-devices and efficiently send quality reports to the system's quality database
- Quality maps built from the database
- Visual overlays generated from the maps that can be viewed in popular earth or map viewers

We make the following contributions in this work: First of all, we propose a system for unsupervised fingerprinting of GPS positioning quality that can visualize positioning quality both over time, in 2D and 3D, and for different types of GPS receivers. Secondly, we evaluate the system to provide evidence that the collected quality maps are accurate, that they remain informative over time and that they capture the differences among GPS receivers. The evaluation takes into account different local environments and their positioning quality conditions (outdoor, indoor, and deep indoor). Therefore, several experimental evaluation were executed in a shopping mall, a cultural history museum and a botanic museum with different motion patterns. Thirdly, we demonstrate how the system enables the design of positioning solutions, mapping between different GPS receivers and visualization of positioning quality over time and in 3D. Fourthly, we propose and evaluate methods for efficient collection of quality reports and show that the methods can significantly decrease the number of quality reports needed when updating quality maps in real-time.

The remainder of this paper is structured as follows: In Section 2, we discuss related work. Subsequently, we introduce the PosQ system in Section 3. Afterwards, we present the results of evaluating fingerprinting of positioning quality in Section 4 and discuss the usage of the system in Section 5. Methods and evaluation results for efficient collection of fingerprints are then presented in Section 6. Finally, in Section 7 we provide summarizing conclusions and discuss directions for future work.

2 Related Work

Related work presented in the following is divided into *estimating positioning quality* and *using information about positioning quality.*

2.1 Estimating Positioning Quality

A GPS receiver will, in addition to position and time, estimate the accuracy of a position fix. The *accuracy* of a position fix intuitively denotes how close the fix is to the correct, but unknown, position [10]. Basic algorithms for estimating the accuracy are described in textbooks such as Misra et al. [13], but receivers generally use proprietary algorithms that take into account information such as the channel-to-noise ratios, signal strengths, number of satellites, *dilution of precision (DOP)* (a measure for the satellites' geometric strength for positioning), range residuals, variations in low level measurements and multipath indicators. PosQ collects such estimates to fingerprint the positioning quality. We use the accuracy estimates produced by the receivers proprietary algorithms because only a subset of the above information can be accessed using the location API's on mobile phones (e.g., Nokia and HTC). Also experimental observations presented by Kjærgaard et al. [8] for indoor environments indicate that the estimation can not be based on a single value such as the *dilution of precision* factor because in their experiments low accuracy was better explained by multi-path and weak-signal effects than the geometric strength of the current satellite constellation.

To predict GPS coverage in cities, Steed [17] proposes a system that can predict the binary information of whether GPS is likely to be available or not. The system uses a city model and satellite orbits to calculate how many satellites are likely to be available in the different parts of a city. The calculations only take line-of-sight conditions into account. Therefore, the system cannot provide the fine-grained positioning quality information in different conditions as pursued in this paper.

To estimate the accuracy of outdoor GSM positioning, Dearman et al. [6] propose two methods: A machine-learning method that uses a linear regression technique on basic signal strength features and a ground-truth method that requires the supervised collection of test data. For indoor WiFi-based location fingerprinting, Lemelson et al. [12] propose several methods to estimate the positioning accuracy based on features of location fingerprints and the online signal strength samples. This paper focuses on GPS positioning, but the system can easily be used to collect information from other positioning technologies.

2.2 Using Information about Positioning Quality

The imperfection of positioning technology and its impact on position-based applications have been studied from several perspectives.

Benford et al. [4] study the problem in the context of a location-based game "Can you see me now?" where runners in the real world chase online players in a fixed game zone. The runners are equipped with GPS-enabled devices for positioning. Based on this experience, they identify five strategies for how to deal with imperfections or "seams" in their vocabulary. The five strategies are to remove, hide, manage, reveal or exploit imperfections. These strategies have also been studied in the context of the location-based game Hitchers based on GSM cell positioning [7]. Benford et al. [4] give evidence that the runners observed imperfections themselves with ongoing playtime and then integrated this knowledge into their gameplay strategies, for instance, by "hiding" in areas of low accuracy, where the runners' positions shown to the online players would "jump" or frequently disappear. To reveal imperfections in their game, Benford et al. propose that the system should distribute knowledge of positioning quality to runners and the technical crew. They propose to provide maps that show areas of good or bad positioning quality. As an example, a visualization for GPS with color maps is given, with data derived from log files.

Dearman et al. [6] have studied the impact of revealing positioning imperfections to quantify if people can take advantage of them. The authors based their conclusions on experiments with visualizing positioning accuracy for GSM positioning as circles. They give evidence that enhancing or replacing the sole target position with a circle, intuitively representing the size of the search area, efficiently supports people in their search which resulted in decreased time until the task was finished.

Oppermann et al. [14] motivate that authoring tools for position-based applications must visualize the imperfections of infrastructures. They propose that authoring tools should provide three layers of information: The physical world layer (e.g. maps and GIS), the infrastructure layer (e.g. visualizations of the infrastructure) and the content layer (e.g. regions, events and assets). They describe an authoring tool built on this model and use it to author a location-aware game named Tycoon based on GSM positioning. To provide the information for visualizing the infrastructure, they drive around the deployment area and map GSM cells. They point out that this method does not scale and argue that a system is needed which collects this information unsupervised during the game and can handle changes over time.

In comparison, this paper proposes PosQ, a system for unsupervised fingerprinting of GPS positioning quality that, both prior to deployment and during use, can provide information about positioning quality. The system represents changes over time, in 2D and 3D, and shows differences among GPS receivers. The system can provide visual overlays that can either be shown to users and administrators or imported into authoring tools to visualize the infrastructure.

3 PosQ - Fingerprinting Positioning Quality

The PosQ system consists of four key elements: GPS receivers that estimate positioning quality, PosQ collectors that efficiently report quality information to a server-side quality database, quality maps built from the database, and visual overlays generated from the maps (see Figure 1). Here, a visual overlay means a transparent information visualization layer enhancing the underlying map. In the following subsections, we describe each of these elements and how they are designed to help meeting PosQ's goals of unsupervised fingerprinting and visualization of GPS positioning quality.

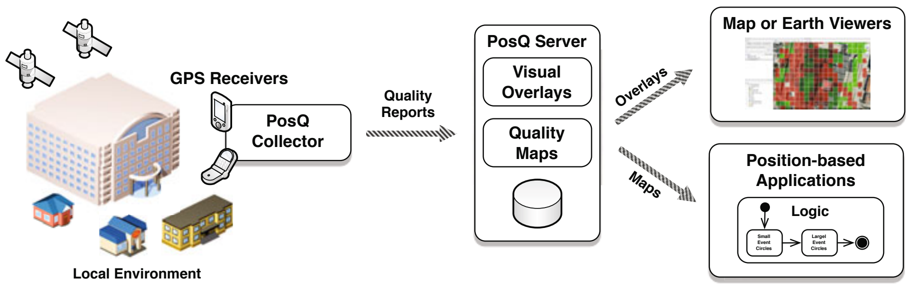

Fig. 1. The four key elements of the PosQ system

3.1 GPS Receivers

As described under related work, most GPS receivers will deliver – in addition to position and time estimates – the accuracy of the position fix. Such estimates are provided both by the built-in GPS receivers in HTC and Nokia mobile phones (available in their standard location APIs) and dedicated GPS receiver from manufactures such as U-Blox (available in PUBX NMEA messages) and Sirf (available using the Sirf binary protocol). PosQ collects such estimates for both the horizontal and vertical accuracy to fingerprint the positioning quality.

3.2 PosQ Collectors

The PosQ collector is a software component deployed on participating devices that collects quality tuples (time, position, horizontal accuracy, vertical accuracy) from either external or internal GPS receivers. The fingerprinting process is therefore fully unsupervised. Depending on the scenario, the information might be stored at the devices until collection is finished or be continuously reported to the server-side quality database. For the latter case, PosQ collectors apply aggregation as well as threshold and area-based updating strategies to limit the number of quality reports. These strategies are described later on.

The collectors can be deployed on the devices as part of a position-based application or as a stand-alone client. In the former case, the collector will enable the user or administrator to judge the positioning quality from visual overlays or for application developers to adapt the logic of their application to the positioning quality in a certain area. It might also be deployed as part of a positioning client, e.g., the client provided by Skyhook Wireless [1]. In the latter case, deploying the collector as a stand-alone client will support administrators in fingerprinting the coverage to judge the quality and know where to augment the positioning with other positioning means to reach a certain coverage and accuracy level.

3.3 Quality Maps

The quality map component provides an expressive and compact interface to the information stored in the quality database. The component aggregates the information stored in the database with regard to location, time stamp and GPS receiver type. Location means that position fixes which are geographically close are combined, because they are likely to have similar error estimates whereas far away positions are independent of each other. Time stamp, because GPS errors are similar within short periods of time. GPS receiver type, because GPS receivers differ in positioning coverage and accuracy. The component also provides meta information about the up-to-dateness of the information.

The component can provide both 2D and 3D quality maps of selected areas. The maps are represented as multidimensional histograms with either two or three dimensions. The histograms are built by aggregating the quality information placed within a histogram bin using the GPS position. The size of the bins is a configurable parameter of the system, but the values should be chosen small enough for the bins to only cover areas with similar positioning qualities. Therefore, for the two horizontal dimensions the size should be chosen to match the size of buildings or building parts whereas for the vertical dimension it should be chosen to match, e.g., the separation of floors.

The quality tuples aggregated by the component will, in high accuracy areas, contain accurate GPS positions and therefore be assigned to the correct bins. However, in low accuracy areas quality tuples might not be positioned within the correct bin and therefore be assigned to a neighboring bin, an effect that we call *pollution* of quality tuples. A *Polluted area* means an area that a calculated position is wrongly mapped to. This pollution of neighboring bins will depend on the accuracy: lower accuracy will increase the polluted area around the true position, but on the other side, an increased area will also lower the absolute impact on each neighboring bin. PosQ counters this problem by design by aggregating quality information over space and time. This means that in high accuracy areas visited by a collecting device, bins will contain many correct tuples. There may also be some polluted tuples, but because the correct tuples outnumber the polluted ones, the bins will be correctly mapped as having high accuracy. To further reduce pollution we include a weighting step in the aggregation that takes into account how jumpy the GPS positions are calculated from the recent history of GPS positions. To aggregate information over time, a weighted average is used that values recent information higher than older information.

To further reduce the pollution problem we have implemented a predictive component following the ideas of Steed [17] that based on information about satellite orbits and a city model estimates for each bin if more than three satellites can be received by line of sight. The information produced by this component is used to mark high accuracy areas prior to the addition of the fingerprinting data thus removing the pollution problem in these areas. We will in the rest of this paper not focus on this component as it is not able to provide detailed quality maps for urban canyons and indoor areas because it does not model signal reflection and attenuation.

In the case that a city model is not available the system might not have any data for a high accuracy area, in this case it could happen that only pollution from a nearby low accuracy area is present. However, firstly, this information will quickly be corrected when a collecting device enters this area and secondly, marking a good area as bad is less a problem than if the system was marking bad areas as good. Within low accuracy areas, the bins might pollute each other, which can slightly skew the values, but never change a low accuracy bin to be shown as a high accuracy one. In other words: Pollution is conservative of accuracy, meaning that it will not falsely upgrade a low-accuracy bin, even if it might downgrade one.

3.4 Visual Overlays

To inspect the positioning quality in an area, PosQ provides visual overlays that color an area depending on either the horizontal or the vertical accuracy of the associated accuracy bin. The overlays are produced in the widely used KML [19] format supported by many map and earth viewers. The map coloring is scaled from pure green for high accuracy (zero meters) over red (above ten meters) to black for low accuracy (forty meters or more) to easily communicate the properties of an area. The visual overlays can both be provided in 2D and 3D and be animated to show timely changes.

3.5 Implementation

The PosQ Collectors have been implemented for several operating systems and run on different types of devices, as listed in Table 1. The PosQ server-side

Table 1. PosQ Collectors implemented for different operating systems, programming languages, devices and GPS receivers

Operating Systems	Languages	Device	GPS		
			Chip	External	Built-in
Windows XP	Java	Asus Eee	U-Blox EVK-5H	•	
Symbian OS	Python	Nokia N95	TI NaviLink NL5350		•
Android OS	Python	Dev Phone 1	Qualcomm MSM7201A		•

quality map and visual overlay components have been written in Java and run in a standard OSGi [2] platform. The communication between the collectors and the server uses TCP connections. To provide easy access for position-based applications, the server-side components are accessible as webservices.

4 Evaluation of Quality Fingerprinting

PosQ's goal is to accurately fingerprint GPS positioning quality. To provide evidence for this ability, we will in this section present results from deploying PosQ in a shopping mall, a cultural history museum and a botanic museum. We will, firstly, argue that the quality maps provided by PosQ are accurate. Secondly, that the quality maps stay valid as the GPS satellite constellation changes over time. Thirdly, that the system can fingerprint the difference in accuracy among GPS receivers. For this presentation, we will focus on 2D quality maps and horizontal accuracy as this is the most common usage of GPS positions.

4.1 Quality Maps are Accurate

To argue that the quality maps provided by PosQ are accurate, we will discuss the results from using the Windows XP PosQ collector with an external U-Blox GPS to fingerprint three areas of a shopping mall: An outdoor area with generally high accuracy (less than three meters), an indoor area with medium accuracy (less than ten meters), and an indoor area with low accuracy (more than ten meters). The visual overlay produced by PosQ from the collected data is shown in Figure 2, together with the collection routes walked in each area. Ideally, in the high accuracy area the bins should be green and sit over the route, in the medium accuracy area the bins should be darker green or light red (worse than ten meters) and be slightly off the route and in the low accuracy area the bins should be mainly red and black (worse than forty meters) and be some bins off the route. As one can see from the figure, these differences indeed show up on the figure as, e.g., in the low accuracy area the bins are less often placed directly over the route as the positioning errors misplace fingerprints in neighboring bins. This even partly pollutes the bins around the medium accuracy area, but none of the medium-accuracy bins is polluted, because the many correctly placed tuples within these bins will outnumber the few polluted ones.

To study the absolute accuracy of the quality maps in greater detail we compare the PosQ quality map values with Root Mean Square (RMS) deviations from ground truth for the bins visited by the collector. However, the used U-blox receiver's documentation does not state which error measure or quantile the error estimates are supposed to match. We have chosen to compare to RMS because it is a standard error measure and is the measure reported by the JSR-179 location API. The RMS errors are calculated from the differences between the GPS positions stored in the collected tuples and manually recorded ground truth positions by the collector. Figure 3 plots in a scatter plot for each visited bin the RMS deviations from ground truth and the corresponding quality map value.

Fig. 2. Visual overlay for a shopping mall as a colored map with scale from bright green (high accuracy) over red to black (low accuracy)

The comparison shows that the map values are for most bins very predictive for the real error. In the low- and medium-accuracy areas, the bin values have a tendency to overestimate the real errors by a couple of meters. The increased deviation for low accuracy area can both be attributed to increased pollution but also that in weak signal conditions receivers will have poorer measurement (e.g., fewer satellites and more fluctuations) to use for the error estimation.

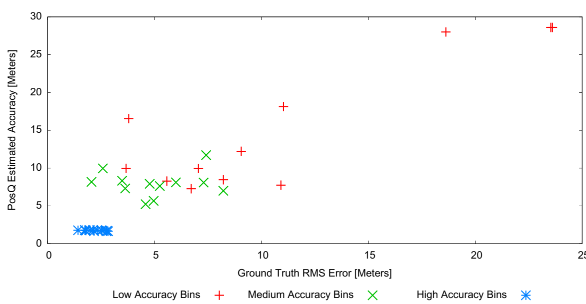

Fig. 3. Comparison of ground truth accuracy with PosQ estimated accuracy for individual bins

4.2 Quality Maps Remain Informative When Satellite Constellation Changes

GPS satellites are moving in their orbits around the earth and therefore the constellation of satellites visible over an area is changing all the time. The orbits of the satellites are designed in such a manner that the same constellation repeats nearly twice a day. For quality maps to be a useful tool, they have to remain informative when the satellite constellation changes. A main argument why they might remain informative is that the impact of the local environment (e.g., indoors because of attenuation) has a high influence on the accuracy. The experimental observations presented by Kjærgaard et al. [8] for indoor environments support this claim, as their results indicate that indoors low accuracy is better explained by multi-path and weak-signal effects than the geometric strength of the current satellite constellation measured as the *dilution of precision* factor.

(a) First Satellite Constellation (b) Second Satellite Constellation

Fig. 4. Visual overlay collected with a G1 device for two different satellite constellations

To support the argument that quality maps remain informative, we have collected fingerprints in a cultural history museum using the Android PosQ collector on a G1 device for two different satellite constellations. We collected the fingerprints on two different days at a different time of day to make sure that the constellations were different. The visual overlays produced by PosQ for each of the two quality maps are shown in Figure 4(a) and 4(b), respectively. Each day we walked the route shown in black on the overlays. The route covers local environment with different conditions such as: outdoors, outdoors in building shadows, indoors in buildings of light building materials (e.g. wood) and buildings of heavy building materials (e.g. bricks and stones).

In both figures one can observe that the G1 receiver is not able to provide complete coverage within this area. If one compares the areas in terms of coverage, the high accuracy areas in the lower left corners and in the middle left part of the overlays correspond. For the low accuracy areas in the upper middle and the upper right corner parts, both overlays mark this area as bad. The placement of the colored bins is a bit offset in the low accuracy areas which is due to the fact that the higher errors will offset the GPS positions in different directions. Therefore, the general trends remain the same, but there are some variations in other areas, where generally the accuracy on the overlay for the second constellation is slightly better than for the first constellation. These variations can both be attributed to changes in the number of people around and the changed constellation. PosQ handles these changes by updating its information as devices reenters an area.

4.3 Quality Maps Capture the Differences among GPS Receivers

The coverage and accuracy of a GPS receiver depend, among other things, on the receiver design, the receiver shielding which protects it from nearby electronic components, and the antenna [13]. Quality maps can capture such differences in GPS receiver performance, for instance, to determine that the same GPS receiver chip mounted in two different phone models might not result in the same accuracy because of differences in antennas or interference from nearby electronic components. To evaluate if the quality maps can capture the differences among GPS receivers, we collected fingerprints using the Windows XP PosQ collector with an external U-blox GPS chip at the same time as with the Android PosQ collector on a G1 device. The devices was carried several centimeters apart to avoid any interference. A limitation of the comparison is that the Android location API does not specify an exact error measure for the accuracy estimates.

The visual overlay produced by PosQ for the quality map is shown in Figure 5. The U-blox receiver has a good coverage in the area and nearly all parts of the walked route are colored, except for some bins in the area with several floored brick houses around the upper middle part. The accuracy is generally high, especially when compared to the accuracy of the G1 for the same constellation in Figure 4(b). If one compares the U-blox and the G1 results, it becomes apparent that even though the levels of accuracy and coverage are different, the same areas have been identified as difficult, for instance, the upper middle part.

5 Applications of PosQ

PosQ can collect quality maps that are accurate and remain valid over time, and it can capture differences in GPS receivers. This opens up possibilities for a number of applications. Firstly, PosQ can be used to design a positioning solution. Secondly, PosQ can help developers adapt their position-based applications to different GPS receivers. Thirdly, PosQ can visualize positioning quality over time and in 3D. We detail these in the remaining of this section.

5.1 Designing a Positioning Solution

In order to deploy a successful position-based application for the cultural history museum, an adequate positioning solution is required. To determine if GPS could be used, we collected fingerprints with PosQ by walking a tour through the museum for 35 minutes as described in the previous section. The result (as one can judge from the overlays in Figure 4 and 5) is that the dedicated U-blox receiver provided positioning in nearly all areas whereas the G1 phone only provided positioning in roughly half of the area. Therefore, a second positioning solution other than GPS is needed to provide coverage in the case of the G1 mobile phone. Because many smart phones today implement GSM and WiFi positioning, these are relevant second options. However, the accuracy of GSM positioning on phones is generally poor [11] and therefore the best secondary

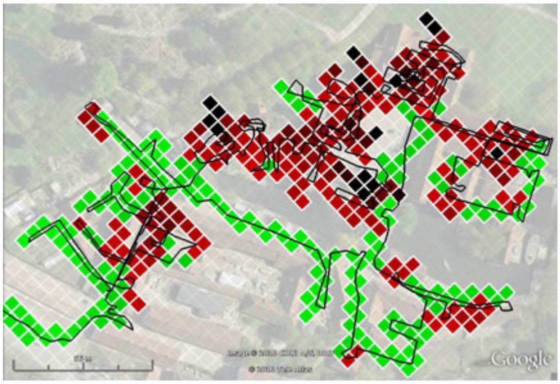

Fig. 5. Visual overlay collected with a U-blox receiver for the second satellite constellation

option is WiFi positioning. As companies providing WiFi positioning (e.g. Skyhook Wireless [1]) allow users to report the location of their own access point, a solution would be to deploy access points in the area and then enter them into such databases.

The cultural museum site is today not covered by WiFi access points and therefore one has to select a placement of access points for positioning. Because the cultural museum shows houses built between 1700 and 1920, the access points must not be mounted in visible locations. Therefore a goal is to add WiFi Positioning where it is needed, using a minimal number of access points. The visual overlays produced by PosQ can thus be used to select an optimal placement of access points. Other positioning options such as bluetooth or IR beacons could be considered using the same approach.

5.2 Mapping between Different GPS Receivers

Many position-based applications define trigger zones which are physical areas where some application logic should be executed as the user enters the area. A developer of such an application would often only design and test on a restricted number of mobile phones. The developer might use the results of Randell et al. [16], who advise that the size of trigger zones should be at least two times the accuracy of the positioning system. To customize the game play, the developer might further increase the size of trigger zones. In this process the developer can use the visual overlays produced by PosQ to decide on the zone sizes in specific areas as also suggested by Oppermann et al. [14]. The problem is: What happens when a user has a phone with another GPS receiver that might either be better or worse than the one used in the designing phase? A solution based on PosQ is to run PosQ on the mobile device and let PosQ compare the fingerprints for the designer's phone with the fingerprints collected by the user's phone to estimate differences in accuracy. We have implemented such a quality mapper extension to PosQ that enables the system to find such mappings.

To test the quality mapper, we used our fingerprints from the deployment of PosQ in a botanic museum and tested with a mapping from the G1 (designer's phone) to a Nokia N95 (user's phone). The case of N95 actually requires the system to find two different mappings as the GPS performance of the N95 depends on whether the keyboard is slid out or not. In the botanic museum case, the keyboard was not slid out, which contributed to the poor coverage of the N95 compared to the G1. The quality mapper compares the values of bins with fingerprinting data from both phones and applies linear regression to find the parameters for a linear model $y = ax$ where y is the values on the user's phone and x the values on the developer's phone. For the given scenario, the mapping parameter a was estimated to be 15.32 which means that the developer's phone is estimated to be fifteen times more accurate that the user's phone. In comparison, when the keyboard is slid out the factor is only 3.71. In an application you would use the factor to scale the size of the trigger zones to adjust to the new phone.

The reason why developers may want to adapt application logic using PosQ instead of the raw estimates from the GPS is that they can not judge the quality of a specific GPS from a single value. By comparing fingerprints between the developer's test phone and the user's phone, you can adapt to the differences between the two phones and not just the current GPS error.

5.3 Visualizing Positioning Quality over Time and in 3D

Visualization of positioning quality is relevant both for end users (as shown by Dearman et al. [6]) and for designers and developers (as argued for by Oppermann et al. [14]). Previous work has considered 2D visualizations for the horizontal accuracy, but, e.g., Oppermann et al. [14] pointed out that also changes over time are relevant. PosQ supports views over time by providing animated visualizations which can show how positioning quality changes as reported by collectors for different points in time.

Because PosQ is based on a 3D quality model and collects information about both the horizontal and the vertical accuracy, the system can visualize the vertical accuracy in 3D. For the cultural museum, a visualization of the 3D model

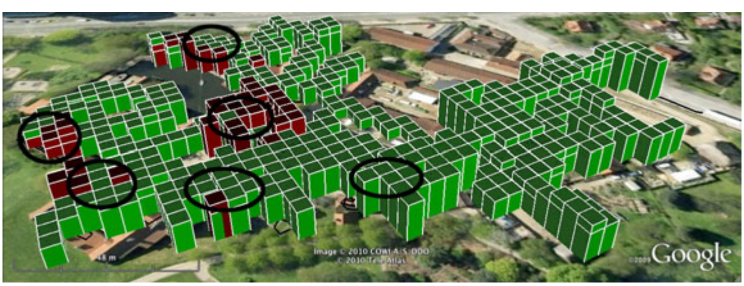

Fig. 6. Visual overlay in 3D for vertical accuracy

is shown with vertical accuracy information in Figure 6. The 3D visualization both provides hints about the curvature of the terrain but also changes for the buildings where the collecting person visited several floors. The places where the vertical accuracy is above ten meters (red to black) is actually co-located in most cases with the places shown with a black circle where the person entered a building and visited several floors within it. The floor changes makes the GPS more uncertain about the height compared to buildings where no floor changes occur.

6 Efficient Collection of Fingerprints

There are several reasons why PosQ should support the efficient updating of quality information in real-time. Firstly, it can be used as motivated by Benford et al. [4] in the context of location-based games to visualize the positioning quality for players and system administrators during play. Secondly, when using positioning for critical decision making, e.g. when the police or fire fighters enter an unknown building to build up a quality map over where positioning is possible and with what quality. In this section, we will focus on 2D quality maps and horizontal accuracy.

It would be problematic to support this by just letting PosQ collectors continuously report all collected information to the quality database. Firstly, the quantity of collected data adds up to 140 MB per hour, which would use a significant amount of bandwidth and result in scalability problems for the quality database as the number of participating devices increases. Sending this amount of data from the mobile devices would also be a problem because of the power consumption of the radio, e.g. the 3G radio in a N95 phone uses around 1.1 watt which if used continuously would decrease the battery life with a factor of twenty [9].

6.1 Updating Strategies

To address this, PosQ collectors implement several smart updating strategies for positioning quality. Firstly, the collectors apply aggregation to tuples to limit the amount of data per update. Secondly, they apply threshold-based updating to only send updates if the quality deviates from former values above a configurable threshold. Thirdly, they apply area-based updating to send an update when the device has moved through a configurable number of bins.

Figure 7 illustrates the steps involved. The PosQ collector demands GPS positioning and uses this initial position to request data for the bins (1). The quality database returns three parameters for each bin to the client: the aggregated accuracy, a time quality and a measurement quality (2). The collector processes a GPS position and the logic determines that a positioning quality update is not needed (3). Later, a GPS position is processed where the client-side logic determines that the used threshold has been exceeded and initiates a positioning quality update (4). The server processes the update and sends back updated bin data (5). Even later, a GPS position is processed where the collector logic determines that the device has moved more that the configured number of bins and initiates a positioning quality update (6). The server processes the update and sends back updated bin data (7).

6.2 Evaluation Setup

We have evaluated these strategies in our shopping mall deployment using the three different areas, as shown in Figure 2: an outdoor area with high accuracy, an indoor area with medium accuracy and medium changes in accuracy, and an indoor area with low accuracy and large changes in accuracy. For each area a path with four corner points was measured. The path in the high accuracy area has a length of 131 meters and passes through 30 bins, in the medium accuracy area the path has a length of 110 meters and passes through 30 bins and in the low accuracy area the path has a length of 168 meters and passes through 36 bins. To know the real position of the GPS receiver during evaluation, the position of the receiver was recorded manually.

Several runs of evaluations were performed along each path to match three motion patterns: (a) walking with stops at all four corner points of a path for one minute each and walking the path twice; (b) continuously walking the path three times with a stop of one minute in the beginning and in the end; (c) standing three minutes at a corner point and then walking to a second corner point to wait three more minutes.

For each evaluation, the PosQ collector and a reference client were run in parallel. The reference collector continuously sends reports to the quality database which means that it provides minimal communication efficiency but maximal information quality. The threshold-based strategy was configured with a deviation threshold factor of 1.5 and the area-based strategy with an area update

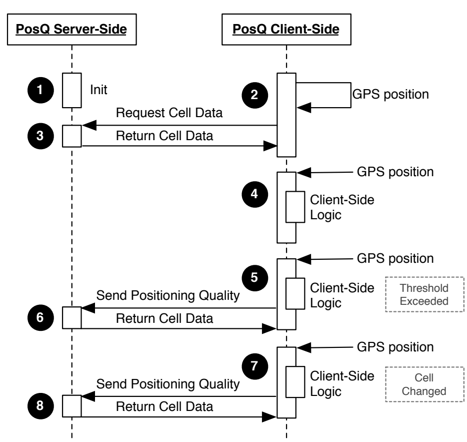

Fig. 7. The steps in the communication between the PosQ collector and the quality database

distance of two bins. These parameters were chosen to test the strategies in a difficult case with tight tracking of the quality in the area due to the small deviation threshold and the short update distance. The goal of PosQ is to use fewer messages than the reference client, while only causing a small or negligible loss in information quality.

6.3 PosQ Decreases the Number of Quality Reports

PosQ is able, when compared to the reference client, to decrease the number of needed quality messages as indicated by the results shown in Table 2 for *motion pattern a*. The number of messages sent by the reference client is linear with time whereas the PosQ client only sends updates when determined by the collector logic. However, because each PosQ update also includes a message with new parameters back to the collector each update is counted as two messages. The PosQ client is in the high accuracy area able to save a factor 9.3 compared to the reference client where as in the low accuracy area it is able to save a factor of 2.9. The difference is both due to that in the high accuracy area there were 47 threshold updates compared to two in the low accuracy area and that in the high accuracy area there were 69 area updates, while only 46 in the low accuracy area. The increased number of updates is due to the more varying accuracy throughout the area as can be seen from Figure 8(b). The results for the medium accuracy area (a factor of 4.8) is in between the two others.

Figure 8 shows the results of the evaluation for *motion pattern a* in both the low and high accuracy areas. Our results indicate that PosQ only impacts the information quality with a small loss compared to the reference client. In both areas, the updates delivered by PosQ tightly match the information reported by the reference client. This is also true for the medium accuracy area (not shown). Figure 8(a) and 8(b) also show when updates are due to bin changes as small arrows. For the low-accuracy area, there were 69 area updates and 46 area updates in the high-accuracy area, which is more than the 37 and 31 area updates necessary, respectively. This could of course be lowered by using a larger area update distance.

The motion patterns have an impact on PosQ performance, because most updates are produced in motion and fewer updates during standing periods. The reason is that the estimated accuracy is more stable when the receiver is

Table 2. Reference compared to PosQ for the number of messages for *motion pattern a* on the routes shown in figure 2

	Accuracy of Area		
	Low	Medium	High
Reference Messages	677	902	893
PosQ Messages	232	188	96
Improvement Factor	2.9	4.8	9.3

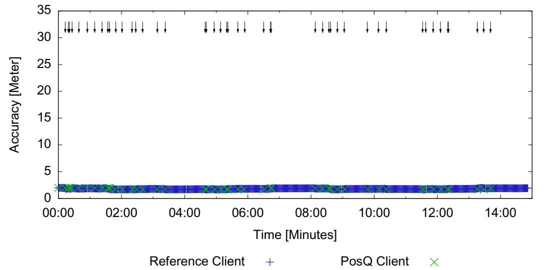

(a) Estimated accuracy in high accuracy area (arrows indicate area updates)

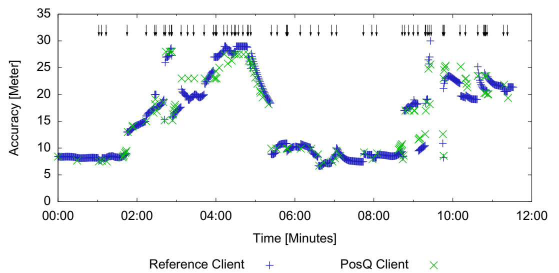

(b) Estimated accuracy in low accuracy area (arrows indicate area updates)

Fig. 8. Results for Positioning Quality

kept motionless and only few erroneous area updates are produced. This can also be noticed from our evaluations for *motion pattern c*, where most updates are produced for the motion-full *pattern b* compared to the nearly motionless *pattern c*. In summary, PosQ is able to decrease the number of needed quality messages with only a small loss in information quality.

7 Conclusion

This paper proposed PosQ, a system for unsupervised fingerprinting and visualization of GPS positioning quality. The system provides quality maps to position-based applications and visual overlays to users and managers to reveal the positioning quality. PosQ reveals the quality in three dimensions: As it changes over time, in 2D and 3D, and for each type of GPS receiver. The system enables unsupervised fingerprinting by running in the background of a GPS enabled device without any need of attention by the device user. It enables people and applications to judge the GPS positioning quality in a local environment with low effort, e.g., in a mixed indoor-outdoor environment in the scope of a city.

Our evaluation of the system provided evidence that the collected quality maps are accurate, that they remain informative over time and that they capture the differences among GPS receivers. The evaluation took into account different local environments and conditions. We demonstrated how the system enables the design of positioning solutions, mapping between different GPS receivers and visualization of positioning quality over time and in 3D. Furthermore, we proposed and evaluated methods for efficient collection of quality reports and showed that the methods can significantly decrease the number of quality reports needed when updating quality maps in real-time. As the PosQ collectors can be deployed on phones as part of applications the only disadvantage for users is that they have to accept the few extra resources that the system uses to realize the benefits of the system.

In our ongoing work we are going to address mainly four issues. Firstly, to integrate the system with other forms of positioning such as WiFi and GSM positioning to fingerprint their positioning quality. Secondly, integrate PosQ with an authoring tool for position-based applications and extend the tool with methods for making use of PosQ's capabilities to provide visual overlays over time and in 3D. Thirdly, study in more detail the problem of pollution of neighboring bins. Fourthly, consider how to integrate privacy protection measures for data contributors.

Acknowledgements. The authors acknowledge the financial support granted by the *Danish National Advanced Technology Foundation* under J.nr. 009-2007-2.

References

1. Skyhook wireless, `http://www.skyhookwireless.com/`
2. Alliance, O.: Open Services Gateway Initiative (2009), Specification download `http://www.osgi.org/Download/Release4V42` (Online, cited February 18, 2010)
3. Bahl, P., Padmanabhan, V.N.: RADAR: An In-Building RF-based User Location and Tracking System. In: Proceedings of the 19th Annual Joint Conf. of the IEEE Computer and Communications Societies (2000)
4. Benford, S., Crabtree, A., Flintham, M., Drozd, A., Anastasi, R., Paxton, M., Tandavanitj, N., Adams, M., Row-Farr, J.: Can you see me now? ACM Trans. Comput.-Hum. Interact. 13(1), 100–133 (2006)
5. Constandache, I., Choudhury, R.R., Rhee, I.: Towards Mobile Phone Localization without War-Driving. In: Proceedings of the Annual Joint Conf. of the IEEE Computer and Communications Societies (2010)
6. Dearman, D., Varshavsky, A., de Lara, E., Truong, K.N.: An Exploration of Location Error Estimation. In: Krumm, J., Abowd, G.D., Seneviratne, A., Strang, T. (eds.) UbiComp 2007. LNCS, vol. 4717, pp. 181–198. Springer, Heidelberg (2007)
7. Drozd, A., Benford, S., Tandavanitj, N., Wright, M., Chamberlain, A.: Hitchers: Designing for Cellular Positioning. In: Dourish, P., Friday, A. (eds.) UbiComp 2006. LNCS, vol. 4206, pp. 279–296. Springer, Heidelberg (2006)
8. Kjærgaard, M.B., Blunck, H., Godsk, T., Toftkjær, T., Christensen, D.L., Grønbæk, K.: Indoor Positioning Using GPS Revisited. In: Floréen, P., Krüger, A., Spasojevic, M. (eds.) Pervasive 2010. LNCS, vol. 6030, pp. 38–56. Springer, Heidelberg (2010)

9. Kjærgaard, M.B., Langdal, J., Godsk, T., Toftkjær, T.: EnTracked: Energy-Efficient Robust Position Tracking for Mobile Devices. In: Proceedings of the 7th Int. Conf. on Mobile Systems, Applications, and Services, pp. 221–234 (2009)
10. Küpper, A.: Location-Based Services: Fundamentals and Operation. Wiley (October 2005)
11. LaMarca, A., Chawathe, Y., Consolvo, S., Hightower, J., Smith, I., Scott, J., Sohn, T., Howard, J., Hughes, J., Potter, F., Tabert, J., Powledge, P., Borriello, G., Schilit, B.: Place Lab: Device Positioning Using Radio Beacons in the Wild. In: Gellersen, H.-W., Want, R., Schmidt, A. (eds.) PERVASIVE 2005. LNCS, vol. 3468, pp. 116–133. Springer, Heidelberg (2005)
12. Lemelson, H., Kjærgaard, M.B., Hansen, R., King, T.: Error Estimation for Indoor 802.11 Location Fingerprinting. In: Choudhury, T., Quigley, A., Strang, T., Suginuma, K. (eds.) LoCA 2009. LNCS, vol. 5561, pp. 138–155. Springer, Heidelberg (2009)
13. Misra, P., Enge, P.: Global Positioning System: Signals, Measurements, and Performance, 2nd edn. Navtech (2006)
14. Oppermann, L., Broll, G., Capra, M., Benford, S.: Extending Authoring Tools for Location-Aware Applications with an Infrastructure Visualization Layer. In: Dourish, P., Friday, A. (eds.) UbiComp 2006. LNCS, vol. 4206, pp. 52–68. Springer, Heidelberg (2006)
15. Patel, S.N., Truong, K.N., Abowd, G.D.: PowerLine Positioning: A Practical Sub-Room-Level Indoor Location System for Domestic Use. In: Dourish, P., Friday, A. (eds.) UbiComp 2006. LNCS, vol. 4206, pp. 441–458. Springer, Heidelberg (2006)
16. Randell, C., Geelhoed, E., Dix, A.J., Muller, H.L.: Exploring the Effects of Target Location Size and Position System Accuracy on Location Based Applications. In: Fishkin, K.P., Schiele, B., Nixon, P., Quigley, A. (eds.) PERVASIVE 2006. LNCS, vol. 3968, pp. 305–320. Springer, Heidelberg (2006)
17. Steed, A.: Supporting Mobile Applications with Real-Time Visualisation of GPS Availability. In: Brewster, S., Dunlop, M.D. (eds.) Mobile HCI 2004. LNCS, vol. 3160, pp. 373–377. Springer, Heidelberg (2004)
18. Varshavsky, A., de Lara, E., Hightower, J., LaMarca, A., Otsason, V.: Gsm indoor localization. Pervasive and Mobile Computing 3(6), 698–720 (2007)
19. Wernecke, J.: The KML Handbook: Geographic Visualization for the Web, 1st edn. Addison-Wesley Professional (November 2008)