Dynamic GPS-position correction for mobile pedestrian navigation and orientation

Jason Daniel Martin[a], Jens Krösche[b], Susanne Boll[c]

[a]Department of Computing Science, University of Oldenburg, Germany,
e-mail: jason.martin@informatik.uni-oldenburg.de
[b]University of Applied Sciences at Hagenberg, Austria,
e-mail: jens.kroesche@fh-hagenberg.at
[c]Department of Computing Science, University of Oldenburg, Germany,
e-mail: susanne.boll@informatik.uni-oldenburg.de

Abstract - Location has become one of the most relevant context parameters for mobile applications. Technology for the positiong such as Global Positioning System (GPS) has been widely adopted for locating the mobile user. However, for mobile applications that use an off the shelf GPS-receiver such as a PCMCIA card attached to a PDA, the correctness of the positioning is yet unsatisfying. Especially pedestrians, walking in urban areas, may experience strong deviations of the detected position from their actual position. In this paper, we present our approach to provide a software-based solution for a more correct position to a mobile pedestrian by still relying on consumer-grade GPS receivers. The results are small-sized, however, promising.

1 Introduction

Over the past several years, there have been significant advances in navigation technologies for the consumer market. This trend has been made possible by developments in the Global Positioning System (GPS), the development of GPS-based navigation applications, and the technological advancement of mobile computing devices. While most navigation applications on the consumer market to date have been developed for automobile navigation, a current area of research and development is that of GPS-based navigation for pedestrians. Pedestrian navigation applications would help pedestrians navigate through unfamiliar cities, for example, and would guide them to various points of interest. The development of software applications for pedestrian navigation is, however, faced with a number of unique obstacles.

The environments in which GPS-based pedestrian navigation applications would most likely be used are urban in nature; for example, a downtown district of a major city. Tall buildings and narrow passageways in these types of environments can present very difficult conditions for obtaining accurate GPS positions. This problem is augmented by the fact that most users of pedestrian navigation systems would use inexpensive, consumer-grade GPS receivers for positioning. Professional-grade GPS receivers, which are far too expensive for the general consumer, often provide GPS error correction technologies, such as hardware-based multipath-reduction, and have powerful antennas capable of obtaining GPS signals in difficult environments. Consumer-grade devices generally do not provide these advanced features. In addition, software-based GPS position correction methods common to automobile navigation applications, such as map-matching, cannot be easily transferred to pedestrian navigation applications. Thus, there is a need for a new GPS-position correction technology that addresses the unique problems of pedestrian navigation in urban environments.

Existing approaches in the area of high quality receivers work on correction on the signal level. Other approaches use additional sources of location information to enhance the position such as Differential Global Positioning System or use additional sensor information such as RFID to increase the precision of the determined position. We also find a research field that exploits further modalities such as video with the location information and combines positioning and and computer vision techniques to improve the received signals correctness. The overall goal of our project, however, was to design, implement, and test a software-based dynamic GPS-position correction technology that addresses the current problems of GPS-based pedestrian navigation in urban environments and that can be integrated into existing pedestrian navigation applications. Two primary properties of the developed system are that (1) it relies solely on a single, consumer-grade GPS receiver for positioning and that it (2) only uses the subset of positioning information that is commonly provided by inexpensive, consumer-grade receivers. The system architecture is based on the Pipes and Filters architectural pattern [BMR+96], which provides a basis for applications which filter an incoming stream of data. Using this pattern, the developed system successively filters and corrects incoming GPS positions via a series of position-correction filters. This paper introduces the developed filtering system as well as the individual position-correction filters in Section 2. In Sections 3, present the test environment and results 4, and discuss the effectiveness of each filter and filter combination in reducing GPS error with respect to pedestrian navigation.

2 System Design and Implementation

Our position correction system takes the original signal as received from the GPS receiver and shifts it through a set of sequential filters. In each filtering step, the system enhances the position by a specific strategy and hand the result over to the next step. The following Section 2.1 introduces this Pipes and Filters architecture before the single filters and their functionality are presented in Section 2.2.

2.1 General System Design

The GPS position correction filtering system is modelled after the Pipes and Filters architectural pattern and is meant to be integrated into an existing pedestrian navigation application. By intercepting incoming GPS positions and passing them through the filter pipeline prior to making them available to the pedestrian navigation application, an improvement in positioning accuracy can be obtained. The general filtering process is illustrated in Figure 1. This flexible and extendable architecture allows any number and combination of filters to be used in the filter pipeline as well as the easy integration of new filters and the customization of the filtering process for various operating environments. The system was implemented using the Java programming language and uses property files to allow the end user or application developer to customize the filter pipeline without having to alter the source code.

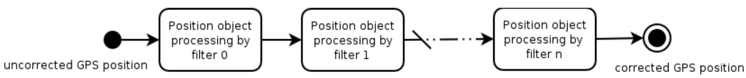

Figure 1: Conceptional filter design

GPS data is most commonly transmitted via the NMEA (National Marine Electronics Association) protocol, in which positioning information is contained in comma-separated ASCII-text sentences. Before this data can be filtered and / or corrected, it must be parsed and packaged into position objects. These position objects can then be passed down the filter pipeline for analysis and correction. Once corrected, they can be provided to the host application for further processing. As part of this project, four position correction filters were developed and implemented: a filter for removing invalid GPS positions, a filter for enforcing pre-defined position quality standards, a map-matching filter specifically designed for pedestrian navigation, and a filter to remove logically inaccurate positions. These filters were conceived, developed, and implemented based on their expected potential to reduce overall GPS error. The following sections introduce and discuss the four position correction filters in more detail.

2.2 Correction Filters

Invalid Position Filter: The goal of the Invalid Position Filter was to identify and remove GPS-positions deemed inaccurate by the GPS receiver. If a GPS receiver deems a position to be invalid, it sets a numeric flag in the "'GGA"' NMEA data sentences to denote this. This information is then included in the resulting position object that is supplied to the filter pipeline. A GPS receiver can invalidate a position for several reasons. For example, a position might be invalidated if there were less than 3 satellites available for the position computation or if the position is too far away from the previous position to be logically valid. Although the criterium for position invalidation varies from receiver manufacturer to receiver manufacturer, it can be reasoned that if a position has been deemed invalid for any reason, it should not be made available to the host application. In this case, the Invalid Position Filter nullifies any position object which is flagged as being deemed invalid by the receiver. By nullifying the position object, it cannot be used by the host application and is disregarded. Once a position object has been processed by the Invalid Position Filter, it is passed down the filter pipeline to the next available filter.

Mask Filter: The goal of the Mask Filter was to remove GPS-positions that do not meet certain quality standards. The NMEA data protocol allows a GPS receiver to transmit information regarding the expected quality of a GPS position. Quality measurements include the Position Dilution of Precision (PDOP), the Signal-to-Noise Ratio (SNR), and the dimension of the position (either 2D or 3D) [DMW05]. A defined threshold level for a specific measurement is often referred to as a "mask". For example, a PDOP of less than 6 is indicates that the spatial geometry of the satellites used to compute the position was very good and that the position is most likely accurate. A high SNR value indicates that the signal strength of a particular satellite is good, which can aid in accurate positioning. In addition, for truly accurate positions, at least four satellites should be used in the position calculation (such a position is known as a 3D position). The Mask filter enables the user to define which quality measurements should be used for filtering (the masks) and what the respective threshold values should be. This information is defined in the properties file, which allows the user or application developer to customize the mask filter as desired. If the incoming GPS position does not meet the quality levels required by the masks, the position object is nullified and passed down the pipeline.

A static nullification of position objects that do not meet the specified masks is, however, undesirable. If GPS conditions are continually poor and the threshold valued are set too high, the situation could arise that every position is nullified. An important aspect of GPS quality measurements is that they represent the statistical accuracy of a position. This means that poor PDOP and SNR values do not necessarily result in an inaccurate position, only that the chances for the position being accurate are low. For these reasons, a dynamic filtering of GPS positions was implemented for the mask filter. The mask filter continually monitors GPS conditions by collecting statistics on the number of positions rejected versus the number of positions accepted by the mask filter. If conditions remain inadequate for some time, the mask filter dynamically adjusts the threshold levels of the masks. If conditions do not improve, the masks are successively disabled, so that the positions are sent down the pipeline regardless of their quality measurements. The motto here is that a bad position is better than no position at all.

Map-Matching Filter: The goal of the Map-Matching Filter was to determine the areas in which a user of a pedestrian navigation system would be travelling and correct any positions that do not adhere to this assumption. As stated before, map-matching methods used for automobile navigation systems cannot be easily transferred to pedestrian navigation system. The reason for this is fairly strait forward; there are no "rules of the road" for pedestrians. Automobile navigation systems can make the assumption that the user must be on a particular road. If a GPS position is received that is near a road, it can be "snapped" back onto the road. On the other hand, pedestrians have complete freedom of movement which means that defining the allowable routes of travel is nearly impossible. If no position correction at all takes place, however, it can lead to confusion and frustration on the part of the user.

Therefore, a different strategy had to be developed. The only assumption that can be made with respect to GPS-aided pedestrian navigation is that users will not be trying to navigate in or through buildings, lakes, or other obstructions. These areas can be thought of as areas in which a pedestrian would not possibly be travelling. In this way, they can be designated as "off-limits" by the Map-Matching filter. If a GPS position is received that is within one of these areas, it can be safely reasoned that the position must be inaccurate. The next step is to project the position back to the edge of the area. Figure 2 shows the downtown district of Oldenburg, Germany and outlines the areas of downtown Oldenburg in which a pedestrian would not be travelling.

Figure 2: Layout of downtown Oldenburg

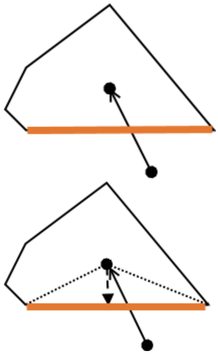

Figure 3: Determining intersected side of Polygon

To define these areas, the Geographic Markup Language (GML) was used. GML is an XML-Schema application that provides an extensible framework for defining and describing geographical objects and features. Any areas deemed "off-limits" by the application developer can be defined as polygons in a GML file which is supplied to the Map-Matching filter at runtime. The Map-Matching filter analyzes each incoming position and checks whether it is within a non-allowable area. If a GPS position is within a defined polygon, then the coordinates of a new position that lies just outside the polygon are computed. The original position coordinates are then overwritten with the corrected coordinates and the position object is sent down the pipeline. The first step in correcting the position coordinates is determining the side of the polygon that was "intersected" by the route of travel. Figure 3 illustrates the corrrection back to this side.

Statistical Correction Filter: The goal of the Statistical Correction Filter was to identify and "smooth out" incorrect positioning. One distinct characteristic of inaccurate GPS positions is a sudden and drastic change in user direction and speed. By identifying such events, the statistical correction filter attempts to adjust the error in speed and direction to produce a "smoother" route of travel.

The algorithm behind the Statistical Correction Filter keeps track of the average distance travelled between GPS positions. If the distance to a new position is more than 25% greater than the average distance between positions, the distance to the corrected position will be the average distance plus 25% of the difference between the average and actual distances. This allows for a slight increase in travel speed for the pedestrian user. In addition, the Statistical Correction Filter analyzes the difference between the last recorded course and the actual course for every GPS position. If the course varies from the previous course by more than 45°, a new course is calculated that lies halfway between the previous course and the actual course. Once the distance and course have been adjusted, the coordinates of the new position are calculated. This new position is then passed down the filter pipeline. The Statistical Correction Filter can only partially eliminate this type of GPS error because in the end, it is impossible to know for sure if the user has really changed his or her direction and / or speed.

3 Implementation and Test

For testing purposes, the position correction filtering system was integrated into the position module of the Niccimon platform [RGB+03, BBK+04b], a framework for individual mobile software development. The different autonomous modules offer basic functionality that is often required within location-based mobile applications. The architecture of the Niccimon platform is illustrated in Figure 4.

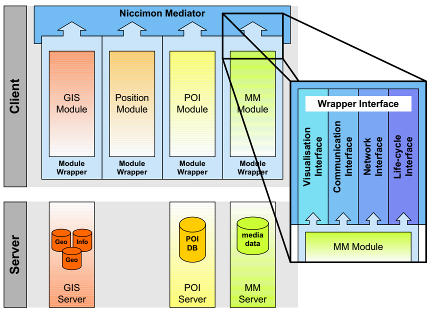

Figure 4: The modular Niccimon platform [BBK+04a]

3.1 Test Environment

To test the developed GPS position correction filtering system in a real-world scenario, a reference route through downtown Oldenburg was decided upon that was characteristic of a path a user of a pedestrian navigation system might take. Downtown Oldenburg is characterized by relatively tall buildings (4-5 stories high) and narrow passageways as shown in Figures 5(a) - 5(d).

(a)

(b)

(c)

(d)

Figure 5: Oldenburg downtown district

The next step was to collect and log raw GPS data along this route. This was done using an inexpensive consumer-grade GPS receiver connected to a PDA. This data served as the "uncorrected" GPS data. After this data was obtained, it was possible to compute the amount the uncorrected GPS data varied from the reference route. When the GPS data varied from the reference route, a polygon could be derived (see Figures 6(a) and 6(b)) whose area could be measured using the following formula [Bey88]:

$$\text{Area} = \sum_{i=0}^{n-1} (x_i y_{i+1} - x_{i+1} y_i)$$

where x and y represent the respective vertex coordinates. The sum of these areas represents the total amount of variance between the GPS data and the reference route.

Next, each filter and filter combination was tested to see which produced the greatest reduction in GPS error. This was done by processing the logged GPS data with all filters individually and then with different filter combinations. The results of filtering were logged and then compared, as with the uncorrected GPS data, with the reference route to determine the amount of variance. If the sum of the polygon areas obtained from using the filtered data was less than that obtained by comparing the uncorrected GPS data to the reference route, it could be concluded that an improvement had been achieved.

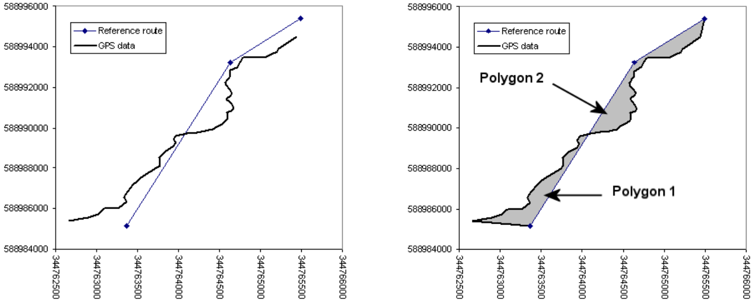

(a) (b)

Figure 6: GPS and reference route data

4 Results of the Position Correction

To automatically determine the area variance between the filtered and non-filtered GPS information and the reference route, a program was developed. This program determines the polygons and their areas as described above, adds the areas together and generates a summary report with detailed analysis information. The first step was to find the area variance of the original GPS data with respect to the reference route. This was determined to be $12,553$ m^2. Thus, the original GPS data contains more than 1.2 hectares of error over the entire data set [1].

The first filter to be tested was the Invalid Position Filter. However, the original GPS data set did not contain any positions deemed invalid by the receiver. This filter would not have corrected the GPS data in any way. The fact, however, that there was over $12,500$ m^2 of error and not one position was deemed invalid, is surprising. In some cases, valid positions were received that registered no satellites in view, which is, of course, impossible. This being said, it is fairly safe to say that even if a few positions were to be deemed invalid by the receiver, the overall improvement in accuracy using the Invalid Position Filter alone as it is currently implemented would be marginal at best. The fact that the Invalid Position Filter could not improve any positions with this GPS data set left 3 Filters to analyze; the Mask Filter, the Statistical Correction Filter, and the Map-Matching filter. These were first tested individually. Then, all possible combinations of the three filters were tested. The results of these tests are summarized in Table 1.

Processing the original GPS data with the Mask Filter alone only improved the accuracy by 28 m^2. Given the fact that there is a calculation error of approximately $\pm 1\%$, it can be concluded that no real change in accuracy was obtained. This was due in part to the fact that many of the inaccurate positions in this data set had acceptable mask values. The mask filter, therefore, enforced the masks throughout the entire data set. This proves the fact that good quality measures (e.g., PDOP, SNR, 3D) do not necessarily result in accurate positions.

Filtering the GPS data through the Statistical Correction Filter alone also did not result in an improvement in accuracy. On the contrary; there was a slight increase in error ($12,700$ m^2).

[1] Almost all of the determined areas contain a slight inaccuracy due to the fact that the formula used assumes that the polygons are non-self-intersecting, which is not always the case. After further analysis, however, this error was determined to be within $\pm$ 1% of the total area.

Filter (Combination) Test	Error (m^2)
Original GPS	12553.09
Mask	12525.00
SCF	12700.68
Map	10298.30
Mask-SCF	12354.19
Mask-Map	10245.94
Map-SCF	10304.57
Map-Mask	10247.60
SCF-Map	10161.47
SCF-Mask	12612.62
Mask-SCF-Map	10194.07
Mask-Map-SCF	11014.84
Map-SCF-Mask	10051.05
Map-Mask-SCF	11267.69
SCF-Mask-Map	10063.17
SCF-Map-Mask	10136.45

Table 1: Results of filter / filter combination tests

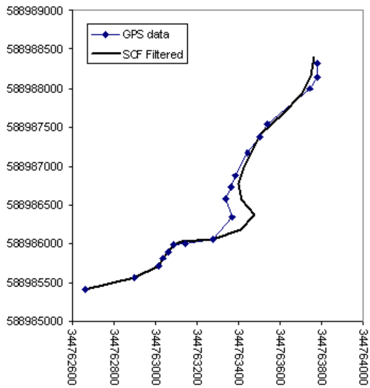

Table 2: Statistical Correction filter

This can be attributed to the fact that sharp changes in direction are "smoothed" out by the Statistical Correction Filter, resulting in potentially larger variance polygons. Inaccurate positions often result in sharp direction changes, which could actually decrease the amount of area variance. This type of behavior, however, does not reflect pedestrian movement and is exactly the type of behavior that the Statistical Correction Filter is designed to correct. Another point to note is that the corrections performed by the Statistical Correction Filter were, on an overall scale, fairly small. This is due to the fact that the GPS data was collected while walking at a slow speed. This resulted in many positions being located very close together. The correction takes place but on a very small scale. To illustrate this, consider Figure 2, which shows the first 25 GPS positions and the correction that took place through filtering with the Statistical Correction Filter. As can be seen, the Statistical Correction Filter identified sharp changes in direction and distance and computed the new positions accordingly.

The Map-Matching filter produced the best results as a stand alone filter, reducing the area variance by $2,254$ m^2. In addition, the positions were corrected to reflect the fact that GPS positions cannot be collected inside of buildings. Figure 7 shows the GPS data after being filtered by the Map-Matching Filter and the polygons that represent the buildings and other areas specified as "off-limits". As can be seen, no positions are located within the specified polygons. However, the situation can arise that the next valid position is "on the other side" of a polygon. In this case, the next position is accepted, which if viewed graphically, causes a line to be drawn across the polygon to the new position. This situation can not be easily prevented because, in the end, the mobile application can never know if a GPS position is truly accurate or not. If this was prevented, and the user really was on "the other side" of the polygon, the system would not accept any positions there. The fact that the system can never know if a position is accurate or not and the fact that it cannot predict future positions limits the amount and type of correction that can logically be applied.

The final step was to compare all possible dual and triple combinations of the three filters

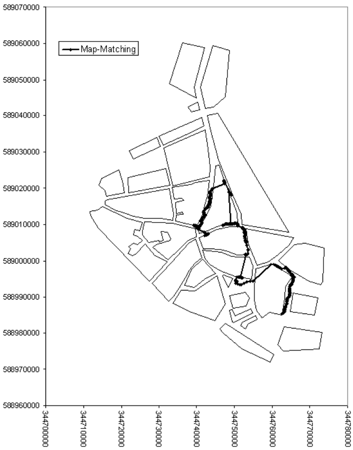

Figure 7: Map-Matching filter

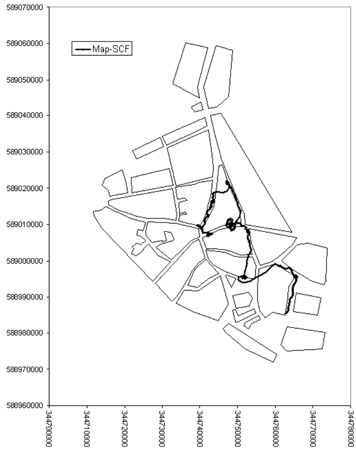

Figure 8: Map-Matching Filter followed by SCF

to determine if they improved accuracy. During analysis, there were two trends that became evident. First, filtering the GPS data with the Statistical Correction Filter and the Map-Matching Filter, in that order, seemed to produce the best results. The best overall reduction in error was achieved with the combination: Mask Filter → Statistical Correction Filter → Map-Matching Filter. The top four filter combinations in error reduction all filtered the GPS data first with the Statistical Correction Filter and then, at some point, with the map-matching filter. The Mask Filter seemed to have little effect on error reduction. The second trend was that filtering GPS data with the Map-Matching Filter before filtering it with the Statistical Correction Filter did not produce good results. The Statistical Correction Filter's goal is to detect sharp changes in direction and distance and modify these to "smooth out" these changes. The Map-Matching Filter, however, produces sharp changes in direction when positions are projected back to the edge of a polygon. The result of the Statistical Correction Filter being further down the pipeline as the Map-Matching Filter is that the sharp changes in direction caused by the Map-Matching Filter are corrected by the Statistical Correction Filter. This ultimately results in a circular-pattern of positions being produced around the projected positions from the Map-Matching Filter, as shown in Figure 8. Thus, placing the Statistical Correction Filter above the Map-Matching Filter in the pipeline will generally increase error and is not recommended.

5 Conclusion

The Map-Matching Filter proved to be the only filter that significantly reduced GPS error. The downtown district of Oldenburg is a maze of narrow passageways between tall buildings. Correcting the incorrect positions in this case meant projecting the positions back onto the narrow passageways, which coincided more closely with the actual path taken. The Invalid Position Filter relies on positions being deemed invalid by the receiver. The test data set, however, contained no invalid positions, although many of the positions were very inaccurate.

Thus, no correction could take place. The Mask Filter did not produce a significant reduction in error, primarily because the delivered GPS-position quality measures do not directly reflect position accuracy, rather they are statistical measures that can be used to make an assumption of the quality of the position calculation. Many of the inaccurate positions in the test data had acceptable position-quality values and consequently, were not rejected by the Mask-Filter. Lastly, the Statistical Correction Filter corrected sharp changes in direction and speed, however, not enough that a user would benefit from it. This resulted from the speed of the user being very slow and many positions being collected over a short distance. The corrections were thus on a very small scale and hardly noticeable to the user.

Two major obstacles presented themselves when developing the position-correction filters. The first obstacle was that the only source of positioning information was a single, consumer-grade GPS receiver. It is nearly impossible to know if a GPS-position is truly accurate or not, because a second positioning source is not available for verification. An exception to this are positions that are obviously inaccurate, hence the success of the Map-Matching Filter. Consequently, position accuracy can only be roughly estimated at best and these estimations have a high rate of error. The Map-Matching and Statistical Correction filters relied on these estimations for correction and thus many corrections were prone to error from the outset. The second obstacle was the fact that it is nearly impossible to make future predictions about where a pedestrian navigator will travel. They do not have to adhere to any rules of movement and thus, differentiating accurate from inaccurate positions, unless it is obvious, becomes much more difficult. In order for this system to produce more significant reductions in GPS error, other technologies, such as the integration of inertial navigation systems and new filter concepts such as Bayes and Kalman filtering [FHL+03] need to be researched. The system itself, however, provides a solid basis for any and all future work.

References

[BBK+04a] BALDZER ; BOLL ; KLANTE ; KRÖSCHE ; RUMP ; MEYER ; SCHERP ; APPELRATH: *Location-Aware Mobile Multimedia Applications on the Niccimon platform.* August 2004

[BBK+04b] BALDZER, Jörg ; BOLL, Susanne ; KLANTE, Palle ; KRÖSCHE, Jens [u. a.]: Location-Aware Mobile Multimedia Applications on the Niccimon Platform. In: *Informationssysteme für mobile Anwendungen (IMA'04).* Brunswick, Germany, October 2004

[Bey88] BEYER, William H.: *CRC standard mathematical tables.* 28. CRC Press, Boca Raton, FL, 1988. – ISBN 0849362180

[BMR+96] BUSHMANN ; MEUNIER ; ROHNERT ; SOMMERLAD ; STAL: *A System of Patterns.* John Wiley & Sons Ltd., 1996. – ISBN 0 471 95869 7

[DMW05] DR. MICHAEL WÖSSNER, Dr. Anja K.: *Wie funktioniert GPS?* Version: March 2005. http://www.kowoma.de/gps/index.htm. Word Wide Web. – Online Ressource, Last checked: 07/ 24/ 2005

[FHL+03] FOX ; HIGHTOWER ; LIAO ; SCHULTZ ; BORRIELLO: Bayesian Filtering for Location Estimation. In: *Pervasive Computing* (2003), July-September, S. 24–33

[RGB+03] RUMP, Norbert ; GERAMANI, Konstantina ; BALDZER, Jörg ; THIEME, Sabine ; SCHERP, Ansgar ; KRÖSCHE, Jens ; MEYER, Jochen: Potentials of pervasive computing in highly interactive workplaces. In: CHA, J. (Ed.) ; JARDIM-GONÇALVES, R. (Ed.) ; STEIGER-GARÇÃO, A. (Ed.): *10th ISPE International Conference on concurrent Engineering: Research and Applications.* Madeira Island, Portugal : A. A. Balkema Publishers, Lisse/Abingdon/Exton(PA)/Tokyo, July 2003, S. 733–739