A Pocket Guide to Indoor Mapping

Pascal Bissig, Roger Wattenhofer, Samuel Welten,
Distributed Computing Group - ETH Zurich,
firstname.lastname@tik.ee.ethz.ch

***Abstract*—In this paper, we present a way to obtain accurate WLAN signal strength maps in indoor environments, without dedicated hardware, and without a time consuming and complicated training process. We need two contributions towards this end. First, we present a novel dead-reckoning technique, to gather accurate user motion estimates. This motion data is combined with information about the signal strength of access points of the wireless infrastructure. Our second contribution lies in the efficient integration of this complementary information into a system that allows for easy mapping and requires nothing but a smartphone. All the required data is gathered from people walking casually around in an area of interest while carrying a smartphone in their pocket.**

***Index Terms*—Spatial signal strength distribution; motion estimation; localization; SLAM, Graph-SLAM;**

I. Introduction

Today, GPS is an essential component of the global information infrastructure. Similarly to the Internet, its applications are affecting many aspects of modern life. Although GPS satellites cover our globe, and current smartphones are equipped with suitable receivers, GPS based localization still has its blind spots, in particular if a user enters a building. Often, navigating inside an unknown building can be just as hard (if not harder!) than navigating outside, in free space.

WLAN signal strength based indoor localization is widely used today. However, for such a system to work, knowledge about the spatial signal strength distribution is required. In this paper, we present a way to obtain accurate WLAN signal strength maps that can be used to localize any WLAN enabled device localization similar to GPS also in indoor environments, without dedicated hardware, and without a time consuming and expensive training process.

Our paper has two main contributions. The first contribution is a novel motion estimation technique that allows us to gather accurate user motion data in an unobtrusive way. The required sensors, a 3-axes accelerometer and a 3-axes gyroscope, can be found in many modern smartphones. Using these sensor measurements, we show how the direction of the leg and hip can be accurately estimated, independent of how the smartphone is placed in the trouser pocket. Based on these estimates, we can track the heading as well as the distance of each step the user takes. This motion model is tailored to fit the requirements of tracking users in indoor environments and therefore does not utilize the (often unreliable) magnetic field to determine the absolute heading. Instead, the change in heading between steps is estimated. We evaluate the performance of this motion model and show that despite our non-restrictive sensor placement and the lack of absolute heading, we still manage to get a qualitatively high motion tracking performance. We also show how the motion model is able to extrapolate from a single configuration parameter to work at different walking speeds. As we will show, our motion data is locally highly accurate, i.e. short walks with a smartphone do exhibit a small relative positioning error.

The positioning error is growing with distance though, so we use information about the signal strength of access points of the wireless infrastructure to globally correct that. Our second contribution lies in the integration of the complementary information – the locally accurate motion data and the globally accurate signal strength data – into a system that allows for easy mapping with nothing but a smartphone in the pocket. More precisely, we obtain additional positioning constraints between locations that the user has visited while walking around in the area of interest. We use the common (least-squares) Graph-SLAM technique to obtain corrected positions for all the signal strength measurements that a user recorded. Compared to the Monte-Carlo technique which is most commonly used for such applications, the Graph-SLAM method leads to maximum likelihood estimates of the signal strength measurement locations. Furthermore, Monte-Carlo approaches have to back-propagate newly gained insights after loop closure; this is not required when using least squares and greatly improves performance when combining data from multiple measurement sessions or users. As a result, the system is capable of recovering accurate relative positions of the signal strength measurements from data that was captured by multiple users walking through an area at different times. We show a qualitative example for the mapping performance in a large university building.

II. Related Work

Crowd-sourcing Localization. Recent work presented by Rai et al. [1] shows the high interest in localization solutions relying on crowd sourced data rather than time consuming training. The introduced system [1] can provide localization information using only crowd-sourced information. Motion estimates and WLAN signal strength measurements are fused using a particle filter approach which also requires a floor plan map. As a result, the locations of the signal strength measurements within the floor plan can be estimated. High localization accuracy is achieved. However, the required floor plan map has to fulfill strict requirements in order for the particle filter to converge to the correct WLAN signal strength measurement locations. The probability density that has to be estimated with a particle has four dimensions which implies a high computational complexity on the mobile device. Our approach does not require a map and the use of efficient sparse

non linear solvers that are readily available, the computational complexity is lower.

Signal Strength Mapping. The problem of acquiring signal strength maps in an unsupervised manner has been approached by Chintalapudi et al. [2]. The system they presented allows WLAN signal strength based localization without annotated training data. No motion data is used in the process. Instead, the relative positions of WLAN Access Points and mobile node locations are modeled as a function of the received signal strengths. The resulting set of overdetermined geometrical constraints is solved using a genetic algorithm. Occasional GPS location measurements are used to obtain location information in the global frame. However, GPS measurements are hard to obtain in indoor environments and may be erroneous. The received signal strengths are the only indication for geometrical distance. Due to non-line of sight effects in indoor environments the complexity of this relationship is beyond the reach of the simple model used in the paper. This shortcoming becomes evident as the localization accuracy drops if the accurate locations of the access points are provided in the training phase. Therefore, the model might be accurate enough to provide localization capabilities, but the estimated map will in reality never converge towards the accurate access point distribution. In our approach, a person's motion is estimated in order to get geometrical distance information. The measured signal strengths are only utilized to recognized previously visited locations. Therefore, even if associations are incorrect, our maps converge to the accurate signal strength distribution because the increasing number of unbiased distance estimations from our motion model.

Motion Estimation. The dead-reckoning method we present in this paper is the enabling factor that allows us to create signal strength maps. Many such methods to estimate human motion using foot-mounted sensors have been described in the literature [3], [4], [5], [6]. These solutions rely on the foot being stationary when in contact with the ground. This zero-velocity interval can be used to counteract the velocity estimation drift caused by integrating the accelerometer signals. An overview of different methods to estimate motion based on foot-mounted sensor measurements have been compared by Skog et al. [7]. Another approach, relying on sensors fixed on a helmet was presented by Beauregard [8]. As our goal is to perform mapping in an unobtrusive way, we desire to measure user motion based on its natural location when users are casually walking around. Therefore, all motion estimation methods that require dedicated hardware such as bodily fixed sensors are useless if we want to allow a wide range of users to contribute in the mapping process. Our motion estimation method only requires a modern smartphone being placed in a trouser pocket which makes it tangible for many people now, and even more so in the future. A motion estimation method that is seemingly similar to the one presented in this paper was described by Blanke and Schiele [9]. This method and a selection of others that are equally nonrestrictive about the positioning of the sensors have been experimentally compared by Steinhoff and Schiele [10]. While the motion direction is estimated using the available sensor measurements, the step length was assumed to be constant and even manually set for each track. The presented evaluation does not allow us to quantitatively benchmark our approach because the evaluation results are distance- and orientation error quantities for each stride. The results are median stride length errors of more than $10\mathrm{cm}$ and median orientation errors of more than 5° whereas we achieve mean localization errors of below 10% of the traveled distance. However, the argument that an estimation error of $10\mathrm{cm}$ per stride is already more than 10% of the traveled distance is highly questionable. Partly due to measurement errors in the ground truth but also because our approach estimates not only step- heading and count but also distance. Also our motion model can cope with different walking speeds without requiring the stride length to be manually configured. In addition to this, our approach does not rely on the use of magnetic field sensors to determine the absolute heading direction. Due to the large magnetic disturbances that may interfere with finding the correct direction towards north, using magnetometers in indoor environments is questionable and may incur large and unexpected errors. As a result, our approach is not able to estimate absolute headings for the steps, but only the change in heading for consecutive steps. Another system that is seemingly similar to ours, but requiring two separate, fixed sensors to estimate motion based on the orientation of the thigh was presented by Lee and Mase [11].

SLAM. In the robotics community, the Simultaneous Localization and Mapping Problem (SLAM) has been an active research topic for over twenty years. Bailey and Durrant-Whyte [12], [13] summarized the most important results and solution ideas. To apply these SLAM solutions, two complementary information sources are required. Firstly, the motion of the measurement device has to be estimated. Secondly, previous locations have to be recognizable if visited again. Both requirements can be met with off-the-shelf smartphones. The known SLAM solutions already have been applied to the indoor mapping problem in several instances. Ferris et al. [14] introduced a method to build signal strength maps in indoor environments based on Gaussian Process Latent Variable models. Motion information is integrated with signal strength observations to estimate the signal strength distribution which can be used for localization. While the approach is able to reconstruct topologically correct maps, the true geometric shape of the building could not be captured. Comparing their reconstructed maps to ours (see Figure 6) clearly indicates the superior mapping accuracy of our approach. Huang et al. [15] presented an adaption of the Graph SLAM method which uses WLAN signal strengths as observations. Pedometer and gyroscope measurements are used to estimate the user's motion. The results show that signal strength measurements are sufficient to recognize previously visited locations. We follow this approach and formulate a sparse non-linear optimization problem based on the motion estimates and signal strength measurements. This has the advantage that many highly efficient solvers are available and can directly be applied to solve the mapping problem. Also this problem formulation allows us to combine motion and observation information from multiple walks or even users. In the publication of Huang et

al. [15], the resulting pose distribution is only shown for a small building that allows for many loop closure constraints. Also, it is unclear how to crowd source the required data as the motion estimation approach is assumed to be known. Our motion estimation approach allows us to present results for larger buildings with sparser loop closure opportunities. Additionally, our approach requires only a smartphone as well as existing wireless infrastructure.

III. Motion

The algorithm we present to track the motion of a user requires a 3-axes gyroscope and a 3-axes accelerometer to be placed loosely in a trouser pocket. Modern smartphones not only contain the required set of sensors, but are likely to be located in a user's trouser pocket. Additionally, the proximity- and light-sensors can be used to determine if the phone might be located in- or outside the pocket. The method is split into three parts which are described in the three following subsections. Firstly we show how the orientation changes of the phone can be estimated using the gyroscopes only. These orientation estimates are then used to track the motion of the users thigh which allows us to find the exact moments in time when the leg reaches its extreme orientations as depicted in Figure 2. We then estimate the length and the direction of each step based on two consecutive orientation extrema.

Orientation. In indoor environments, the earth magnetic field is heavily disturbed by power cables, concrete reinforcements and such. In addition, magnetic disturbances are caused by varying electrical currents in the sensor frame itself. To complicate the use of these magnetic field sensors even more, infrequent recalibration of the sensors occurred without notification. Since these errors are hard to accurately model, we decided to not use the magnetic field sensor to counter the heading drift caused by the gyroscopes. We found that the heading drift caused by the gyroscopes is nearly time independent and can be effectively corrected with the observation model presented later on. Therefore, we only use the gyroscope measurements $(\omega_x, \omega_y, \omega_z)$ to estimate the orientation quaternion $\mathbf{q}$ which can be computed using the Hamilton product:

$$\mathbf{q_t} = \mathbf{q_{t-1}} \cdot (1, \frac{\omega_x}{2}, \frac{\omega_y}{2}, \frac{\omega_z}{2}) \tag{1}$$

This orientation estimate relates the sensor frame to an arbitrarily chosen coordinate frame which has an offset to the earth coordinate frame that is governed by the gyroscope drift. While the orientation estimate drift leads to errors in the motion estimates, knowledge of the absolute orientation in the earth coordinate frame is not required.

Rotation Based Step Detection. The following step detection and estimation is based on the assumption that a smartphone that is placed in a trouser pocket approximately follows the motion of the thigh. We use this assumption to infer the evolution of the orientation of the thigh and the axis around which it rotates (the hip axis). We then use these vectors to find orientation extrema which we use to determine step direction and distance. The most relevant vectors that are used in the following discussion are shown in Figure 2. In the following discussion, all vectors are expressed in sensor coordinates. The acceleration of the earth's gravitational field $\mathbf{g}$ is transformed into the sensor coordinate frame using the orientation quaternion $\mathbf{q}$ from Equation 1. In a first step, the

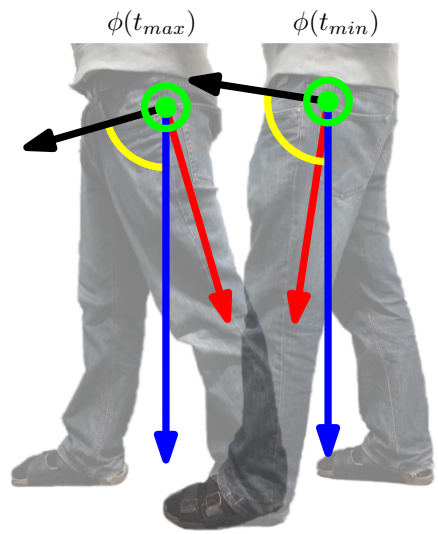

Fig. 2. The most important vectors used to find $\phi(t)$ are shown for a minimum of $\phi(t)$ and a maximum of $\phi(t)$ respectively. The leg direction $\mathbf{L}$ is shown in red, the hip direction $\mathbf{H}$ is shown in green (points out of the image), the earth gravitational force $\mathbf{g}$ is shown in blue, the inclination angle $\phi(t)$ is shown in yellow and the vector used to find the inclination $\mathbf{L} \times \mathbf{H}$ is drawn in black.

direction of the leg $\mathbf{L}$ in the sensor coordinate frame can be estimated using a low-pass filter:

$$\mathbf{L} = L_C \cdot \mathbf{L} + (1 - L_C) \cdot \mathbf{g}$$

The cutoff frequency $L_C = \sqrt{\mathbf{L} \cdot \mathbf{g}}$ is chosen to allow the leg estimate to quickly converge whenever the device orientation relative to the thigh changes. During normal use, $\mathbf{L}$ is oscillating around $\mathbf{g}$ and the leg estimate is not significantly converging towards $\mathbf{g}$. The normalized estimate $\frac{\mathbf{L}}{|\mathbf{L}|}$ can be used to determine the rotation axis between the leg and the direction of the gravitational force:

$$\mathbf{r} = \mathbf{g} \times \mathbf{L}$$

Based on the rotation axis $\mathbf{r}$, we estimate the direction of the hip $\mathbf{H}$ using a low-pass filter:

$$\mathbf{H} = H_C \cdot \mathbf{H} + (1 - H_C) \cdot \mathbf{r} \tag{2}$$

The cutoff frequency $H_C = |sin(|\mathbf{r}|)|$ is chosen to increase the influence of the rotation axis $\mathbf{r}$ on the hip estimate $\mathbf{H}$ if the angle between $\mathbf{L}$ and $\mathbf{g}$ grows. The reason for this is the fact, that the rotation axis can be determined with higher accuracy if the angle between the two vectors is larger. In case the rotation axis is pointing into the opposite direction of the hip estimate, the rotation axis is reversed ($\mathbf{r} = \mathbf{r} \cdot \mathrm{signum}(\mathbf{r} \cdot \mathbf{H})$) before applying the low-pass filter in Equation 2. The hip axis $\mathbf{H}$ converges to one of the two (opposing) main rotation axes of the leg. In the absence of absolute heading information, the two convergence possibilities of $\mathbf{H}$ deliver equal estimation

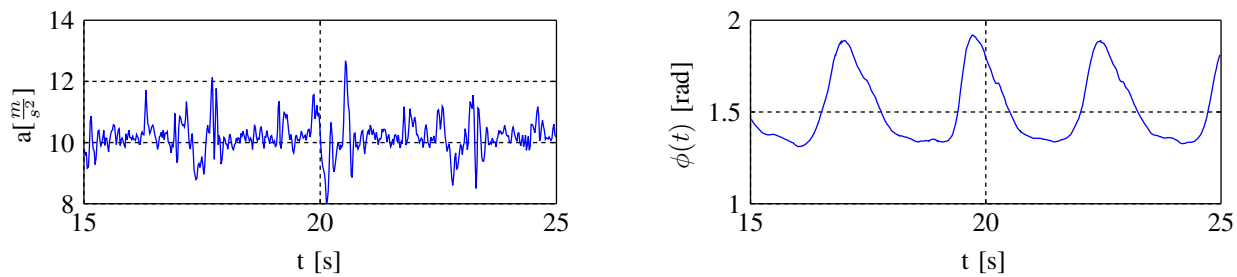

Fig. 1. A comparison of the accelerometer magnitude (left) and inclination angle ϕ (right) during several steps on a threadmill. The treadmill setting for the measurements was set to $1\frac{km}{h}$. Clearly the inclination ϕ is suitable to detect steps.

performance and results. The hip- and leg-direction estimates are then used to find $\phi(t)$:

$$\phi(t) = \arccos\left((\mathbf{L} \times \mathbf{H}) \cdot \mathbf{g}\right) \tag{3}$$

The evolution of $\phi(t)$ is used to find the points in time where the the thigh is maximally displaced. Figure 1 visualizes $\phi(t)$ and for comparison also shows the evolution of the acceleration magnitude signal $a(t)$. In $phi(t)$ the minima and maxima are easy to detect whereas in the accelerometer signal, it is extremely hard to extract isolated steps. Each minimum and maximum in $\phi(t)$ corresponds to the end of the last step and the beginning of the next. With accelerometer based step detection methods, counting steps gets increasingly difficult at low walking speeds, as the foot impact gets less articulated in the accelerometer signal. In addition to this, $\phi(t)$ allows us not only to count steps, but also to find the orientation extrema of the leg swing which allows us to accurately estimate the step length and direction.

Step Estimation. The direction and length of each step is estimated based on the smartphone orientations recorded in the minima and maxima of $\phi(t)$. The change in orientation between minima and maxima can be expressed as axis $\mathbf{a}$ and angle α. Because absolute heading information is not available, $\mathbf{a}$ can be used as the walking-direction. The angle α can be used to estimate the step length.

$$\mathbf{s} = \frac{\mathbf{a}}{|\mathbf{a}|} \cdot c \cdot \sin\left(\frac{\alpha}{2}\right) \tag{4}$$

The constant c is user dependent and has to be configured.

IV. Observation

In addition to the motion data, the smartphone is able to record the evolution of received signal strength indicators (RSSI) for all visible access points. In indoor environments, inferring distance from received signal strengths is infeasible. Non-line of sight effects and antenna imperfections lead to a spatial signal strength distribution which cannot be captured in a function that depends on the distance between sender and receiver. This is why we use these signal strength measurements not to infer physical distance, but to recognize previously visited locations only. We achieve this recognition by comparing two landmarks L_1 and L_2 using the following signal space distance measure:

$$\begin{aligned} s(L_1, L_2) = & \frac{1}{|M_1 \cap M_2|} \sum_{e=M_1 \cap M_2} |M_1(e) - M_2(e))| \\ & + \frac{1}{|M_1 \setminus M_2|} \sum_{e=M_1 \setminus M_2} |M_1(e) - l_{min})| \\ & + \frac{1}{|M_2 \setminus M_1|} \sum_{e=M_2 \setminus M_1} |l_{min} - M_2(e))|. \end{aligned} \tag{5}$$

The sets of visible access points are denoted as M_1 and M_2 respectively. Access points that report an RSSI lower than l_{min} are neglected. In addition, access points that are only visible in the other landmarks are considered to be received with signal strength l_{min}. Two landmarks assumed to be captured in the same location (associated) if their distance measure $s(L_1, L_2)$ is lower than a threshold s_{th}.

V. Fusion

The complementary characteristics of the motion- and observation association constraints can be exploited to counteract the divergence of integrating motion estimates using the observation associations. THe similar to finding the equilibrium point of a system of springs and masses. The visited locations correspond to the masses and are described as point in space as well as current motion heading $x_i = (x, y, \phi)$. The motion and association constraints correspond to the springs. The stiffness of the springs corresponds to the confidence level of the estimates. Whereas the association constraints are modeled as springs with equilibrium length zero and do not constrain the heading difference, the motion constraints are modeled as springs with equilibrium length equal to the estimated step length that also constrain the change in heading ϕ between consecutive poses.

$$x_i = f_i(x_{i-1}, u_i) + w_i. \tag{6}$$

The pose x_i is linked to the previous pose x_{i-1}, using the sensor readings u_i. The model uncertainty is captured in the Gaussian noise term w_i. Landmark associations are described as follows:

$$0 = ||x_{i_k} - x_{j_k}|| + v_k. \tag{7}$$

This captures the fact that if we associate two landmarks, we expect them to be recorded in spatially close locations. The Gaussian noise term v_k may vary depending on the association

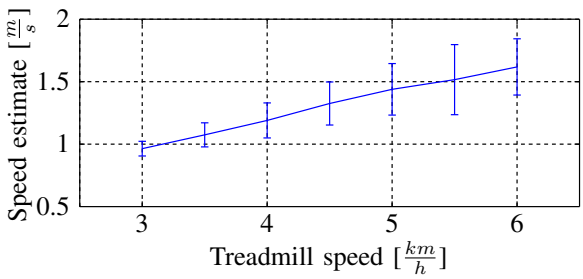

Fig. 3. Comparing the motion model speed output to the ground truth (Treadmill speed setting). Clearly, the motion model output linearly depends on the actual walking speed which means that when properly configured, the motion model will work at different walking speeds.

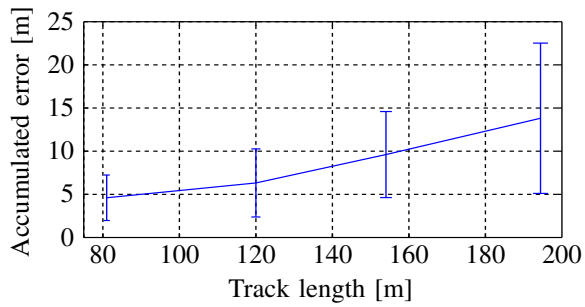

Fig. 4. Relating the traveled distance with the resulting motion model integration error.

quality where k addresses one specific association between two poses x_{i_k} and x_{j_k}. The optimization problem is defined as follows:

$$\Theta^* = \underset{\Theta}{\operatorname{argmin}} \Big[\sum_{i=1}^{M} ||f_i(x_{i-1}, u_i) - x_i||^2_{\Lambda_i} + \sum_{k=1}^{K} ||x_{i_k} - x_{j_k}||^2_{\Sigma_k} \Big]. \tag{8}$$

For simplicity, the notation $||e||_\Sigma = e^\intercal \Sigma^{-1} e$ was used. The solution is the set of poses Θ that minimizes the given cost function based on the constraints obtained from the motion- and observation model. We solve the sparse non-linear optimization problem using iSAM [16].

Compared to a particle filter based approach, this non-linear least squares problem allows us to overcome the lack of absolute heading information as well as heading drift without increasing the computational complexity. Note that a particle filter would need to sample the joint probability function of spatial location, heading offset and heading drift. Adding dimensions to the probability distribution causes the number of particles to grow exponentially.

VI. Experiments and Results

The following experiments were carried out using a Samsung Nexus S smartphone. In a first step, the motion- and observation models are evaluated separately because the motion model will deliver user specific performance whereas the observation model will be largely user independent.

Motion. Firstly the speed estimation is compared to the actual walking speed. A single user was walking on a treadmill whose speed setting was used as the ground truth walking speed. The results shown in Figure 3 indicate a linear dependance between actual speed and motion model estimate. This is the desirable outcome since it indicates that the motion model can be configured for a specific user with only one parameter and deliver accurate speed estimates over a variety of walking speeds. The variances shown for each treadmill speed setting indicates a large variance of of the speed estimates even though the treadmill- and therefore walking speed was very constant during the experiment. The results shown in Figure 3 were used to find the parameter c from Equation 4.

The distance estimation was evaluated with eleven people that were asked to walk down a 51 meter long straight hallway. The motion model output had a mean of 51.7 meters and a standard deviation of 4.4 meters. This result indicates, that the motion model may be adapted for different users. To evaluate the motion model accuracy in a more realistic setting in which people were walking through hallways, doors and bends, the same eleven people were asked to walk along four predefined tracks with increasing lengths (51m, 81m, 120m, 154m, 195m). All the tracks ended in their starting point which means, the localization error after each track is captured in the distance between start- and estimated end-point. Figure 4 shows the increasing localization error caused by accumulating motion estimation errors. As expected, the localization error increases as the traveled distance grows. A large portion of the error is caused by the drifting heading estimate which could not be corrected using the magnetic field sensors. Although the closed tracks facilitate evaluating the localization error, note that this evaluation scheme does not capture the fact that the motion model might systematically over- or underestimate the step lengths for different users. However, combined with the results shown in Figure 3 and the distance measurement accuracy of the straight hallway experiment we conclude that the motion model works for a variety of people walking at different speeds only requiring to be configured at one speed, or even only using the users body height.

Observation. Observations are associated to each other if their mutual signal space distance is below a given threshold. Therefore, we require the signal space distance to be low for two spatially close fingerprints. On the other hand, comparing spatially distant fingerprints should lead to a high signal space distance. To evaluate how well this requirement is met by the smartphone WLAN observations, we carried out the following experiment. We collected a large number of fingerprints by walking at constant speed through office building hallways. The hallways are are forming a rectangle (50m x 10m) which is traversed twice during the experiment. The distance measure for all the pairs of recorded fingerprints are shown in Figure 5. Note that the distance measure is commutative and therefore, the distribution is symmetric. Also, elements that are close to the diagonal indicate signal space distances of two fingerprints that have been recorded within a short period of time. The farther away from the diagonal, the larger the time span between the two recorded fingerprints. Because the experiment

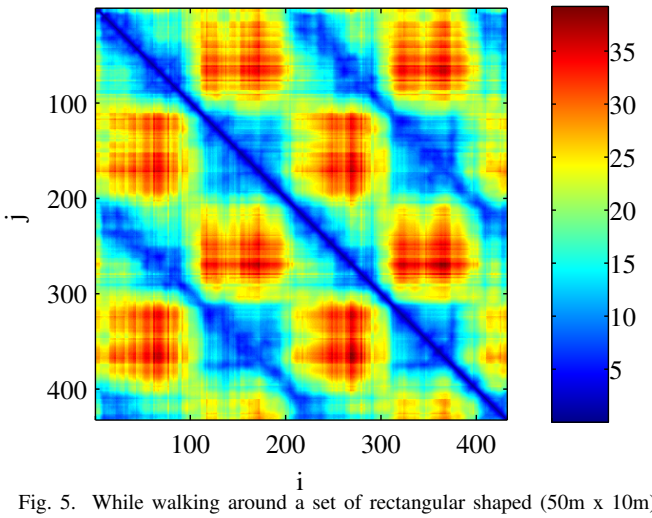

Fig. 5. While walking around a set of rectangular shaped (50m x 10m) set of hallways twice, roughly 420 fingerprints were collected. This figure shows the fingerprint signal space distance for all pairs i, j of recorded fingerprints.

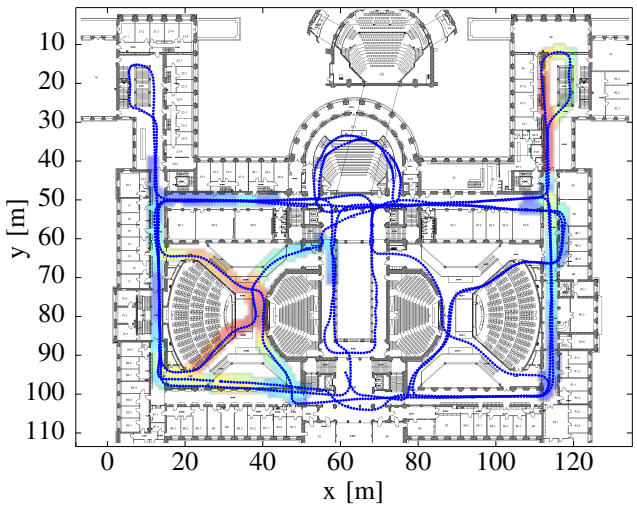

Fig. 6. Fingerprint distribution for ETH Zurich main building (HG). Fingerprint locations are indicated as blue dots. Signal strengths for selected access points are shown in blue to red colors

was conducted while walking at constant speed, this also means that the time span in which two fingerprints were recorded linearly translates to a spatial distance. Since the goal of this experiment is to get an idea of how well our signal space distance measure is able to distinguish fingerprints that are far from fingerprints that are close, it is desirable to have low distance measures along the diagonal, but high distance measures the further away from the diagonal. In addition to the main diagonal, two secondary diagonals which are originate from the second round through the rectangular set of hallways. Similar to the main diagonal, we desire low distance measures close to the secondary diagonals but high distance measures the further away we get. Clearly the secondary diagonals are less clear than the main diagonal especially, the off diagonal elements are less distinguishable from the diagonal elements. The four by four checkerboard pattern is a result of the rectangular track and shows that fingerprints recorded in the two opposite long hallways have a large signal space distance. However, within one hallway, the spatial distance of fingerprints with small signal space distances can grow large. Due to this shortcoming, erroneous associations have to be expected. In addition, the similarity between fingerprints within one hallway will impede localization performance.

Fusion. Fusing motion- and association data by non-linear optimization leads to a map of fingerprints. The quantitative evaluation of the map quality is difficult as long as no localization scheme is used to measure localization performance based on the generated fingerprint map. Figure 6 shows the fingerprint distribution resulting from the fusion step for the ETH main building. The recorded fingerprints are drawn as blue dots. The received signal strengths for two distinct (non overlapping) access points are indicated in blue to red colored areas around the corresponding fingerprints. The relative fingerprint positions are the result of the fusion step. The floor plan as well as the rotation and translation between floor plan and fingerprint distribution is not automatically obtained but added by hand. The fingerprint locations obtained in the fusion step resemble the true shape of the building with with few exceptions (intersection on the right side).

References

[1] A. Rai, K. K. Chintalapudi, V. N. Padmanabhan, and R. Sen, "Zee: zero-effort crowdsourcing for indoor localization," ser. Mobicom. ACM, 2012, pp. 293–304.

[2] K. Chintalapudi, A. Padmanabha Iyer, and V. N. Padmanabhan, "Indoor localization without the pain," ser. MobiCom, 2010.

[3] E. Foxlin, "Pedestrian tracking with shoe-mounted inertial sensors," *Computer Graphics and Applications, IEEE*, vol. 25, no. 6, pp. 38–46, 2005.

[4] S. Godha and G. Lachapelle, "Foot mounted inertial system for pedestrian navigation," *Measurement Science and Technology*, vol. 19, no. 7, p. 075202, 2008.

[5] S. Beauregard, "Omnidirectional pedestrian navigation for first responders," ser. WPNC, 2007.

[6] L. Ojeda and J. Borenstein, "Non-GPS navigation for emergency responders," 2006.

[7] I. Skog, J.-O. Nilsson, and P. Handel, "Evaluation of zero-velocity detectors for foot-mounted inertial navigation systems," ser. IPIN, 2010.

[8] S. Beauregard, "A helmet-mounted pedestrian dead reckoning system," ser. IFAWC, 2006.

[9] U. Blanke and B. Schiele, "Sensing location in the pocket." ser. UbiComp, 2008.

[10] U. Steinhoff and B. Schiele, "Dead reckoning from the pocket - An experimental study," ser. PerCom, 2010.

[11] S.-W. Lee and K. Mase, "Activity and location recognition using wearable sensors," *Pervasive Computing, IEEE*, vol. 1, no. 3, pp. 24 – 32, 2002.

[12] H. Durrant-Whyte and T. Bailey, "Simultaneous localisation and mapping (SLAM): Part I," *Robotics Automation Magazine*, vol. 13, no. 2, pp. 99–110, 2006.

[13] T. Bailey and H. Durrant-Whyte, "Simultaneous localization and mapping (SLAM): Part II," *Robotics Automation Magazine*, vol. 13, no. 3, pp. 108–117, 2006.

[14] B. Ferris, D. Fox, and N. Lawrence, "WiFi-SLAM using Gaussian process latent variable models," ser. IJCAI, 2007.

[15] J. Huang, D. Millman, M. Quigley, D. Stavens, S. Thrun, and A. Aggarwal, "Efficient , generalized indoor WiFi GraphSLAM," ser. ICRA, 2011.

[16] M. Kaess, A. Ranganathan, and F. Dellaert, "iSAM: Incremental smoothing and mapping," *Transactions on Robotics*, vol. 24, no. 6, pp. 1365–1378, 2008.