

Received-Signal-Strength-Based Logical Positioning Resilient to Signal Fluctuation

Thomas Locher, Roger Wattenhofer, Aaron Zollinger
{lochert@student, wattenhofer@tik.ee, zollinger@tik.ee}.ethz.ch
Computer Engineering and Networks Laboratory
ETH Zurich, Switzerland

Abstract

Positioning based on received signal strengths from base stations is highly sensitive to the effects of signal attenuation, reflection, and scattering. Moreover, experimental measurements show that received signal strengths fluctuate significantly over time due to noise and interference. Unlike most of the related work our approach does not rely on the presence of a priori information about base station positions, objects influencing radio signal propagation, and signal propagation characteristics. Instead of trying to compute the current position defined by coordinates, the goal of our system is to infer a user's present logical position. In addition, our system particularly differs from previous work in the chosen statistical approach, which explicitly copes with signal strength fluctuation over time and allows control over system accuracy by means of specification of confidence intervals.

1 Introduction

Positioning based on received signal strength patterns is difficult [8]. Scattering and reflection of electromagnetic waves can strongly distort the signals [6]; shielding severely attenuates the signal strength. A single person can alter the signal strength by up to -3.5dBm [2]. What is more, when measuring signal strengths, a high level of measurement noise is to be expected, especially when using standard hardware [9]. Moreover noise and interference can cause the received signal strengths to fluctuate significantly over time. Nevertheless there is a strong correlation between the distance of a base station and its signal strength [1], which can be exploited by a suitable positioning system.

Most of the existing approaches to positioning based on received signal strengths assume to have knowledge about base station positions and use trilateration or triangulation methods. Both trilateration—determining the position of a

device by measuring *distances* to base stations with known positions—and triangulation—position inference based on measured *angles* to base stations—are however highly sensitive to errors caused by signal attenuation, reflection, and scattering.

Consequently positioning maps are in most cases not sufficient if they only contain information about base station positions. Satisfactory results are hardly achievable unless these maps include additional information about factors and conditions influencing signal propagation, such as walls, buildings, or mountains. Even if such information is available, however, modeling their effect is highly complex, error-prone, and in most practical cases impossible. In this paper we choose an approach that—instead of requiring a priori information about such architectural or environmental conditions—is based on a relatively short, simple, and effortless calibration of the system and is therefore more suitable for most practical purposes.

Unlike most of the related work our approach is exclusively based on signal patterns received from base stations without knowledge about base station positions. As a result, our system does not aim at approximating *coordinates* of the current position of the mobile user, but finding the most probable *logical* location, such as a room, floor or area, among the set of previously learned logical locations. For that purpose several measurements are performed at each logical location and the measured signal strength patterns are recorded in a *calibration phase*. This information is then used in an *operation phase* to infer the currently most probable logical location.

However, even if errors induced by insufficient modeling of signal attenuation, reflection, and scattering can thus be alleviated, the problem remains that noise and interference give rise to significant signal fluctuation over time. Our experimental measurements show that this fluctuation is significant and cannot be neglected.

Due to the aforementioned correlation between the distance of a base station and the measured signal strength, a simple procedure to estimate the mobile user's current loca-

tion is to create a straightforward mapping between signal strength patterns and locations. However, since the signal strengths fluctuate over time, a statistical approach manages to capture the characteristics of the locations more accurately. Several other approaches applying different statistical schemes with a varying degree of complexity have been proposed. Our system model distinguishes itself from these proposals in that our statistical approach explicitly accounts for high fluctuation of signal strengths. Furthermore, a confidence interval can be specified in order to specify or measure the quality of the result.

Another aspect in which our system differs from related work, is that the signal strength of a base station measured at a specific position is assumed to be normally distributed over time. A series of measurements have been carried out and suggest this assumption.

The paper is organized as follows: The subsequent section discusses previous approaches to the problem of positioning based on received signal strengths. Section 3 presents our experimental measurements with respect to signal strength fluctuation and its modeling. Section 4 describes the statistical approach chosen in our system for both the calibration and the operation phases. Section 5 discusses experimental results gained as well as problems and limitations experienced with our implementation of the system. Section 6 concludes the paper.

2 Related Work

The goal of the majority of the systems presented in this section is to determine the mobile user's physical position. There is one system, called Nibble [3] forming an exception in that it is similar to our approach, also working with logical locations; yet the system models differ. Nibble's modular probabilistic approach for inferring location uses Bayesian networks. A Bayesian network is a graphical representation of a joint probability distribution that explicitly declares dependency relationships between random variables in the distribution.

As mentioned above, in the systems presented in the following paragraphs, the physical position of the mobile user is approximated. Therefore, the geometry of the surrounding area has to be known to a certain extent. Even the knowledge of the exact positions of all base stations is often a prerequisite. Due to the correlation between signal strength and distance, one simple approach used in the RADAR system [2] to solve the positioning problem is to build up a so-called radio map, which stores position-signal strength pairs. In order to approximate the user's current position, the average of the k nearest neighbors is returned. A problem of this approach is the choice of k . A variation of this approach is the joint clustering technique [12]. A cluster is a set of positions where signals from exactly the

same base stations are received. This set of base stations is denoted the cluster key. Once the right cluster is found in the operation phase, Baye's theorem is used to determine the probability of each location within the cluster. The critical aspect here is the choice of the dimension of the joint distribution. An advantage is on the other hand the use of a probability distribution for the signal strengths. It has been stated that an optimal strategy must consider the probability distribution of the signal strengths and that taking the average of several signals received reduces the error [11].

Another system, called GPPS [9], uses a maximum likelihood estimator. Gaussian process models are built for the distribution of the signal strengths, using the Matérn kernel function. The resulting maximum likelihood estimator is returned as the solution, that is, there is no error bound. The positions of all the base stations ought to be known, otherwise they are approximated.

A totally different approach is used in the LEASE system [7]. This infrastructure-based system uses a small number of stationary emitters and sniffers in order to locate the user. The sniffers collect information about the user and the stationary emitters. Afterwards, the collected information is forwarded to the location estimation engine (LEE), which knows the positions of the stationary emitters. A similar system called Palantir has been proposed [4]. In this system, the floor of a building is divided into grids and the signal strength in the middle of each grid is approximated again using sniffers, which are the main component of this system. A nearest neighbor search is performed in order to approximate the user's current physical position.

Another interesting idea is the combination of various localization techniques [5]. Since it is not clear whether a single positioning algorithm can find the optimal solution, several different methods are combined in order to achieve more accurate results. The contributions of all methods have to be weighted to minimize erroneous information.

3 Experimental Received-Signal-Strength Measurements

This section discusses signal strength fluctuation over time. First, experimental measurements show that the variance over time of signals received from base stations are significant and therefore deserve special attention. Second, our measurements suggest the assumption that signals received over a certain period of time may be modeled as normally distributed.

To illustrate the problem of signal fluctuation, signals were measured using the WRAPI Wireless Research API [10] with standard IEEE 802.11b Wireless LAN network adapters. The following table shows the signal strength pattern measured at one specific sample location (60 measurements over a period of 6 seconds, signal strengths in mW):

Base station	Average	Variance
00-07-50-D6-00-FC	1.640386e-7	1.664545e-15
00-07-50-D6-01-0F	2.299025e-5	8.428953e-11
00-07-50-D6-02-F8	2.384267e-9	3.631800e-19
00-07-50-D6-02-FA	2.760481e-6	1.499120e-12
00-07-50-D6-03-CB	4.656942e-9	1.664351e-18

Merely 15 minutes later the following pattern was determined at exactly the same position:

Base station	Average	Variance
00-07-50-D6-00-FC	1.762277e-7	4.949886e-15
00-07-50-D6-01-0F	2.219300e-5	1.057022e-10
00-07-50-D6-02-F8	3.073024e-9	3.603822e-18
00-07-50-D6-02-FA	1.809146e-6	7.098445e-13
00-07-50-D6-03-CB	4.461042e-9	3.218017e-18

While the average signal strength differs from the older measurements by a factor of at most 1.5 (fourth entry in above base station table), the difference in variance reaches a factor of 10 (third base station). Fluctuation becomes even clearer when the period of time between those measurements is increased to one day:

Base station	Average	Variance
00-07-50-D6-00-FC	1.481586e-7	8.501779e-16
00-07-50-D6-01-0F	1.974024e-5	1.851989e-11
00-07-50-D6-02-F8	1.798881e-8	7.769301e-17
00-07-50-D6-02-FA	6.498956e-7	5.176300e-14
00-07-50-D6-03-CB	2.202195e-8	1.167230e-16

In this case, even the average signal strength differs by a factor of up to 7.5, while the difference in variance reaches a factor of 214 (third base station). It has to be noted that there are differences between base stations. The signal strengths of some base stations appear to be more stable than those of others. For example, the third base station in the above tables is responsible for the highest factors as far as both the average and the variance is concerned.

As a consequence to the observation that signals received from base stations can fluctuate significantly, the question arises how signal strengths measured over time are distributed. Figure 1 shows sample histograms of signal values received from a base station with weak signals and strong signals, respectively, both measured over a shorter and a longer period of time, respectively. Not only the shapes of the resulting histograms depicted but of most of our measurements suggest that received signal strengths of a base station can be modeled as normally distributed over time.

4 System Model

In our system model, the positions of the base stations are not required. In the so-called *calibration* phase, the system learns the logical mapping between base station identifiers, their signal strengths and the logical locations. The

system memorizes this mapping and uses it in the *operation* phase, when the current logical position is to be determined.

The system allows for the specification of a confidence interval. The result of the system is the logical location that best matches the received signals. If the measured signals are not compatible with any logical location, considering a specific confidence interval, then no solution is to be returned. Alternatively, all logical locations that are compatible with the given confidence interval can be returned together with the probability that the mobile user currently is located at the respective logical location.

4.1 Calibration Phase

Let M denote the mobile user and b_i denote the i th base station as seen by M . b_i is a pair (a_i, s_i) , where a_i is the identifier of the base station and s_i is the received signal strength (in dBm or mW). M collects data in order to characterize a logical location L . For that purpose, M performs k measurements at that logical location.

A measurement is a set of pairs $b_i=(a_i, s_i)$, that is, a measurement has the form $S^{(j)} = \{b_1, b_2, \dots, b_n\}$, where $j \in [1, k]$. n is the number of different base stations from which signals are received. Technically, n is not known until the k th measurement has been carried out. Each measurement probably receives a signal from one or more base stations that another measurement does not. If there is no signal of b_i in the j th measurement, b_i however occurs in another measurement taken at the same location, then $s_i^{(j)}$ is set to 0.¹

The characteristics of L have to be determined utilizing this data. First, \bar{S} is calculated, which for each a_i stores all the average signal strengths \bar{s}_i , that is, $\bar{S} = \{(a_1, \bar{s}_1), \dots, (a_n, \bar{s}_n)\}$, where $\bar{s}_i = \frac{1}{k} \sum_{j=1}^k s_i^{(j)} \forall i \in [1, n]$. Either all b_i are considered, or the dimension is reduced by excluding those base stations from which signals were rarely received in those k measurements. Assuming that such a reduction has been performed, let there be $m \leq n$ base stations from which signals have been received sufficiently often. Because of noise and interference, it is not sufficient to consider only \bar{S} as the characteristics for L . In order to take these factors into account, a probabilistic approach appears to be appropriate. The normal distribution is used as an approximation of the distribution of the signal strengths over time. For that purpose \bar{S} has to be extended to include the variance v_i of each of its associated base stations, that is, $S^* = \{(a_1, \bar{s}_1, v_1), \dots, (a_m, \bar{s}_m, v_m)\}$ and $v_i = \frac{1}{k-1} \sum_{j=1}^k (s_i^{(j)} - \bar{s}_i)^2 \forall i \in [1, m]$ where S^* is the extended form of \bar{S} . Not only does S^* describe the average signal strengths, but also the variance of the individual

¹This is only reasonable if signal strengths are measured and stored in mW and not in dBm.

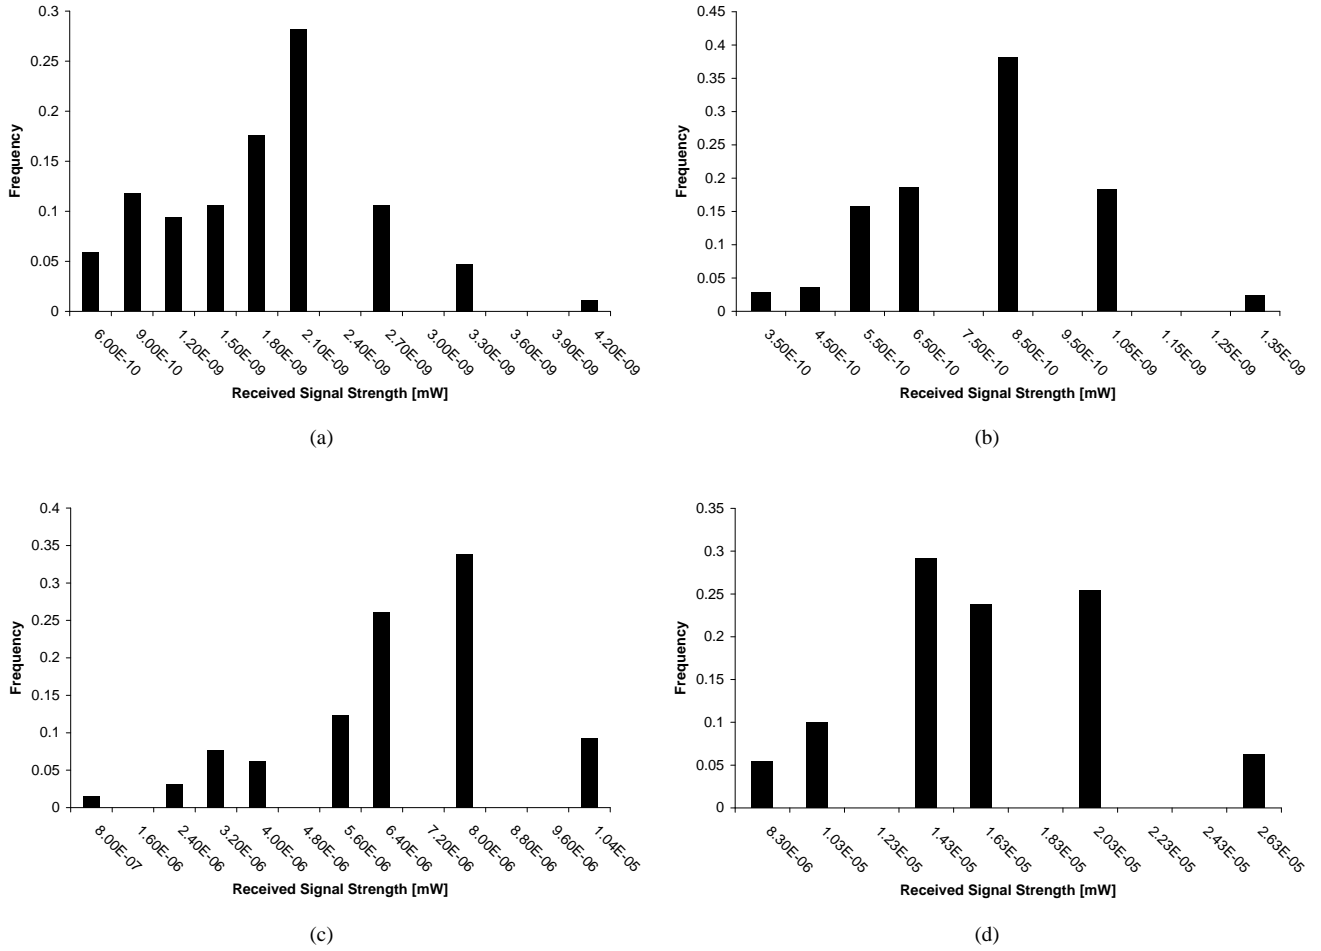


Figure 1. Histograms of measured signal strengths of a base station with weak signals measured over a period of 3 minutes (a) and one hour (b) and of a base station with strong signals measured over the same periods (c and d).

contributions; therefore S^* captures the characteristics of a particular location more precisely since—as illustrated in Section 3—the fluctuation of signal values of different base stations can vary significantly.

4.2 Operation Phase

In the operation phase, the system is interrogated. In order to derive M 's current logical location, a fresh set of measurements $S = \{b_1, b_2, \dots, b_d\}$ is acquired. Let $\mathcal{L} = \{L_1, L_2, \dots, L_p\}$ be the set of all logical locations known to the system and let n_i be the dimension of L_i , that is the number of base stations constituting the characteristics of L_i . In addition, let S_i^* denote the characteristic set of L_i .

Given that all measurements are probabilistic, it is de-

sirable to find the logical location $\hat{L} \in \mathcal{L}$ that maximizes the probability of M being at this location, as opposed to being at any of the other locations of \mathcal{L} . Naturally, \hat{L} also has to respect the given confidence interval. As mentioned, alternatively all locations that comply with the confidence interval together with their likelihood could be given as the result.

In order to achieve these goals, the value Z_i is calculated for each logical location of \mathcal{L} according to the formula

$$Z_i := \sum_{j=1}^{n_i} \frac{(\bar{s}_{ij} - s_j)^2}{v_{ij}},$$

where \bar{s}_{ij} and v_{ij} denote the average signal strength of the j th base station of the logical location L_i and the vari-

ance thereof, respectively; s_j denotes the (averaged) signal strength obtained from the same base station in the new measurements (the canonical order of S being adapted to S_i^*). If there is a particular base station b_j that appears in S_i^* , but not in S , then s_j is simply set to 0.² If, in contrast, there is a base station appearing in S , but not in S_i^* , then it is ignored since it is not part of the characteristics of L_i . Alternatively, the characteristic signal strength of this base station could be set to 0; however, filtering certain base stations in the calibration phase could not be applied in that case.

Assuming that the signal strength of a base station received at a logical location is normally distributed, the individual terms of Z_i are squared normal variables. Hence Z_i is χ^2 -distributed with n_i degrees of freedom. L_i can only be considered M 's current location, given a confidence interval p , if Z_i is at most γ_i . γ_i denotes the value for which it holds that the integral of $\chi_{n_i}^2(t)$ from 0 to γ_i is equal to p , where $\chi_{n_i}^2(t)$ is the χ^2 -probability density function with n_i degrees of freedom. In order to obtain the value γ_i , the inverse cumulative density function is used:

$$\begin{aligned} \text{inverseCumulativeDensity}(n_i, p) &= \gamma_i \\ \iff \int_0^{\gamma_i} \chi_{n_i}^2(t) dt &= p. \end{aligned}$$

Given the Z_i for all logical locations, it is straightforward to determine the best logical position. Moreover, by choosing the $L_i \in \mathcal{L}$ for which Z_i is minimal, we obtain the solution that not only minimizes the least squares error weighted with the inverse of the variance but also constitutes the maximum likelihood estimator, which is shown in the remainder of this section.

The chosen L_i minimizes the least squares error by definition. It remains to be shown that it is at the same time the maximum likelihood estimator. For that purpose, the characteristic set S^* of one logical location is considered. Since there is a canonical order imposed on the elements of S^* , it is possible to omit the identifiers in each element for simplicity. Furthermore, all the average signal strengths and the variances thereof are renamed in order to emphasize their statistical meaning, that is, S^* has the form $\{(\mu_1, \sigma_1^2), \dots, (\mu_n, \sigma_n^2)\}$. According to our assumption about the distribution of the signal strengths, the probability density function for the measured signal strength value x given μ_i and σ_i is³

$$\forall b_i: \quad \text{prob}(x|\mu_i) = \frac{1}{\sqrt{2\pi}\sigma_i} e^{-\frac{(x-\mu_i)^2}{2\sigma_i^2}},$$

²Again, this is only reasonable if signal strengths are stored in mW and not in dBm.

³For simplified representation, the expression $\text{prob}(x|\mu_i)$ denotes that both μ_i and the corresponding variance are given.

where the base stations b_i are the base stations that constitute the characteristics of the logical location with the characteristic set S^* . The formula can be extended to include all those base stations together:

$$\text{prob}((x_1, \dots, x_n)|(\mu_1, \dots, \mu_n)) = \prod_{i=1}^n \left(\frac{1}{\sqrt{2\pi}\sigma_i} e^{-\frac{(x-\mu_i)^2}{2\sigma_i^2}} \right).$$

The definition of $\text{lik}((\mu_1, \dots, \mu_n))$, the likelihood of the currently measured location to be the logical location with (μ_1, \dots, μ_n) , is useful in the analysis:

$$\text{lik}((\mu_1, \dots, \mu_n)) = \text{prob}((x_1, \dots, x_n)|(\mu_1, \dots, \mu_n)).$$

The maximum likelihood estimator maximizes $\text{lik}((\mu_1, \dots, \mu_n))$:

$$\begin{aligned} \max\{\text{lik}((\mu_1, \dots, \mu_n))\} &= \\ &= \prod_{i=1}^n \left(\frac{1}{\sqrt{2\pi}\sigma_i} \right) \max \left\{ \prod_{i=1}^n e^{-\frac{(x-\mu_i)^2}{2\sigma_i^2}} \right\} \\ &= \prod_{i=1}^n \left(\frac{1}{\sqrt{2\pi}\sigma_i} \right) \max \left\{ e^{-\sum_{i=1}^n \frac{(x-\mu_i)^2}{2\sigma_i^2}} \right\}. \end{aligned}$$

It holds that

$$e^{-\sum_{i=1}^n \frac{(x-\mu_i)^2}{2\sigma_i^2}}$$

is maximal if and only if

$$\sum_{i=1}^n \frac{(x_i - \mu_i)^2}{\sigma_i^2}$$

is minimal. This means that the maximum likelihood estimator is in fact obtained by minimizing the weighted error in the least squares sense.

Finding all logical locations that comply with the given confidence interval p and computing their likelihood is straightforward: All logical locations L_i for which $Z_i \leq \gamma_i$ holds are possible solutions.

5 Experimental Results

In order to assess the feasibility of our approach and to evaluate the quality of the results thereby achieved, we implemented our system using standard off-the-shelf hardware. In particular we employed a standard IEEE 802.11b network adapter in combination with the Wireless Research API [10] to acquire signal strength patterns from Wireless LAN base stations.

The experimental results can be grouped into two parts. The first part, which is treated in Section 5.1, focuses on the performance of the system presented in this paper. In particular, the effect of parameters such as the number of measurements performed and the confidence interval are analyzed. In the second part, problems and limitations of positioning with WLAN are described.



Figure 2. Floor plan of the building where the described system evaluation has been carried out. Base station positions are indicated by crosses; room H02 is on the next higher floor.

5.1 Accuracy of the System

The measurements to determine how accurately the system is able to infer the correct position included one experiment where the system was presented five logical locations to be learned, in particular five different rooms in one building. Four of them are on the same floor, while the fifth room (H02) is on the next higher floor (cf. Figure 2). For every room, 60 measurements were performed in the calibration phase with an interval of 100 milliseconds between consecutive measurements. Operation phase measurements were started immediately after calibration of the complete setting. The different operation experiments described in the following were taken during a period of approximately 2 hours. In the first experiment, 10 measurements with the same interval of 100 milliseconds were performed in the operation phase. The confidence interval was set to 80%. Since the system never displayed an incorrect location—it either displayed the correct answer or it concluded that no location matched the new measurements—it is sufficient to list how often it found the right solution:

Room	Correct	No Match
G05	75%	25%
G21	50%	50%
G26	75%	25%
G31	100%	0%
H02	100%	0%
Total	80%	20%

Since the system never returned a bad result, the confidence interval could be increased in order to allow more inaccurate measurements to be considered and thus reduce the frequency of “no matches.” However, in doing so, the set of possible locations increases and the system starts yielding erroneous results. There is clearly a tradeoff between obtaining a result in most cases and allowing the system to make mistakes. Therefore, the appropriate setting of the confidence interval is crucial. As an extreme example, in the following experiment, the confidence interval has been set to 100%, that is, every logical location is considered to be a possible location. On this account, it is only reasonable to retrieve the best matching location. The system then always returns the location with the lowest least-squares error:

Room	Correct	No Match
G05	70%	30%
G21	90%	10%
G26	70%	30%
G31	60%	40%
H02	80%	20%
Total	74%	26%

Room	Correct	Wrong
G05	90%	10%
G21	90%	10%
G26	100%	0%
G31	100%	0%
H02	100%	0%
Total	96%	4%

The accuracy can be improved by prolonging the operation phase. In the second experiment, 20 measurements were performed in the operation phase with an interval of 1 second between successive measurements:

Surprisingly, the best matching location is in most cases the right location, indicating that the confidence interval can be set to a high value without rendering the localization scheme inoperative.

5.2 Problems and Limitations

Besides high signal fluctuation over time, the signals measured at different positions even in the same room can also vary strongly. In the same room, the following pattern was obtained at a different position a few seconds after the first measurements displayed in Section 3:

Base station	Average	Variance
00-07-50-D6-00-FC	8.590684e-9	2.350362e-17
00-07-50-D6-01-0F	9.444324e-6	5.584365e-12
00-07-50-D6-02-F8	2.700181e-9	1.164344e-18
00-07-50-D6-02-FA	1.877516e-5	1.999351e-10
00-07-50-D6-03-CB	2.871098e-7	1.592052e-14

The signal strengths deviate from the former measurements to such an extent that it is scarcely reasonable to consider the two patterns the characteristics for the same logical location. Apparently—at least in an indoor setting—signal propagation depends on edificial conditions to a degree that prohibits formulation of one logical location to represent a whole architectural unit, such as a room. It appears that consequently the approach presented in this paper can in some cases require more finely grained subdivision of logical units to become reasonably accurate.

A further limitation of the system lies in the fact that the quality of the results depends highly on the constellation of the mobile user and the base stations. Also moving objects influencing signal propagation, such as vehicles, are not accounted for. All these issues however form a difficulty for any system based on received signal strength patterns, independently of the chosen approach.

6 Conclusion

Positioning based on received signal strength patterns is difficult not only due to radio signal shielding, reflection, and scattering. Our experimental measurements show that signals received from Wireless LAN base stations can also fluctuate significantly over time. However, this effect is shown to be plausibly modeled as received signal strength values following a normal distribution.

In contrast to most of the related work which aims at computing the mobile user's current position *coordinates* based on a priori knowledge about base station positions—a process highly sensitive to signal shielding, reflection, and scattering—our approach exclusively operates with received signal strength patterns. Consequently our system tries to match the user's current position with previously learned *logical* locations. Furthermore, our system distinguishes itself from previously proposed solutions in the chosen statistical approach, explicitly accounting for high signal fluctuation over time and allowing for a tradeoff be-

tween result accuracy and error rate by specification of confidence intervals.

In particular, our experimental system evaluation shows that, in order to achieve high accuracy, many measurements ought to be performed in the calibration phase. By repeatedly adding measurements after a longer period of time and by also prolonging the operation phase, the problem of signal fluctuation over time can be controlled to a reasonable extent.

Acknowledgements

We would like to thank Jürgen Bohn for valuable discussions and for helping us get WRAPI to work.

References

- [1] P. Bahl and V. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. *IEEE Infocom*, 2000.
- [2] P. Bahl, V. Padmanabhan, and A. Balachandran. A software system for locating mobile users: Design, evaluation, and lessons. *MSR Technical Report*, February 2000.
- [3] P. Castro, P. Chiu, T. Kremenek, and R. Muntz. A probabilistic room location service for wireless networked environments. *Proceedings of UbiComp*, 2001.
- [4] S. Ganu, A. Krishnakumar, and P. Krishnan. Infrastructure-based location estimation in WLAN networks. *IEEE Wireless Communications and Networking Conference (WCNC)*, 2004.
- [5] Y. Gwon, R. Jain, and T. Kawahara. Robust indoor location estimation of stationary and mobile users. *IEEE Infocom*, 2004.
- [6] H. Hashemi. The indoor radiation propagation channel. *Proceedings of the IEEE*, 81(7):943–968, 1993.
- [7] P. Krishnan, A. Krishnakumar, W.-H. Ju, C. Malloes, and S. Ganu. A system for lease: Location estimation assisted by stationary emitters for indoor RF wireless networks. *IEEE Infocom*, 2004.
- [8] T. Roos, P. Myllymäki, H. Tirri, P. Misikangas, and J. Sievänen. A probabilistic approach to WLAN user location estimation. *The Third IEEE Workshop on Wireless LANs*, 2001.
- [9] A. Schwaighofer, M. Grigoras, V. Tresp, and C. Hoffmann. GPPS: A Gaussian process positioning system for cellular networks. *Advances in Neural Information Processing Systems*, 16, 2004.
- [10] Wireless research application programming interface WRAPI. <http://ramp.ucsd.edu/pawn/wrapi>.
- [11] M. Youssef and A. Agrawala. On the optimality of WLAN location determination systems. *Communication Networks and Distributed Systems Modeling and Simulation Conference*, 2004.
- [12] M. Youssef, A. Agrawala, and A. Shankar. WLAN location determination via clustering and probability distributions. *IEEE International Conference on Pervasive Computing and Communications (PerCom)*, 2003.