

Key Factors for Position Errors in 802.11-based Indoor Positioning Systems

Thomas King, Thomas Haenselmann, and Wolfgang Effelsberg

Technical Report
Department for Mathematics and Computer Science
University of Mannheim
TR-2007-003

{king,haenselmann,effelsberg}@informatik.uni-mannheim.de

Abstract. Indoor positioning systems based on 802.11 and fingerprints offer reasonably low position errors. We study the key factors for position errors by systematically investigating (1) the number of access points, (2) the number of samples in the training phase, (3) the number of samples in the position determination phase, and (4) the setup of the grid of reference points. Further, we squeeze out the best of the positioning system by selecting advantageous values for these parameters. For our study, we utilize a test environment with a size of about 312 square meters that is covered with 612 reference points arranged in an equally spaced grid.

1 Introduction

During recent years, we have seen considerable improvements in down-sizing computer hardware and in increasing the capacity of rechargeable batteries, as well as the advent of wireless networks for the mass markets. These technologies allowed manufacturers to build mobile devices that have a similar performance as desktop computers had several years ago. The benefit of these mobile devices can be leveraged by so-called *location-based services*: Applications that act differently depending on the position of the user or, even better, pro-actively offer location-dependent information to the user. Location-based services are currently a hot topic in research, and are considered to be a promising market.

Nowadays, the *Global Positioning System* (GPS) [1] is the predominant outdoor positioning system. Though GPS works well in many outdoor scenarios, it suffers from obstacles such as skyscrapers creating shielded street canyons or walls and ceilings blocking the radio signals indoors. One of the most promising technologies that could be an equivalent to GPS for indoor applications are *802.11-based positioning systems* [2] [3]. Lately, 802.11 hardware is readily available and installed nearly everywhere where people live and work. Another important fact is that 802.11 can be used for communications as well as for positioning purposes at the same time. Even better, almost all modern PDAs, cellphones and laptops are capable to communicate with the 802.11 infrastructure because they are shipped with built-in 802.11 hardware.

The most promising 802.11-based positioning systems utilize the so-called *fingerprint* approach [2]. This technique comprises two stages: An offline training phase and an online position determination phase. During the offline phase, the signal strength distributions are collected from access points at pre-defined reference points in the operation area. They are stored in a table together with their physical coordinates. An entry in this dataset is called a fingerprint. During the position determination phase, mobile devices sample the signal strengths of access points in their communication range and search for similar patterns in the fingerprint data. The best match is selected, and its physical coordinates are returned as the position estimate.

Recent research has mainly focused on algorithms that compute the best match (e.g., [4] [5] [6]). Although, the authors of these papers provide experimental results and compare their own work to existing approaches, they neglect an in-depth analysis of the impact of different factors for position errors. To our knowledge, this paper is the first to present a detailed analysis of key factors causing position errors. The questions we seek answers to are the following:

- How does the number of access points influence position errors?
- What is the impact of the number of training set samples on position errors?
- What is the impact of the number of online samples on position errors?
- How does the grid spacing and starting point of the grid of reference points contribute to position errors?
- What is the lower bound of the average position error achievable with a 802.11-based positioning system if all parameters are set to the best possible values?

Answers of the above questions have implications on the planning, deployment and administration of 802.11-based positioning systems. Furthermore, our analysis will also be helpful for the research area of position determination algorithms.

We use our test environment on the entire floor of an office building on the campus of the University of Anonymity to carry out our study. Although we have worked with only one test environment, the consistency of parts of our results (e.g., number of samples) with results published by other researchers indicates that the conclusions we draw are indeed meaningful. Further, we selected the position determination algorithm proposed by Haeberlen et al. [6] because it shows the best performance. However, we also performed tests with other algorithms (e.g., [2] [7]) and our spot checks indicate that the results are also applicable to these algorithms.

Our results are as follows:

- The number of access points is a primary factor in determining position errors.
- For the training phase, 20 samples at each reference point are sufficient.
- For the number of samples in the position determination phase, no single value can be determined. The trade-off here is about improved position errors

and the time required to calculate a position fix. Thus, for a positioning system running in tracking mode, a high frequency of position updates is required and hence we recommend three samples. Otherwise, 15 samples lead to the best position errors.

- Although, a grid spacing of 0.5 meters leads to the best results, the amount of time required to collect the data for the training phase is hardly bearable. Again, we have to trade position error against time. So, we recommend a grid spacing between 1.0 and 2.5 meters. An operator can select a grid spacing in this range depending on the amount of time he is willing to spend and the position accuracy he is expecting. To find a suitable starting point for a given grid spacing we provide an algorithm.
- We observed a bottom line of 2.0 meters for the average position error, even if we select advantageous values for the parameters.

The rest of the paper is organized as follows. The next section (Sec. 2) presents the related work. In Sec. 3 we describe our experimental setup and methodology in-depth. Subsequently, Sec. 4 presents a detailed analysis of various values for particular parameters we have identified as key factors for position errors. In Sec. 5, we discuss the implications that can be drawn from the results we observed. We conclude the paper in Sec. 6.

2 Related Work

In their preliminary work Bahl et al. proposed the first 802.11-based indoor positioning system [2]. In this paper, the authors provided a few experimental results and mainly focused on position determination algorithms. A few months later, Bahl et al. released a technical report [8] that offers additional experimental results and more position determination algorithms. Although the authors provide results from a second test environment they mainly focus on a tracking algorithm.

In contrast to the algorithms presented by Bahl et al., Castro and Muntz came up with the idea of using probabilistic algorithms [3]. In [7] and [6], two groups from Rice University have embraced this idea and proposed two probabilistic algorithms. The first algorithm requires a histogram of signal strength samples at each reference point resulting in huge piles of data. In their second approach, the histograms are replaced with Gaussian distributions to alleviate the burden of handling large amounts of data. Furthermore, the Gaussian approximation makes their system more accurate.

Moustafa and Agrawala show in a mathematical analysis that probabilistic approaches outperform deterministic position determination algorithms [9]. Further, Moustafa et al. propose different algorithms for the position determination and for the tracking of users [10] [4].

Although all these papers mainly focus on the algorithms to determine the position of users and most of them provide experimental results, none of them systematically investigates the key factors leading to position errors.

3 Experimental Setup and Measurement Methodology

In this section, we first briefly describe our experimental environment (Sec. 3.1). We then present the hardware and software setup (Sec. 3.2). Subsequently, we report how we collected the data used in the experiments. Finally, we describe the overall experimental methodology (Sec. 3.4).

3.1 Local Test Environment

We deployed the positioning system on the second floor of an office building on the campus of the University of Anonymity. The operation area is nearly 15 meters in width and 36 meters in length, covering an area of approximately 312 square meters. The floor plan of the operation area is shown in Fig. 1. The large hallway in the left part of the map is connected by two narrow hallways that are separated by rooms such as a copier room, an archive and a kitchen. The rooms depicted on both sides of the narrow hallways are mainly used as offices, and due to access restrictions they could not be included into the operation area.




Fig. 1. Floor plan of the local test environment. The operation area is painted in gray. The blue markers represent the offline reference points and the purple markers show the randomly selected online points. The access points are marked by orange circles and an orange ring.

3.2 Hardware and Software Setup

Initially, the test environment was covered by one Linksys / Cisco WRT54GS and two enterasys RBT-4102-EU access points administered by the computer center of our university. We additionally installed 11 access points: Two D-Link DWL-G700AP, three NETGEAR WG102, and six Linksys / Cisco WRT54G

access points. All access points support 802.11b and 802.11g. Except of one enterasys access point, all access points are located on the same floor as our operation area. This particular enterasys access point is placed on a lower floor, however, it covers the operation area completely. The position of this access point is marked by an orange ring and the positions of the other access points are marked by orange circles (see Fig. 1).

As a client, we used a Lucent Orinoco Silver PCMCIA network card supporting 802.11b. This card was plugged into an IBM Thinkpad R51 running Linux kernel 2.6.13 and Wireless Tools 28pre. To collect signal strength samples, we implemented a framework that contains two parts: A library that cooperates with the network card driver to perform scans and capture internal driver information [11], and an easy-to-use application that stores these information in a file together with additional data such as the physical position and a timestamp. Further, the application configures the library to select a scan frequency and scan technique for the signal strength measurements. For our experiments we used active scanning. Active scanning is defined in the 802.11 standard [12] and it is a technique to find a suitable gateway to the Internet by measuring the signal strength of access points within communication range.

From the driver our library collects the following information for each device that replies to an active scan:

- MAC address of the device
- received signal strength
- noise level
- mode of the device (i.e. access point or ad-hoc)
- frequency used for the communication

Although only the MAC address, mode and received signal strength values are required by 802.11-based positioning systems, we stored the additional information for further analysis and debugging purposes.

3.3 Data Collection

The grid of reference points applied to the operation area includes 612 points with a spacing of 0.5 meter (see the blue markers in Fig. 1). During the offline phase, we collected 110 signal strength samples at each reference point, resulting in 72,600 samples in total. We spent over ten hours to collect all the data, however, we want to point out that for a productive deployment of a positioning system 20 signal strength samples and a grid with grid spacing of 1.5 meters will be sufficient (see Sect. 4), cutting down the expenditure of time to less than half an hour.

For the online phase, we randomly selected 83 coordinates. The only condition to select a point inside the operation area as a online point is that it is surrounded by four reference points of the grid. Again, we collected 110 signal strength samples for each online point, leading to 9,460 samples in total. In Fig. 1 the online points are marked by purple dots.

3.4 Experimental Methodology

Metrics and Parameter Space. For our experiments, we consider a two-dimensional operation area. We define *position error* as the Euclidian distance between the real physical position and the estimated position. Based on this definition, we consider two metrics during our experiments:

- average position error
- standard deviation of the position error

The former metric is also called accuracy, the latter is sometimes named precision. Both metrics are important because users need highly accurate and precise position estimates.

We have identified the parameter space for our measurements as follows:

- *Number of access points:* To study the impact of the number of access points, we vary the number of enabled access points between one and 14.
- *Number of training set samples:* The time required to collect the training set can be approximated by the number of reference points times the number of signal strength samples at each reference point. To lower the deployment burden of fingerprint-based positioning systems, time requirements should be minimized. For this, it is mandatory to know how many samples at each reference point are required during the training phase to produce stable position estimates. Therefore, we varied the number of signal strength samples from one to 110.
- *Number of online samples:* The number of online samples required to calculate a position estimate determines the time how often position updates are available to the user. Typically, a wireless network card requires at least 250 milliseconds to perform an active scan, so, the time between two position updates is a multiple of 250 milliseconds, depending on the number of samples used. For this, we also varied the number of online samples between one and 110.
- *Grid spacing:* As previously mentioned, the time required to collect the training set depends on the number of reference points. For a given operation area, the number of reference points depends on the grid spacing and the starting point of the grid. If the grid spacing is doubled, the number of reference points is approximately square rooted. The grid of reference points that covers our operation area has a grid spacing of 0.5 meters, allowing us to vary the grid spacing between 0.5 and 4.0 meters in 0.5 meter steps.
- *Starting point of the reference grid:* As mentioned in the last item, the number of reference points also depends on the starting point of the reference grid. Especially, in obstacle indoor areas, different starting points might lead to various ways the operation area is covered with reference points. To study the impact of the starting point, we varied the starting point for grids with a spacing larger than 0.5 meters.

Experiments. To investigate each parameter of the parameter space we use the data we have collected as described in Sect. 3.3. We developed a software-suite called Loceva [13] to switch off different values of particular parameters, so that we are able to quickly emulate various scenarios. This approach allows us to study scenarios that could otherwise hardly be investigated due to the enormous amount of time it would take to carry them out.

We define a *basic experiment* that is used as a basis for the subsequent studies. If a study of a particular parameter requires an extension of the basic experiment, the changes are described in the according section. The basic experiment is defined as follows:

- Nine access points are used.
- For each reference point, 20 offline samples are randomly selected out of the 110 samples.
- For the online phase, three samples are randomly chosen from the 110 samples for each coordinate.
- A grid spacing of 1.0 meter is applied.
- The starting point of the grid is set 0.5 meters north (relative to the point of origin).

This basic experiment is repeated 1000 times to achieve statistically stable results. We now present the various experimental results.

4 Experimental Results

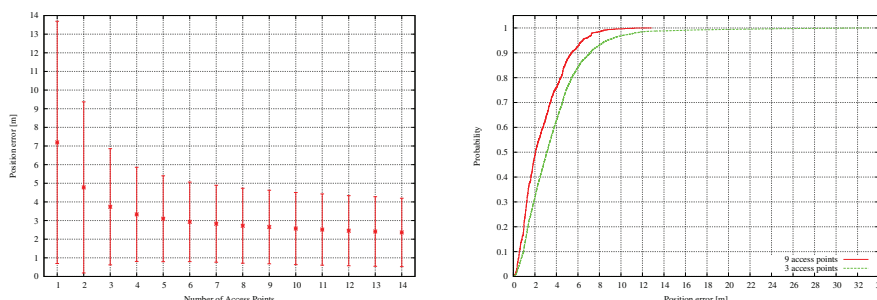
In this section, we first present the impact of the number of access points on position errors (Sec. 4.1). Subsequently, we discuss the influence of the number of samples in the training phase as well as in the position determination phase on position errors. In Sec. 4.4, we present the results of the experiments with various grid spacings and starting points. The best is squeezed out of the positioning system in the last subsection.

4.1 Number of Access Points

To investigate the effect of different numbers of access points on position errors, we extend the basic experiment by varying the number of access points between one and 14. For this, we randomly select the particular number of access points out of the 14 access point covering the operation area.

Figure 2(a) shows the average position error and its standard deviation with respect to the number of access points. As expected, the average position error decreases with an increasing number of access points. Furthermore, we see a marginal utility for each access point added. For instance, the position error drops from about 7.2 meters to about 4.8 meters if the number of access points is increased from one to two. This corresponds to a reduction of more than 2.4 meters or 33 percent. If the number of access points is further increased the

reduction is about one meter (from 4.77 to 3.74 meters). If we make a large step and add the 10th access point we see that the average position error is reduced by about only nine centimeters (from 2.65 to 2.56 meters). Furthermore, not only the average position error decreases, but also does the standard deviation. For example, the standard deviation is about 4.6 meters when two access points are used; it is reduced by about 1.4 meters to 3.1 meters in case three access points are available. And again, the standard deviation shows a similar diminishing utility as seen by the average position error.



(a) The effect of the number of access points on position errors.

(b) The cumulative distribution function of position errors for 3 and 9 access points.

Fig. 2. The impact of the number of access points on position errors.

From the literature we know that in areas of the developed world where people live and work it is common to see on the average three access points [14] [15]. So, with three access points, an average position error of about 3.7 meters is achievable. If we are trying to reduce the average position error by 33 percent (that corresponds to 1.2 meters), at least an additional six access points are required. Unfortunately, in the scenarios described in the literature, such a high number of access points is quite uncommon, however, in some environments such as multi-story buildings or universities, it is typical to see dozens of access points. Even better, access points are quite cheap nowadays, allowing a positioning system operator to deploy additional access points just for the matter of positioning accuracy.

To get a deeper understanding of what the position error distribution for the three and nine access points scenarios look like, we printed the cumulative distribution function in Fig. 2(b). From this figure we see that in 95 percent of all cases the position errors are smaller than 6.5 meters if nine access points are used. If only three access points are utilized, position errors are smaller than 8.7 meters in 95 percent of all cases. It is important to note that not only the average position error is of interest if we compare positioning systems. Important is also the length of the tail of the distribution. Thus, we see that the largest position error is about 12.85 meters for nine access points and 33.23 meters for three access points.

For the basic experiment we have chosen nine access points. Nine access points are more than what we usually encounter, however, we selected this number because we think operators may install extra access points for the sake of position accuracy. We selected the eight access points located in the hallways because this is the place where network operators usually install these devices. The ninth access point is the one installed in the large office in the south-west part of the map. However, the location of the access points is of minor importance because our results show that a particular selection of the nine out of the 14 access points influences position errors only slightly (about a few centimeters).

4.2 Training Set Size

In the training phase, an operator walks from reference point to reference point and collects signal strength samples. Therefore, two factors mainly determine the time required to collect the training data: The number of reference points and the number of offline samples taken at each reference point. The impact of the former on position errors is discussed below. The latter is the objective of this section.

Usually, active scanning is used to collect signal strength samples of access points within communication range. A typical wireless network interface requires 250 milliseconds to complete a scan. Thus, decreasing the number of samples required at each reference point directly reduces the time required to gather the training data.

To see how the training set size affects the positioning accuracy we conducted an experiment that extended the basic experiment by varying the number of offline samples. We start with one offline sample. Next we select five samples. Further, we increase the samples in steps of five up to 110.

As we see from Fig. 3(a), the average position error drops from 4.36 meters to 2.78 meters in case the number of offline samples is increased from one to five. Further, we see a constant decrease of the average position error if the number of offline sample is increased. However, at around 20 samples the marginal utility of adding five additional samples is less than one centimeter, or in other words the average position error only decreases by less than one centimeter if another five samples are added. Not only the average position error is saturated at around 20 samples, also the standard deviation decreases only slightly in case additional samples are collected. This being said, we see that taking more than 20 offline samples is not worth the effort. This is why, we selected 20 samples for the training set size of the basic experiment.

4.3 Online Set Size

In this section, we focus on the online phase and analyze the impact of the number of signal strength samples on position errors. This is interesting to know because the number of samples determines the time required to calculate a reliable position estimate. As mentioned in the previous section, a common 802.11 wireless network card takes 250 milliseconds to measure the signal strength of access

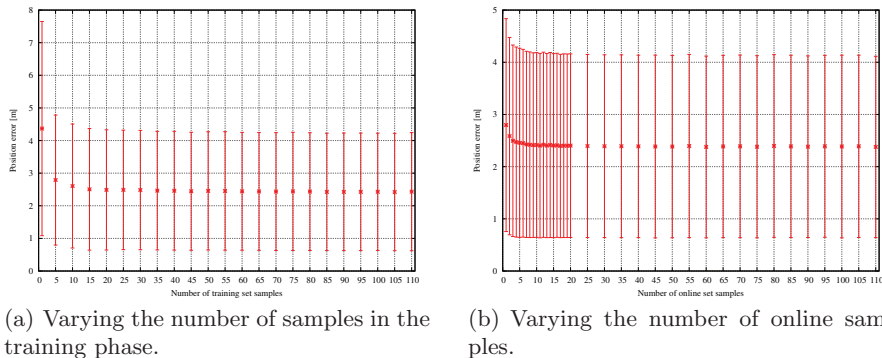


Fig. 3. Average and standard deviation of position errors depending on the number of online and offline samples, respectively.

points within communication range and hence the time to calculate a position fix is a multiple of 250 milliseconds.

The basic experiment is extended in such a way that we vary the number of online samples from one to 110 for this experiment. In the range between one and 20 we investigate every single step, whereas we use an increment of five in the range of 20 to 110.

Figure 3(b) depicts position errors with respect to the number of online samples. We present average and standard deviation because we want to see if a large number of online samples impacts these two measures. Unfortunately, the standard deviation is more or less unaffected by the number of samples. With one signal strength sample the standard deviation is about 2.03 meters and it drops to 1.82 meters if the signal strength is sampled four times. In the range of four to 110 samples, the standard deviation varies only between 1.82 and 1.73 meters.

We see a similar behavior with the average position error. With only one signal strength sample the average position error is about 2.79 meters, but it drops down to 2.49 meters if three signal strength samples are collected. If the number of signal strength samples is further increased the average drops only slightly and we see a diminishing marginal utility. For instance, with 20 samples an average position error of about 2.4 meters is achievable and with 110 samples the position error is on average 2.37 meters.


For the default online set value in the basic experiment, we have to make a trade-off. On one hand, we are interested in the best possible position estimate and on the other hand we want position updates as often as possible. Especially, if the positioning system is running in tracking mode, position updates should be offered quite frequently. For this, we select three samples as the default value because it trades time against position error in such a way that waiting another 250 milliseconds improves the average position error by only about two centimeters whereas the last scan improved the error by about nine centimeters.

4.4 Grid Setup

As stated in Sect. 4.2, one of the factors that determines how much time it takes to collect the training data is the number of reference points. In the literature, most authors utilize an equally spaced grid of reference points to cover the operation area. This makes the whole process quite easy and does not require any further operator interaction. An equally spaced grid is defined by a starting point, the grid spacing, and the angle of the grid alignment. To simplify the scenario we assume that the grid is aligned in the same way as the building that comprises the operation area.

Although the grid spacing is a relevant factor, it is also important to pay attention to the starting point of the grid. Especially, in indoor scenarios we usually face a lot of obstacles (e.g., cabinets, tables or locked rooms) that fragment a floor into subareas. Therefore, the starting point of a grid may determine the size of the area a reference point is associated with. For instance, our test area contains three connected hallways. The small rooms (the copier room, the kitchen, and the archive) in the middle of the virtual large hallway chop it into two narrow hallways that are linked by spaces between the rooms. So, if we overlay such fragmented areas with a grid that utilizes a given grid spacing, it may occur that different numbers of reference points can be deployed depending on the starting point of the grid. Figure 4 illustrates an example of two grids with a grid spacing of 2.0 meters using different starting points. The grid in Fig. 4(a) covers the operation area with 51 reference points, whereas the grid in Fig. 4(b) comprises only 19 reference points. The starting points for these grids are shifted 0.5 meters north, 1.0 meters east and 1.5 meters north, 0.5 meters east, respectively.

An investigation of the impact of the starting point and the grid spacing on position errors is the subject of this section.



(a) This grid contains 51 reference points. The starting point is moved 0.5 meters north and 1.0 meters east compared to the point of origin.

(b) This grid contains 19 reference points. The starting point is moved 1.5 meters north and 0.5 meters east compared to the point of origin.

Fig. 4. Two grids of reference points with a grid spacing of 2.0 meters but different starting points.

Starting Point. We define the point of origin by selecting the bottom left reference point of the 0.5 meter spaced grid (see Fig. 1). Based on this point of origin and the 0.5 meters spaced grid of 612 reference points we derive grids with different starting points and grid spacings between 1.0 and 4.0 meters. For instance, four different 1.0 meter spaced grids can be created by selecting only every other reference point and by moving the starting point 0.5 meters north, east or both.

In the following we selected a grid spacing of 2.0 meters because this scenario can easily be described and shows a tendency that is valid for all other grid spacings as well. With such a grid spacing, 16 starting points can be selected. Depending on the starting point, different numbers of reference points can be applied to the operation area. Table 1 lists the various starting points and the corresponding number of reference points.

Table 1. This table shows the number of reference points for different starting points of 2.0 meter spaced grids. The first value of the starting point column represents the north offset, the second value the east offset.

Starting point	# reference points	Starting point	# reference points
0.0, 0.0	41	1.0, 0.0	34
0.0, 0.5	26	1.0, 0.5	21
0.0, 1.0	50	1.0, 1.0	43
0.0, 1.5	47	1.0, 1.5	42
0.5, 0.0	44	1.5, 0.0	33
0.5, 0.5	30	1.5, 0.5	19
0.5, 1.0	51	1.5, 1.0	42
0.5, 1.5	48	1.5, 1.5	41

From the table we see that the number of reference points varies between 19 and 51, or in other words, depending on the starting points of the grid up to 37 percent of the maximum number of reference points are available. Examples for these extremes are depicted in Fig. 4(a) and Fig. 4(b), respectively.

Although a small number of reference points means that the time required for the data collection phase can be reduced, we expect position errors to increase because each reference point is responsible for a larger region of the operation area. In the following, we investigate this question. For this, we extend the basic experiment by selecting a grid spacing of 2.0 meters and select starting points as listed in Table 1. Figure 5 shows the average position error grouped by the number of reference points. The average position error is between 2.71 and 3.23 meters. 2.71 meters are achieved in case of 50 reference points and 3.23 meters in case of 43 reference points. Furthermore, from the graph we see that the average position error slightly improves if the number of reference point is increases. This tendency is not strictly consistent. For instance, the case of 43 reference points is an outlier.




Fig. 5. Average position error vs. number of reference points.

If we look at different grids of reference points and how they fit into the operation area we see that for some grids there are large areas that are not covered by any reference point. Other grids cover the operation area more "smoothly". For instance, in Fig. 4(b) we see that the two horizontal hallways are not covered at all by any reference points whereas in Fig. 4(a) the operation area is more evenly covered. If we further relate the average position error to the smoothness of arrangement of reference points it follows that smoother grids achieve better position errors. For example, the smooth grid of Fig. 4(a) achieves an average position error of 2.87 meters in contrast to the rough grid of Fig. 4(b) that achieves on 3.17 meters on the average.

It is relatively easy for a human to decide which scenario shows the smoother arrangement of reference points if two scenarios are given. To let the computer take the same decision, we conceived the following algorithm: Randomly select 312 points inside the operation area and add up, for each of these points, the distance to its closest reference point. The scenario with the smaller result is the one with the smoother grid arrangement. We selected 312 points for our 312 square meter operation area. We achieved great results by sticking to the one point for one square meter rule during our tests.

This algorithm can be used to determine which starting point leads to the smoothest grid and therefore to a small average position error. We verified the practicability of this approach by letting the algorithm select the starting point of the smoothest grid and then we compare the average position error of this grid with the other grids. For all grid spacings between 1.0 and 4.0 meters, the algorithm selected a starting point that leads to a grid that achieved at least the third best average position error.

In many real-world deployments there will be no pre-defined starting points to choose from (as in our case). The operator can manually select a few possible starting points and use the approach described above to find the best suitable starting point for the operation area in question.

Grid Spacing. In this section, we focus on the grid spacing and how it impacts the precision and accuracy of the positioning system. For this, we vary the grid spacing of the basic experiment from 0.5 meters stepwise by 0.5 meters until

a spacing of 4.0 meters is reached. For all grid spacings larger than 0.5 meters we applied the technique described in the previous section to select a suitable starting point.

From Fig. 6 we see that with a grid spacing of 0.5 meters an average position error of 2.33 meters and a standard deviation of 1.73 meters is achievable. In case a grid spacing of 4.0 meters is applied the average position error goes up to 2.97 meters and the standard deviation increases slightly to 1.98 meters. Each time the grid spacing is increased by 0.5 meters the average position error increases between 3 and 11 centimeters. This is interesting to notice because if we assume a perfect positioning system that always finds the closest reference point then the average position error should increase 19 centimeters each time the grid spacing is increased by 0.5 meters¹. We call this error the *inherent position error*. Real-world positioning systems are usually not perfect and this is why a second kind of error adds to the position error: the error caused by selecting a reference point that is not closest to the user's real position. This second part is dubbed the *real position error*. Coming back to the observed position errors, it follows that the real position error decreases if the grid spacing is increased. Or in other words, the positioning system is getting better in finding the closest reference point if the grid spacing is increased.




Fig. 6. Average and standard deviation of position errors w.r.t grid spacing.

Figure 7 shows the average signal strength at each reference point for one access point. We selected grid spacings of 0.5 and 2.0 meters to exemplify why it is getting easier for the positioning algorithm to find a closer reference point if the grid spacing is increased. From the figures we see that the number of at least two reference points that share the same average signal strength is decreasing if the grid spacing is increased. If we count only the reference points that do not share their signal strength value with other reference points, we see that in our sparse example nine such reference points can be found. This corresponds to 18 percent of all reference points. In contrast, the example with the 0.5 meter spaced grid contains only five such reference points or 0.008 percent of all reference points.

¹ Let x and y be random variables $\in [0, \dots, a = \frac{\text{gridspacing}}{2}]$ then the average positioning error can be determined by: $E(\sqrt{x^2 + y^2}) = \frac{1}{a^2} \int_0^a \int_0^a \sqrt{x^2 + y^2} dx dy$

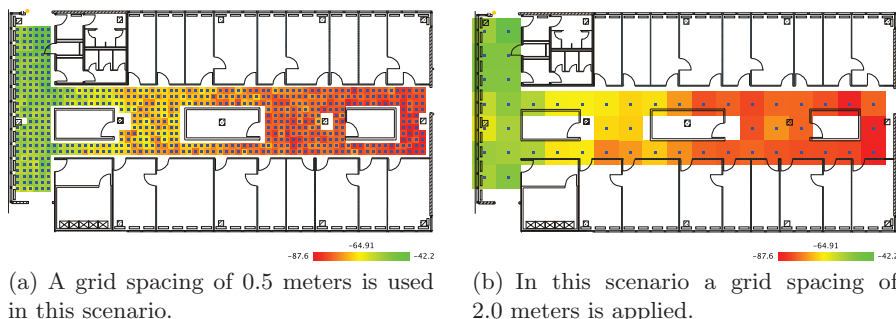


Fig. 7. Two maps with different grid spacings show the average signal strength at each reference point for one access point. The position of the access point is marked by an orange circle.

As already mentioned, the number of reference points is one factor that mainly determines the time require to collect the data in the offline phase. From the previous section we know that the exact number of reference points depends on the shape of the operation area as well as on the inside the operation area. However, for an operator it might be interesting to get a rough estimate of the number of reference points that might cover the operation area. Especially in large deployments this information is helpful to assess the time requirements for the training phase. To simplify the calculation we assume that the operation area is rectangular and we omit obstacles that should not be covered by reference points. These assumptions are valid because a rectangle can be drawn around every shape an operation area might have and omitting obstacles does not increase the actual number of reference points. Instead, the actual number of reference points would be smaller and so our approach yields to a upper bound. Let l be the length of the rectangle approximating the operation area and w its width. Furthermore, let d be the selected grid spacing of the equally spaced grid of reference points. Then the maximum number of reference points m that cover this area can be calculated as:

$$m = \left(\frac{l}{d} + 1\right) * \left(\frac{w}{d} + 1\right)$$

For instance, if we approximate our operation area of 312 square meters by a rectangle of 36 times 9 meters and select a 2.0 meter spaced grid the area can be covered with at most 105 reference points. From Table 1 we know that in practice the maximum number is 51. The difference here is caused by obstacles and a non-rectangular operation area.

For the basic experiment we selected a grid spacing of 1.0 meter because this is the size used by other researchers in the literature (e.g., [5] [8]). For the starting point, we used our aforementioned algorithm to find the smoothest grid resulting in a starting point moved 0.5 meters north compared to the point of origin.

4.5 Best Case Scenario

In this section, we squeeze the best out of the positioning system. For a few application areas, it sounds quite feasible to deploy a few extra access points and to spend a few more minutes for the collection of the training data, especially if in return the accuracy and precision of the positioning system increases. To exploit the potential of 802.11-based positioning systems, we selected all the parameters that produced the best results in the previous sections: 110 samples for the offline set, 110 samples for the online set, a grid spacing of 0.5 meters, all 14 access points.

As expected, the average position error as well as the standard deviation achieve the best results presented so far: 2.06 meters on the average with a standard deviation of 1.65 meters. Compared to best results presented in the previous sections the average position error drops by about 27 centimeters, corresponding to an improvement of 14 percent (see Sec. 4.4). The standard deviation improves by about five percent, from 1.73 to 1.65 centimeter.

In Fig. 8 we present the cumulative distribution function of this experiment because we are interested in the 95th percentile and in the long tail of the distribution. As you see, we achieve position errors of less than 5 meters in 95 percent of all cases. Furthermore, the maximum position error is less than 6.5 meters. Compared to the results presented in the previous sections, we see that with this experiment, the long tail of the position error distribution can be reduced by more than 5.5 meters (see Sec. 4.1). This is an interesting and important result which makes 802.11-based positioning systems more robust.




Fig. 8. Cumulative distribution function for the best case scenario.

5 Discussion

After we have presented the key factors for position errors in the previous sections, we want to discuss the interesting and surprising implications that are caused by this results. First, we are going to stress the importance of the number of access points for 802.11-based positioning systems. Second, we will discuss the bottom line of the average position error we have observed.

Our results in Sec. 4.1 already show the importance of the number of access points. We want to emphasize their importance here again. Although the other parameters of the parameter space are also important factors to influence position errors, they all come with the drawback that decreasing position errors by using these parameters increases the amount of time required to get the system to work considerably. Instead, adding an extra access point to the operation area requires only a fixed amount of time for the installation. During the training phase as well as during the position determination phase, an additional access point does not influence the time requirements in any way. Increasing the number of access points leads to a decrease of the average position error and its standard deviation or in other words makes the position system more accurate and robust. These two facts make extra access points highly appealing for performance improvements of 802.11-based positioning systems.

In all of our experiments we have seen a “hidden” bottom line for the average position error of 2.0 meters. Even if we select the most advantageous values for the parameters of the positioning system, we are not able to under-run the lower bound of 2.0 meters. This is consistent with the results published by other researchers (e.g., [2] [8] [4]). From this, we draw the conclusion that an average position error of 2.0 meters is the lower bound for 802.11-based positioning systems.

To further improve position errors of 802.11-based positioning systems researchers came up with the idea of sensor-fusion [14]. This means that additional sensors such as Bluetooth or a digital compasses are used in combination with 802.11. First publications proof that the average position error can be reduced to 1.65 meters in case a digital compass is used [5]. However, this approach lessens the advantage that every 802.11-enabled device can be used for positioning purposes out of the box.

6 Conclusions

In this paper, we have presented a measurement study of key factors for position errors for 802.11-based positioning systems. Our results show that for the training phase 20 samples at each reference point are enough. For the position determination phase at least three samples should be selected to achieve reasonable position errors. If the positioning system is not used to track users, better results are achievable with 15 samples. How the grid of reference points should look like cannot be definitely said because there is a trade-off between position errors and time. Therefore, we recommend a grid spacing between 1.0 and 2.5 meters depending on the time the operator is willing to spend to gather data for the training phase. To find a suitable starting point for the reference grid, we have presented an algorithm. The number of access points is a great means to improve position errors for many reasons. Access points are cheaply available nowadays and can be easily installed. Further, the number of access points does not influence the time requirements of the training and position determination phase.

All our experiments as well as the results presented by other researchers show that on average position errors of less than 2.0 meters is not achievable. Thus, we draw the conclusion that there is a "hidden" bottom line of 2.0 meters for the average position error that cannot be under-run.

References

1. Kaplan, E., Hegarty, C., eds.: *Understanding GPS: Principles and Applications*. Second edn. Artech House Incorporated (2005)
2. Bahl, P., Padmanabhan, V.N.: RADAR: An In-Building RF-Based User Location and Tracking System. In: *Proc. IEEE InfoCom*. (2000) 775–784
3. Castro, P., Muntz, R.: Managing Context Data for Smart Spaces. *IEEE Personal Communications* (2000) 44–46
4. Youssef, M., Agrawala, A.: The Horus WLAN Location Determination System. In: *Proc. ACM MobiSys*. (2005) 205–218
5. King, T., Kopf, S., Haenselmann, T., Lubberger, C., Effelsberg, W.: COMPASS: A Probabilistic Indoor Positioning System Based on 802.11 and Digital Compasses. In: *Proc. ACM WinTECH*. (2006) 34–40
6. Haeberlen, A., Flannery, E., Ladd, A.M., Rudys, A., Wallach, D.S., Kavraki, L.E.: Practical Robust Localization over Large-Scale 802.11 Wireless Networks. In: *Proc. ACM MobiCom*. (2004) 70–84
7. Ladd, A.M., Bekris, K.E., Rudys, A., Marceau, G., Kavraki, L.E., Wallach, D.S.: Robotics-Based Location Sensing using Wireless Ethernet. In: *Proc. ACM MobiCom*. (2002) 227–238
8. Bahl, P., Padmanabhan, V.N.: Enhancements of the RADAR User Location and Tracking System. Technical Report MSR-TR-2000-12, Microsoft Research, Microsoft Corporation One Microsoft Way Redmond, WA 98052 (2000)
9. Youssef, M., Agrawala, A.: On the Optimality of WLAN Location Determination Systems. In: *Proc. CNDS*. (2004)
10. Youssef, M., Agrawala, A., Shankar, A.U.: WLAN Location Determination via Clustering and Probability Distributions. In: *Proc. IEEE PerCom*. (2003) 143–150
11. King, T., Kopf, S.: Loelib - A Location Library. Website: <http://www.informatik.uni-mannheim.de/pi4/lib/projects/loelib/> (January 2007)
12. Institute for Electrical and Electronics Engineers, Inc.: ANSI/IEEE Standard 802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. Website: <http://standards.ieee.org/getieee802/> (1999)
13. King, T., Kopf, S.: Loceva - A Location Evaluation Tool. Website: <http://www.informatik.uni-mannheim.de/pi4/lib/projects/loelib/loceva.html> (January 2007)
14. LaMarca, A., Chawathe, Y., Consolvo, S., Hightower, J., Smith, I., Scott, J., Sohn, T., Howard, J., Hughes, J., Potter, F., Tabert, J., Powledge, P., Borriello, G., Schilit, B.: Place Lab: Device Positioning Using Radio Beacons in the Wild. In: *Proc. IEEE PerCom*. (2005) 116–133
15. Bychkovsky, V., Hull, B., Miu, A., Balakrishnan, H., Madden, S.: A Measurement Study of Vehicular Internet Access Using In Situ Wi-Fi Networks. In: *Proc. ACM MobiCom*. (2006) 50–61