Chapter 6

Experimental Results

6.1 Introduction

In the previous chapters different aspects related to tracking and localisation using Radio Frequency (RF) sensors were presented. This chapter provides simulation and experimental results to underline the findings of the previous chapters. It aims to show the previously described effects and to provide an understanding of the localisation quality that can be achieved using RF sensors. It presents the results in three separate sections:

- Sensor modelling: This aspect is demonstrated with a sensor model example using data obtained from real geo-referenced measurements. The example is also used to illustrate the change in the sensor model occurring when compressing it. Furthermore, the effect of antenna mounting height is illustrated with real data sensor models.
- 1D tracking: Several simulation examples using a single-sensor or a multiple-sensor setup highlight the performance of the three filtering types when tracking the position of an agent in one dimension. The simulation examples are well suited to demonstrate certain effects as they occur as a consequence of the filter type used.

  Experimental tracking with real data is also shown to provide an insight into actual RF based tracking performance.
- 2D tracking: Results of range-only tracking and localisation in two dimensions are presented, using real data in a multi-sensor setup. Simulations are employed to highlight the effects as they are occurring in such a setup.

6.2 Sensor modelling results

6.2.1 Example of a sensor model using real data

The basis for the implementation of a sensor model as presented in this thesis is geo-referenced data. Received Signal Strength Indicator (RSSI)-distance pairs, as shown in figure 6.1 (top, left) are logged for the maximum sensor range of 20 m.

Approximating the exponential component of the logged data leads to a curve as shown in black in figure 6.1 (top, right). It can be seen that the exponential approximation may be improved through a better choice of the parameters. Nevertheless, this plot also illustrates how different real data can be, when compared to an exponentially decaying approximation such as the nth power model. Obviously such an approximation is insufficient for high accuracy localisation.

The approximation of the RSSI-distance dependency as based on the sum of the exponential approximation and the Fourier based approximation as shown in figure 6.1 (bottom, right) shows that accurate modelling of the initial data is possible using this approach. The original data points are well represented by this interpolation.

The antenna pattern obtained from experimental data, illustrated in figure 6.1 (bottom, left), shows that the sensor under consideration exhibits a directional behaviour in its propagation pattern. It has an increased gain in the azimuth direction of about 110 degrees, while the gain is decreased in the azimuth direction of about 290 degrees.

Figure 6.2 shows the sensor likelihood function corresponding to a measurement of -90 dBm. This is obtained when implementing the sensor model using the previously described components for the distance dependency and the angular dependency. It can be clearly seen that this sensor likelihood function is not omnidirectional, but exhibits a directional behaviour. Furthermore, it is possible to see that this sensor model exhibits radially strong spatial oscillations as they are introduced through the oscillations of the signal mean (figure 6.1 (bottom, right)).

The model is shown in Cartesian coordinates and it is important to notice that, for a distance $r > 20$ m from the sensor location, the model is no longer valid. This can be seen in the corners of the plots, where the sensor likelihood function is not valid.

6.2.2 Compression

As explained in the previous chapter, the Fourier based approximation lends itself to compress the sensor model, or individual sensor measurements. Figure 6.3 and figure 6.4 both show two sensor likelihood functions corresponding to the measurements of -84 dBm and -94 dBm. The likelihood functions were obtained from the sensor model as shown in the

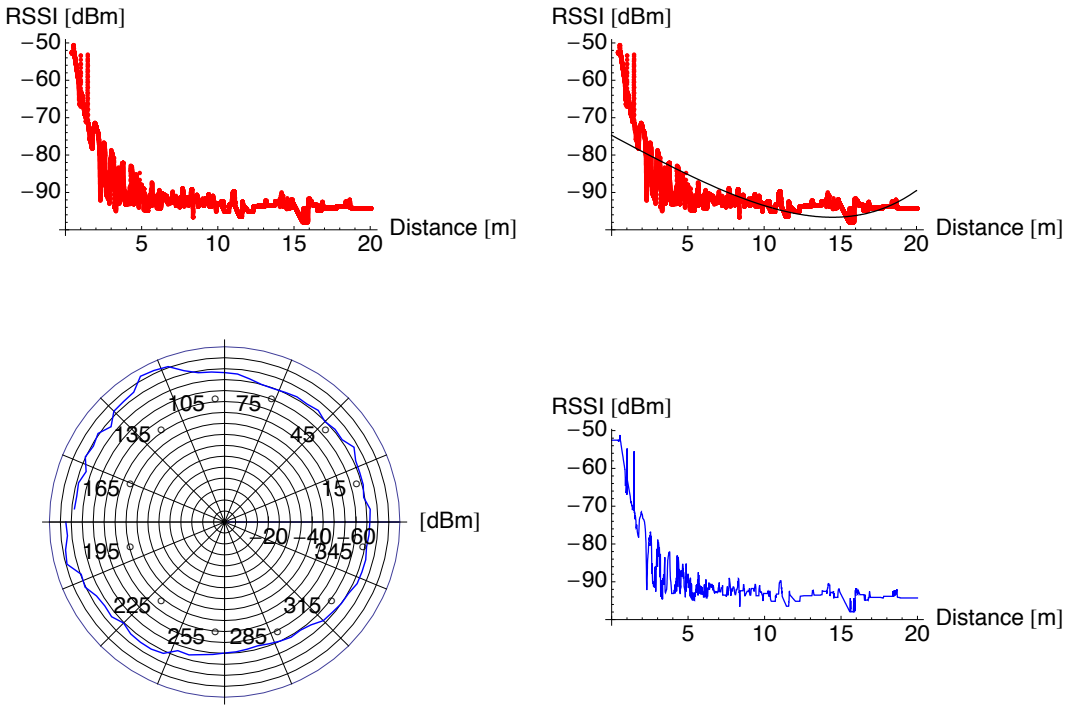

Figure 6.1: Sensor model example with real data (individual components) — The training data (top, left) is fitted with an exponential approximation (top, right). Using the Fourier based description the data points can be approximated very accurately (bottom, right). Finally the angular behaviour of the sensor is also obtained from experimental data (bottom, left).

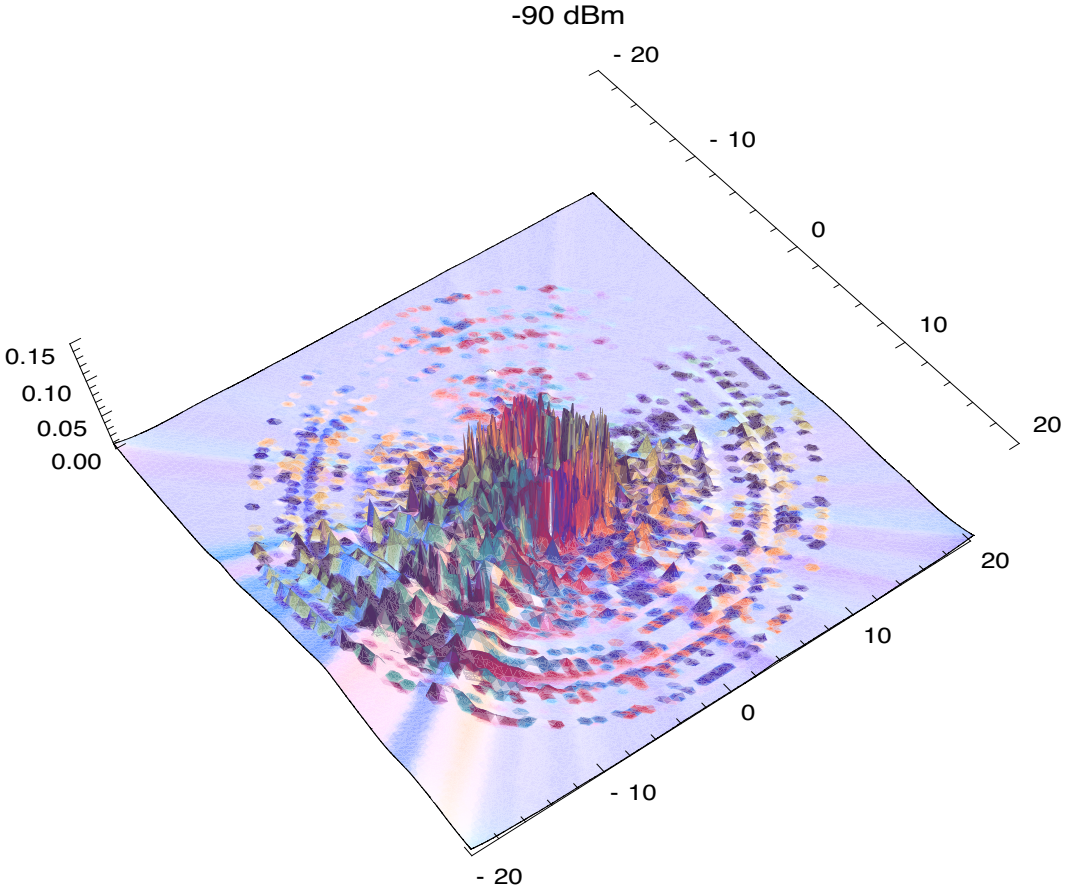

Figure 6.2: Sensor model example with real data (final model) — Combining the individual components shown in figure 6.1 the final sensor model is obtained. A sensor likelihood function for a measurement of -90 dBm is shown. The sensor is located in the centre at (0,0).

previous subsection. Figure 6.3 shows the likelihood functions obtained with all Fourier coefficients, whereas figure 6.4 shows a highly compressed version of the same likelihood functions. The compression effect can be seen as a smoothing of the surface of the sensor likelihood function. For the measurement of -84 dBm the resemblance between the uncompressed and the compressed version of the sensor likelihood function can be seen. This is harder (if not impossible) to judge visually for the measurement of -94 dBm. Here the Kullback-Leibler (KL) measure could give an indication when comparing two compressed models to the initial uncompressed model, but as already mentioned, the value obtained from the KL does not translate into a meaningful figure for the localisation process. It does not determine in which areas the compressed model will perform as well as the original one and in which areas it will not. At least with this simple method there are still a number of simulations and also experiments necessary to confirm that a compressed sensor likelihood function performs satisfactorily in a given application.

For the compressed models, especially at the border of the sensor model, artefacts occur around ranges $r \approx 20$ m. This can be seen in the figures where the compressed sensor likelihood functions suddenly exhibit strong oscillations which are not present in the uncompressed likelihood functions. These artefacts have to be dealt with as they are an obvious sign that high levels of compression can alter the sensor likelihood function significantly.

In the figures the number of Fourier coefficients has been reduced from approximately 20000 for the uncompressed model (figure 6.3) to less than 500 Fourier coefficients for the compressed version (figure 6.4). This is a significant reduction and may prove useful when computing power, or communication bandwidth is limited and measurements (or Probability Density Functions (PDFs)) are required to be transmitted. It is important to stress that compression should only occur to the extent that the compressed sensor likelihood function is still useful for the application.

6.2.3 The influence of antenna mounting height

It was previously presented that the antenna mounting height will affect the signal strength readings. This is demonstrated in the experimental data plotted in figure 6.5 for three different transmitter heights. These figures show the sensor likelihood function for one common given reference angle with the likelihood encoded in a grey scale. This representation is useful to evaluate the effect of mounting height.

For very close distances between transmitter and receiver the effect of signal attenuation is very prominent if the transmitter and receiver are vertically separated. The closer they are mounted vertically, i.e. they have the same or similar height, the less the attenuation effect is present as the Line-of-Sight (LOS) ray will be almost unaffected by the antenna

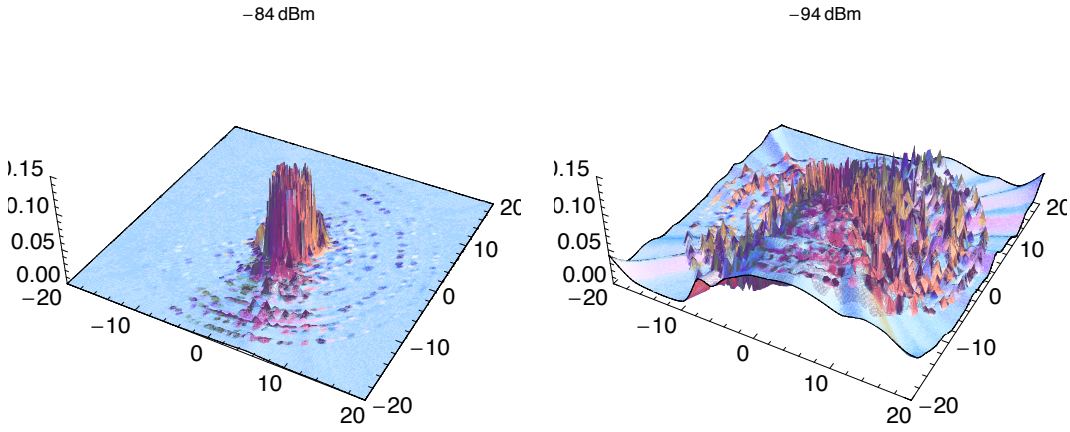

Figure 6.3: Model compression (uncompressed model) — Two measurement likelihood functions for measurements of -84 dBm and -94 dBm are shown. The likelihood function for a measurement of -84 dBm shows high likelihoods for close distances, whereas the likelihood function for -94 dBm is more distributed over the maximum sensing range of 20 m.

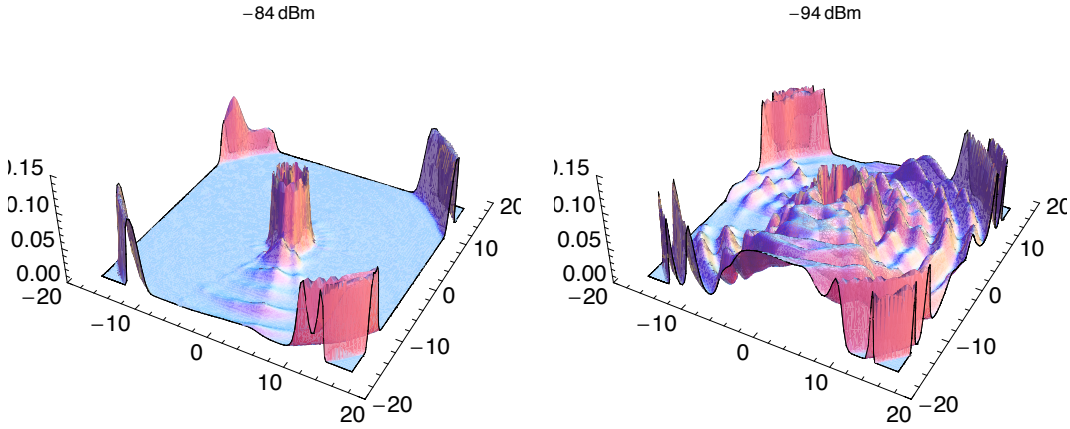

Figure 6.4: Model compression (compressed model) — The likelihood functions for the same sensor and the same measurements from figure 6.3, but obtained with a reduced number of parameters (Fourier coefficients), show a much smoother behaviour. Note that the higher likelihoods towards the maximum sensor range are artefacts stemming from the compression and have to be dealt with as they strongly misrepresent the original sensor model.

gain. This is consistent with the radiation pattern of the $\lambda/4$ antenna, which is used in the hardware for the experiments.

If the separation distance is increased, the effect stemming from the mounting height diminishes. This is because in such a geometrical arrangement the paths of the rays for all three cases become more and more similar. The sensor likelihood functions become more similar and the antenna radiation pattern becomes less relevant for this type of antenna.

A clever arrangement of the fixed sensors in an application, which makes use of the effect of the antenna mounting height on the behaviour of the sensor likelihood functions, can be used to eliminate areas where multiple hypotheses may occur frequently. By adjusting the mounting height of a sensor, its model is changed which leads to different sensor likelihood functions. It should be noted that multiple sensor observations are combined in the localisation and tracking process to obtain a position likelihood. This position likelihood will have more or less modes (hypotheses) depending on the arrangement of the sensor installation locations.

6.3 1-dimensional tracking results

Most of the main effects relating to the three filtering algorithms under consideration can be best shown in 1 dimension. In this case these effects are not masked by other effects that only occur in 2 dimensions. In this section simulations are used to illustrate certain points, and experimental results are used to validate the main contributions.

In the parts showing the simulated results a target is tracked using either one or three *ideal* RF sensors with the sensor model based on the deterministic two-ray model. Starting at a distance of 5m, the distance between transmitter (target) and receiver is increased at a constant rate using the constant velocity model as the process model. The motion occurs along one axis only and the position of the stationary receiver is assumed to be at position (0) for the single sensor case and at the spatially different positions (0), (2) and (4) in the multi-sensor case.

For the experimental results corresponding to the 1-dimensional tracking, the geo-referenced data were taken along a radial track shown in figure 6.6(a). An independent data set was used to evaluate the performance when tracking. This track is shown in figure 6.6(b). In the experimental dataset the observations of the individual sensors occur asynchronously.

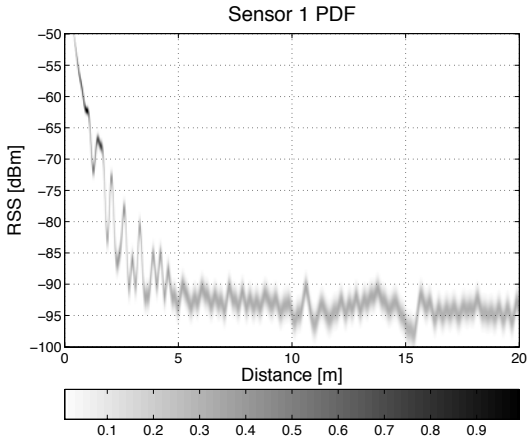

(a)

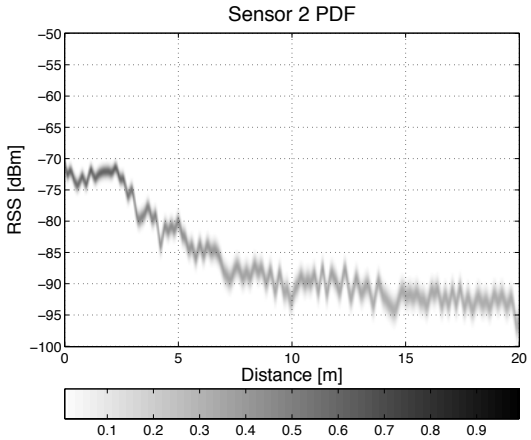

(b)

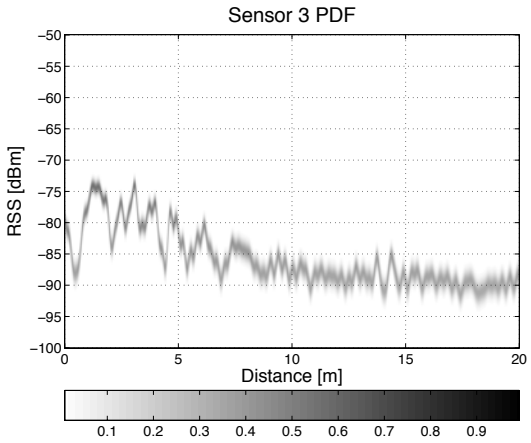

(c)

Figure 6.5: Sensor models for three different receiver heights — Three different sensor models as they are obtained when the receiver height is varied and the transmitter height is kept constant at 2m. (a) 2m mounting height, (b) 2.5m mounting height and (c) 3m mounting height. The influence for short distances can be clearly seen, with a strong signal drop for the higher antenna mounting height.

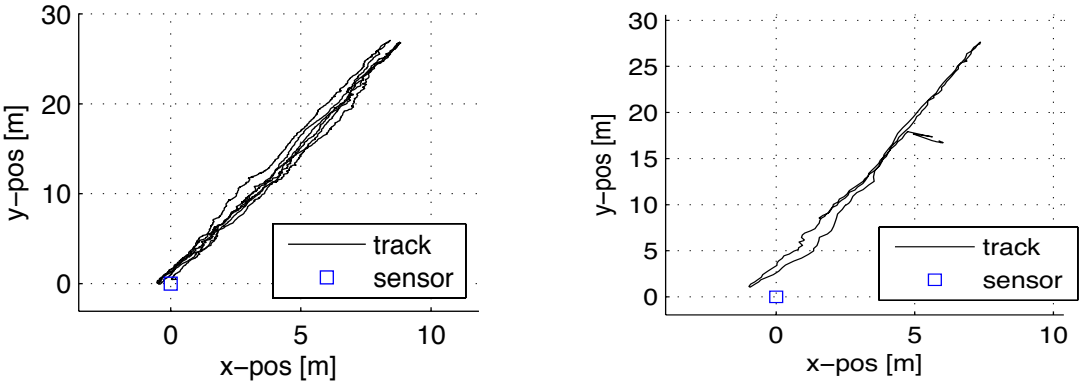

Figure 6.6: Experimental tracking with RF sensors — multi-sensor case (experimental setup): (a) the track used in the sensor modelling process and (b) the track used for the experimental tracking. The three sensors used are all mounted at the same location (0,0) but at different heights which results in a different sensor model for each sensor.

6.3.1 Kalman Filter tracking

Example 1 — Single sensor tracking (simulation)

In the theoretical considerations in subsection 4.4.2 the difficulty (or impossibility) of approximating multi-modal sensor measurements using a uni-modal Gaussian distribution suitable for the Kalman Filter (KF) was discussed. Subsection 5.2.2 showed how the nth power model will, most of the time, estimate distances as either too short or too long.

In this example a single target is tracked using one RF sensor. As the KF cannot handle the multi-modal measurements resulting from this sensor model, the measurements are approximated either based on the corresponding nth power model or by calculating the first two moments of the multi-modal observation distribution as discussed in the previous chapter. Both sensor models, the two-ray model and the nth power model, were shown previously and can be seen in figure 4.4.

Figure 6.7 shows the results when using the nth power model. It can be seen that the mismatch between the model where the simulation observations are generated (two-ray model), and the model where the KF observations are approximated (nth power model), has a serious impact on filter performance. The filter is consistently overconfident in its covariance estimate due to the artificially chosen measurement covariance.

It is also very interesting to see that, as expected, the behaviour of the true sensor model with its high frequency spatial oscillations for small separation distances is reflected in the KF estimate. These oscillations occur because of the effect described in subsection 5.2.2.

If the measurements are approximated by calculating the mean and the variance of the

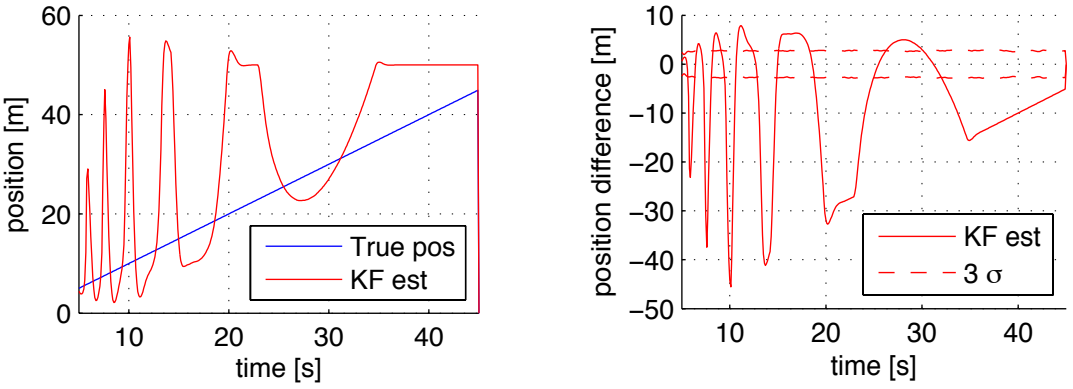

(a) (b)

Figure 6.7: Tracking with the KF and the nth power model — 1 sensor case (simulation): (a) True position, i.e. distance between transmitter and receiver and the KF estimate. (b) Difference between the true position and the KF estimate with corresponding 3σ bounds.

true multi-modal observation distribution, the results are similar to the previous case. As figure 6.8 shows this measurement approximation also exhibits oscillating behaviour, though at a smaller magnitude than for the nth power model approximation. What can also be seen is that the state covariance now better reflects the true values. This is due to the fact that the variance for the observation is calculated anew for every observation, and this is reflected in the state covariance.

In both cases the tracking results are very poor and confirm the theoretical considerations from chapter 4 and chapter 5.

Example 2 — Multi-sensor tracking (simulation)

The multi-sensor tracking case is interesting for two reasons. Firstly, it is of interest to see whether multiple sensors can alleviate the severe failure mode of the single sensor tracker. Secondly, in the real application an agent will usually be seen by more than one tag reader as it is necessary to have more than one sensor in spatially different locations to be able to localise in 2 dimensions.

Figure 6.9 and figure 6.10 show that the result is similar to the previous example using only a single sensor. Again, the estimate oscillates around the true value. For this result though, spatial mixture frequencies occur, which is the effect of having more than one sensor arranged at spatially different locations. The comparison with the single sensor example also shows that the magnitude of the deviations is decreased.

Again, the approximation based on the calculation of the mean and the variance of the

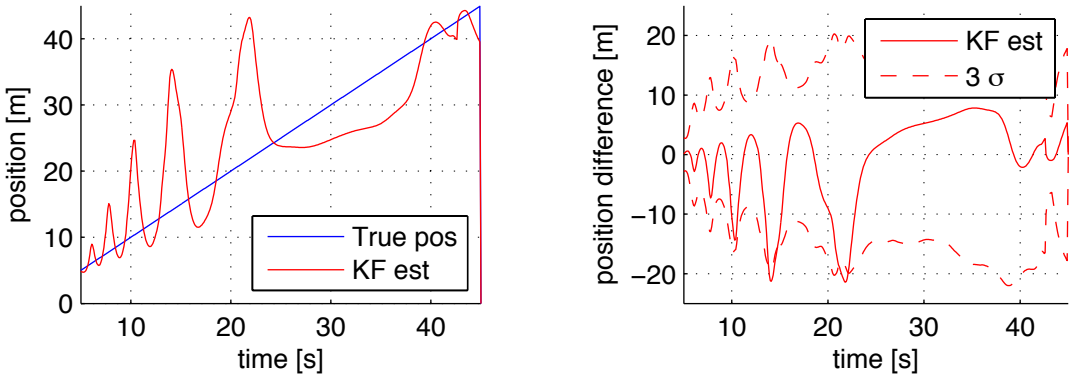

(a) (b)

Figure 6.8: Tracking with the KF and calculating the mean and variance of the multi-modal observation distribution — 1 sensor case (simulation): (a) True position, i.e. distance between transmitter and receiver and the KF estimate. (b) Difference between the true position and the KF estimate with corresponding 3σ bounds.

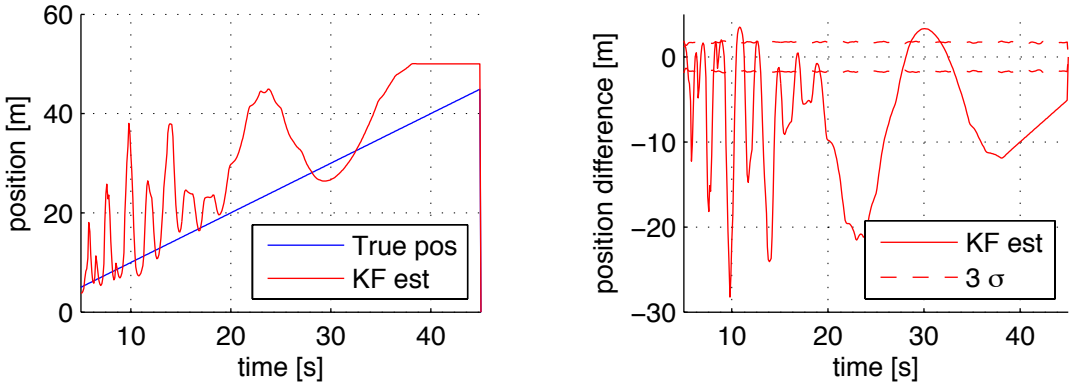

(a) (b)

Figure 6.9: Tracking with the KF and the nth power model — multi-sensor case (simulation): (a) True position, i.e. distance between transmitter and receiver and the KF estimate. (b) Difference between the true position and the KF estimate with corresponding 3σ bounds.

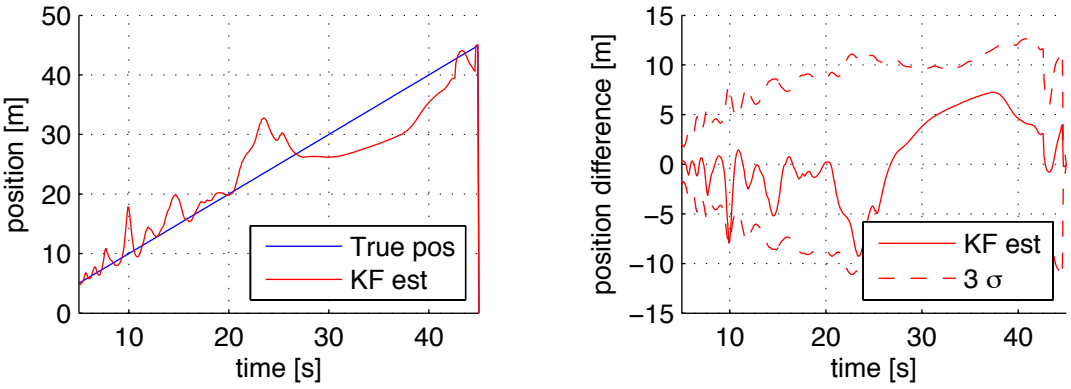

(a) (b)

Figure 6.10: Tracking with the KF and calculating the mean and variance of the multi-modal observation distribution — multi-sensor case (simulation): (a) True position, i.e. distance between transmitter and receiver and the KF estimate. (b) Difference between the true position and the KF estimate with corresponding 3σ bounds.

measurement likelihood is slightly better than the nth power model. In this case also, the state covariance is consistent with the behaviour caused by the difference between the true position and the estimated position.

It seems that the use of multiple RF sensors is beneficial to the tracking and localisation process. Nevertheless the behaviour is in both cases very similar and the tracking result is not satisfactory.

Example 3 — Multi-sensor tracking (experiment)

This example shows experimental tracking using the KF to underline the findings obtained through simulation. The experimental setup is as described at the beginning of this section. Figure 6.11 shows the deviation of the estimate from the true distance, while figure 6.12 shows the true and the estimated track. The deviation in figure 6.11 oscillates in an almost deterministic manner. The zero crossings of the deviation can be allocated to two distinct positions (areas) on the track. The first area is very close to the sensor locations at (0,0) (marked by the circled two) and the second area is about midway between the tracks (marked by the circled one). When looking at the measurement approximations for the KF, the reasons for the behaviour of the deviation become apparent.

Typical measurement approximations are shown in figure 6.13 and figure 6.14. Figure 6.13 portrays approximations that are typical for larger separation distances between transmitter and receiver (the area marked by the circled one in figure 6.11). Figure 6.14 shows

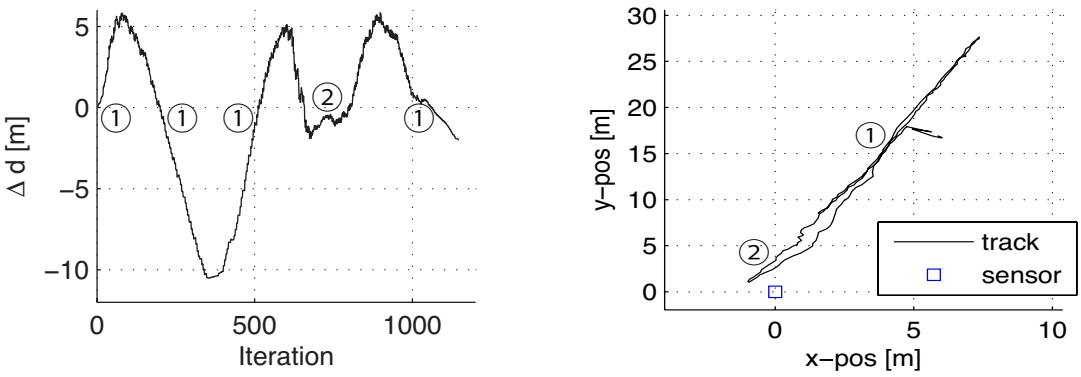

(a) (b)

Figure 6.11: Experimental tracking with RF sensors — multi-sensor case for the KF (2D experimental results): (a) The deviation as obtained through the tracking process and (b) the track used for the experimental tracking. The circled numbers in (a) are points where the deviation is zero or very small and correspond to the locations with the circled numbers in (b). Note that (b) shows the trajectory used to record the experimental data. The estimated trajectory is shown in figure 6.12.

the measurement likelihood for short distances (the area marked by the circled two in figure 6.11).

For larger distances, the underlying multi-modal distribution is widely spread as shown in figure 6.13. The corresponding Gaussian approximation will have its mean somewhere in the middle of the underlying distribution. This happens for a relatively large area. For many of the measurements for larger distances, the Gaussian approximation of the multi-modal distribution looks similar. This means that the mean of this distribution is in almost the same location, marked with the circled one. As a result of this, the position is estimated to be almost stationary in that location. As the agent in reality keeps moving, this leads to an increase in the deviation when the agent moves away from this point, and to a decrease in the deviation when the agent gets closer.

Figure 6.14 shows that for small distances, the multi-modal distribution is peaked around short distances. This results in a comparatively better approximation with a Gaussian distribution. Similarly to the previous situation, a dominant point (or area) develops where the mean of most of these Gaussians approximating the measurements lies. In this case though, this point is closely located to the sensor location at position (0,0). Therefore, the state estimate tends to be at or around the point marked with the circled two.

This tracking example using real data clearly documents that the KF is not suited for an application with such highly non-Gaussian measurement likelihood functions. The systematic failures expected were demonstrated with simulations and experimental results.

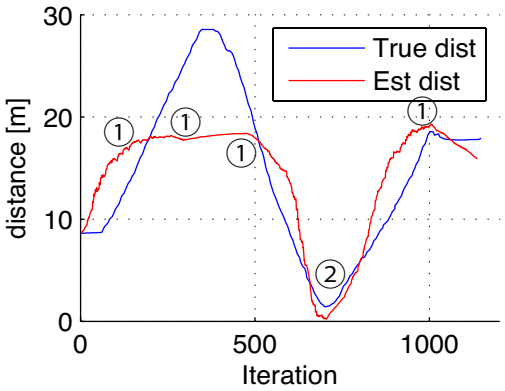

Figure 6.12: Experimental tracking with RF sensors — multi-sensor case for the KF (experimental result - estimated track): The true and the estimated track. The circled numbers correspond to the numbers in figure 6.11.

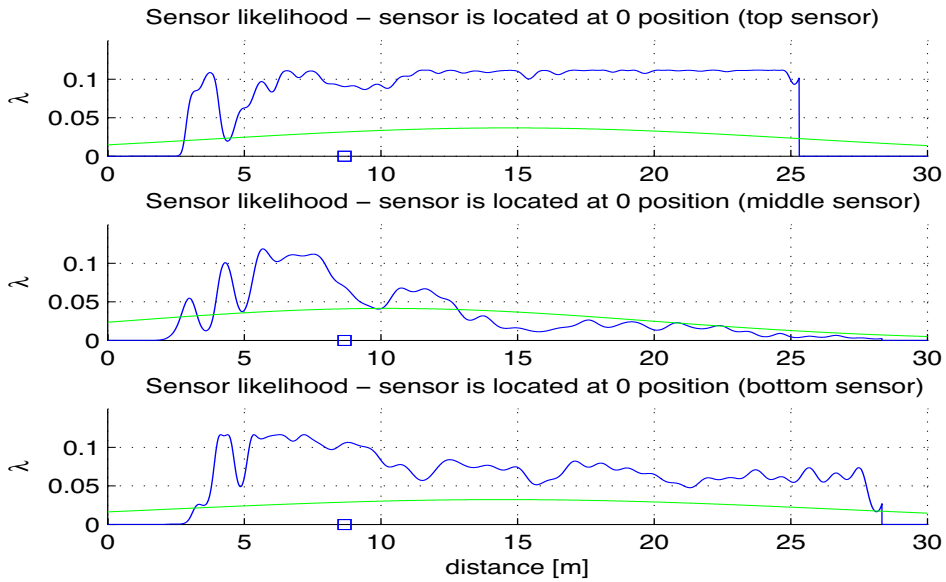

Figure 6.13: Experimental tracking with RF sensors — multi-sensor observations for the KF (large distances): Three measurement likelihood functions as obtained from the three sensors for large distances. For a wide range of distances between the transmitter and the receiver the multi-modal distribution (blue line) is very wide spread. The resulting Gaussian approximation (green line) is very similar with very small variations in the mean. This represents the true location very badly.

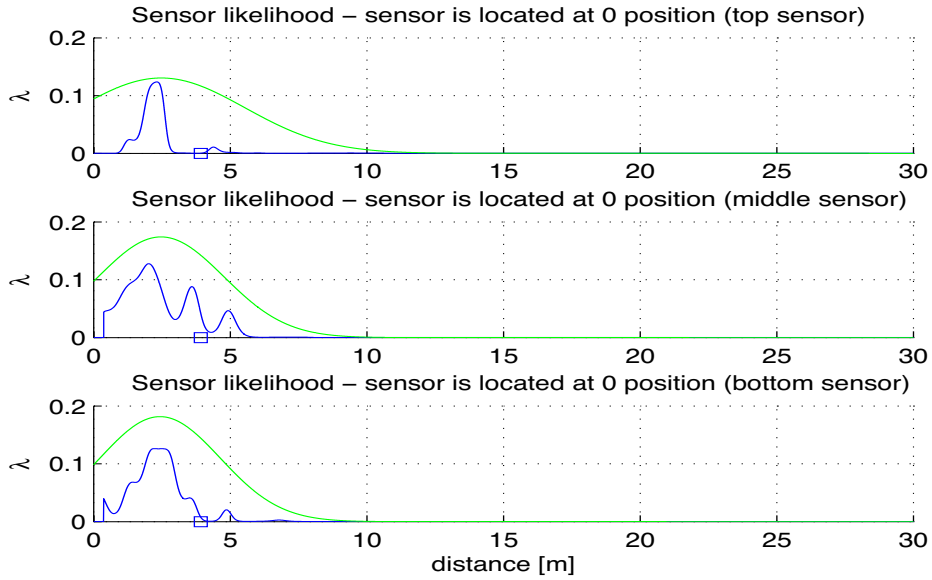

Figure 6.14: Experimental tracking with RF sensors — multi-sensor observations for the KF (small distances): Three measurement likelihood functions as obtained from the three sensors for small distances. In this case the multi-modal distributions (blue line) are peaked around small distances which leads to the Gaussian approximations (green line) centred close to (0).

This is clearly due to the critical mismatch between the Gaussian approximation and the multi-modal distribution which corresponds to a real measurement likelihood function.

As the KF is obviously not suited for tracking and localisation using RF sensors, it is not considered later on in the section dealing with 2-dimensional range-only tracking and localisation.

6.3.2 Particle Filter tracking

Example 4 — Multi-sensor tracking (simulation)

It was shown in chapter 4 that the Particle Filter (PF) can work with multi-modal distributions. Indeed, this sample based approach can be used to approximate distributions of any shape. The following example demonstrates the importance of the process model for the PF when used with the multi-modal observations as they occur with RF sensors. The simulation data correspond to the three-sensor simulation data set used previously.

This example uses the constant velocity model and compares two cases. The first case is characterised by high process noise, and the second case by low process noise. All other parameters are the same. Having high process noise can be interpreted as being unsure

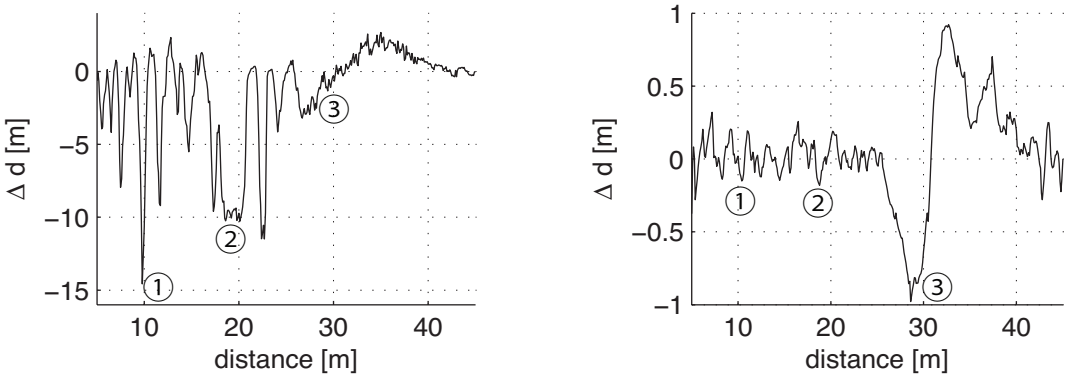

Figure 6.15: Tracking with the PF — multi-sensor tracking (simulation): (a) The difference between the true position and the estimated position as obtained with high process noise and (b) when the process noise is very small. The circled numbers correspond to the state PDFs shown in figure 6.16 to figure 6.18.

about the quality of the process model. This is a reasonable assumption when trying to model the motion of people, as it is hard to predict (model) their behaviour. Low process noise is used when the process model describes the behaviour of the modelled process well.

In PFs, especially with highly multi-modal and peaked state distributions, the added process noise is often increased so that the particles can *explore* the distribution. This approach avoids the problem of particles getting stuck in one peak. Such behaviour is desired to ensure that particles can move away from hypotheses which do not correspond to the true state, and move towards the true state.

The comparison is done by looking at the difference between the true state and the estimated state. This difference, together with some examples of what the state PDF looks like at particular points in the estimation process, explains why PFs need a good process model. This is important when the PFs are used in conjunction with highly multi-modal observations.

In figure 6.15, the result of the 1-dimensional PF tracking is shown. Figure 6.15(a) shows the result for high process noise (which is probably closer to the situation of modelling a person's motion) and figure 6.16(b) illustrates the results of very low process noise. Overall, it can be seen that the tracker performs well for the case with small process noise, with only small deviations from the true state. As soon as the process noise is increased to match the behaviour of a person, the quality of tracking diminishes.

For three specific locations along the track the effects of the process noise are analysed. The circled numbers in figure 6.15 correspond to the state PDFs shown in figures 6.16 to 6.18. The left figures (a) correspond to the high process noise case and the right figures (b)

to the low process noise case. The following three points describe the effects occurring at the corresponding three locations:

1. The high process noise leads to particles being spread and moved away from the true position. During subsequent resampling, even more particles move away from the true position as the measurement likelihood function is very peaked around the true position, and spread out for larger distances. A very peaked PDF is hard to represent with particles, and often only a very limited number of particles with high weights are included in such areas. Figure 6.16(a) shows that for the true position only a small number of particles with high weights are used to represent the distribution, and that many particles have moved away from the true position.

   In the case of low process noise (see figure 6.16(b)) the state PDF represents the state very well. The peak is located at the true position. This is an indication that the spread of the particles for the case of high process noise is indeed due to the process noise and not a consequence of the observation likelihood functions.

2. Figure 6.17 confirms the previously described situation. Again we see a tendency for the particles to move away from the true state towards larger distances, i.e. areas where the observation likelihoods are more spread and not so highly peaked. In this example the distribution for the case with low process noise, figure 6.17(b), is marginally broader than before, but still provides a good representation of the true state.

3. For larger distances the influence of the process noise diminishes or is hidden by the fact that the true position is in the distance range where the particles would have been moved to. This is shown in figure 6.18.

So, overall it can be seen that in a PF application with higher process noise for the constant velocity model (corresponding to a person's motion with all the possible erratic changes in direction and velocity), and using multi-modal observation likelihood functions typical for RF sensors (with high spatial oscillations for small distances and lower spatial oscillations for larger distances), PF range estimates tend to be biased towards larger distances.

It has to be noted that in the example presented here the PF state estimate was calculated as the weighted mean of the particles and that this mean, especially in the high process noise case, is not a accurate representation of the true state. The PDF in contrast represents the state as a distribution, and in all three examples has particles at the true location. So if the PDF is interpreted to represent all possible locations, it can be used to show the possible locations of the target that is tracked, even in the case of using high process noise in conjunction with the constant velocity model.

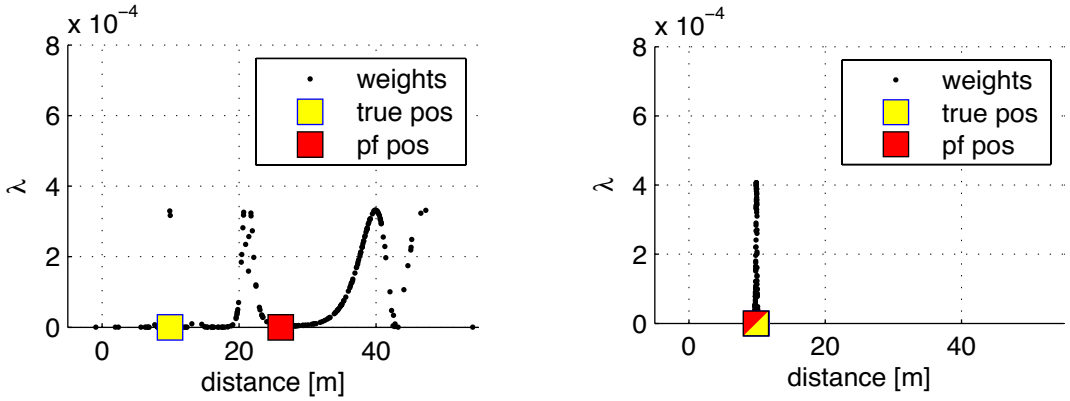

Figure 6.16: PF multi-sensor tracking (simulation — situation 1): The true state is at a distance of 10 m. (a) High process noise lets many particles *explore* the state PDF. (b) A good process model leads to the particles concentrated at the true position. The symbols in the figure are on top of each other.

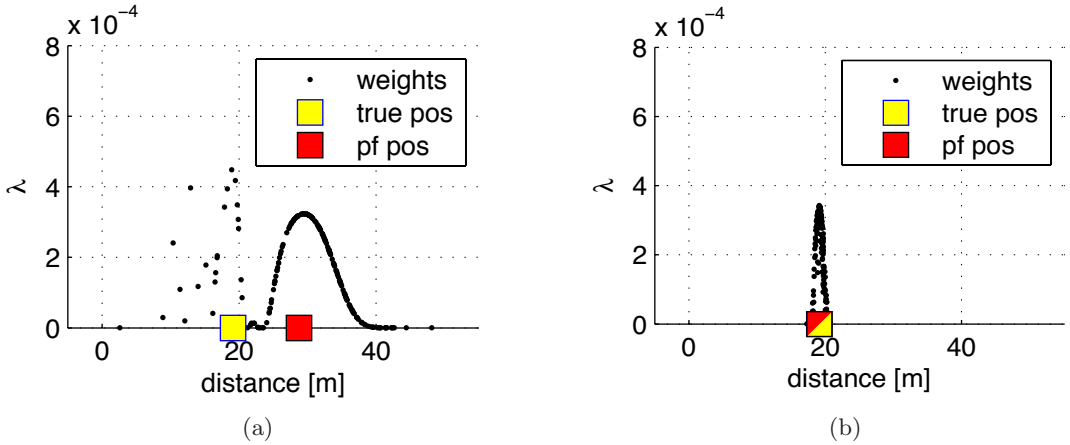

Figure 6.17: PF multi-sensor tracking (simulation — situation 2): The true state is at a distance of 20 m. (a) Similar to the figure 6.16(a) very few particles are at the true position.(b) The estimated position is close to the true position. The symbols in the figure are again on top of each other.

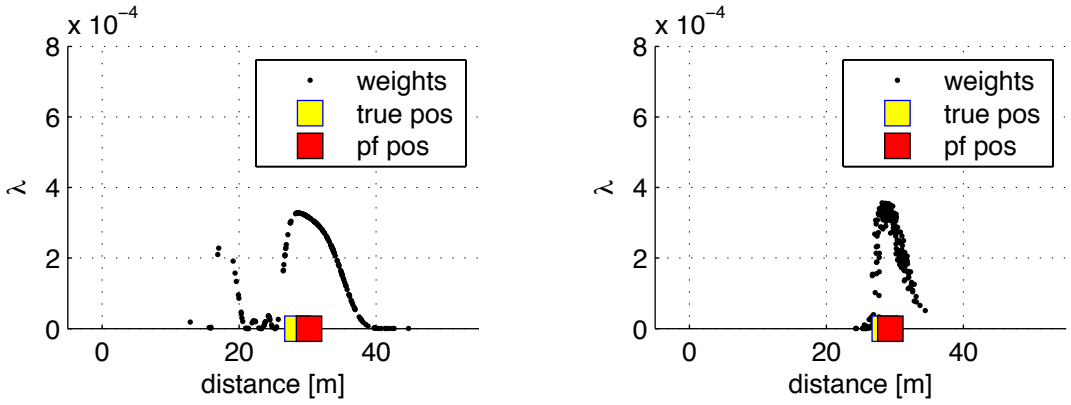

Figure 6.18: PF multi-sensor tracking (simulation — situation 3): The true state is at a distance of 30 m. (a) and (b) For larger distances the estimate of the position is good also in the case of high process noise.

Example 5 — Multi-sensor tracking (experiment)

The experimental result when tracking the position of a target in 1 dimension provides confirmation for the findings of the previous simulation examples. The position of the target can be tracked quite well with the PF as shown in figure 6.19. If the process noise is small (figure 6.19(a)) then the deviation between the true position and the estimated position is also smaller than in the case of higher process noise (shown in figure 6.19(b)). This underlines the importance of a good process model. Nevertheless the parameters used in figure 6.19(b) with the high process noise correspond more to a person, than the very well tuned case shown in figure 6.19(a). Figure 6.20 shows the estimated tracks corresponding to the two described situations.

In both cases a systematic deviation is present at the same point in time, which is indicated by the circled two. This deviation occurs for large distances towards the maximum range of the sensors, as can be seen in figure 6.21.

We first analyse the situation indicated by the circled one, where the deviation is minimal. Figure 6.22 shows a typical example of the state PDF in this region. The particles explore the PDF to both sides of the true location, and the PDF is peaked at the true location. If a resampling step occurs, many particles will be *moved* to the area of the true location as the weights of the particles in this area are high, and as there are already a high number of particles present. This concentration of particles at the true position will lead to a further concentration of particles in a resampling step.

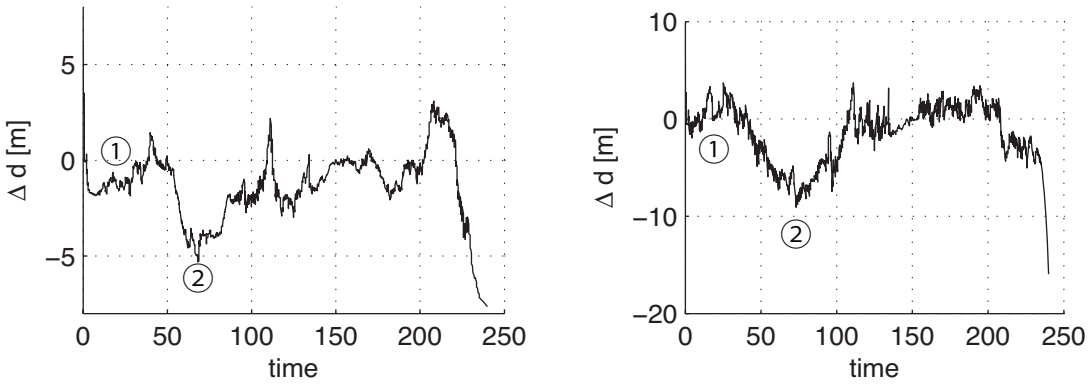

Figure 6.19: Experimental tracking with RF sensors — multi-sensor case for the PF (experimental result): (a) The deviation as obtained through the tracking process when running the filter with low process noise and (b) deviation as obtained through the tracking process when running the filter with high process noise. The circled numbers in both figures correspond to the locations with the circled numbers in figure 6.21.

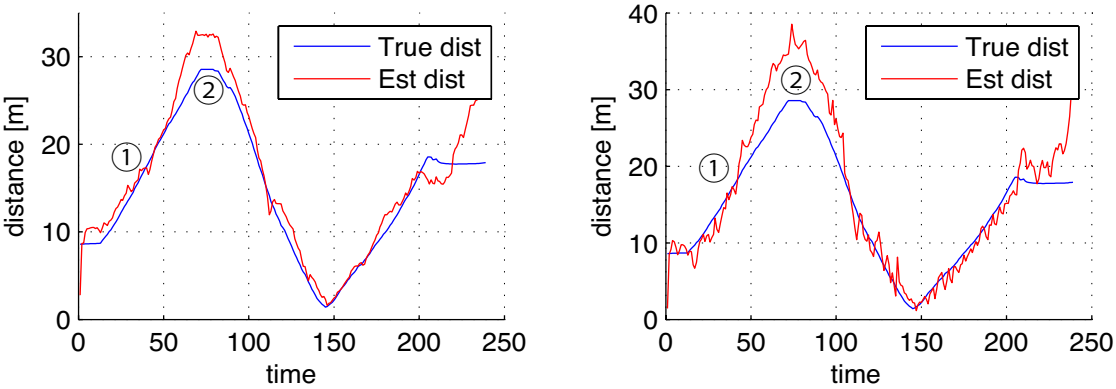

Figure 6.20: Experimental tracking with RF sensors — multi-sensor case for the PF (experimental result - estimated track): (a) The estimated track as obtained through the tracking process when running the filter with low process noise and (b) the estimated track as obtained through the tracking process when running the filter with high process noise. The circled numbers in both figures correspond to the locations with the circled numbers in figure 6.21.

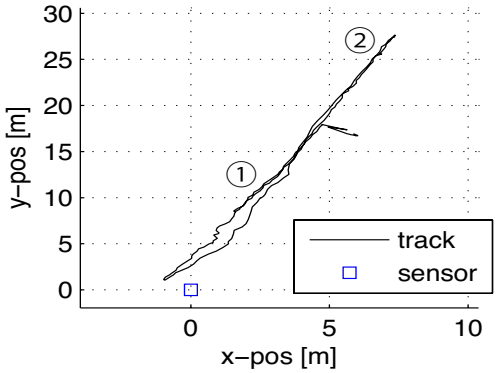

Figure 6.21: Experimental tracking with RF sensors — multi-sensor case for the PF (analysed locations): The two locations (areas) that are analysed with respect to the deviation of the tracker are indicated by the circled numbers. Note that this plot shows the trajectory used to record the experimental data.

For the second case, marked with the circled two, where the deviation from the true position occurs, an instance of the PDF is shown in figure 6.23. Here the true location is almost at the maximum sensing range of the sensors. For large distances the observation likelihood functions tend to be spread. This is reflected in the state PDF, which is also spread over a wide range with similar high probabilities in this area. The weighted mean, often used to represent the most likely state, does not provide a good result in this case as it is not near the true location.

With PFs, and especially in cases where the state PDF tends to be multi-modal, the weighted mean approach has to be used carefully so that incorrect conclusions are not drawn. Indeed it is the shape of the distribution given by the particle location and their corresponding weights which, most of the time, has to be interpreted as a whole. Any reduction to the weighted mean, or any other measure may come at the risk of a complete misinterpretation of what the distribution actually represents.

The examples using simulation and experimental results clearly highlighted that although the PF approach is capable of handling the multi-modal observation likelihood functions, it is not straight forward enough to be used as the resulting multi-modal state PDFs have to be interpreted carefully.

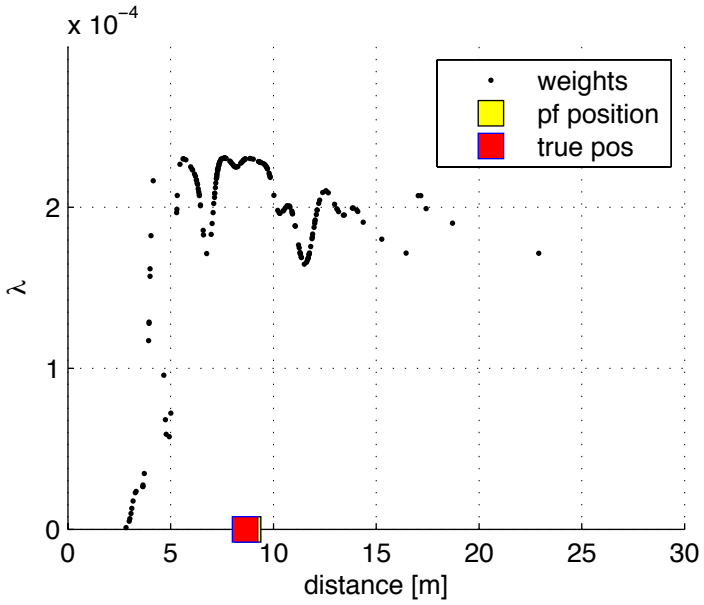

Figure 6.22: PF multi-sensor tracking (experiment — situation 1): This figure depicts the area indicated by the circled one in figure 6.21. The weighted mean (yellow square) corresponds quite well to the true location. Also the PDF is peaked in the area of the true location.

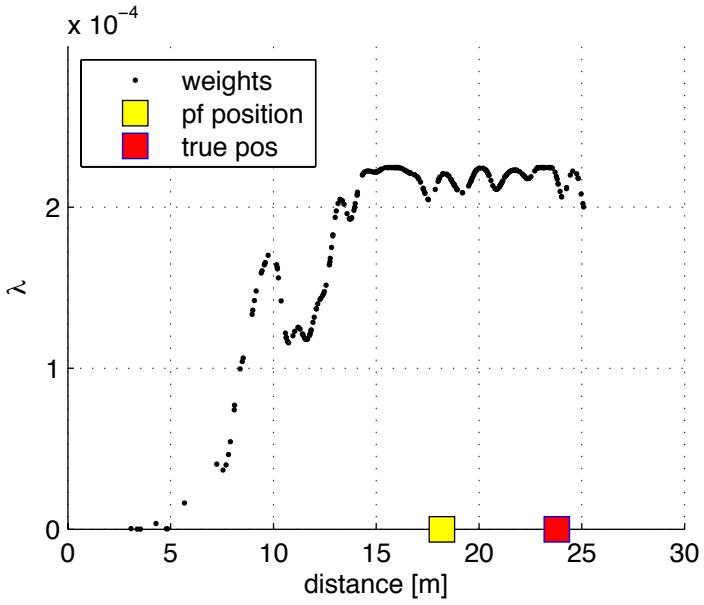

Figure 6.23: PF multi-sensor tracking (experiment — situation 2): This figure depicts the area indicated by the circled two in figure 6.21. The weighted mean (yellow square) does not correspond to the true location anymore, but the PDF is peaked in the area of the true location.

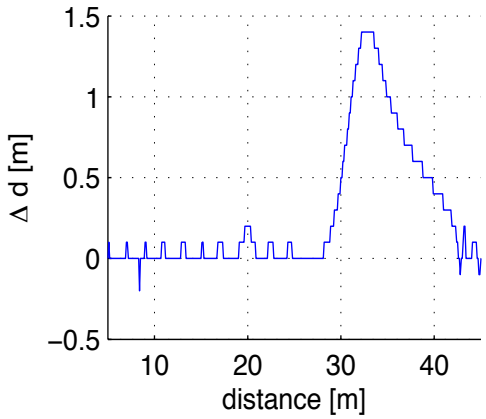

Figure 6.24: Tracking with the HF — multi-sensor tracking (simulation): The deviation between the true state and the estimated state is, for short separation distances, very small. Often only with a grid size of 0.1 m.

6.3.3 Histogram Filter tracking

The final probabilistic tracking technique investigated with 1-dimensional tracking examples is the Histogram Filter (HF). The motion in the trackers for the simulation and the experimental results are implemented as a convolution of the state PDF with a 1-dimensional Gaussian distribution.

Example 6 — Multi-sensor tracking (simulation)

The simulation using multiple sensors provides an exceptionally good result for short to medium distances, with only small deviations at larger distances. In this example the estimated position used to calculate the deviation from the true position is obtained as the Maximum Likelihood (ML) estimate of the state PDF. The deviation of the estimated state from the true state is minimal as shown in figure 6.24. Most of the time the actual deviation is only about the grid size of 0.1 m. This result encourages the use of the ML estimate to obtain the position of the target in HF applications using multi-modal observations modelled in the way described in this thesis. It also hints at the possible suitability of the HF for an application with such sensors. The deviation increases only for larger distances. The reason for this will be explained later on.

The HF is the tracking technique where the motion of the target is modelled through the convolution of the state PDF with a Gaussian kernel. Figure 6.25 shows an example of how this convolution affects the shape of the PDF and smoothes it.

Figure 6.26 shows typical observation likelihood functions as they are obtained for larger

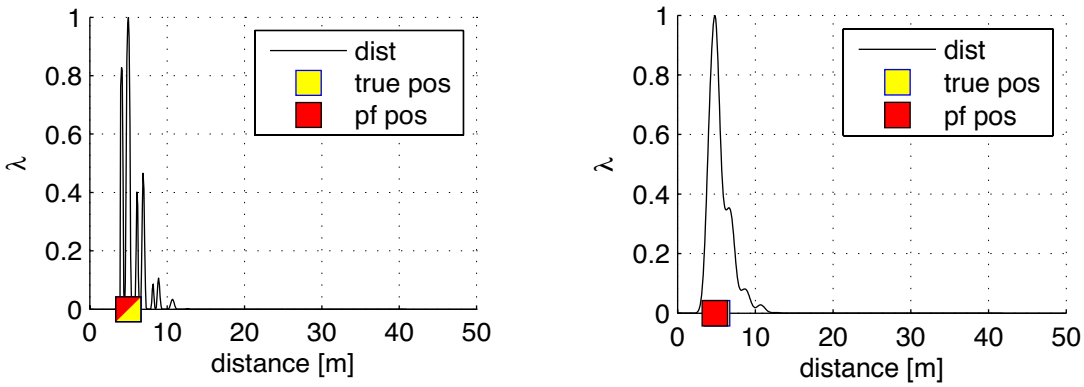

(a) (b)

Figure 6.25: HF multi-sensor tracking (simulation — convolution of the state PDF): (a) shows an example of the state PDF just after a measurement is incorporated. (b) shows the state PDF after a period without observations. The convolution with the Gaussian kernel smoothes the PDF.

distances. It can be seen that these likelihood functions are spread in the area where the true position of the target is located (indicated by the black line). The product of the three likelihood functions is also not heavily peaked (see figure 6.27 (top)), nevertheless the ML is very close to the true position. This is the reason for the deviation of the estimated state for larger distances.

The updated state PDF is obtained as the product of the previous state with the observation likelihood functions. As the observation likelihood functions are very broad in the area of the peak of the state PDF, only small changes to the peak of the previous state PDF are introduced. This results in a *lag*. In figure 6.24, the deviation first increases and then decreases. This corresponds to the observation likelihoods getting initially more spread, and then becoming more and more peaked again. To summarise, a spread observation likelihood function can lead to a lag, while a peaked one gives a more responsive behaviour.

The simulations in 1 dimension showed that the HF might be well suited to the tracking problem as formulated in this thesis. It has the advantage over the KF in that it can use the multi-modal distributions, and does not suffer from the PF problem where the particles may end up in the wrong part of the PDF if the process noise is too high. The drawback of the HF is the potential lag between the true and the estimated position.

Example 7 — Multi-sensor tracking (experiment)

From the simulation we saw that the HF seems to be well suited for the application with multi-modal observation functions. This section investigates the filter operation when work-

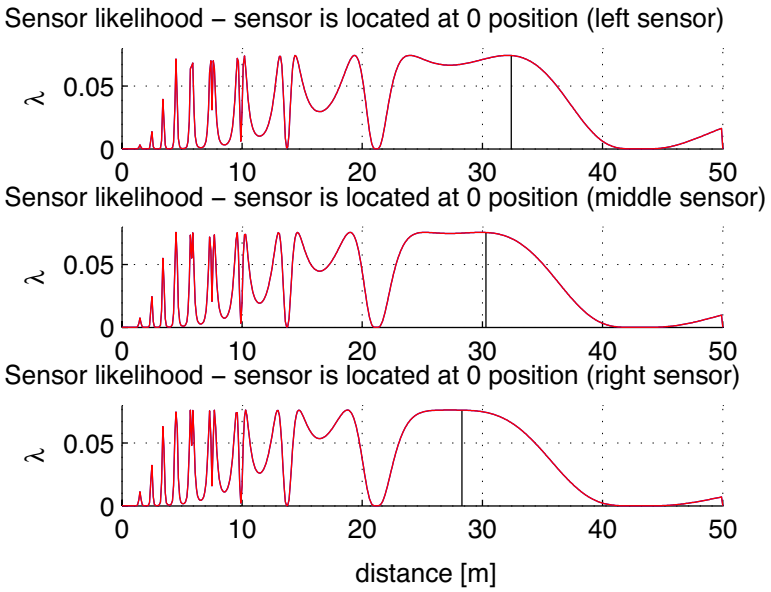

Figure 6.26: HF multi-sensor tracking (simulation — observation likelihoods): Three measurement likelihood functions as obtained from the three sensors for large distances. For a wide range of distances between the transmitter and the receiver the multi-modal distribution is very often wide spread. The vertical black line indicates the true position of the target with respect to the sensors' position.

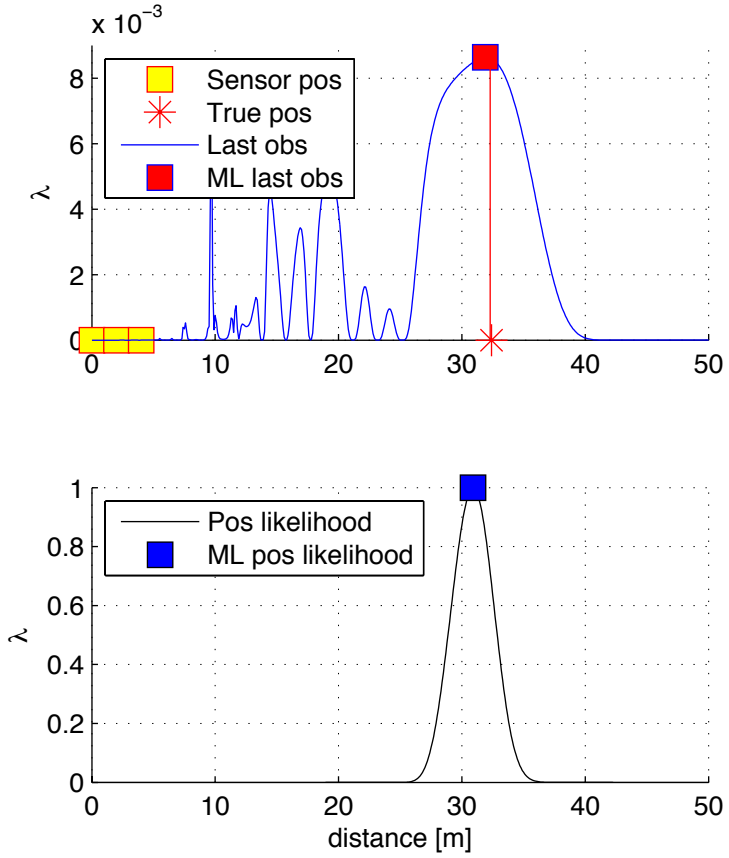

Figure 6.27: HF multi-sensor tracking (simulation — state PDF): (top) The product of the three measurement likelihood function as shown in figure 6.26. Also shown are the positions of the three sensors (yellow square), the true position of the target (red star) and the ML of the product of the three observation functions. Note how the ML estimate of the product of the observation functions does not coincide with the true position, but is a little bit to the left. (bottom) shows the state PDF after incorporating the measurements. The ML estimate of the state PDF lags even more behind the true position.

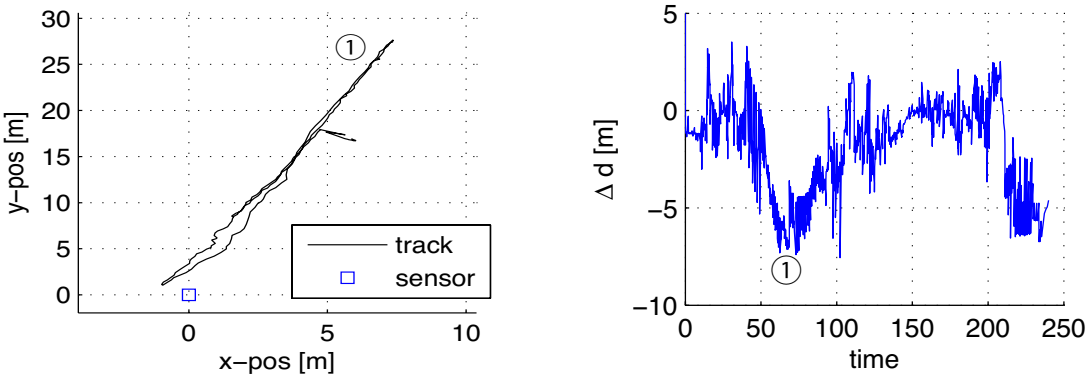

(a) (b)

Figure 6.28: Tracking with the HF — multi-sensor tracking (experiment): The deviation between the estimated and the true state is small, but the same systematic deviation from the simulations occurs again.

ing with experimental data. In figure 6.28(b) the deviation between the true state and the estimated state is shown. For the whole experiment the deviation is small and again a systematic deviation occurs, indicated by the circled one. The explanation for the deviation here is the same as given in the previous simulation. It is due to the fact that the observation likelihoods are broad and not peaked at large distances.

The examples using the PF and the HF employed two different approaches to obtain the position of the target. The weighted mean presented in the PF examples appears to be less suitable for this task than the ML estimate shown in the previous HF examples. Nevertheless, in both cases the PDF appears to represent the state quite well.

6.4 2-dimensional tracking results

This section presents 2-dimensional tracking results using RF sensors. As mentioned in chapter 5, multiple sensors are needed to resolve the bearing ambiguities inherent to range-only sensors. Therefore, all results in this section are obtained using multi-sensor setups. Furthermore, the HF approach has been adopted to present the results. In the previous section it has been demonstrated that the KF is completely unsuitable for this problem, and that the PF might be an alternative to the HF.

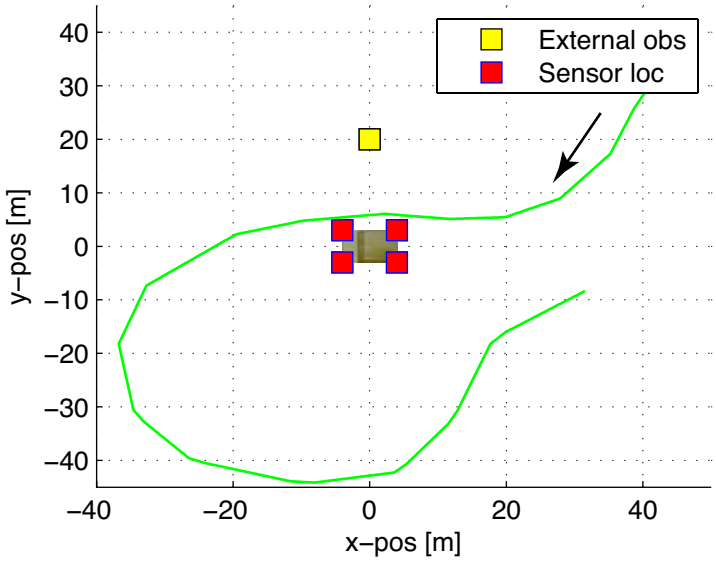

Figure 6.29: 2D simulation setup — The track used in the simulations is shown with the arrow indicating the direction along which the agent moves. In the centre the truck can be seen with the 4 sensors attached at its corners. Also shown is the location of an external observer that can provide additional external information.

6.4.1 Simulation

A simulation example shows 2-dimensional range only tracking. As in the previous 1-dimensional examples, the sensor model used is a two-ray model and the motion is modelled through the use of a convolution with a 2D Gaussian kernel. The directionality of the antenna pattern is neglected and an omnidirectional radiation pattern is assumed. The position estimate is obtained as the ML estimate of the 2-dimensional state PDF. As will be seen, the ML estimate of the state PDF will often be the correct location, though there are instances where this estimate will be wrong.

The setup for the simulation is shown in figure 6.29. In the middle of the figure the truck is shown with four RF sensors attached, one at each corner. Also shown is the track of the agent and an arrow indicating the direction along which the agent moved.

Observations are obtained from a sensor as soon as the agent is within the maximum sensing range of a sensor. This leads to the situation that sometimes only one sensor is observing, sometimes two, three or all four. Also, the generated measurements are corrupted by noise. The tracker is not initialised with a specific position, i.e. initialised close to the true position.

The most relevant effects are described and analysed fully in this section.

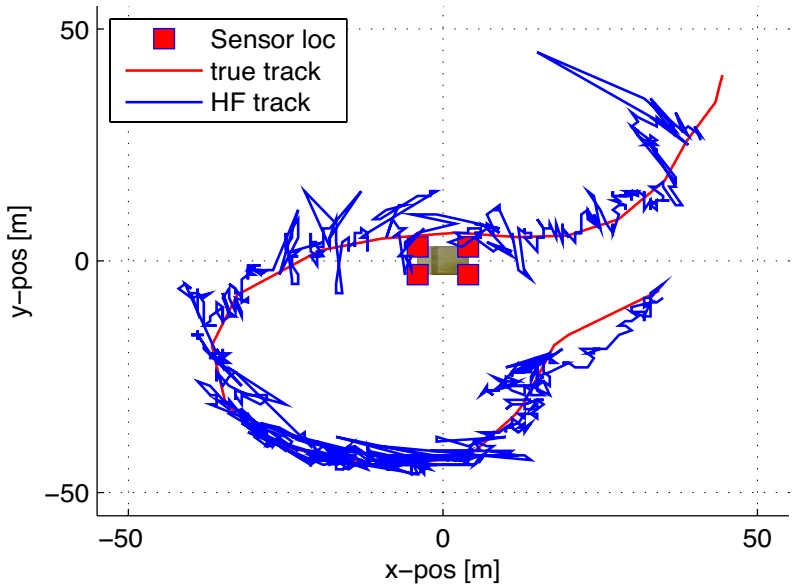

Figure 6.30: 2D simulation results (track) — The estimated track is shown in blue.

Results

Figure 6.30 shows an overview of the estimated track. As the tracker is not initialised with a specific position close to the true position, it takes a number of measurements until the estimated position is close to the true track. An interesting point to note is that most of the changes to the estimated track are tangential to the line connecting the truck with the true track. This is consistent with the ranging-error based failure modes explained in chapter 5. With these failures the position estimate moved in a tangential manner when the estimated range of one sensor was incorrect. The tangential movements occur more often than radial ones as the observation likelihood functions are better defined radially than tangentially.

Figure 6.31 depicts the deviations between the estimated and the true position. It can be seen that the tracker starts up with a large deviation, but soon this deviation decreases significantly and stays at a low level.

Five different situations are now analysed. In the corresponding figures, the PDF is shown on the left, while the measurement is shown on the right.

Startup — one sensor observing

After startup, only the sensor closest to the agent makes observations. A typical observation is shown in figure 6.32, with the state PDF as obtained after multiple observations from the same sensor.

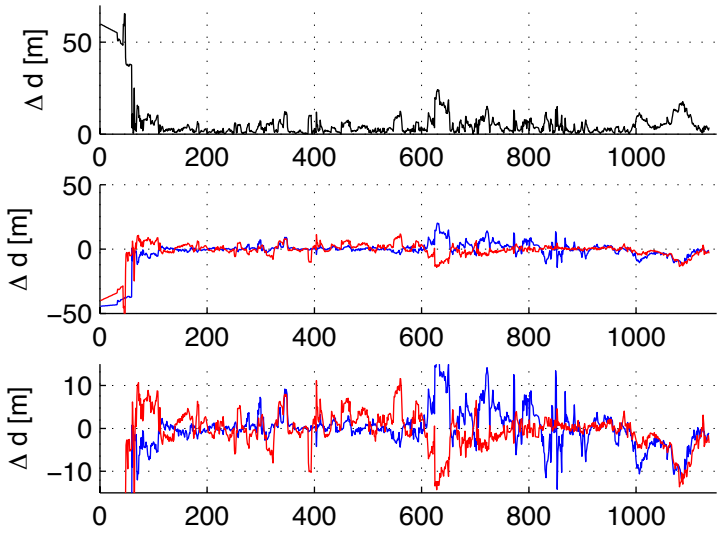

Figure 6.31: 2D simulation results (deviations) — (top) The overall distance deviation between estimated and true position. (middle) The deviations of x-position (blue) and y-position (red) separately and (bottom) as before, but y-axis magnified.

Two important points are well illustrated here. The first is the effect of discretisation. Ideally, both the measurement likelihood function and the state PDF would be absolutely rotation-symmetric. As the HF discretises both, some information is lost, leading to a state PDF as shown here where the probabilities for a given distance are not the same, although they should be. The second point to notice is that although the ML estimate is in the wrong position due to the discretisation issue, the measurement likelihood function and the state PDF are consistent with respect to the true position, i.e. in both cases the true position is in areas of high probability.

Startup — two sensors observing

As more sensors make observations simultaneously the measurement likelihood functions become more non-symmetrical, something that is desired as it is needed to reduce the number of hypotheses in the state PDF. The initial state PDF was rotation symmetric because only one sensor observed previously, but as soon as a second sensor also makes observations the combined measurement likelihood functions become axially symmetric. If these are now combined with the rotation-symmetric state PDF from the previous example, the number of hypotheses is collapsed to only two. Here again the effect of the discretisation and numerical calculations can be seen. If the state PDF and the measurement likelihood

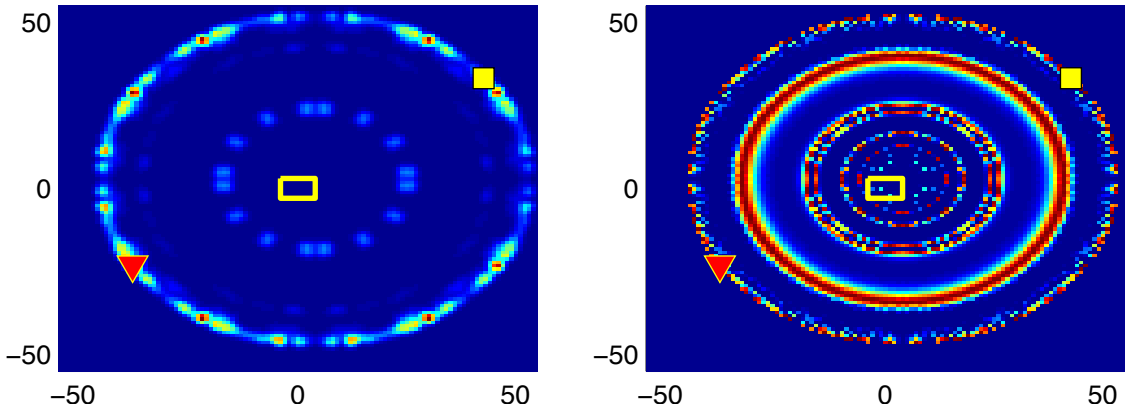

Figure 6.32: Startup — one sensor observing: As the measurement likelihood functions are rotation-symmetric, the resulting state PDF is too. The yellow □ shows the true position, the red ▽ the estimated position, and the yellow rectangle indicates the position of the sensors.

functions were described in a closed form the state PDF would be perfectly symmetrical and the ML estimate would return not one, but both locations as the ML estimate. Due to the numerical implementation the ML estimate returns one location, which unfortunately is not on the true location. Nevertheless, the state PDF shows high probability where the true location is. This is shown in figure 6.33.

Startup — three sensors observing

A third additional sensor now provides the information necessary to be able to resolve the last ambiguity. The combined measurement likelihood function from the three sensors is now non-symmetrical as the sensors are not aligned co-linearly. This combined measurement likelihood function, when fused with the previous axially-symmetric state PDF, removes the secondary hypothesis and resolves the last ambiguity. Now the true state and the estimated state are very close together. Figure 6.34 shows this situation.

Multiple hypotheses

As just shown, the use of multiple sensors observing at the same time is able to reduce the number of hypotheses to one only. This however is only a possibility, and is dependent on the alignment of the sensors, the measurements they make and the shape of the prior state PDF. Multiple hypotheses may still be present or emerge after the fusion of a combined measurement likelihood function. In figure 6.35 such a situation is shown. The prior state PDF exhibited one mode only and after the incorporation of the combined measurement likelihood function obtained from multiple sensors, three hypotheses are present. The situ-

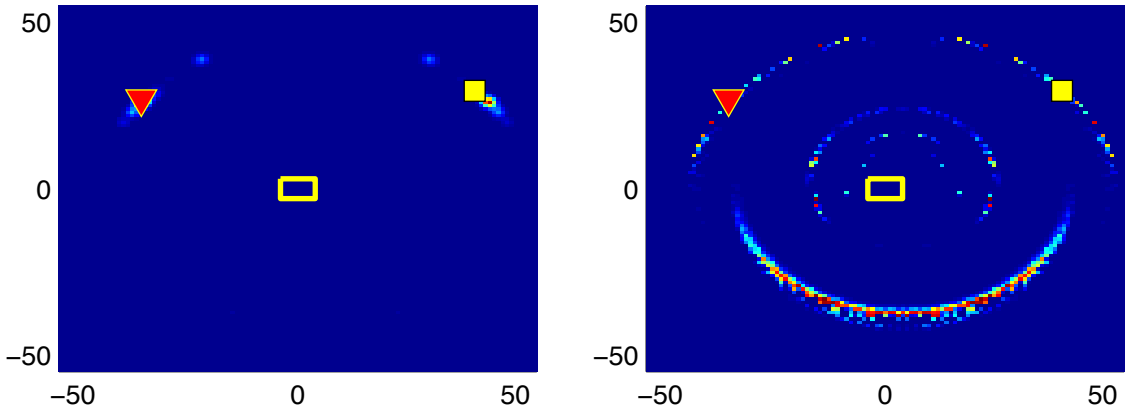

Figure 6.33: Startup — two sensors observing: As soon as a second sensor also makes observations the number of hypotheses is reduced. Note how the combined measurement likelihood function is now axially symmetric. The yellow □ shows the true position, the red ▽ the estimated position, and the yellow rectangle indicates the position of the sensors.

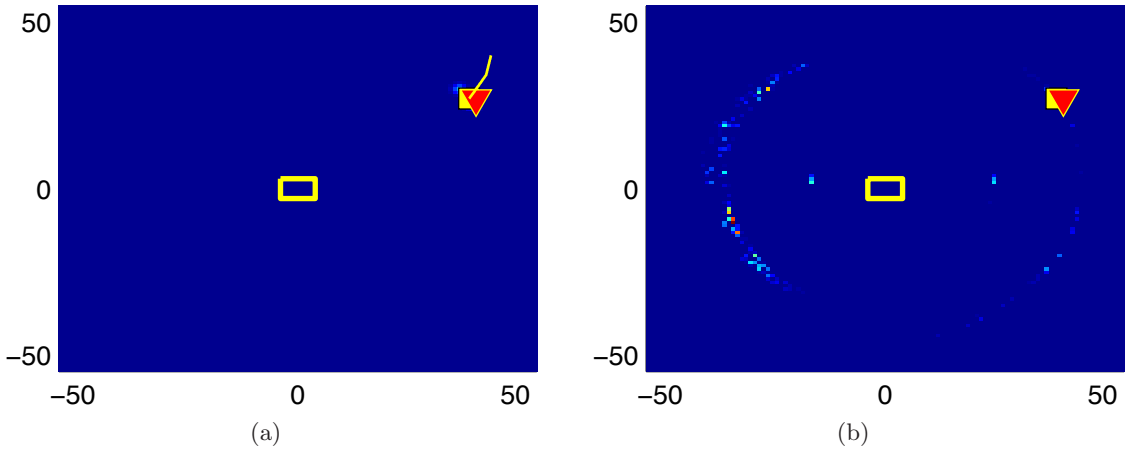

Figure 6.34: Startup — three sensors observing: Three sensors are able to collapse the number of hypotheses to one hypothesis only. The combined measurement likelihood function is not symmetric anymore. The yellow □ shows the true position, the red ▽ the estimated position, and the yellow rectangle indicates the position of the sensors.

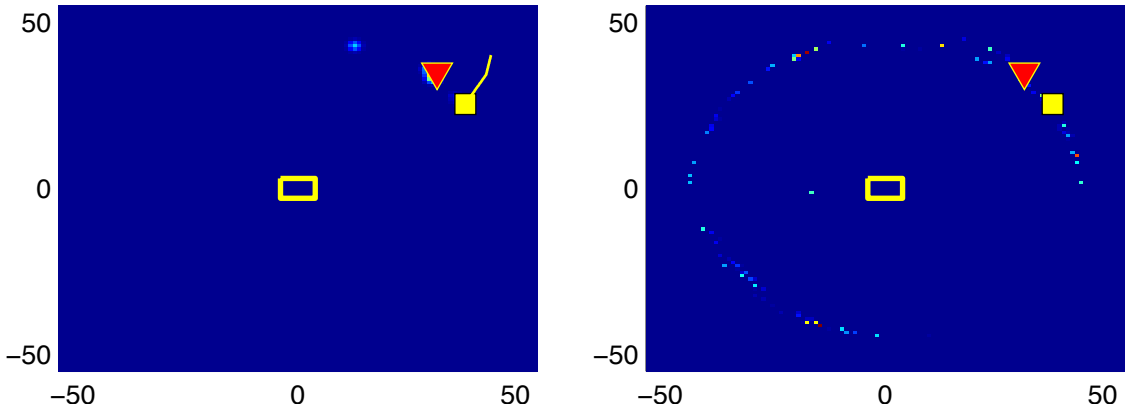

Figure 6.35: Multiple hypotheses: Although multiple sensors are observing the agent and the combined measurement likelihood function is not symmetric multiple hypotheses in the state PDF appear. Note that the hypothesis under the the true position is hidden by the yellow square, there is actually some probability mass at the correct position. The yellow □ shows the true position, the red ▽ the estimated position, and the yellow rectangle indicates the position of the sensors.

ation is further worsened by the fact that the ML estimate is not the true position anymore. In the simulations performed, the occurrence of multiple hypotheses with a possible shift of the ML estimate from the true position was temporally very limited.

Ranging error

In the previous chapter the effects of ranging errors were elaborated. An example as it occurred in the simulation is presented in figure 6.36. This example is obtained as the combination of four sensor measurements. A tangential shift occurs in the combined measurement likelihood function, but in this case it does not affect the state PDF as it is not a systematic effect, but occurs randomly with different magnitudes and direction.

The simulation result presented here shows the performance of 2-dimensional range-only tracking with RF sensors using multi-modal observation likelihood functions. Overall the behaviour of the tracker is as expected and exhibits the phenomena as described earlier in this thesis. Also, the deviations occurring are not too large. It was shown however, that sometimes the ML estimate does not provide the true position, and that in such situations the entire state PDF must be interpreted.

6.4.2 Experimental results

This subsection shows two 2-dimensional RF-RSSI range-only tracking examples obtained using experimental data from outdoor environments. The first example shows a four-fixed-

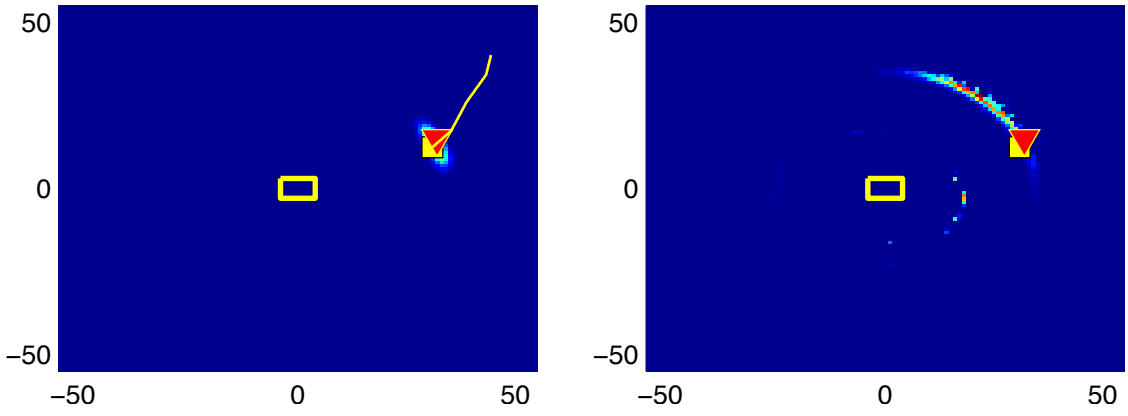

Figure 6.36: Ranging error — As the measurements are corrupted by noise, combined measurement likelihood functions that are tangentially shifted can occur. If this happens sporadically, the state PDF may not be greatly affected and the estimated and true position still coincide well. The yellow □ shows the true position, the red ▽ the estimated position, and the yellow rectangle indicates the position of the sensors.

sensor setup tracking one agent without the use of an external observer, while the second example shows how measurements obtained from an external observer can be incorporated to aid the localisation process. In these two examples the directional behaviour of the antenna is incorporated, and the sensor observations are made out of sequence. The reference position for the ground truth was obtained using Real Time Kinematic (RTK) Global Positioning System (GPS).

2D tracking without external observer

This example shows two dimensional tracking similar to the previous simulation. The setup for this experiment is shown in figure 6.37. Four sensors observe one agent walking along a straight line away from the sensors, and then back towards them.

The simulation showed that it is not reasonable to expect a smooth estimated track with range-only sensors, such as RF RSSI devices. Figure 6.38 shows the estimated positions obtained as the ML estimate of the complete state distribution. Once again, it becomes apparent that the ML estimate will often return an unsatisfactory estimate in such a localisation application. Nevertheless, it is worthwhile analyzing this figure, as it confirms some of the conclusions made previously in this thesis.

The estimated track deviates in the previously described tangential fashion, especially for larger distances. It is interesting to note that the estimated positions for very far distances are completely incorrect and do not follow the track at all. Instead, the estimated positions branch off tangentially. For medium and shorter ranges, the majority of the estimated

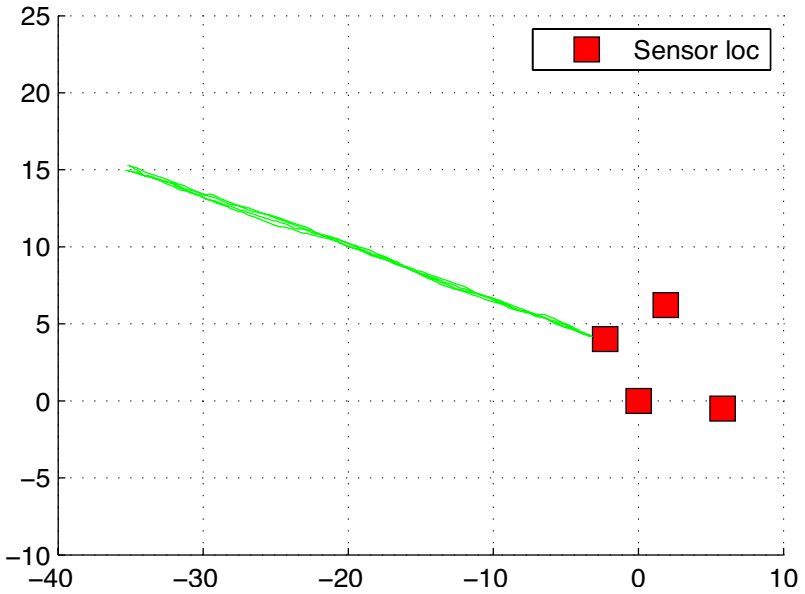

Figure 6.37: 2D experimental setup (1^{st} example) — The track used in this experiment is shown. In the centre the 4 stationary sensors can be seen.

positions are quite close to the ground truth. Furthermore, the arrangement of the four stationary sensors as used in this experiment is capable of determining the bearing to the agent. The majority of the ML position estimates are on the correct side of the stationary sensors, where the true track is located.

The deviations are shown in figure 6.39. For a significant part of the experiment the overall deviations are well below 10 m. Two spikes with increased error of more than 20 m occur. These points correspond to the areas of maximum separation between fixed sensors and agents, and these are the areas where the tangential deviations occur.

Five different typical situations as they occurred during the experiment are now discussed in detail. The figures show the combined measurement likelihood function on the right (labelled (b)), and the state PDF on the left (labelled (a)). The fixed sensors positions are indicated by the red squares, while the ground truth is shown using the Δ, and the estimated position using the ∇ symbol.

Close distance: For close distances the estimation works well. Figure 6.40 and figure 6.41 show two typical, but different measurement situations for these distances. Both measurement likelihood functions are correct for these ranges, but exhibit different characteristics. The measurement likelihood function in figure 6.40 is fairly compact. In contrast, the measurement likelihood function in figure 6.41 is more spread. The true position is located inside of high likelihood within this function, indicating that

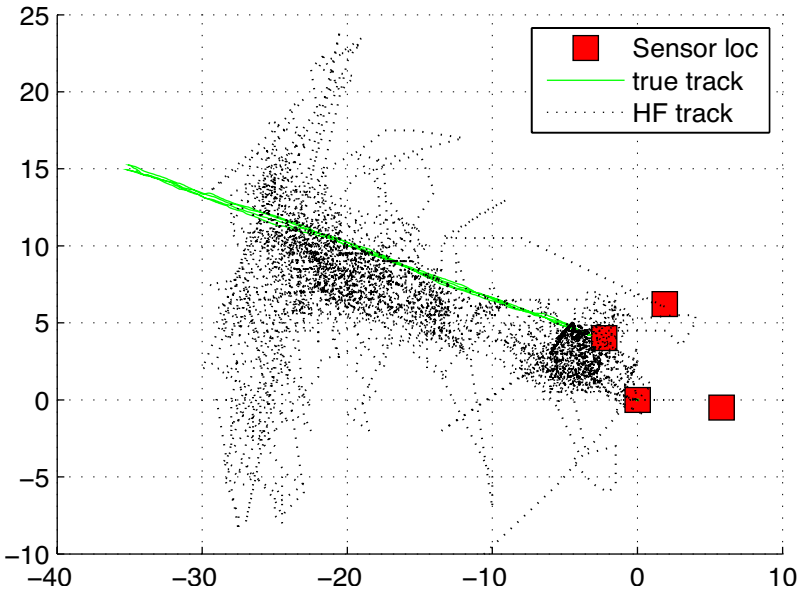

Figure 6.38: 2D experimental results (1^{st} example — track) — The estimated track is shown in black. It can be seen that for large distances the estimated position often moves tangentially to the true track, which is radial with respect to the fixed sensors.

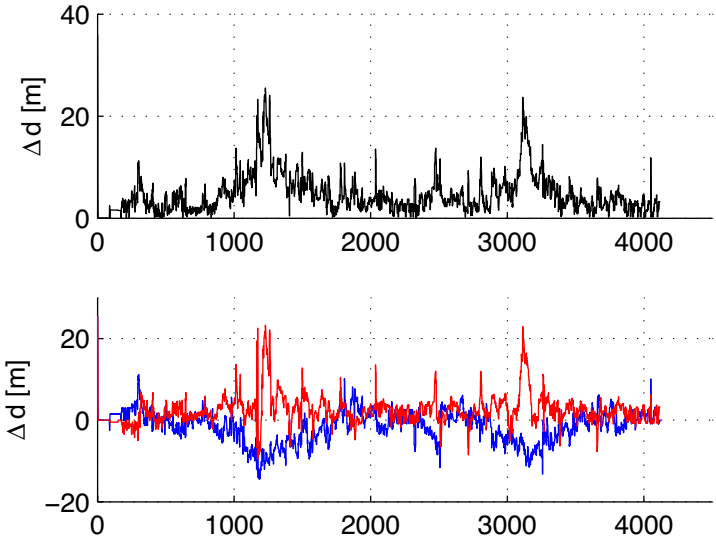

Figure 6.39: 2D experimental results (1^{st} example — deviations) — (top) The overall distance deviation between estimated and true position. (bottom) The deviations of x-position (blue) and y-position (red) separately.

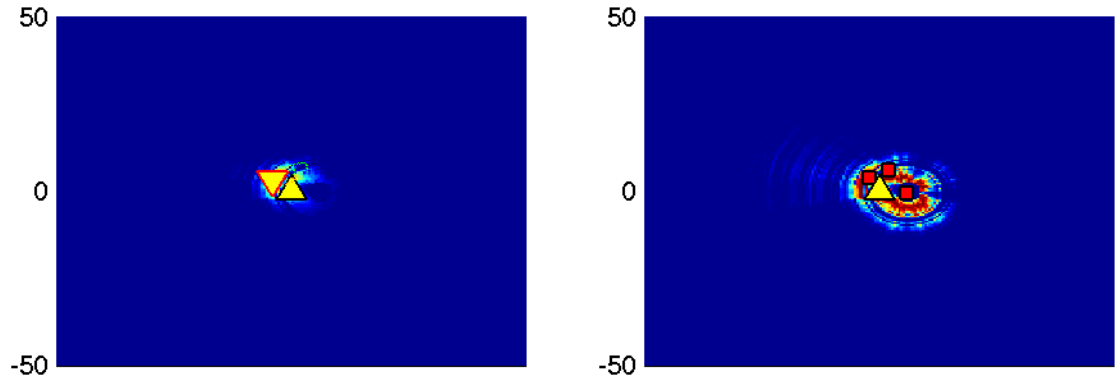

(a) (b)

Figure 6.40: 2D experimental results (1^{st} example — situation one at close distance): The agent is very close to the fixed sensors and the obtained observation likelihood function is very *compact*. Estimated and true position are close together.

the observation might have been correct. This is difficult to assess for such multi-modal behaviour. In the same figure the directionality of the radiation pattern can be seen fairly well.

Medium distance: For medium separation distances the tracking continued to work well. A typical observation likelihood function is presented in figure 6.42. The likelihood mass is concentrated in the area where the true position is and the state PDF represents the situation in a realistic manner.

Incorrect observation: Occasionally, incorrect observations occur due to the noisy properties of the RF channel. Figure 6.43 shows an example of this situation. In this case the measurement likelihood function shows increased likelihood mass in areas other than the true location. If such measurements occur at a very low rate compared to correct measurements, then their influence on the state PDF may be negligible. In the case depicted here, the wrong measurement was a singular event and did not heavily influence the state PDF. Although the probability mass of the state PDF is not as concentrated as in the previous cases, it still is located in the correct region, and in this case the ML estimated position is still close to the ground truth.

Far distance: Large separation distances were identified as problematic in the 1-dimensional examples and the 2-dimensional simulation. The 2-dimensional experimental result confirmed this, with the estimated track deviating tangentially from the true track for large distances. This example clearly presents the multi-modal nature of the RF sensors used, which must be considered. When comparing the observation likelihood function occurring for this large separation distance, shown in figure 6.44, with the

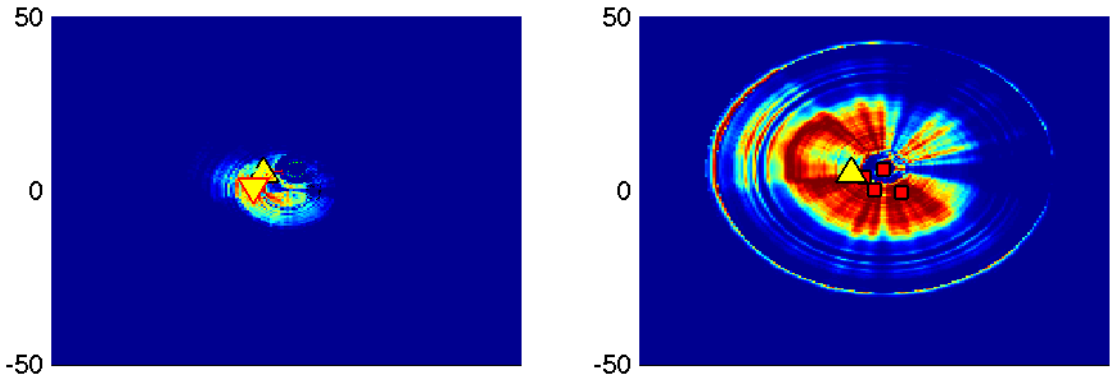

(a) (b)

Figure 6.41: 2D experimental results (1^{st} example — situation two at close distance): The agent is still close to the fixed sensors but the obtained observation likelihood function is now spread, with significant *mass* spread to the sides. Estimated and true position are again close together.

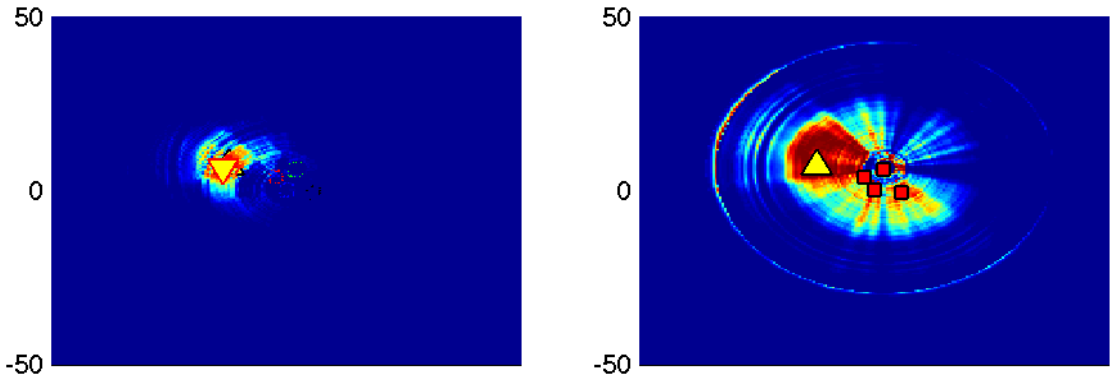

(a) (b)

Figure 6.42: 2D experimental results (1^{st} example — situation at medium distance): The agent has moved further from the fixed sensors, and the observation likelihood function shows very high likelihood concentrated in the area where the true position is. Estimated position and true position are very close together (the symbols overlap).

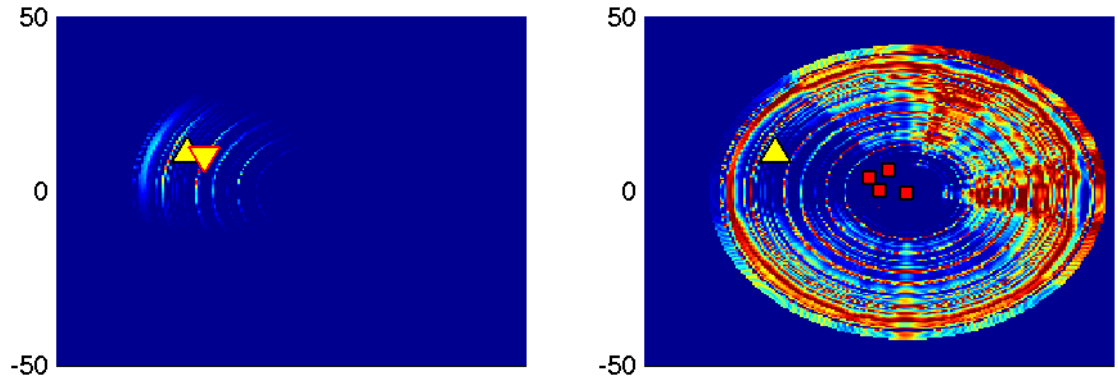

(a) (b)

Figure 6.43: 2D experimental results (1^{st} example — wrong observation): Again for a medium distance, but this time the measurement likelihood function does not correspond to the true location. This probably occurred because of a noise-corrupted measurement. The likelihood mass is spread in areas far from the true position. Nevertheless, the state PDF still represents the real situation well.

one for close distance, as shown in figure 6.41, it can be visually established that they are identical. So in summary, for two extremely different separation distances the same measurement value must have occurred. This aligns well with the theoretical considerations obtained through the two-ray model, and is in contrast to the nth power law, which does not allow for the same measurement value to occur at two different distances. This is a classic example of how a real RF sensor behaves.

This measurement is correct for the separation distance at which it was obtained. This can be seen from the fact that the ground truth is in an area of high likelihood, though it is a very concentrated (small) area around the true position.

When considering the development of the state PDF it can be seen that the incorporated measurement divides the previously singular hypothesis into two separate hypotheses. Nevertheless, the ML estimate is still close to the true position.

Far distance with hypothesis split: For very far distances, the estimated track showed a tangential deviation from the ground truth. A typical observation likelihood for this case is shown in figure 6.45. Although the measurement is correct, with high likelihood at the true position, a significant part of the area covered by the sensor measurement also has high likelihood mass. For this large separation distance the frequency of this type of measurement is high. Many measurement likelihood functions similar to this one will be obtained and incorporated into the state PDF. Ultimately this results in a split hypothesis, moving the estimate tangentially away from the true location. As this split is symmetrical, it is hard to predict to which side the ML estimate will

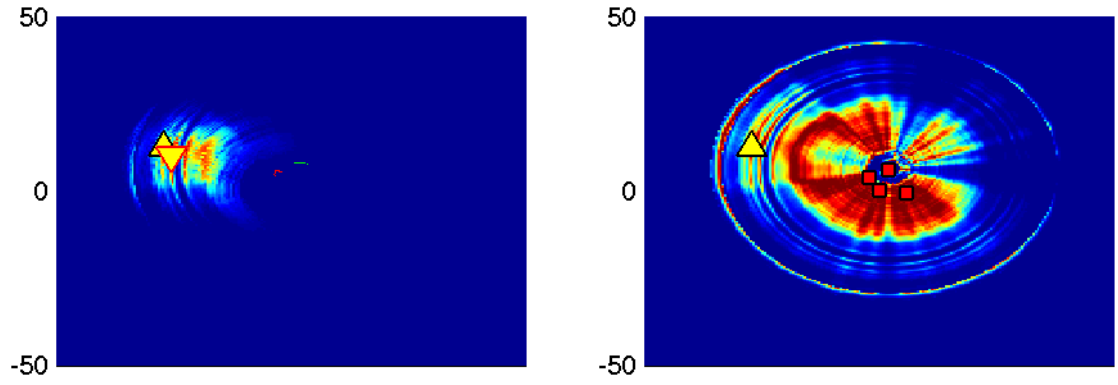

(a) (b)

Figure 6.44: 2D experimental results (1^{st} example — situation at far distance): The measurement likelihood still has likelihood mass at the true position, but this is now reduced. The true and the estimated position still coincide well.

move, and as the estimated track shows for this region, the ML estimate tends to jump between both newly developed hypotheses.

This experimental example showed that 2-dimensional RF RSSI range-only based localisation and tracking is possible. The results are promising, especially for shorter and medium separation distances, with a systematic deviation that has to be dealt with for larger separation distances. The issues detected in the simulation were confirmed to exist in the real application. It was also shown that for sporadically occurring incorrect measurements the tracking result is not significantly affected. Furthermore, the ML estimate to obtain the estimated position performed well, except for the case of large separation distances.

2D tracking with external observer

This example shows how an external observer can help the localisation process. The external observer provides additional RF-RSSI range measurements. In this example only the process of obtaining the external measurement and incorporating it is shown.

Figure 6.46 shows this process, with the measurement likelihood function in the left and the state PDF in the right column. The fixed sensor positions are indicated by the red squares, while the ground truth is shown using the Δ symbol, and the estimated position uses the ∇ symbol. The position of the external observer is marked with a $\square$. The circled numbers describe the sequence of occurrence:

1. Initially the estimated state PDF and the ML estimate represent the true location with a deviation.

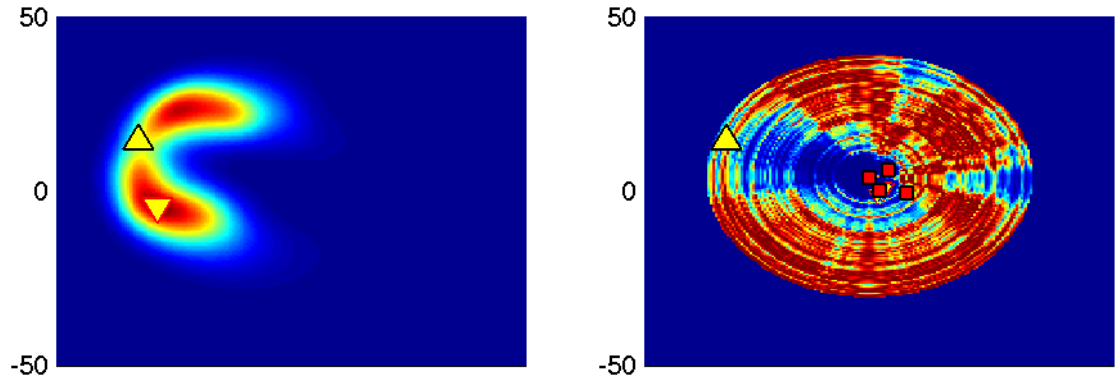

(a) (b)

Figure 6.45: 2D experimental results (1^{st} example — situation at very far distance): In this case for far distances, the measurement likelihood still has some likelihood mass at the true position. But in this case two hypotheses have formed, now misrepresenting the true occurring situation.

2. A subsequent measurement occurs by one of the internal sensors. Note the peaked characteristic for such internal measurements.

3. The measurement, after incorporation, moves the previous hypothesis even further away from the true location.

4. After some time without observations, the estimated location is still in the same position. The convolution of the state PDF with a Gaussian kernel to incorporate the possible motion of the agent has smoothed the state PDF.

5. The external observer makes a measurement of the agent. As the position of the external observer is known only with uncertainty, the measurement likelihood function is obtained as the convolution of the measurement likelihood function as obtained by the external observer with a Gaussian kernel depicting the uncertainty about the position of the external observer. This convolution smoothes the likelihood function significantly.

6. After incorporation of the measurement from the external observer, the state PDF and the ML estimate have moved closer to the ground truth location. Also note that in this case the resulting state PDF is still smooth and not peaked as it is after incorporation of an internal measurement.

In this example the use of external range information was shown. The quality of this information depends significantly on the quality of the position information of the external

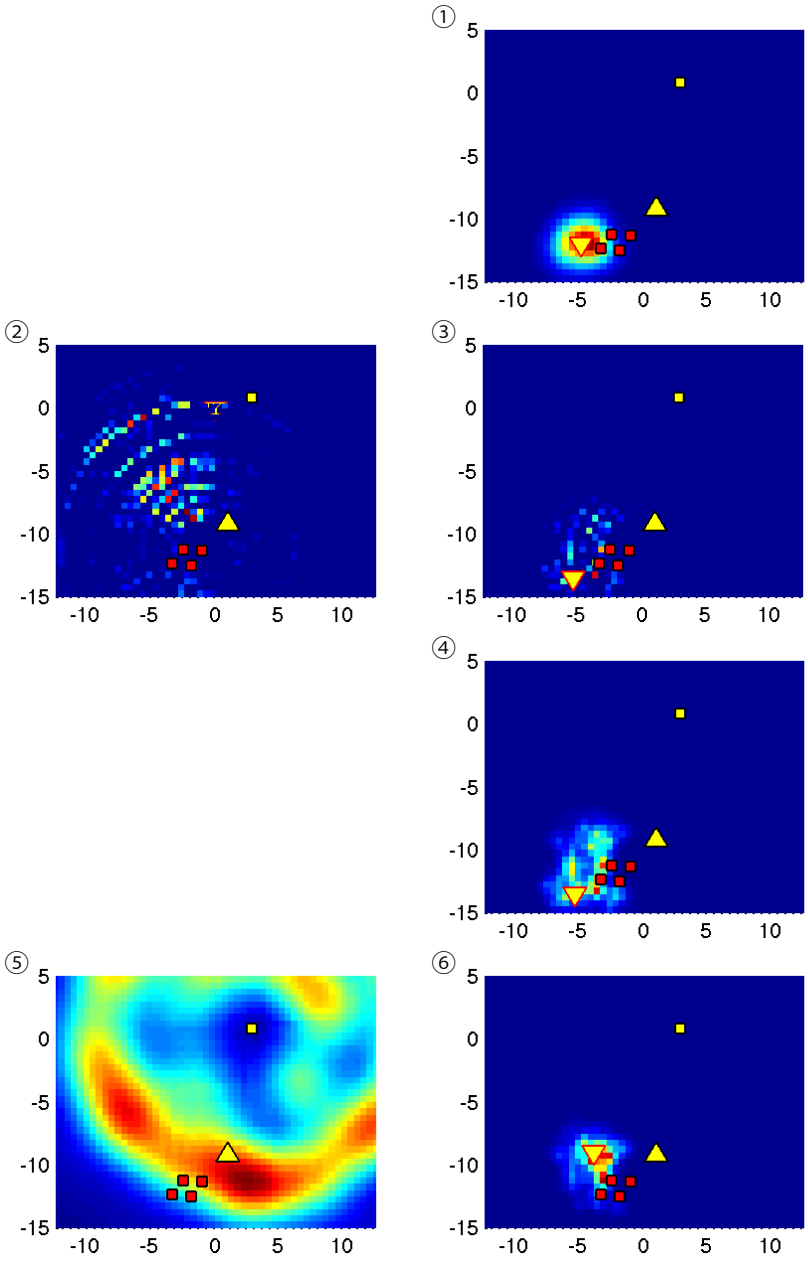

Figure 6.46: 2D experimental results (2^{nd} example) — The left column shows the observation likelihood functions, and the right column the development of the state PDFs. The circled numbers indicate the order of occurrence. Refer to the text for the explanation.

observer, as shown in chapter 5. Nevertheless, this information was shown to be useful for correcting the estimated position as with any other measurement.

6.5 Summary

This chapter showed 1-dimensional and 2-dimensional RF RSSI based range-only tracking and localisation performance with both simulation and experiments. The use of a probabilistic filtering technique capable of handling the multi-modal nature of RF measurements is absolutely necessary as the examples using the KF showed. The PF and HF in contrast are capable of handling such measurements and also of representing the corresponding, often multi-modal and non-Gaussian, state distributions.

The *noise susceptible* nature of range-only based localisation with its effects could be seen in the simulations and also in the experiments.

A systematic deviation for large separation distances was found and has to be considered. This deviation is due to the nature of the measurement likelihood functions for large distances which allows for large changes in distance for small measurement variations. A way to deal with this issue would be to use another geometrical arrangement (large baseline) of the fixed sensors. Nevertheless, for the targeted application this is not an option as the sensors cannot be mounted at large baselines. This issue could be potentially alleviated through the use of external information obtained from other agents that have good knowledge of their own position.

The results obtained experimentally encourage further work, especially in a real environment, i.e. with the sensors attached to trucks in a mining environment. This is important as the radiation pattern will be heavily influenced by the mounting location and the outdoor environment.