Information theoretic framework for the optimization of UWB localization systems

Information Theoretic Framework for the Optimization of UWB Localization Systems

Anthony Schenck*†‡, Edwin Walsh*†, Jonas Reijniers*, Ted Ooijevaar†, Risang Yudanto†, Erik Hostens†, Walter Daems*†, Jan Steckel*†

*FTI CoSys Lab, University of Antwerp, Belgium
†Flanders Make Strategic Research Centre, Belgium
‡anthony.schenck@uantwerpen.be

***Abstract*—In this paper we will propose a method to evaluate the performance of a certain ultra-wide band fixed anchor configuration in complex indoor environments by making use of the mutual information as the performance metric. Furthermore we will introduce an incremental algorithm that will determine an optimized anchor configuration for complex indoor environments. By making use of heuristics we are able to ensure that the time required to complete the algorithm is feasible on commercial grade computers, even for large-scale floor plans.**

***Index Terms*—Information Theory, Ultra Wide Band Ranging, Indoor Positioning Systems**

I. Introduction

Ultra wide band (UWB) ranging has become a popular technology used for indoor localization [1]–[3], due to its precision, low power consumption and its wide availability with products from DecaWave and UbiSense among others. Due to the fact that the data provided by these sensors is in the form of ranges, they can be easily transformed into a position using trilateration. The placement of the fixed anchors of a UWB-system has a significant impact on the achievable accuracy of the localization system [4]. Many different techniques have been proposed to determine an optimal placement of fixed anchors. Some have employed analytical methods such as minimizing the mean square error (MSE) by determining the optimal distance between anchors [2]. Others have used geometric approaches in order to minimize the geometrical dilution of precision (GDOP) between the anchors [4]. Genetic algorithms (GAs) have also been used to determine an optimal anchor placement configuration [5].

In this paper, we will approach this optimization problem by making use of the amount of mutual information as the performance metric [6], [7] using a heuristic algorithm for the placement of the fixed anchors. This work is based on the mutual information model presented in [8] and which has been applied for RSS-based localization in [9]. In this paper we adapt the previously introduced model to a UWB-based localization system. As we will demonstrate further in this paper, there is a need to optimize the computational algorithm for large-scale optimization problems as we target industrial applications with large floor plans and a high number of fixed anchors. For the optimization we resort to heuristic optimization techniques instead of meta-heuristics, as these general techniques tend to have difficulties in efficiently solving these large-scale non-linear problems with tight constraints [10], as found in UWB anchor placement.

The rest of this paper is organized as follows: In Section II we describe the mathematical model of UWB ranging measurements using a probabilistic approach. Section III explains the real world measurements we performed to tune our simulation parameters. In Section IV we will expand the likelihood model of Section II to information theory. Section V explains the algorithm we used to optimize the anchor placement as well as the heuristics we have used to speed up the calculations. Lastly, in Sections VI and VII we will present the results of our optimization algorithms and our conclusions.

This work has been carried out within the framework of Flanders Make's ICON project 'LOCATOR' (LOCalization system for Accurate Tracking and navigation for autonomous OpeRation, HBC.2016.0469) funded by the agency Flanders Innovation & Entrepreneurship (VLAIO) and Flanders Make. Flanders Make is the Flemish strategic research center for the manufacturing industry.

II. Indoor localization with UWB ranging

In order to evaluate the positioning of a set of anchors in a certain environment, we start from a probabilistic approach to localization: we calculate the posterior probability distribution of the sensor position $\vec{p}$. To achieve this, we generate measurement ranges between a certain point in the environment where a mobile node can be situated, and one or more anchors. Figure 1 demonstrates one of these measurements. Anchors A_1, A_2, A_3, A_5 and A_6 determine their range with regard to the measurement location $\vec{p}$. The range measurement of anchor A_4 is blocked by an obstacle and provides no data for the measurement. The mobile node's position in the environment can be defined by:

$$\vec{p} = \begin{bmatrix} x_p & y_p & z_p \end{bmatrix}^T \tag{1}$$

The anchor locations can be defined in the same way:

$$\vec{a}_i = \begin{bmatrix} x_i & y_i & z_i \end{bmatrix}^T \tag{2}$$

A range measurement between a set of anchors and a position is defined by the range between the position and the anchors ($\vec{r}_i$) with the addition of noise (η) and a certain bias (δ):

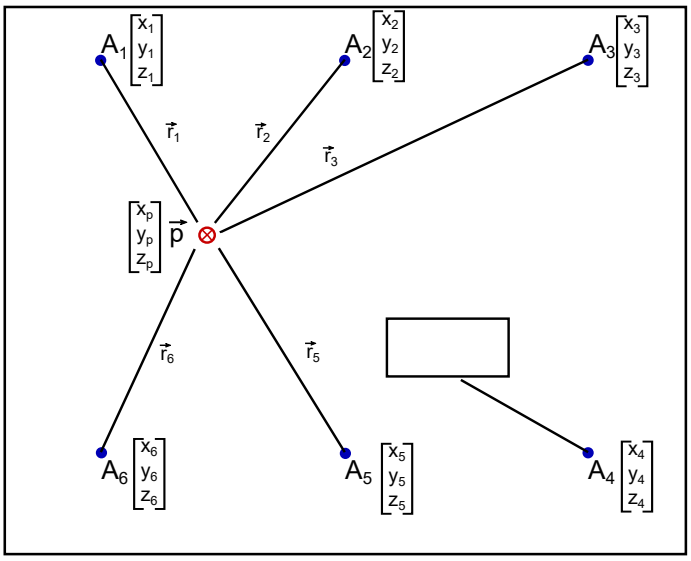

Fig. 1. Example of a ranging measurement. $A_1, ..., A_6$ are anchor positions, $\vec{p}$ is the location where the measurement is being taken. A_1, A_2, A_3, A_5, A_6 have line-of-sight and contribute to the result with there respective ranges r_i, while the range measurement of A_4 is being blocked by an obstacle.

$$\vec{m} = \begin{bmatrix} \vec{r}_1 \\ \vec{r}_2 \\ \vdots \\ \vec{r}_n \end{bmatrix} + \begin{bmatrix} \vec{\eta}_1 \\ \vec{\eta}_2 \\ \vdots \\ \vec{\eta}_n \end{bmatrix} + \begin{bmatrix} \vec{\delta}_1 \\ \vec{\delta}_2 \\ \vdots \\ \vec{\delta}_n \end{bmatrix} \tag{3}$$

with

$$\vec{r}_i = \|\vec{a}_i - \vec{p}\|_2 \tag{4}$$

and

$$\eta \in \mathcal{N}[0, \sigma^2] \tag{5}$$

We select the noise from a normal distribution with zero-mean because we assume the bias of a range measurement to be known and stable for relatively long time periods and can therefore be removed for our purposes. In Section III we will elaborate on this assumption. Based on this range measurement, we define the likelihood that the generated ranges $\vec{m}$ are measured at a certain position $\vec{p}_i$:

$$\mathcal{L}(\vec{m}|\vec{p}_i) = \Gamma \cdot \exp\left[-\frac{1}{2}(\vec{m} - T_{\vec{p}_i})\Sigma_{\vec{p}_i}^{-1}(\vec{m} - T_{\vec{p}_i})^T\right], \tag{6}$$

with Γ the Gaussian normalization factor:

$$\Gamma = (2\pi)^{-l/2}|\Sigma_{\vec{p}_i}|^{-1/2} \tag{7}$$

l the number of anchors, $T_{\vec{p}_i}$ the calculated, noise-free expected ranges to the anchors from position $\vec{p}_i$ and $\Sigma_{\vec{p}_i}$ a constant covariance matrix. In the computational implementations we have foreseen that the covariance matrix Σ can be a function of the pose $\vec{p}$, which allows to model parts of the environment with a lower signal-to-noise ratio (e.g. due to the presence of electromagnetic interference from rotating machinery). Figure 2 shows the likelihood distribution for a single position in the center of a room with one, two, three and four anchors respectively, providing the measurement ranges. For each position, the entirety of the discretized environment has to be evaluated with the generated ranges. In our simulation model, we have to discretize the environment using a selectable grid size. This gives us a series of 2D-voxels that correspond to the mobile node positions in the environment.

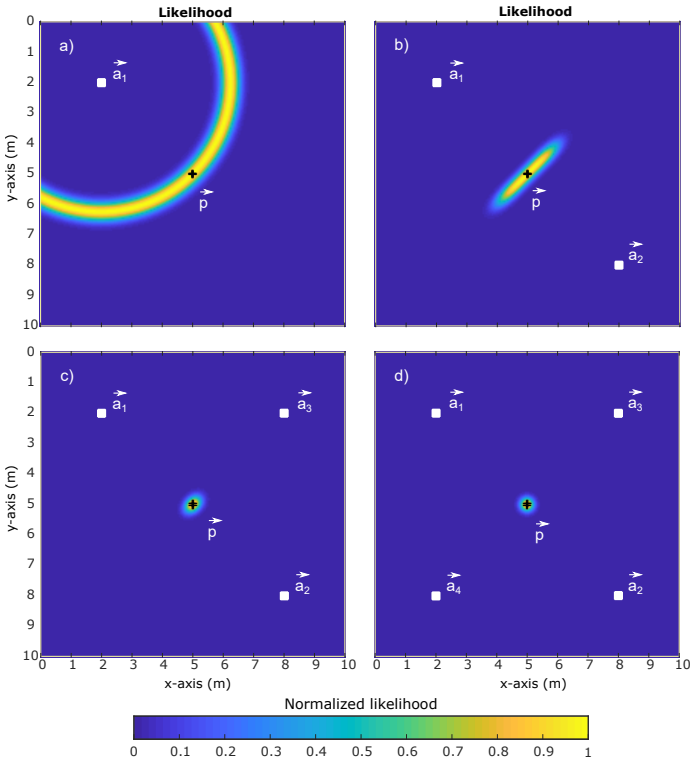

Fig. 2. Influence of anchor configurations on the likelihood distribution for a single measurement location. The white squares represent the anchors while the black cross is the measurement location.

Using the likelihood distribution of (6), we can calculate the probability density function (PDF) of a certain 2D-voxel being occupied by the mobile node as:

$$\mathcal{P}(\vec{p}_k|\vec{m}) = \frac{\mathcal{L}(\vec{m}|\vec{p}_k) \cdot \mathcal{P}(\vec{p}_k)}{\sum_{\vec{p}_k} \mathcal{L}(\vec{m}|\vec{p}_k) \cdot \mathcal{P}(\vec{p}_k)}, \tag{8}$$

with $\mathcal{P}(\vec{p}_k)$ being the prior distribution for the mobile node. Using this prior distribution, we introduce a-priori information about the mobile node position such as sections that cannot be occupied by the mobile node (due to obstacles) or the boundaries of the environment. We can calculate the likelihood model with different levels of a-priori information. The first level is to use no a-priori information and only constrain the surface area of the environment. The second level is to only make use of the outline of the floor plan by means of a mask. The third level can expand this mask to also contain inaccessible areas. The fourth level can make full use of the knowledge of the floor plan and determine LOS conditions in

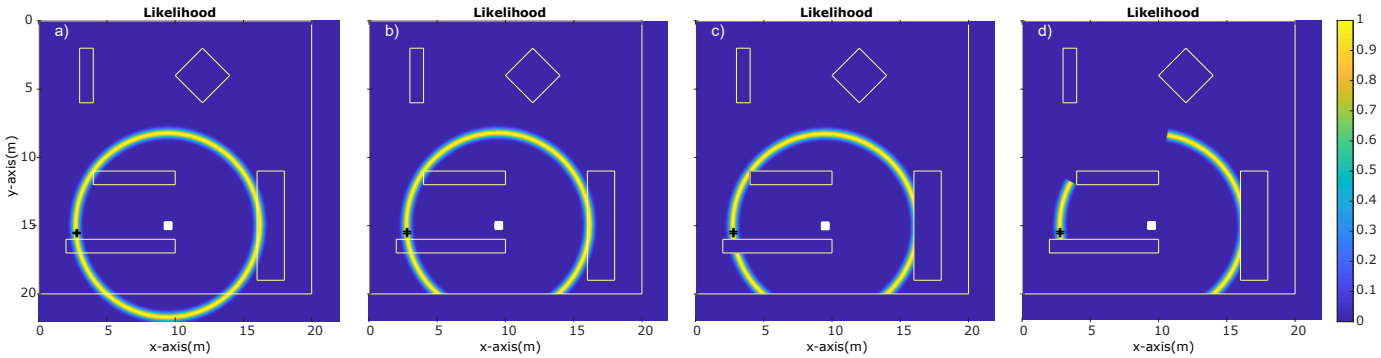

Fig. 3. Difference in the amount of a priori information used. (a) shows the likelihood distribution when no information about the environment is used in the measurements. In (b) the measurements are limited by the outline of the floor plan. (c) adds additional floor plan information by leaving out inaccessible areas within the floor plan. Finally in (d) the measurements are only generated for 2D-voxels that have LOS conditions.

advance. In this situation, when a position has NLOS of one or more anchors, the likelihood for this position increases. This happens because we then know the mobile node can not be at a position that does have coverage of these anchors.

Figure 3 demonstrates the influence of using different types of a-priori information on the posterior probability distributions. Figure 3 (a) shows the situation where we do not use any a-priori information. The likelihood distribution is also calculated outside the floor plan, which is of no interest to us, increasing the time required for the calculations. In figure 3 (b) the outline of the floor plan is used as a mask to limit the likelihood distribution to our region of interest. Figure 3 (c) has additional masking information about areas inside the floor plan that are inaccessible, for instance a warehouse rack or a wall. Lastly, figure 3 (d) makes full use of the knowledge of the environment and only calculates the likelihood distribution for those 2D-voxels that have LOS with the anchor. In the remainder of the paper, we use situation c) for all calculations. This way, we do not increase the likelihood at a certain mobile node position due to certain anchors having NLOS of that position.

III. Measurements to tune the model

The UWB anchor placement optimization relies heavily on the sensor model introduced in Section II. Therefore, it is crucial that this model adequately describes how the range measurements are affected in various scenarios (i.e. combinations of configuration and environmental conditions). There are a number of ranging errors arising from time delays in the hardware that are hard to physically describe in detail (e.g. the effect of the antenna design). For this reason, we have carried out numerous experiments to motivate the selected model. We will first discuss the UWB setup that has been used to gather the data, and the different scenarios that are at play in a typical application. All data was collected in static conditions: we assume that the UWB ranging measurements experience little influence of the Doppler effect from the movement of the mobile nodes at the typical low speeds in practical applications.

A. UWB setup

The UWB system used in this work is the DecaWave EVB1000 evaluation board, comprising of a DW1000 radio chip, a micro-controller from ST Microelectronics (STM32F105 ARM Cortex M3) and an external UWB antenna. The DecaWave DW1000 is a single-chip UWB transceiver based on the 802.15.4-2011 standard. It can operate in 6 frequency bands with a center frequency between 3.5 GHz to 6.5 GHz and provides data communication with data rates of 110 kbps, 850 kbps and 6.8 Mbps. The micro-controller communicates with the radio chip via a Serial Peripheral Interface (SPI). The evaluation board provides a micro USB port to communicate with a PC.

B. Scenarios

The DecaWave UWB system used a time of flight based two-way ranging (TWR) protocol to estimate the distance between two modules, the anchor and the mobile node. The effect of a wide range of scenarios has been analyzed, comprising of:

- line of sight (LOS) versus various forms of non-line of sight (NLOS) conditions,
- the selected center frequency (channel),
- the selected preamble length,
- range distances between anchor and mobile node,
- mobile node and anchor orientations,
- mobile node and anchor height.

In this paper, we have only cropped out the effect of LOS versus NLOS conditions, as the other effects are not yet incorporated in the design framework. The influence of the other scenarios will be evaluated and implemented into the framework in the future. The other variables were set to a center frequency of 6.5 GHz, a preamble length of 64 symbols, a range distance of 10 m, 0 degree anchor and mobile node orientations, and a height of 1.5 m.

The NLOS conditions were introduced by obstructing the transfer path using sheets made of plasterboard, metal, EMF absorbent material, and wood, in a first set of experiments (Figure 4, left). In a second experiment a concrete wall (Figure

Fig. 4. The test setup with NLOS conditions introduced by (left) a plasterboard sheet and (right) a concrete wall used to investigate their effect on the UWB ranging error.

4, right) was used to obstruct the line of sight between the anchor and the mobile node.

For each test 1000 samples were acquired. The range error for each sample has been determined by comparing the UWB based range with a ground truth range measurement using a mobile laser distance measurement device. The results obtained for the various NLOS conditions have been compared with the results obtained for the LOS condition.

C. Results

Analysis of a wide set of test scenarios revealed that the range errors are approximately normally distributed for both the LOS and the NLOS conditions as shown in Figure 5 (a) and (b). Contrary to what is sometimes found in the literature [1], we have observed little skewness or multi modality in the distributions of our test measurement results. The fact that we do not observe this skewness is because we performed our tests with a different measurement setup than in [1]. In [1] the combined range errors of sixty one different measurement locations are shown, resulting in the skewness due to NLOS measurements affecting the results. Whereas our results come from performing the same measurement 1000 times at the same location. Figure 5 (c) and (d) shows that the NLOS conditions mainly affect the bias, but have a minor effect on the standard deviation of the range error. The large deviation of the bias of the concrete wall is due to fact that this test was performed with an actual wall, while the other tests were performed with movable panels. The biases are positive for the NLOS conditions, but vary significantly depending on the specific NLOS condition.

From these results it is concluded that the NLOS conditions result in large bias error variations, that are hard to predict in real situations. Therefore, the bias errors for NLOS conditions will not be taken into account in the model.

IV. Information Theoretic Model

Now that we have established our likelihood model and have determined the parameters for our simulations, we can use them to determine the Shannon entropy in the environment. We use the amount of entropy in the environment as the base for our metric to determine if a measurement of a mobile

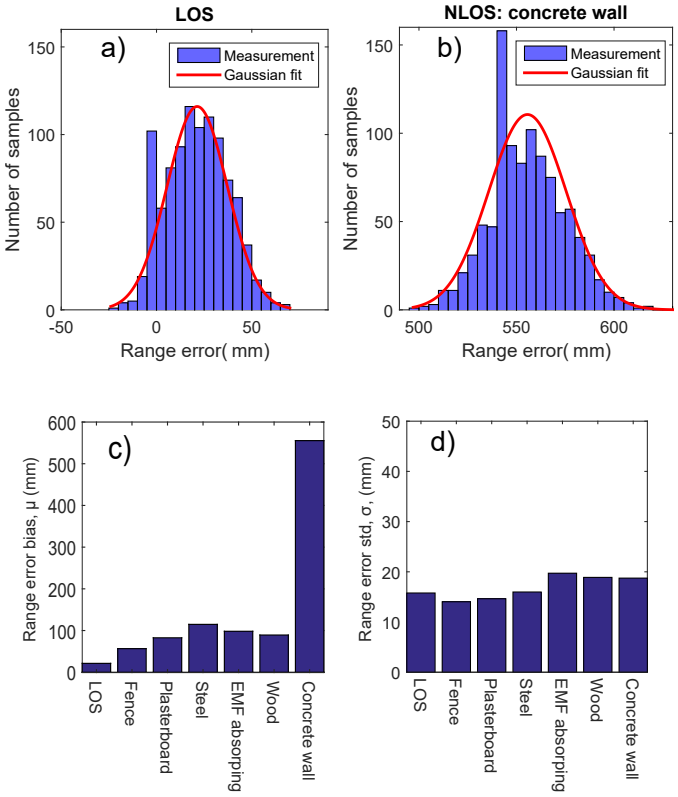

Fig. 5. (a) Range error distributions for LOS and (b) NLOS measurements at a range distance of 10 m showing the approximately normal distributions for both measurements. (c) Bias and (d) standard deviation of the range error at a range distance of 10 m, showing an increased positive bias for NLOS conditions, whereas the standard deviation is barely affected.

node on a certain 2D-voxel will produce a reliable localization. When a certain location $\vec{p}_k$ has a low amount of entropy, there exists little ambiguity on the position estimate resulting from a measurement in this location. This will result in a higher localization precision. We can calculate the entropy in each possible 2D-voxel in the environment $\vec{p}_k$ by using the result of (8):

$$H(\vec{p}_k|\vec{m}_k) = -\sum_{i=1}^{n} \mathcal{P}(\vec{p}_i|\vec{m}_k) \cdot \log_2 \left[\mathcal{P}(\vec{p}_i|\vec{m}_k)\right] \qquad (9)$$

with n the amount of 2D-voxels in the environment. Please note that for each value of i, the PDF is calculated for each 2D-voxel in the environment. This PDF depends on the individual noise realization that is drawn from the noise distribution. To take into account the full noise distribution instead of a single noise realization, we approximate the true value of the Shannon entropy using a Monte Carlo approximation. A new $\vec{m}_k$ is generated M times which is then averaged in order to obtain a single estimate for the Shannon entropy:

$$\langle H(\vec{p}_k)\rangle = \frac{1}{M} \cdot \sum_{j=1}^{M} H(\vec{p}_k|\vec{m}_{k,j}) \qquad (10)$$

In our model, we make use of the mutual information as the performance metric. This can be easily determined by subtracting the entropy from the maximum amount of information that could possibly be present in the environment:

$$\langle I(\vec{p}_k; \vec{m}_k) \rangle = \log_2(n) - \langle H(\vec{p}_k) \rangle \tag{11}$$

Figure 6 illustrates the Monte Carlo approximation to determine the amount of information at a certain 2D-voxel, as well as the effect of having more or less anchors in LOS. The measurement location M_1 has LOS to anchor A_3. The likelihood model will be evaluated M times for this point using the Monte Carlo approximation, which will result in an estimate for the mutual information for this position. The process is repeated for all other possible mobile node positions in the environment. The measurement location M_2 is in LOS of both A_2 and A_3, which leads to a higher amount of information due to the posterior distribution of this measurement being a lot more focused. Higher amounts of information implies that more anchors have LOS to the mobile node location. We can quantify the total amount of information in the environment as follows:

$$I_{tot} = \sum_{k=1}^{n} \langle I(\vec{p}_k) \rangle \tag{12}$$

V. Optimization of anchor placement using information theory

We have established a method to evaluate a certain anchor configuration by using mutual information. Now we want to use this approach to optimize an anchor configuration. To accomplish this we will use an incremental algorithm, which makes use of the overall distribution of mutual information. To evaluate a certain anchor configuration, we will take both the total amount of information as well as the minimum information amount in a floor plan into account. To be able to optimize the placement of anchors, we start with defining a configurable grid on which the anchors can be positioned. Next, we iterate over these positions with the first anchor, and eventually place it at the position where the total amount of information in the environment is maximal (see equation (12)). In the next step, we iterate over the remaining anchor positions, calculating the total amount of information of the two anchors combined. We again select the anchor that adds the highest amount of information. Once a certain maximum threshold value has been reached in a 2D-voxel, we will clip it to this value when calculating the total amount of information. We do this to limit the influence of small increases of information in many 2D-voxels, versus larger increases in fewer 2D-voxels. We keep repeating this process, adding one anchor at a time, until a predefined minimal information threshold has been reached in each mobile node location $\vec{p}_k$.

To determine these thresholds, we use a floor plan with the same dimensions and grid size as the floor plan for which we

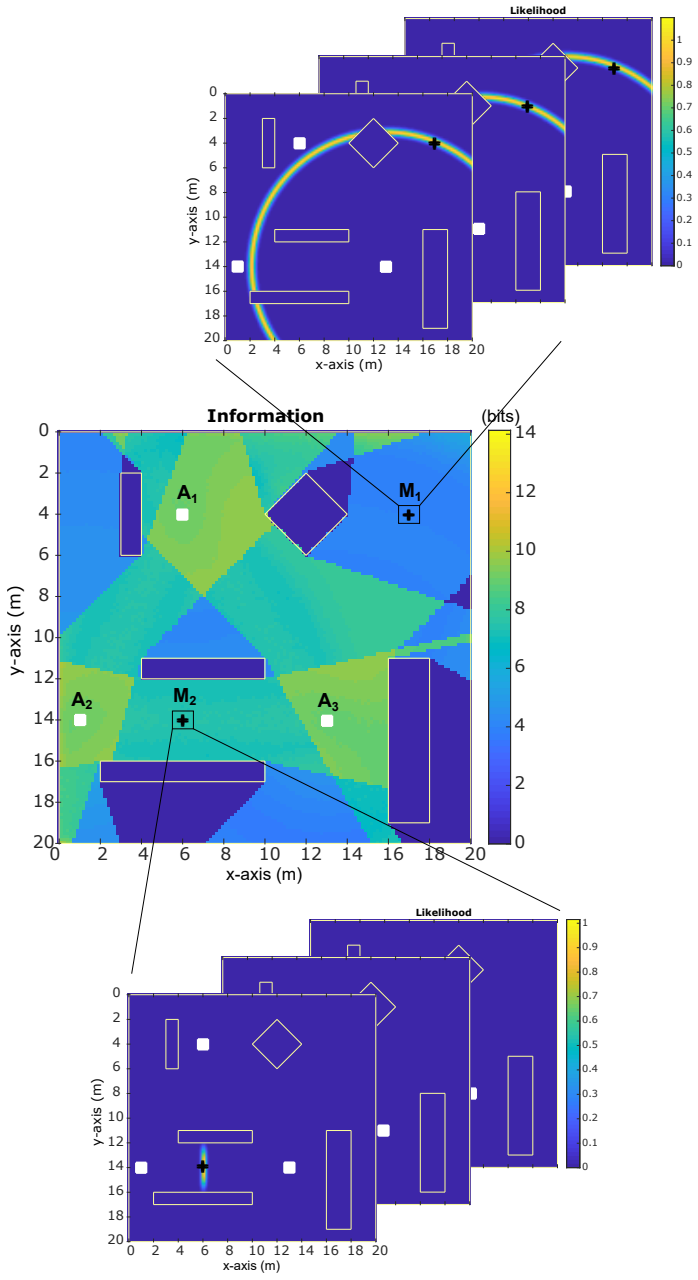

Fig. 6. For each 2D-voxel, the likelihood function is calculated multiple times and averaged in a Monte Carlo approximation. Measurement M_1 only has LOS of anchor A_3, resulting in a lower amount of information. This applies to all the 2D-voxels in the area that are colored light blue. Measurement M_2 is in LOS of both A_2 and A_3, resulting in a higher amount of information. The dark blue areas are 2D-voxels that are not in LOS of any of the anchors (or are inaccessible) resulting in zero information.

want to perform the optimization, without any obstacles. We place three anchors equidistantly from each other on a circle, with the 2D-voxel we want to evaluate as the center of this circle. Figure 7 (a) demonstrates this setup. Figure 7 (b) shows the same setup with five anchors, which we use to determine the information cut-off threshold. Once a 2D-voxel has reached this amount of information, we ignore any additional increase.

Due to the large amount of calculations needed to execute this algorithm, it is unfeasible to use it for anything but the smallest floor plans. Our final scope is to be able to optimize the anchor placement in environments of 100m$\times$200m covered by hundreds of anchors. The algorithm described above would take months to complete. We need to calculate the likelihood distribution of every 2D-voxel multiple times for the Monte Carlo simulation and for each likelihood distribution of a 2D-voxel, we need to evaluate the entire floor plan. This is unfeasible for our targeted real world industrial situations, using commercial grade computers or even modest computation clusters. For this reason we changed the above algorithm to reduce the number of calculations, by reducing the search space for the placement of the next anchor, and by reducing the amount of anchor locations we test.

We start by randomly selecting a 2D-voxel from 1% of all the 2D-voxels with the lowest amount of information. Next, we choose all the anchor positions that have LOS of this 2D-voxel. From these candidate anchor locations, we randomly select twenty. We evaluate these twenty anchor locations and add the one that adds the most amount of information to our anchor configuration, the same way as described in the previous algorithm. We then evaluate the entire environment with the current anchor configuration, and test if we have reached the predetermined minimal information threshold. If not, we select a new 2D-voxel from the lower 1% of information and repeat the selection process. We decided to select from the lowest 1% to add some amount of randomness to the algorithm while still optimizing for areas with low amounts of information. In our experiments we tested with different amounts of anchors to evaluate and found that with twenty we had a good trade-off between speed and performance. These values can change depending on the situation of the environment.

This optimized, greedy algorithm reduced the computation time by a factor of fifteen, with only a very small trade-off in the quality of the end result. We can see a slightly slower increase in the amount of information with each additional anchor. To further reduce the computation time, we added another optimization while calculating the mutual information. Instead of discretizing the floor plan with a small grid size, we make use of two different grid sizes: a course (1m) and a fine grid size (5mm). When we want to calculate the mutual information for a large-sized 2D-voxel, we perform the calculations on the finer, local grid. We do this only locally, in the area of the large-sized 2D-voxel, always with the same amount of 2D-voxels. We demonstrate this method in figure 8 which equates to setting a circular Bernoulli prior in equation (8). Therefore we possibly lose the multi modal character of the likelihood function. However, in most applications there will be some form of sensor fusion or recursive tracking and by varying the radius of the finer, local grid we are able to account for the effectiveness of the sensor fusion algorithm. This method of calculating will not have a large impact on our result, since the exact value of information is less important to us: we are more interested in the relative increase of the mutual information. The required run-time for an optimization problem is reduced greatly (by a factor greater than 3500 for a floor plan of 20m$\times$20m with a grid size of 5mm) by using this method, as we will demonstrate in the following section.

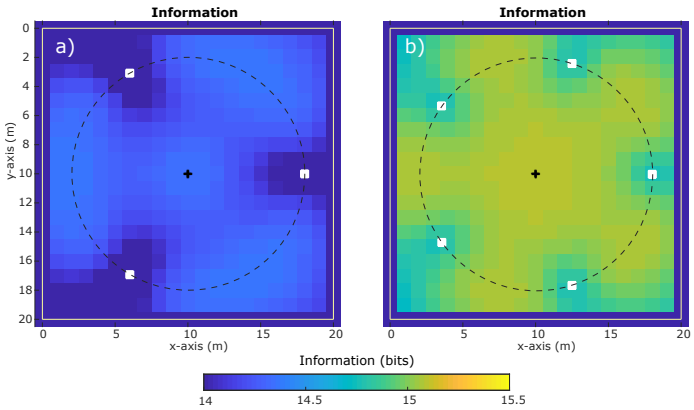

Fig. 7. Setup to determine the thresholds. The white squares represent the anchors. The floor plan has the same size with the same grid sizes as the floor plan for which we want to perform the optimization, without any obstacles. We place the anchors equidistantly from each other on a circle, with the center of this circle the 2D-voxel we want to evaluate ((10, 10) in this case). Figure (a) shows the configuration to determine the lower information threshold with three anchors while figure (b) shows the configuration to determine the maximum information threshold. Note that the color scale has been adjusted for emphasis.

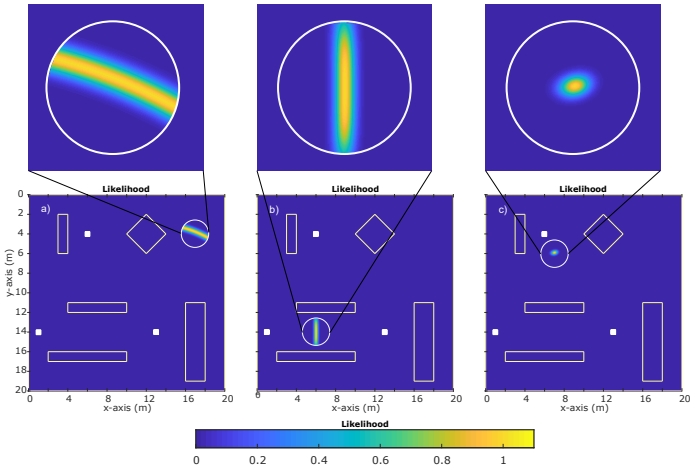

Fig. 8. Example of the use of two different grids. The white squares represent the anchor locations, the white circle is added to accentuate the area of the fine grid size 2D-voxels, for which the information is calculated. Compared to figure 2 we predominantly lose information in the case where we have LOS of only one anchor. Figure (a) shows the situation where only one anchor has LOS. In figure (b) there are two anchors with LOS and finally figure (c) has three anchors with LOS.

VI. Results

In all of the subsequent simulations, we have used the following parameters: a large grid size of 1m, a small grid size of 5mm, a spacing of 0.5m between the anchor locations, a standard deviation on the generated measurement ranges of 20mm and five Monte Carlo samples. The anchor locations were limited to be placed alongside a wall to reduce complexity. We used a mask of the inaccessible areas of the floor plan as a priori information as explained in figure 3 (c). The calculations were performed on a workstation with an Intel I9 7960X CPU with 128GB of RAM in MATLAB R2017b, parallelized to use 16 cores.

To determine the performance of our algorithm, we created a floor plan of 20m $\times$ 20m. In the floor plan we placed multiple obstructions to add to the complexity of the environment. Figure 9 shows the layout of our test environment. We used our incremental algorithm with variable grid size to be used as a benchmark for our speed optimization with a sub selection of anchors. We did not evaluate this problem with the absence of the variable grid size heuristic, as this approach would have taken over a month to complete. Figure 9 shows the result of the incremental algorithm, after placing fifteen anchors. It took approximately thirty hours to reach this result. The resulting information amounts can be seen in table I. To compare the performance of our optimized algorithm with the benchmark, we ran our optimized algorithm twenty times, with the same parameters as the benchmark. The result can be seen in figure 10. Figure (a) compares the total amount of information present in the environment with an increasing number of anchors. The blue line represents the results of the benchmark algorithm while the red dashed line is the mean of the results of the twenty optimized runs, with the error bar the standard deviation of these twenty runs. Figure (b) shows the minimum amount of information present in the environment for increasing numbers of anchors. The blue line is the minimal value with the benchmark algorithm while the red line is the mean of minimal values of the twenty optimized runs, with the error bar the standard deviation between these twenty runs. From these graphs we can conclude that our optimized algorithm performs very well as shown by the data in table I. We do notice in our results that it is difficult to obtain high information levels in narrow spaces, e.g. at the bottom right of Figure 9. This occurs due to GDOP since there is no room for the anchors to be far enough apart.

	Benchmark (bits)	Optimized (bits) (avg of 20 runs)
Max. info	16.19	16.13
Min. info	13.03	13.25
Mean info	15.40	15.33
Total info	4342.64	4323.69

TABLE I
COMPARISON OF THE INFORMATION AMOUNTS BETWEEN THE TWO ALGORITHMS AFTER PLACING FIFTEEN ANCHORS

As an additional test, we asked ten colleagues to place fifteen anchors to the best of their abilities in the same environment to compare their performance to our algorithm. We compared their results with the data of the twenty runs of our optimized algorithm by using a right-tailed Wilcoxon rank sum test [11]. The result of this test showed that our algorithm performed better than the test subject at a 1% significance level for the total, mean and minimum information levels.

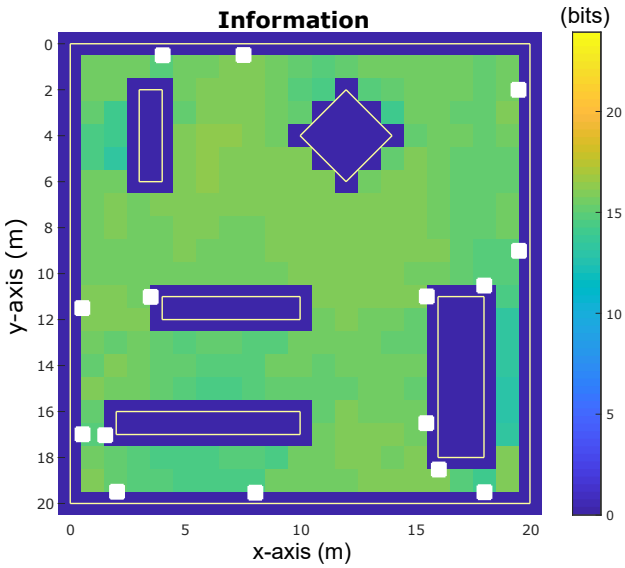

Fig. 9. Result of the incremental algorithm with variable grid size after placing fifteen anchors. The minimal information value is 13.03 bits. The maximum amount of information is 16.19 bits. The anchor locations were limited to be against a wall to limit the complexity. It took approximately 30 hours to calculate. The white squares represent the anchor locations.

VII. Conclusion

In this paper we have introduced a method for evaluating an anchor configuration in a UWB localization setting using an information-theoretic framework. We also introduced a heuristic algorithm to optimize the beacon placement, based on the mutual information distribution in a complex environment. To arrive at a solution that is computationally feasible, we have optimized the algorithm from a computational point of view, which resulted in a reduction in run-time for complex environments from approximately thirty hours to around two hours. In the future we will validate this approach in large-scale ($>$ 100m $\times$ 200m) industrial settings with many obstacles. We will experiment with different approaches in our anchor selection process to further increase the performance of our algorithm.

References

[1] A. R. Jiménez and F. Seco, "Comparing Ubisense, BeSpoon, and DecaWave UWB location systems: Indoor performance analysis," IEEE Transactions on Instrumentation and Measurement, volume. 66, no. 8 pp. 2106 – 2117, 2017.

[2] S. Monica and G. Ferrari, "UWB-Based Localization in Large Indoor Scenarios: Optimized Placement of Anchor Nodes," IEEE Transactions on Aerospace and Electronic Systems, Volume 61, no. 2, April 2015.

[3] J. Choliz, M. Eguizabal, A. Hernandez-Solana and A. Valdovinos, "Comparison of algorithms for UWB indoor location and tracking systems," IEEE Vehicular Technology Conference (VTC Spring) Yokohama, Japan , pp. 1-5, 2011.

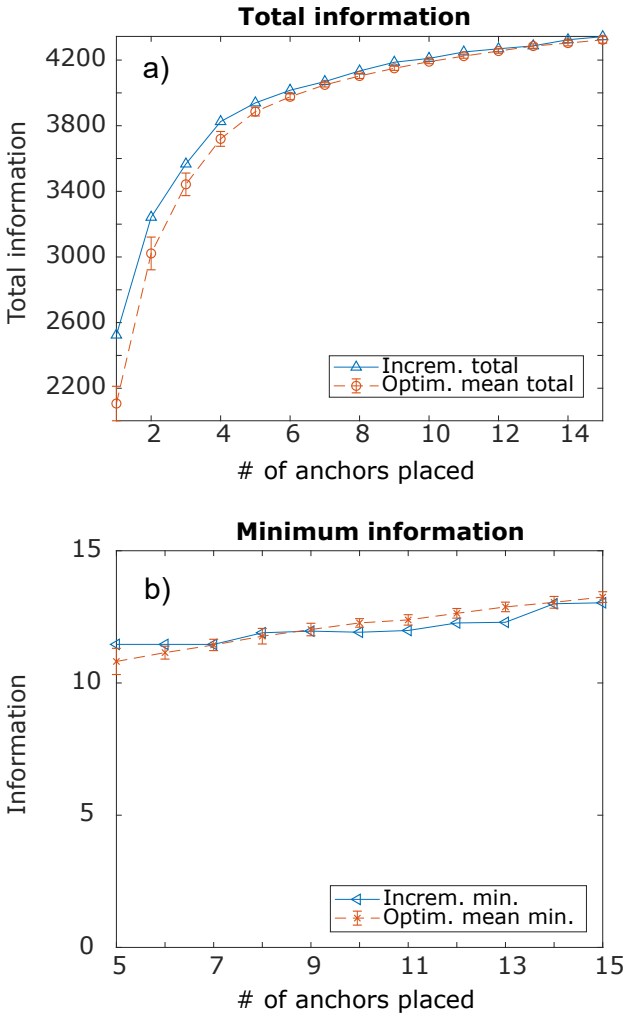

Fig. 10. Comparison between our benchmark algorithm and our optimized algorithm. Figure (a) compares the total amount of information present in the environment with an increasing number of anchors. The blue line represents the results of the benchmark algorithm while the red dashed line is the mean of the results of the twenty optimized runs, with the error bar the standard deviation of these twenty runs. Figure (b) shows the minimum amount of information present in the environment for increasing numbers of anchors. The blue line is the minimal value with the benchmark algorithm while the red line is the mean of minimal values of the twenty optimized runs, with the error bar the standard deviation between these twenty runs.

[4] K. Piwowarczyk, P. Korbel and T. Kacprzak, "Analysis of the influence of radio beacon placement on the accuracy of indoor positioning system," Proceedings of the 2013 Federated Conference on Computer Science and Information Systems Krakow, Poland, pp. 889-–894, November 2013.

[5] T. Leune, T. Wehs, M. Janssen, C. Koch and G. von Cöln, "Optimization of wireless locating in complex environments by placement of anchor nodes with evolutionary algorithms," 18th Conference on Emerging Technologies & Factory Automation (ETFA) Cagliari, Italy, September 2013.

[6] C. E. Shannon, "The mathematical theory of communication," University of Illinois Press, 1963.

[7] A. Krause, A. Singh and C. Guestrin, "Near-optimal sensor placements in Gaussian processes: theory, efficient algorithms and empirical studies," Journal of Machine Learning Research 9, pp. 235-284, February 2008.

[8] J. Reijniers, D. Vanderelst and H.Peremans, "Morphology-induced information transfer in bat sonar," Physical Review Letters Volume 105 Issue 14, pp. October 2010.

[9] D. Laurijssen, J. Steckel and M. Weyn, "Antenna Arrays for RSS Based Indoor Localization," IEEE SENSORS 2014 Proceedings, pp. 261-264, Nov 2014.

[10] D. F. Jones, S. K. Mirrazavi and M. Tamiz, "Multi-objective meta-heuristics: An overview of the current state-of-the-art," European Journal of Operational Research Volume 137, Issue 1, pp. 1-9, 2002.

[11] S. A. Glantz, "Primer of biostatistics," 6th edition, McGraw-Hill Education, 2005