Distributed Localization in Cluttered Underwater Environments

Muzammil Hussain
Computing Laboratory
University of Oxford
Oxford, OX1 3QD, United Kingdom
Muzammil.Hussain@comlab.ox.ac.uk

Niki Trigoni
Computing Laboratory
University of Oxford
Oxford, OX1 3QD, United Kingdom
Niki.Trigoni@comlab.ox.ac.uk

ABSTRACT

Mapping and exploration of closed fluid tanks, such as nuclear waste storage tanks and water sewage treatment ponds, by a swarm of underwater robots requires them to be aware of their positions, as they traverse the depths of the tank. This paper considers the challenges faced in localizing robots using reference-based localization in such clutter-prone environments. Clutter affects localization in two ways: it blocks reference communication signals, effectively preventing them from reaching the more remote regions in the tanks, and, second, it leads to the formation of non-line-of-sight (NLOS) distance measurements. In this paper, we consider the application of distributed localization techniques to cluttered environments and assess their performance. We also investigate how to place robots in a multi-hop network to improve the accuracy of position estimation in the far-fetched regions of the tank, when distributed localization is used.

Categories and Subject Descriptors

H.4 [**Information systems applications**]: Miscellaneous; C.3 [**Special-purpose and application-based systems**]: Real-time and embedded systems

General Terms

Algorithms, Performance, Measurement, Experimentation, Design

Keywords

Localization, underwater, wireless sensor networks, non-line-of-sight, distributed localization, bellman-ford, discrete event modelling

1. INTRODUCTION

Localization is defined as the process of estimating the location of a point in n-dimensional space. Localization is especially crucial in wireless sensor networks since the rationale behind deploying them in the first place is to accomplish tasks based on a spatial dimension in the environment - be it collecting sensor measurements or actuating sensors based on certain conditions. In other words, for processing sensor data, it is important to ascertain *where* the sensor readings are coming from. For example, in a vineyard where sensor networks are used to monitor temperature and humidity, it is important that the deployed sensors know their own positions so that they can stamp the recorded measurements with the location of occurrence. Similarly in a deployment using mobile sensors, such as the NASA Mars Tumbleweed project, the mobile sensors should be aware of their positions in order to perform their exploration and mapping work.

Most localization techniques are based on the use of *references*. References are location-aware nodes that assist other *unlocalized* nodes to estimate their positions. Localization can be broadly classified as *single-hop* localization and distributed *multi-hop* localization. Single-hop localization occurs when direct communication links exist between the references and the unlocalized node. For example, wireless sensors, equipped with GPS receivers, can localize themselves with the help of GPS satellites. Distributed localization [10, 11, 9] is employed when an unlocalized node cannot be reached by a sufficient number of references.

Reference-based localization operates by estimating distances between the references and the unlocalized node. When there exists a direct, unobstructed path between the two, the corresponding distance measurement still suffers from a small measurement error, which is traditionally modelled as a Gaussian distribution. However, in the presence of obstacles between a reference and the unlocalized node, reflected, non-line-of-sight (NLOS) distance measurements are formed. Such distance estimates have large positive errors and cause a substantial degradation in localization accuracy.

In this paper, we evaluate the performance of various localization techniques in cluttered underwater environments. Our aim, in particular, is to localize a swarm of submersible robots (like the prototype shown in Fig. (1c)), as they explore enclosed tanks. A typical scenario is shown in Fig. (1a). Here, we intend to send a *sensor* robot to explore/map distant areas of the tank. *Anchors* are installed on the tank surface and are used for positioning the underwater robots. In order to assist the sensor robot to localize in the specified *sensing* area, a number of *localizer* robots are deployed.

We perform an empirical study to compare the performance of various localization techniques in the presence of NLOS distance measurements errors. We study how the

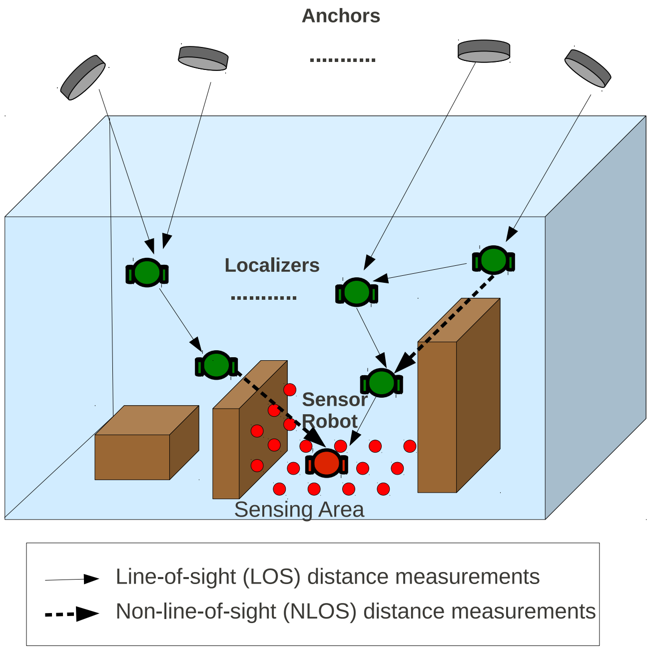

(a) Localization of the sensor robot in the (discretized) sensing area. Anchors, at the surface, and deployed localizer robots enable the localization of the sensor robot. The cluttered environment gives rise to NLOS measurements.

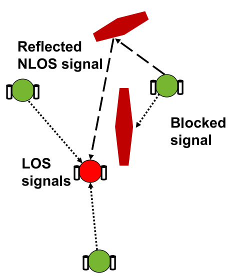

(b) Line of Sight (LOS) and Non Line of Sight (NLOS) measurements

(c) Submersible underwater robot prototype (Courtesy: Simon Watson, University of Manchester)

Figure 1: Localization in cluttered underwater environments

anchor communication range and the magnitude of NLOS errors impact the localization accuracy of single-hop and multi-hop localization techniques. We show that although DV-Distance offers the most promise when compared to other competing techniques in environments with large NLOS errors, its accuracy is prohibitively low when localizers are placed randomly in the cluttered environment. We thereby proceed to propose an algorithm, OPTPLACDVDIST, that computes placements for localizers in the cluttered environment, taking into account the obstacle topology as well as the communications ranges of anchors and localizers. We show that the resulting localizer arrangements improve the localization error of DV-Distance by an order of magnitude, thus making DV-Distance a practical and viable localization technique for cluttered NLOS environments.

The remaining of the paper is organized as follows : Section (2) briefly discusses factors affecting localization error

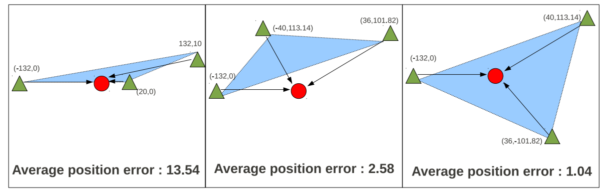

(a) Geometry 1 (b) Geometry 2 (c) Geometry 3

Figure 2: Effect of reference geometry on localization error. The line-of-sight (LOS) measurements are contaminated with Gaussian noise $\mathcal{N}$(0,1). The position error is an average over 10000 samples.

in reference-based localization. Section (3) evaluates the performance of various single-hop and multi-hop localization techniques in cluttered environments. Section (4) introduces the idea of optimal placement of localizer robots in cluttered environments to improve the accuracy of DV-Distance. Section (5) reviews related work, and Section (6) presents concluding remarks and ideas for future work.

2. LOCALIZATION ERROR

In this section, we enumerate the various factors that affect localization error, particularly in reference-based localization. The relationship between the various factors is also briefly discussed.

The error in reference-based localization is dependent on *three* factors :

1. **Distance measurement errors** : These are errors in the distance estimates between the unlocalized node and references. These are classified as line-of-sight (LOS) and, conversely, non-line-of-sight (NLOS). LOS errors are typically modelled as Gaussian of small magnitude [15], while NLOS errors can follow arbitrary, non-negative distributions [3].

2. **Reference position errors** : These are position errors of the references themselves. They affect localization error, in that they cause a discrepancy in the measured distances to references and the reference positions. Typically, reference errors arise in distributed localization, where previously localized nodes are used to localize other (neighbouring) nodes, like in iterative localization [11, 10, 1].

3. **Geometry of references** : Reference geometry affects localization error, in that it magnifies the errors in references positions and distance measurements. Conversely, if there is absolutely no error present, the geometry of the references will not have any effect. Levanon [6] shows that the optimal reference geometry, yielding the least magnification of distance errors in the localization error, is obtained when the references are arranged as a regular polygon with the unlocalized node in the center. The impact of reference geometry on localization error is illustrated in Fig. (2).

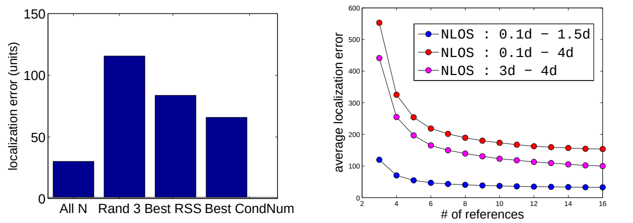

(a) Single-hop localization (b) Effect of minimum number of references

Figure 3: Localization error in single-hop localization with NLOS distance measurements. Each data point is an average of 10000 samples.

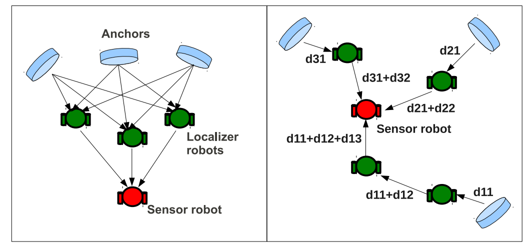

(a) Iterative localization (b) DV-Distance

Figure 4: Distributed localization techniques - iterative localization and DV-Distance - can localize a sensor robot in the remote regions of a tank.

3. LOCALIZATION IN CLUTTERED ENVIRONMENTS

In this section we investigate the performance of existing localization techniques in cluttered NLOS-prone environments. First, we look at single-hop localization and empirically evaluate the accuracy it can achieve in the presence of NLOS errors. We then consider two distributed localization techniques, iterative localization [11, 1] and DV-Distance [9, 12], and evaluate their performance in NLOS-prone conditions. For simplicity, throughout the paper, we consider a two-dimensional setting, as the results are analogous in higher dimensions, and reduce the sensing area to a single point. We estimate position using a linear system of equations (itself derived from the non-linear set of equations for trilateration).

3.1 Single-hop localization

In order to evaluate the performance of single-hop localization in NLOS environments, we run the following simulation: We consider the unlocalized node (e.g. a sensor robot) to be situated in the center of a 100×100 area, with 16 references arranged around the square perimeter, as shown in Fig. (5). The unlocalized node can obtain only NLOS distance measurements to the references. The NLOS bias is modelled as a distance-dependent, uniformly random variable. Thus, the NLOS distance measurement is:

$$d_{NLOS} = d + \mathcal{U}(u_{min}, u_{max}) * d \quad (1)$$

where u_{min} and u_{max} are the bounds of a uniform distribution and d is the Euclidean distance between a reference and the unlocalized node.

We compare the localization error of four variants of single-hop localization:

1. **All N**: where all 16 references are considered.
2. **Random 3**: where only three (randomly chosen) references are used for localization.
3. **Best RSS**: where the position with the least Residual Sum of Squares (RSS) (of the least squares position estimate) is chosen.
4. **Best Condition Number**: where the position with the least condition number of the linear system of equations (used for trilateration) is chosen.

For the third and fourth variants, we consider all possible subsets of the total number of references - which makes them computationally expensive. Fig.(3a) illustrates the results, when the NLOS distance measurement errors are drawn from $\mathcal{U}(0.1d, 1.5d)$. Using all the available references provides the most accurate position estimates (on average). Though Albowicz et al. [1] claim that RSS can be a good indicator of NLOS range measurements, RSS does not reliably reflect the effect of large NLOS errors in the position estimate when the majority of the range measurements are NLOS. Minimizing the condition number also proves to be an ineffective technique for obtaining smaller localization errors. This is because, even with a low condition number (in other words, a good reference geometry), the localization error still depends on the magnitude of the NLOS errors. In other words, large NLOS errors will impact localization error regardless of the geometry of the references. Fig. (3b) shows that, besides the magnitude of NLOS errors, the variability in their magnitudes also has a significant impact on the localization error. For instance, we see that there is a significant increase in position error when the NLOS bias is drawn from $\mathcal{U}(0.1d, 4d)$ compared to $\mathcal{U}(3d, 4d)$. This is due to two factors: the geometry of the references in this setup, and the (linear) least squares technique used for position estimation. In the presence of a good reference geometry, localization error increases, not only with the magnitude, but also with the variance of distance measurement errors.

3.2 Multi-hop localization

We are now in a position to investigate the application of distributed localization techniques in NLOS environments. For our simulations, we consider two distributed localization techniques: iterative localization; and DV-Distance. Additional *localizer* robots are used to assist a sensor robot to determine its position, as previously illustrated in Fig. (1a).

Iterative localization uses previously localized robots as *pseudo-anchors* to localize other robots in their neighbourhoods. For example, in Fig. (4a), localizer robots are first localized by anchors and then in turn localize the sensor robot. Thus, we distinguish between two types of references - anchors and pseudo-anchors. Basic iterative localization, introduced by Savvides et al. [11], does not allow a node to refine its location once it is localized. Improvements to iterative localization by Albowicz et al. [1] and Liu et al. [7] allow a localized node to relocalize itself, based on a metric which quantifies the localization error (for example, RSS [1]). Here, when a localized node receives a new reference advertisement, it will use the metric to ascertain whether the quality of the new reference is better than its

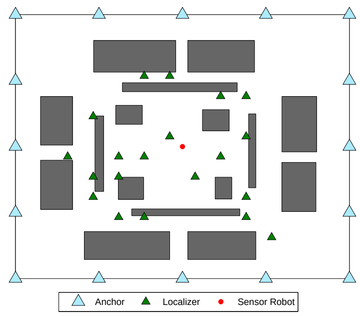

Figure 5: Clutter Topology

current set of references. If it is not, it ignores the new reference. Even if it decides to relocalize with the new reference, it will accept the subsequently calculated position, only if the metric for this new position is better than that for the current position. It is thus imperative for the metric to accurately reflect actual (unknown) position error, otherwise previous, more accurate, position estimates will be replaced, in subsequent re-localizations, by estimates with larger errors. In the paper, we use the RSS-metric enhanced version of iterative localization.

DV-Distance presents an alternative method for localizing the sensor robot. Here, the cumulative distance to the anchors is disseminated, using distributed Bellman-Ford algorithm, by intermediate localizer robots to the sensor robot. For example, in Fig. (4b), localizer robots propagate anchor advertisements, together with cumulative distances to the anchors, towards the sensor robot. The sensor robot uses the anchor positions and cumulative distances to the anchors to calculate its own position.

3.3 Evaluation in NLOS-prone environments

We use Prowler, a MATLAB-based discrete event simulator (DES) [13], to perform wireless sensor network simulations in a cluttered setting. In the paper we have used a simple inverse-square attenuation model as the obtained results are analogous with other existing propagation models. We have modified the inherent Noisy Disk model [15] of the simulator to implement NLOS distance measurements. In other words, the condition for two nodes to communicate, with obstacles between themselves, is:

$$d + n_{NLOS} < R_c$$

where n_{NLOS} denotes the simulated NLOS error, d is the Euclidean distance between the two nodes and R_c is the transmission radius of the transmitting node. (Under the Noisy Disk model, the two nodes can communicate if $d < R_c$, with the Gaussian error being added later on). Anchors may have a different communication range than that of localizers. We denote the former as R_{anch} and the latter as R_{loc}.

Fig. (5) shows the topology of an obstacle-prone environment, used in the simulations. Here, the anchors are situated on the periphery, and the localizers (represented by the small green triangles) are randomly distributed in the area amongst the clutter (represented by the grey objects). The sensor robot, situated in the center, does not have an unobstructed view of any of the anchors. We assume that, when two nodes do not have any clutter between them, the line-of-sight (LOS) distance measurement, d_{LOS}, is:

$$d_{LOS} = d + \mathcal{N}(0, \sigma^2) \qquad (2)$$

where d is the Euclidean distance between the two nodes and σ is the standard deviation of the Gaussian noise. We set σ to 1 throughout all the simulations. When clutter blocks the direct path, the resulting non-line-of-sight (NLOS) measurement is modelled as in Eqn. (1). We perform simulations using two uniform distributions, $\mathcal{U}(0.1d, 1.5d)$ and $\mathcal{U}(0.1d, 4d)$, for modelling NLOS errors. The communication range of localizers is fixed at 25 units.

We compare three localization techniques: single-hop localization, where the sensor robot will only be able to localize with NLOS-contaminated distance measurements from anchors; iterative localization with a RSS-based re-localization criterion; and, DV-Distance, which estimates cumulative distances to reachable anchors. Iterative localization and DV-Distance make use of localizers, which are randomly placed in the area. For consistency, identical random localizer placements are used to compare the two techniques. For single-hop localization, obviously, no localizers are used.

Fig. (6) compares the localization error obtained with the three localization techniques. We implicitly show a measure of successful localizations of the sensor robot, in the same graphs, by removing any data point for which the localization success ratio is not 100% (in other words, where the sensor robot was not able to localize in *every* simulation run for that particular configuration). When NLOS errors are drawn from $\mathcal{U}(0.1d, 1.5d)$, single-hop localization provides the best position accuracy, but only for large values of anchor communication range. Key factors for the good performance of single-hop localization are the good reference geometry provided by the peripheral anchors and the large number of references used (as discussed previously in subsection (3.1)). However, the performance of single-hop localization degrades for large NLOS errors, drawn from $\mathcal{U}(0.1d, 4d)$, as shown in Fig. (6b). The minimum value of R_{anch} for which at least three anchors are able to communicate directly with the sensor robot can be computed as (d_{nr} + $MAX(d_{nlos})$), where d_{nr} is the distance between the sensor robot and the third nearest reference and $MAX(d_{nlos})$ is the maximum possible NLOS distance error in the simulation. For example, when the NLOS error distribution is $\mathcal{U}(0.1d, 1.5d)$, the R_{anch} threshold is $(50 + 1.5 \times 50) = 125$.

In case of iterative localization, when R_{anch} increases such that the anchors can directly reach the sensor robot, we find that the localization error increases significantly. The reason is that, now, not only are anchors used for localization, but also localizers. The localizers can have large errors in their own positions. Thus, in scenarios where anchors can reach the sensor robot directly, using iterative localization does not have any benefits.

We find a more complex relationship between the localization error and the anchor communication range, in the case of DV-Distance. As R_{anch} increases, the sensor robot has access to an increased number of anchors, but those to which the cumulative distances bear larger (NLOS) errors. Unlike, the version of iterative localization we use here, DV-Distance always uses all reachable anchors to localize. This, in turn, has a detrimental effect on the localization accuracy. On the other hand, DV-Distance is able to localize a sensor robot in most scenarios, even for small R_{anch} values. The reason for this is that DV-Distance merely requires the

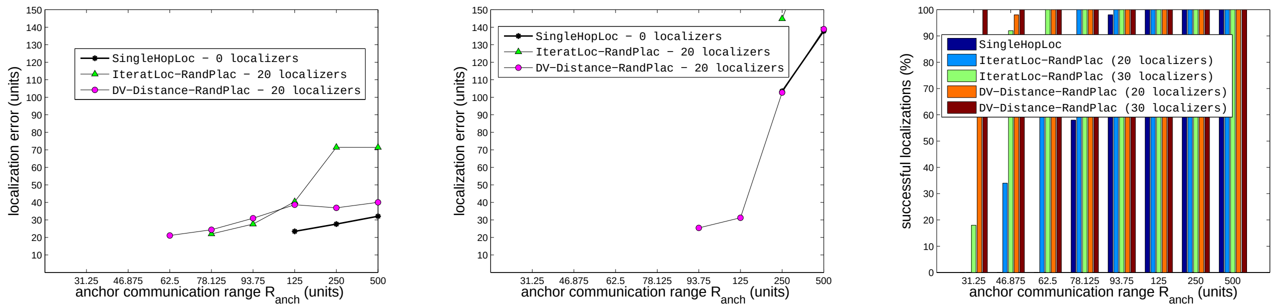

(a) Localization Error (20 localizers, NLOS bias $\mathcal{U}(0.1d, 1.5d)$) (b) Localization Error (20 localizers, NLOS bias $\mathcal{U}(0.1d, 4d)$) (c) Localizations successes (NLOS bias $\mathcal{U}(0.1d, 1.5d)$)

Figure 6: Comparison of localization techniques, varying the anchor communication range R_{anch}. Each data point is an average of 50 independent simulation runs. In case of iterative localization and DV-Distance, 50 distinct localizer topologies were used for each data point.

localizers to propagate the cumulative distance information further towards the sensor robot. Localizers need not necessarily localize themselves to be able to participate in the multi-hop localization process. Consequently, DV-Distance does not require a large number of anchors, and even three would suffice. Iterative localization, on the other hand, requires a larger number of localizers as well as anchors to *propagate localization* towards the sensor robot. Fig. (6c) shows that using a larger number of localizers improves the chances of localizing the sensor robot, for both DV-Distance and iterative localization. Moreover, one can conclude from the figure that DV-Distance succeeds in localizing the sensor robot more frequently than iterative localization for a given number of localizers, even for small values of R_{anch}.

We can conclude that, in the presence of very large NLOS errors, using distributed localization techniques and, in particular, DV-Distance with a small R_{anch}, can provide better localization accuracy than single-hop localization.

4. PLACEMENT OF LOCALIZERS IN CLUTTERED ENVIRONMENTS

In the previous section, our empirical study showed that DV-Distance typically outperforms single-hop localization and iterative localization in cluttered environments with large NLOS errors. Unlike iterative localization, DV-Distance is only affected by range measurement errors, and not by propagating reference position errors. Nevertheless, the localization error of DV-Distance is still very high (greater than 20 units in a 100x100 area), which makes it impractical to employ in cluttered environments. In this section, we focus on DV-Distance and investigate two possible ways of reducing its localization error: 1) by increasing the number of localizer nodes; and 2) by carefully placing localizer nodes in the cluttered area.

4.1 Impact of Localizer Density

Let us first consider whether we can improve the performance of DV-Distance by increasing the number of localizer nodes in the area. We run a simulation using the cluttered topology shown in Fig.(5). The number of anchors is fixed to 4, one on each edge of the square perimeter. The communication range of the anchors is fixed at 300 units, enabling the sensor robot to be reached from all the anchors even with the maximum amount of NLOS bias, which itself is modelled as $\mathcal{U}(0.1d, 1.5d)$, where d is the Euclidean distance between two nodes.

Fig. (7a) shows that we can significantly improve the localization accuracy of DV-Distance by deploying more localizers in the cluttered area. As we increase the number of localizers, the DV-paths become straighter and hence shorter, and the errors in the estimated distances between anchors and sensor robot decrease. However, a very large number of localizers is required (greater than 100) in order to achieve localization accuracy that is comparable to that achievable in clutter-free environments. Notice in the same figure, that in clutter-free environments where range measurement errors are small and normally distributed ($\mathcal{N}(0, 1)$), the localization error increases slightly with the number of localizers. The reason is that negative errors from the Gaussian distribution often result in cumulative distances that underestimate the real distances between sensor robot and anchors.

Fig. (7c) shows the proportion of LOS and NLOS DV-paths in the cluttered environment. An LOS DV-path is a multi-hop path between a anchor and the sensor robot, in which *all* component distances are line-of-sight. Otherwise the DV-path is classified as a NLOS path. We see that the proportion of LOS paths increases with the density of localizers, but it stops increasing after 100 localizers. Initially, as we add localizers, more edges tend to become LOS, and DV-paths tend to be shorter if they use LOS edges only. As the number of localizers increases further (>100), short NLOS edges are created that have a small positive bias; using at least one of these edges in a DV-path reduces its zig-zag shape, and thus its length. Hence, more DV-paths tend to make use of at least one small NLOS edge. This is supported by Fig. (7b), which shows that the distance error of NLOS DV-paths reduces with the number of localizers and approaches that of LOS DV-paths. Fig. (7d) shows that, independent of whether DV-paths consist of LOS edges only, or a mixture of LOS and NLOS edges, they tend to become longer in hops, as we increase the number of localizers.

While we have shown that increasing the number of localizers in the cluttered environment improves the localization accuracy of DV-Distance, the required number of localizers, to see a marked improvement, is prohibitively large. Fig. (7a) shows that we need at least 50 localizers, placed randomly in the area, in order to obtain a localization error of

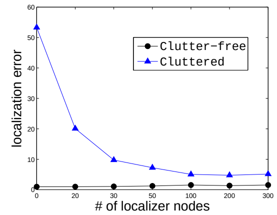

(a) Localization error with varying localizer density

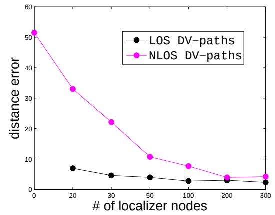

(b) Distance error of DV-paths

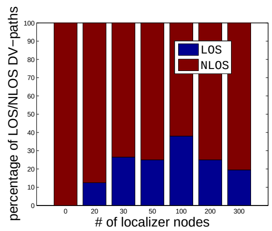

(c) Proportion of LOS/NLOS DV-paths

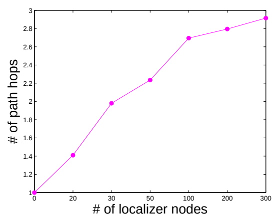

(d) Number of hops of the DV-paths

Figure 7: Effect of localizer density on the localization accuracy of DV-Distance in cluttered environments. Each data point is an average of 100 sample runs.

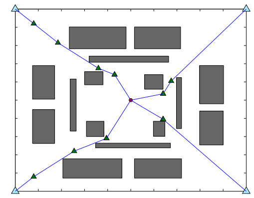

(a) Anchor communication range : 500 units

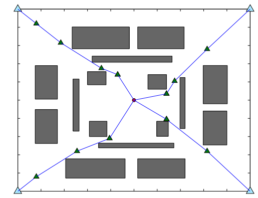

(b) Anchor communication range : 31.25 units

Figure 8: Optimal localizer placements computed by OPTPLACDVDIST, for varying anchor communication ranges. Resolution is set to 625 (25x25) points.

less than 10 units. We show, in the following subsection, that it is possible to achieve small localization errors with a much smaller number of localizers, if we *place them in carefully chosen positions* in the cluttered environment.

4.2 Impact of Localizer Placement

In this subsection, we explore how to deploy a relatively small number of localizers in pre-determined positions in the cluttered environment so as to reduce the localization error of DV-Distance. We propose an optimal placement algorithm for localizers, namely OPTPLACDVDIST. The input to the algorithm, a graph $\mathcal{G}$, is formulated considering the environment area size, the clutter topology, the available number of localizers and the communication range of localizers and anchors. We divide the cluttered area into a grid, and consider each point of the grid that is not covered by an obstacle as a possible localizer position. This set of points, together with the points representing the positions of anchors and the sensor robot, form the vertices of $\mathcal{G}$. The edges of $\mathcal{G}$ are based on the communication ranges of the anchors and localizers. For example, we draw an edge between an anchor and a potential localizer position, if their distance is smaller that the anchor communication range. Similarly, we add an edge between two localizer positions if their distance is smaller than the localizer communication range. The weight of each edge depends on whether it intersects with an obstacle. The weight of a LOS edge, which does not intersect with obstacles, is set to its Euclidean length d. The weight of a NLOS edge, which intersects with one or more obstacles, is greater than its length d. In our simulations, it is set to $c * d$, where $c = 1.75$.

The optimal placement algorithm, OPTPLACDVDIST, is a modification of the Bellman-Ford algorithm, which produces the shortest paths from a given source, to all other nodes in a weighted graph. In our case, we are interested in the shortest paths between the position of the sensor robot and each of the anchor positions. Our proposed algorithm is optimal under the assumption that we allocate localizers evenly across anchors; for example, if we have 20 localizers and 4 anchors, at most 5 localizers will be allocated to form a DV-path between the sensor robot and each anchor node. Given the limited number k of localizers per DV-path, we have modified the Bellman-Ford algorithm to search for shortest paths between the sensor robot and each of the anchors that have fewer than $k + 1$ hops.

Fig. (8) shows optimal DV-paths produced for two different anchor communication ranges. Up to 4 localizers are allocated to each anchor in both cases. However, not all of them are being used in practice. Only 10 out of 16 localizers are used when the anchor communication range is large (500 units) and 12 out of 16 localizers are used when the anchor communication range is small (31.25 units). In this experiment, the localizer communication range is set to 25 units.

We are now in a position to evaluate the performance of DV-Distance when it uses carefully placed, instead of randomly placed localizers. Fig. (5) shows how the error in distance estimates between sensor robot and anchors depends on the number of localizers, the grid resolution and the localizer communication range. Fig. (9a) shows that for a given grid resolution, the distance error initially decreases with the number of localizers and, beyond a certain point, it remains constant. Moreover, increasing the grid resolution significantly improves the accuracy of distance estimates. For example when 3 localizers are used per anchor, a 4x4 grid resolution yields a distance error of about 50 units, a 7x7 grid reduces the error to 20 units, and a 13x13 grid reduces it further to 4 units. However, the algorithm becomes more computationally-intensive for higher resolutions. Fig. (9b) shows how the errors in distance estimates between the sensor robot and anchors are reduced when we carefully place localizers in positions computed by the OPTPLACDVDIST algorithm, as opposed to placing them randomly in the cluttered environment. For example, if we have 8 localizers at our disposal (i.e. 2 per anchor), and we place them randomly, we get distance errors of about 50 units, whereas if we place them using our modified Bellman Ford algorithm, we obtain errors of less than 5 units (with 13x13 or 25x25 grid resolutions). Finally, Fig. (9c) illustrates the impact of the localizer communication range on the distance error. As we increase the localizer communication range, we have

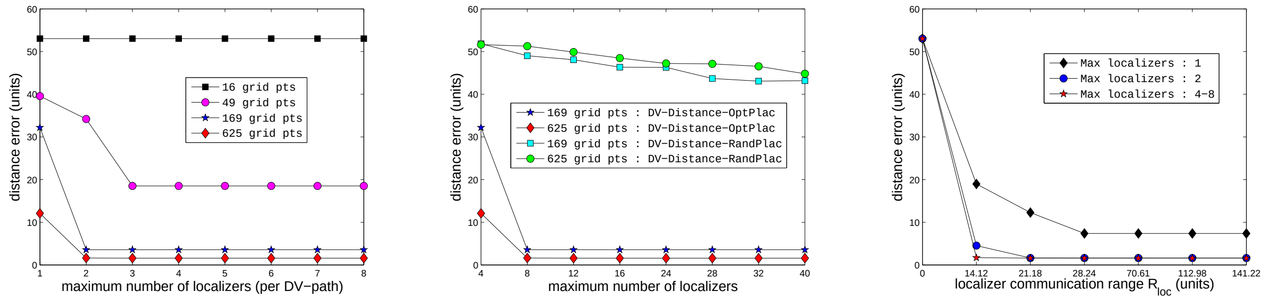

(a) Effect of grid resolution

(b) Comparison of distance errors with random and optimal localizer placement

(c) Effect of localizer communication range

Figure 9: Error in the estimated distance between sensor robot and anchors in DV-Distance with optimal localizer placements computed using OPTPLACDVDIST. The distance error decreases as we increase the localizer density, the grid resolution and the localizer communication range.

more flexibility on the placement of localizers, thus leading to potentially shorter DV-paths, and thus smaller distance errors.

So far, we have shown that carefully placing localizers in a cluttered environment allows us to estimate more accurately the distances between sensor robot and anchors. We now show how it also considerably increases the accuracy in localizing the sensor robot. Fig. (10) compares the localization accuracy of single-hop localization, which does not make use of localizers, iterative localization with random localizer placement, DV-Distance with random localizer placement and DV-Distance with optimal localizer placement. The considered optimal placements are shown in Fig. (8). In case of DV-Distance with optimal placement, we use only 4 anchors for our simulations, while for others, 16 anchors are used, as discussed in the previous section. We use the clutter topology shown in Fig. (5) with the anchor communication range varied from 31.25 to 500 units. The localizer communication range is set to 25 units. Notice that the localization error of DV-Distance with optimal localizer placement is one order of magnitude smaller than that of alternative schemes, even when the latter employ 30 localizers. Moreover, when

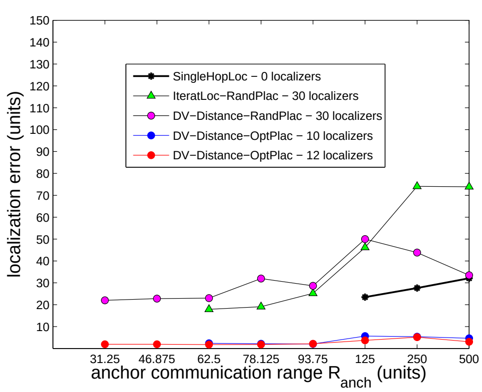

Figure 10: Comparison between the error of localization techniques using random localizer placement and optimal localizer placement. Each data point is an average of 50 simulation runs.

the optimal localizer placement with 12 localizers is used, the sensor robot can be localized in the cluttered environment for all considered values of R_{anch}. Thus, our localizer placement algorithm results in reasonably low errors (about 5 units in a 100x100 area), allowing the practical application of DV-Distance in cluttered environments.

5. RELATED WORK

Distributed localization can be sub-divided into two variants : hop-based localization and iterative localization. Langendoen et al [5] provides a good comparison between Hop-Terrain [10], APS [9] and n-Hop multilateration [12]. The paper proposes to differentiate two stages in distributed localization : initial position estimation and position refinement. In this paper, we considered DV-Distance and a version of iterative localization that uses Residual Sum of Squares (RSS) for error management. In case of DV-Distance, subsequent re-localizations generally tend to improve the accuracy of the position estimate, unless a DV-path with a large error is received. In case of iterative localization with RSS-based error management [1], the position can only be improved when the RSS can correctly differentiate between distances with large NLOS errors or references with large errors. Albowicz et al. [1] claim that the RSS value is a good indicator of the presence of NLOS errors. However, this is based one the assumption that only a small minority of the distances are NLOS in nature. Our simulations have shown that RSS is not a good indicator of NLOS distances when they make a majority of the available distances. Liu et al. [7] propose an error management method of iterative localization, based on Reliable Least Squares (RLS), which uses matrix-reconditioning to improve the position estimates even when the references have a poor geometry. However the technique will not work for non-Gaussian NLOS scenarios as the optimality of the underlying RLS technique is based on the assumption that both references and distances suffer from Gaussian errors. Moore et al. [8] propose to control the propagation of error in iterative localization by constraining the reference geometry. However, this comes at the expense of reducing the number of successful localizations throughout the network.

Whitehouse et al. [14] compare the robustness of various distributed localization algorithms using an empirically derived distance error model [15]. They demonstrate that iterative localization techniques result in much larger localization errors than their counterparts. However they attribute these larger errors only to the propagation of references errors. The primary reason for large errors in iterative localization is the incidence of poor reference geometry in arbitrary topologies. DV-Distance, while suffering from larger distance errors (when random topologies are considered), can avoid this via an error-free optimally placed set of references.

A large amount of research has been been done in the area of optimal reference placement for single-hop localization. A number of techniques propose using the condition number of the observability matrix and the Geometric Dilution of Precision (GDOP) as metrics to minimize while searching for the optimal placement of sensors around a target [2, 6]. The Cramer Rao Bound (CRB) is used as a cost function, while minimizing localization error, in a number of works. Jourdan et al. [4] present an optimal sensor placement algorithm, RELOCATE, based on the minimization of a cost function, the Position Error Bound (PEB) [3]. The paper considers the scenario where distance measurements can be corrupted by NLOS biases and proposes that for a given set of references, their optimal placement is a trade-off between the reference GDOP and the NLOS errors.

To the best of our knowledge, this is the first work that deals with optimal placement of localizer nodes in a multi-hop distributed localization setting, where distance measurements are affected by NLOS errors.

6. CONCLUSION AND FUTURE WORK

In this paper we considered the novel application of distributed localization techniques in cluttered environments to mitigate the effects of non-line-if-sight (NLOS) distance measurements. We start by empirically comparing the performance of single-hop and distributed multi-hop localization techniques for a given clutter topology. We conclude that, in the presence of very large NLOS errors, using distributed localization techniques with a small anchor communication range is preferred to single-hop localization. Moreover, DV-Distance typically outperforms iterative localization, because the latter suffers from pseudo-anchor position errors that are often magnified by their bad geometries. Despite the benefits of DV-Distance, it still required a very large number of localizers to yield a reasonably low localization error. In order to make DV-Distance practical in cluttered environments, we proposed placing localizers at carefully selected positions taking into account the clutter topology. Our localizer placement algorithm allowed DV-Distance to reduce its localization error by an order of magnitude compared to when DV-Distance uses randomly placed localizers.

Currently, our localizer placement algorithm allocates a fixed number of localizers to each DV-path. In the future, we plan to improve our algorithm to allow a localizer to be shared across DV-paths. In our current work we have assumed that NLOS bias is a distance-dependent uniformly distributed random error. We would like to further investigate the performance of the localization techniques evaluated in the paper in the presence of empirically validated NLOS error distributions. Finally, in future work, we would like to relax the assumption that we have full knowledge of the clutter topology, and also consider the practical issue of deploying localizers to their selected positions.

7. ACKNOWLEDGEMENTS

This work is supported by the EPSRC grant EP/F064209/1 on Actuated Acoustic Sensor Networks for Industrial Processes (AASN4IP). We would like to acknowledge the support from Sarfraz Nawaz and Simon Watson.

8. REFERENCES

[1] J. Albowicz, A. Chen, and L. Zhang. Recursive position estimation in sensor networks. In *ICNP 2001*, pages 35–41. Citeseer.

[2] T. Hegazy and G. Vachtsevanos. Sensor placement for isotropic source localization. *Lecture notes in computer science*, 2003.

[3] D. Jourdan, D. Dardari, and M. Win. Position error bound for uwb localization in dense cluttered environments. *IEEE Transactions on Aerospace and Electronic*, 00:3705–3710, 2008.

[4] D. Jourdan and N. Roy. Optimal sensor placement for agent localization. *ACM Transactions on Sensor Networks (TOSN)*, 4(3):13, 2008.

[5] K. Langendoen and N. Reijers. Distributed localization in wireless sensor networks: a quantitative comparison. *Computer Networks*, 43(4):499–518, 2003.

[6] N. Levanon. Lowest GDOP in 2-D scenarios. *IEE Proceedings: Radar, Sonar and Navigation*, 147(3):149–155, 2000.

[7] J. Liu, Y. Zhang, and F. Zhao. Robust distributed node localization with error management. In *MobiHoc*, pages 250–261. ACM New York, NY, USA, 2006.

[8] D. Moore, J. Leonard, D. Rus, and S. Teller. Robust distributed network localization with noisy range measurements. In *Sensys*, pages 50–61. ACM New York, NY, USA, 2004.

[9] D. Niculescu and B. Nath. Ad hoc positioning system (aps). In *IN GLOBECOM*, pages 2926–2931, 2001.

[10] C. Savarese, J. Rabaey, and K. Langendoen. Robust positioning algorithms for distributed ad-hoc wireless sensor networks. In *USENIX technical annual conference*, volume 2. Monterey, CA, 2002.

[11] A. Savvides, C. Han, and M. Strivastava. Dynamic fine-grained localization in ad-hoc networks of sensors. In *MOBICOMM*, pages 166–179. ACM New York, NY, USA, 2001.

[12] A. Savvides, H. Park, and M. Srivastava. The n-hop multilateration primitive for node localization problems. *Mobile Networks and Applications*, 8(4):443–451, 2003.

[13] G. Simon, P. Volgyesi, M. Maroti, and A. Ledeczi. Simulation-based optimization of communication protocols for large-scale wireless sensor networks. *2003 IEEE Aerospace Conference, 2003. Proceedings*, 3, 2003.

[14] K. Whitehouse and D. Culler. A robustness analysis of multi-hop ranging-based localization approximations. In *IPSN*, page 325. ACM, 2006.

[15] K. Whitehouse, C. Karlof, A. Woo, F. Jiang, and D. Culler. The effects of ranging noise on multihop localization: an empirical study. In *IPSN*, page 10. IEEE Press, 2005.