Revisiting the Cramer-Rao Bound for Localization Algorithms

S. Dulman and P. Havinga
Pervasive Systems, EEMCS
University of Twente – The Netherlands

A. Baggio and K. Langendoen
Parallel and Distributed Systems, EEMCS
Delft University – The Netherlands

I. Introduction

Acquiring position information by means of ad-hoc networks and in particular wireless sensors networks (WSNs) received a lot of attention in the past years. Survey works, such as [1], [2], show a large number of techniques/algorithms that can be used to solve the localization problem.

The techniques used are often borrowed from other fields of science and modified to fit the context of wireless sensor networks [3–5]. In order for results established in other fields of science to hold for the problem at hand, particular care must be taken to ensure that the assumptions are still valid. Even the slightest mismatch in the underlying assumptions could render the well-known techniques useless and lead to wrong results.

In this paper, we address the usage of *lateration* [6] and the associated *Cramer-Rao Bound* (CRB) [3], concepts borrowed from GPS localization [7]. Via a series of counter-examples, we show how these concepts fail to deliver the expected results when applied to the field of WSNs. The goal of this paper is to bring forward the idea that a foundation based on geometrical considerations – rather than estimation theory – should be employed when studying the basic mechanisms and boundaries for localization in WSNs.

II. The localization problem

Let the position of a node in the 2D plane be written as z. The position error is defined as the euclidean distance between the real (z) and estimated position ($\tilde{z}$) of a node, $e_i = \|z, \tilde{z}\|$. We call an *anchor* a node that acquired *exact* position information (anchor i has position a_i). A node without knowledge of its position is simply called *node*.

Assume a scenario with three anchors (blue circles in Fig.1(a)) and a node positioned in the central area (red circle). Assuming the exact distances $\overline{d_i}$ between the node and anchors are known, the node is able to infer its position $\tilde{z}$ by intersecting the three circles centered at a_i with radii $\overline{d_i}$. If errors are added to the distance estimates, simulating measured distances d_i, no position in the 2D plane is likely to satisfy the distance constraints. Solving the localization problem then takes the form of choosing the estimated position that minimizes a certain metric. Lateration [6] is one of the most popular techniques and works as follows:

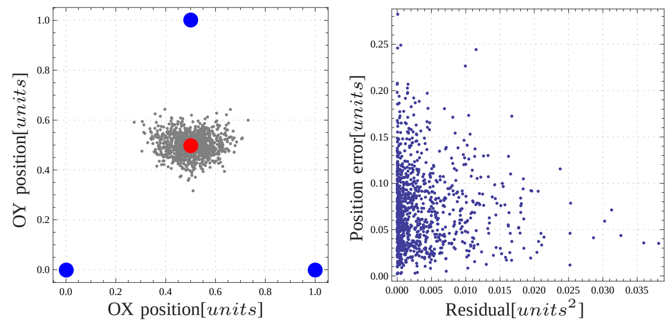

(a) Network topology (b) Position error vs. residual val.

Fig. 1. Localization in the presence of noise

$$\tilde{z} = \arg_{\tilde{z}} \min \sum_{i=1}^{n} \omega_i \left(d_i - \tilde{d}_i \right)^2 \quad (1)$$

where $\tilde{d}_i$ are the distances between the computed position and the anchors, and ω_i a set of weights. When varying the noise in the measured distances and using equal weights, we obtain a cloud of estimated positions as shown in Fig.1(a) in gray.

Tracing back the lateration technique we come across the concept of Cramer-Rao Bound applied to the localization problem [3]. CRB defines the lower bound on the precision of a localization estimator. The articles [3], [8], [9] solve the localization problem step-wise as follows:

- **Basic concept**: known noise distributions are applied to distance measurements;
- **Theoretical step**: determine CRB to obtain an idea of the achievable accuracy of unbiased position estimators;
- **Algorithm**: from the formulation of CRB determine the metric to be minimized in order to obtain a position (e.g. using the lateration procedure);
- **Post processing**: use coefficients such as Geometric Dilution of Precision [10] (GDOP) to correct issues not captured by the estimation theory (such as the geometry of anchor deployment).

It is *assumed* that minimizing the sums in equation (1) leads to better position estimates. In some works the *residual value*, i.e. the minimum value for the sum, serves as an indicator on how good the algorithm performs. In Fig.1(b), we plot the residual value versus the position error for each of the

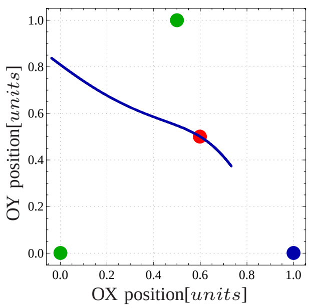

(a) Computed position trajectory

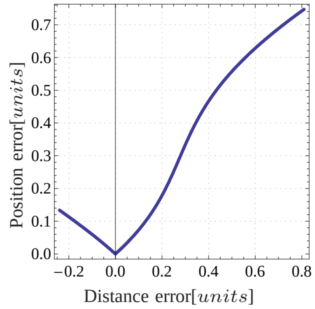

(b) Position errors vs. distance error

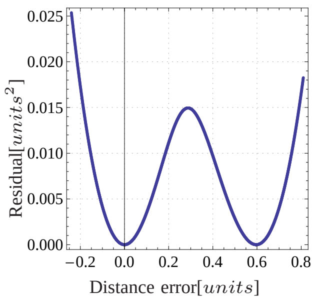

(c) Residual vs. distance error

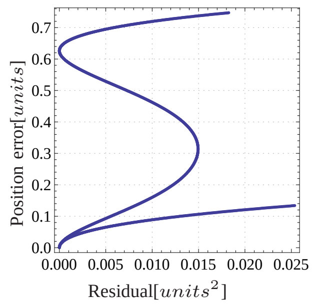

(d) Position error vs. residual

Fig. 2. Lateration behavior for ”regular” anchor setup

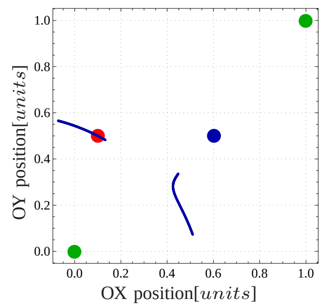

(a) Computed position trajectory

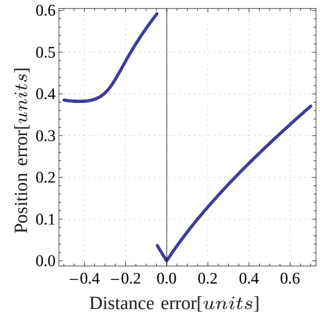

(b) Position errors vs. distance error

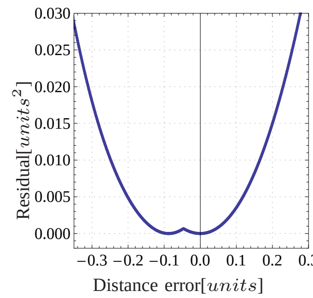

(c) Residual vs. distance error

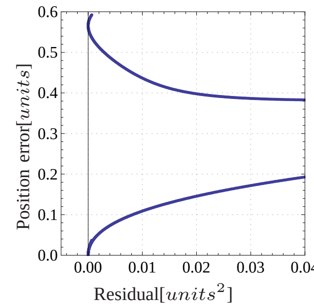

(d) Position error vs. residual

Fig. 3. Lateration behavior for ”hard” anchor setup

gray points in Fig.1(a). We expected a curve passing through the origin to validate the assumption that a residual of 0 corresponds to a position error of 0. The results in Fig.1(b) show that position error and residual value are not correlated as assumed.

Considering the same anchor scenario as before and a node placed at coordinates $(0.6; 0.5)$ (red dot in Fig.2(a)), we performed the following experiment: we kept the distances to two anchors constant and equal to the real distances (no noise) and only varied the distance towards the bottom-right anchor from Fig.2(a). The blue line in Fig.2(a) shows the set of resulting positions computed by the lateration procedure. Fig.2(b) shows the position error versus the induced distance error. For a distance error of 0 – when the true distance is fed to the algorithm – the real position is returned. In other cases as expected, the larger the distance error is, the larger the position error becomes. Fig.2(c) shows the value of the residual versus the induced distance error. The residual function exhibits two local minima: one corresponding to the real position (distance error equals zero) and one corresponding to the symmetrical position of the node with respect to the first two anchors (in green in Fig.2(a)).

III. Residual value behavior

Fig.2(d) shows the position error versus the residual value. This highly nonlinear graph shows two curves passing through the origin. The curves can be explained by the fact that in the residual expression the influence of the distance error enters always as a squared factor. Underestimating a distance or overestimating it leads to different behavior in the lateration procedure (as shown in Fig.2(b)). This leads to the semiplane $x \in (-\infty, 0)$ being folded over the semiplane $x \in (0, \infty)$. The zigzag shape of the upper curve is a direct consequence of the residual shape in Fig.2(c).

Non linearity is not the only problem at hand: discontinuities occur as well in the case of specific anchor topologies. For example, in Fig.3 we ran a similar experiment with the difference that we placed the three anchors in an almost collinear configuration. When varying the distance to the middle anchor, shown in blue in Fig.3(a), one can notice that the computed position ”jumps” at a certain moment over the line defined by the almost collinear anchors. The highly nonlinear behavior remains in the residual representation in Fig.3(c) and the nonlinear curves in the graphs in Fig.3(b) and Fig.3(d) exhibit discontinuities as well.

IV. Cramer-Rao Bound in localization

We repeated the experiment described in [9] to study the behavior of CRB (see Fig.4). Assume three anchors a_i and a node z placed at the origin of the system of coordinates. The anchor a_1 is fixed and a_2 and a_3 are rotating around z, while maintaining the same distance towards z. In Fig.4, the red and blue circles represent their trajectories. The goal is to compute the CRB for various angles ϕ_2 and ϕ_3 that a_2 and a_3 make with the horizontal axis in order to capture the effects of anchor geometry on localization error.

The results of this experiment are presented in Fig.5. The axes OX and OY represent the angles ϕ_2 and ϕ_3 and are graded directly in π. The spikes, rising to $+\infty$, represent discontinuities in CRB (the axis OZ was cropped). For improved clarity, we shifted the axes with 0.5π to clearly show the four discontinuities and we represented CRB rather than 1/CRB [9]. CRB shows discontinuities in the cases

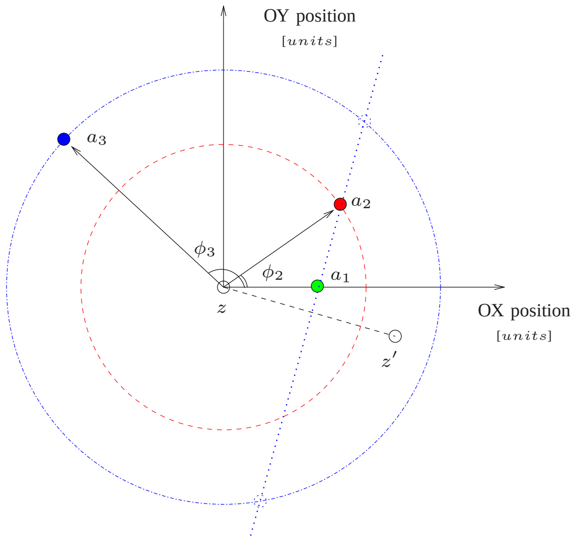

Fig. 4. Circular deployment of anchors

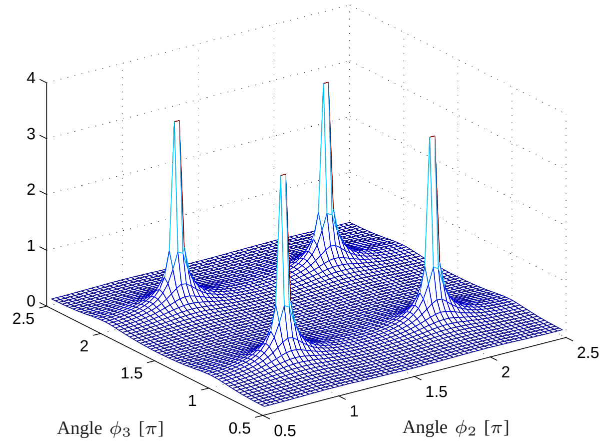

Fig. 5. Cramer-Rao Bound

where collinearity occurs, that is with the pairs of angles $(\phi_2, \phi_3) \in \{(\pi, \pi), (\pi, 2\pi), (2\pi, \pi), (2\pi, 2\pi)\}$.

There is a subtlety not described in [9]: the collinearity situation for which the CRB goes to infinity involves *the node as well* and not only the anchors. Although the findings we cited above intuitively seem correct *they are false*. In the case of three collinear anchors in the 2D plane shown as a blue dotted line in Fig.4 and given the distances between the node and the anchors $\overline{d_i}$, there is an uncertainty on which side of the line the node resides: at z or at its mirrored position z'. This uncertainty is not infinite. We define as a metric the *flipping uncertainty* equal to $\|z, z'\|$. The flipping uncertainty goes to zero when the node gets very close to the line determined by the collinear anchors. This implies that the only case in which a node can compute a position given a set of collinear anchors is when the node resides on the *same line* as the anchors. This contradicts the insight given by CRB.

Fig.4 allows us to make a second important observation. For every position of a_2, there exist two possible positions in which the line determined by a_1 and a_2 intersects the trajectory of a_3. This means that for every position a_2, there exist two possibilities to place a_3 such that all three anchors are collinear. The relationship between a_1, a_2 and a_3 basically reduces to the intersection of the line (a_1, a_2) with the circle centered at z having radius $\|z, a_3\|$. Graphically the result is presented in Fig.6. A point is placed in this picture for each situation in which the angles ϕ_2 and ϕ_3 lead to a collinear situation for the three anchors. The OZ coordinate measures the flipping uncertainty. Although Fig.6 might look similar to Fig.5, attention should be paid to the values on the OX and OY axis – for the points where CRB goes to infinity, the flipping uncertainty is actually 0.

The last two observations lead to the conclusion that CRB does not indicate all the troublesome anchor configurations. Worse than that, it indicates infinite uncertainty in the only situations in which positions can be computed, that is when all the anchors and the node are collinear.

As a final remark on CRB, it is worth noting that it gives a boundary on the *variance* rather than on the *mean value* of an unbiased estimator. This topic has been previously addressed as *accuracy versus precision* [11].

Previous work proposing CRB as a mechanism to determine the uncertainty of position computation shows that under certain assumptions the effects of anchors geometry can be computed as a separate coefficient, known as GDOP [7]. The expression of GDOP is actually directly derived from the expression of CRB under the following simplifying assumptions: the parameters of distance estimates (mean and variance) are considered equal thus the final formula takes only angles into account. Neither of these simplifications holds in sensor networks as opposed to a satellite system – mean values of distances of a node towards various anchors can be of different orders of magnitude and the variances are proportional to the actual distances. Furthermore, as CRB fails grasping the characteristics of the underlying geometric setup, GDOP is of little use in our case.

We conclude that while it makes sense to use CRB in the context of GPS where the distances between a node and the anchors are very large and the amount of error on the distances is assumed insignificant, the method cannot be applied to WSNs.

V. Towards an Explanation

We make abstraction of the exact procedure of computing a position. We assume the coordinates of the n anchors a_i are known. We place a node at a position z. Let f be the function that translates a_i and z into a set of distances $\overline{d_i}$. Let f^{-1} denote the function through which a_i and $\overline{d_i}$ are mapped back into the position z. The function f^{-1} denotes the perfect one-hop distance-based localization algorithm running with noiseless data. In practice, exact distances are not available. We thus feed into the function f^{-1} noisy measurements d_i and

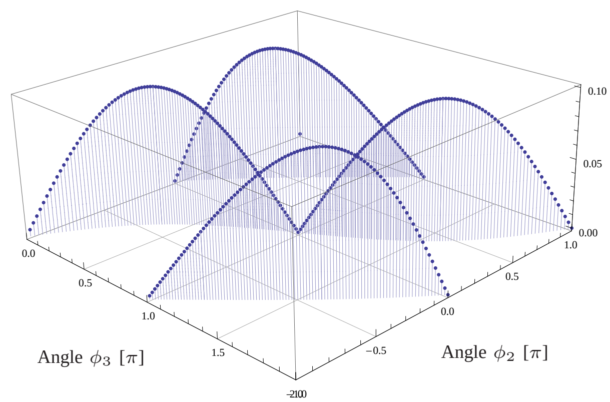

Fig. 6. Flipping uncertainty

receive a computed position $\tilde{z}$. The new distances between $\tilde{z}$ and a_i, written as $\tilde{d_i}$, slightly differ from d_i, compensating for the noise.

For the case of collinear anchors, given anchor positions a_i and distances $\overline{d_i}$, two positions z, z' can be computed (see Fig.4). Note that a_i and $\overline{d_i}$ alone do not offer enough information to distinguish between the two cases, thus f^{-1} is undefined (a function associates only one output value to a given input). The default definition of f^{-1} leads to undefined behavior over the input domain, even when *the real distances* $\overline{d_i}$ are available. This is one of the reasons why CRB does not hold: f^{-1} is assumed always defined and equal to the real position of the node. In order to account for the flipping uncertainty, triggered by anchor geometry *and* by measurement noise as in Fig.3(a), f^{-1} needs to be redefined to output a single value also in the case collinearity (or in the case of only one or two anchors being present).

Another observation we can make is that a ”stable” geometrical structure that leads to only one cluster of possible positions for a node in the presence of small amounts of noise, can become ”unstable” for larger amounts of noise. By unstable we understand that the resulting positions will be part of several geographic clusters (see Fig.3(a)). This means that there is a relation between the shape of the geometrical setup (positions of the node and anchors) and the amount of tolerable noise.

VI. Conclusions and Future Work

In this paper we showed that the one-hop distance-based localization mechanism has *geometry* as its foundation and not *measurement noise* as CRB assumes. The localization boundaries should thus be explored first from a geometrical perspective and then complemented by knowledge of the noise characteristics. We have supported this claim with a series of examples showing the limitations of the current approach.

Based on this argumentation, we propose a radical change in the way in which the localization problem is to be addressed:

- **Basic concept**: geometrical setup (positions of the anchors and of the node);
- **Theoretical steps**: define a function f^{-1} mapping anchor positions and distances to an estimated position; determine the geometrical boundaries for maximum allowed errors;
- **Algorithm**: from the formulation of geometrical boundaries determine the metric to be minimized to obtain a position (leading to a ”geometric” lateration procedure);
- **Post processing**: explore new metrics for the positioning error (euclidean distance is considered the default one – it is nevertheless a one-dimensional metric that cannot express all the characteristics of a higher-dimensional phenomenon).

This description is also the base for our future work, in which we wish to explore the trade-offs between geometrical setup and amount of noise as well as define a boundary under which noisy measurements will lead to a set of clustered positions. Our final target is to provide a clear formulation of the achievable boundaries of localization algorithms and to offer a new metric to be minimized taking the geometry of the setup and the measurement noise into account.

References

[1] G. Mao, B. Fidan, and B. Anderson, “Wireless sensor network localization techniques,” *Comput. Netw.*, vol. 51, no. 10, pp. 2529–2553, 2007.

[2] S. Gezici, “A survey on wireless position estimation,” *Wirel. Pers. Commun.*, vol. 44, no. 3, pp. 263–282, 2008.

[3] N. Patwari and al., “Locating the nodes: cooperative localization in wireless sensor networks,” *IEEE Signal Process. Mag.*, vol. 22, no. 4, pp. 54–69, 2005.

[4] L. Hu and D. Evans, “Localization for mobile sensor networks,” in *Proc. MobiCom 2004*. New York, NY, USA: ACM, 2004, pp. 45–57.

[5] R. Bischoff and R. Wattenhofer, “Analyzing connectivity-based multihop ad-hoc positioning,” in *Proc. PerCom 2004*, 2004, pp. 165–174.

[6] D. Niculescu and B. Nath, “Ad hoc positioning system (aps),” in *Proc. IEEE Global Telecommunications Conference GLOBECOM ’01*, B. Nath, Ed., vol. 5, 2001, pp. 2926–2931 vol.5.

[7] A. Savvides, W. Garber, R. Moses, and M. Srivastava, “An analysis of error inducing parameters in multihop sensor node localization,” *IEEE Trans. Mobile Comput.*, vol. 4, no. 6, pp. 567–577, 2005.

[8] D. Torrieri, “Statistical theory of passive location systems,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-20, no. 2, pp. 183–198, 1984.

[9] M. Spirito, “On the accuracy of cellular mobile station location estimation,” *IEEE Trans. Veh. Technol.*, vol. 50, no. 3, pp. 674–685, 2001.

[10] J. Zhu, “Calculation of geometric dilution of precision,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 28, no. 3, pp. 893–895, 1992.

[11] K. Langendoen and N. Reijers, “Distributed localization in wireless sensor networks: a quantitative comparison,” *Comput. Netw.*, vol. 43, no. 4, pp. 499–518, 2003.