Tradeoff Between Estimation Performance and Sensor Usage in Distributed Localization Problems

Juan Liu and Qingfeng Huang
Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, CA94304, USA.
email: {qingfeng.huang, juan.liu }@parc.com

Abstract: In distributed sensor networks, the tradeoff between application performance and resource consumption is a fundamental issue. In this paper, we analyze the tradeoff in the context of localization problems. In particular, the localization performance is characterized using the Cramer-Rao lower bound which reflects the performance limit in the family of unbiased estimators. For a variety of distance-sensitive sensing models, the bound comes in closed-form. In view of the tradeoff, we design a circular incremental inclusion scheme for efficient sensor tasking and propose to use the theoretical Cramer-Rao lower bound as a guiding heuristic. Our results reveal that there is a performance limit for collaborative sensing, and the expansion of sensing group size incurs fast diminishing returns on performance. The study results in new insights and provides general guidelines for sensor tasking in practical applications.

Keywords: Cramer-Rao lower bound, localization, sensor networks, performance analysis.

Biographical notes: Juan Liu received her Ph.D. degree in Electrical Engineering from the University of Illinois at Urbana-Champaign in 2001. She is now a researcher in Palo Alto Research Center (PARC). Her research interests include signal processing, statistical modeling, detection and estimation, network routing, and their applications to distributed sensor network problems.
Qingfeng Huang received his D.Sc. degree in Computer Science from Washington University in St. Louis in August 2003. He has published papers in multiple areas including mobile computing, sensor networks, intelligent transportation systems, neuroscience, and quantum physics. His current interest include algorithms and middleware for mobile computing and sensor networks. He is currently a research scientist at Palo Alto Research Center.

1 Introduction

Distributed sensor networks are often severely constrained in resources such as on-board battery and communication bandwidth. How to design an application achieving good performance without over-tasking the sensors poses a significant challenge. In this paper, we focus on localization problem and investigate the tradeoff between localization performance and resource consumptions. Using localization as an example, we seek to answer questions such as the following:

- How should one characterize the application performance? Is there a limit on performance?
- How do sensors contribute to the application performance? And what is the cost incurred to task these sensors?
- Given a sensor tasking scenario, how does the performance vary with resource consumption?
- Can one design a good sensor tasking strategy to optimize application performance at moderate resource costs?

These problems are fundamental in understanding the tradeoff, and will be helpful to the design of scalable distributed sensing systems.

We first briefly introduce the localization problem, which is a canonical example of sensor network applications. Assume a signal source is in a sensor field and its signal is detected by many sensors. The signal source location $\mathbf{l_t} = (x_t, y_t)$ is unknown. The goal is to obtain an estimate $\hat{\mathbf{l}}_t$ as close as possible to the true $\mathbf{l_t}$ based on the sensor measurements. This localization problem often occurs in target tracking, where the signal source is a moving target. Or it can be used in node location service, in which the signal sources are the sensor nodes sending out beacons for localization. In both cases, we refer to the signal source as target. For this problem, there is a body of literature on how localization can be achieved. A variety of algorithms have been proposed, covering a wide spectrum from deterministic to statistical, from simple linear methods to sophisticated optimization approaches. A few representative examples include the distributed least-square based method [1], sequential Bayesian method [2], and node localization through multi-lateration[3]. In this paper, we do not focus on developing any detailed localization algorithm, but spend our effort on analyzing general performance limits. In particular, the focus is on the analysis of cost (in terms of communication) and benefit (in terms of performance improvement) of tasking individual or group of sensors.

Assessing the exact performance of an estimator may be difficult; extensive simulations may be involved. In this paper, we use Cramer-Rao lower bound (CRLB) to characterize localization performance. CRLB represents the best estimation performance in the family of unbiased estimators, and is asymptotically achievable [4]. In essence, it reflects how easy or difficult it is to estimate an unknown parameter (signal source location in this case) from a set of observations. CRLB is a general metric; it is independent on the particular algorithm that is used, hence allows us to characterize performance without worrying about algorithmic details. Related work in [3] also uses CRLB for analyzing localization performance and has looked into the effect of node geometry and density on localization results. In this paper, we treat node layout as random (unlike the deterministic assumption in [3]) in return for more general applicability. We also explain the diminishing performance improvement observed in many localization experiments: as more sensors are tasked, localization performance is only marginally improved. In this paper, we justify this intuition analytically and quantify the sensor contribution. This provides a mechanism to study the performance-cost tradeoff. Our analysis demonstrates that given a sensor network with a certain density, depending on the sensing model, there may exist a limit for localization accuracy. Furthermore, the return on (communication and computation) investment (ROI) diminishes as performance approaches the limit. This result provides guideline for sensor tasking: querying sensors at increasingly higher communication cost will not improve the performance beyond the asymptotic limit. This limit analysis is novel.

The paper is organized as follows. Sec. 2 quantifies individual sensor contribution to localization accuracy. For two commonly used sensing models, namely range sensors and directional sensors, the contribution, measured as Fisher information (FI), is computed. Sec. 3 investigates the collective performance of multiple sensors. This measures the limit of collaborative sensing. Based on the heuristics gained from the CRLB analysis, we design in Sec. 4 a dynamic circular incremental inclusion (CII) scheme for sensor tasking. Sensor groups are formed based on distance using a geographical multicast communication scheme. Sec. 5 explores more implications of CRLB results and the performance-cost tradeoff. Sec. 6 provides some discussion about the applicability of the techniques illustrated in this paper. Sec. 7 concludes the paper.

2 CRLB for individual sensors

Assume that a target is located at $\mathbf{l_t}$. A sensor measurement z with observational $p(z|\mathbf{l_t})$ conveys some information about the target location. To quantify the contribution, we use the Fisher information matrix (FIM) defined as follows:

$$I = E_{p(z|\mathbf{l_t})}\left[(\nabla_{\mathbf{l_t}} \log p(z|\mathbf{l_t})) \cdot (\nabla_{\mathbf{l_t}} \log p(z|\mathbf{l_t}))^T\right] \qquad (1)$$

As Fisher information is expressed in terms of the gradient (and implicitly the curvature) of the observational model $p(z|\mathbf{l_t})$ with respect to the unknown parameter, it reflects the model sensitivity. The more sensitive the observation model, the larger the Fisher information, hence the more accurate an estimator may get.

The Cramer-Rao lower bound theorem states that the covariance of any unbiased estimator is lower-bounded by the inverse of the Fisher information matrix [4], i.e.,

$$Cov(\hat{\mathbf{l}}_t(\mathbf{z})) \geq I^{-1}. \qquad (2)$$

The notation of $A \geq B$ means that the matrix $A-B$ is positive semi-definite. The right hand side of (2) is known as the CRLB. It characterizes the best possible performance in the unbiased estimator family, and is asymptotically achievable. Note that CRLB is independent of any particular estimation schemes, but rather reflects the performance limit imposed by the observational model $p(z|\mathbf{l_t})$ alone. In other words, it shows how easy or difficult it is to estimate a set of parameters from the observation z.

2.1 Sensor Models

Now we introduce a generic sensing model. For a sensor located at $\mathbf{l} = (x, y)$, we assume that its measurement of signal from $\mathbf{l}_t$ takes the form

$$z = f(\mathbf{l}_t, \mathbf{l}) + n. \qquad (3)$$

The function $f(\mathbf{l}_t, \mathbf{l})$ is a scalar function representing the measurement in the noiseless case. In practical situations, it is contaminated by some noise n. For simplicity of illustration, we assume that the noise is independently Gaussian distributed with zero mean and variance $\sigma^2(\mathbf{l}_t, \mathbf{l})$. The analysis in this paper can be extended to non-Gaussian or correlated noise.

The sensor measurement form (3) is very general. It covers a large variety of sensing models. Here we list two commonly seen examples. From this point on-wards, we

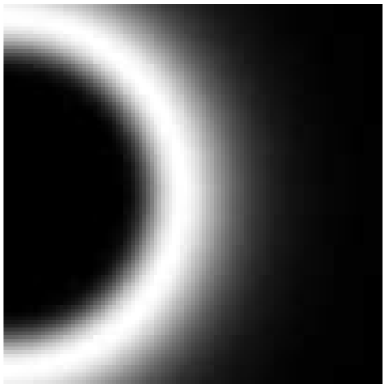

(a) range sensor

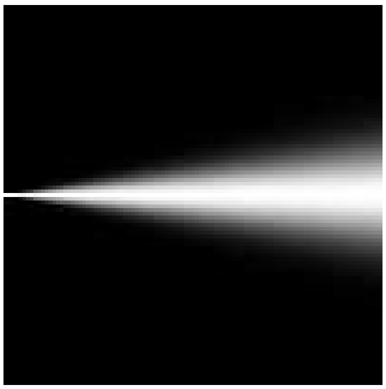

(b) direction sensor

Figure 1: Observational model.

use the notation $r = |\mathbf{l_t} - \mathbf{l}|$ for the distance between the sensor and the target, and $\Delta x = x_t - x$ and $\Delta y = y_t - y$ for the displacement in x and y directions respectively.

1. Range sensors. These sensors infer about target distance based on the received signal strength (RSS). Assuming that the target is a point signal source emitting constant energy and the signal propagation is isotropic, the RSS takes the approximate form

$$f(\mathbf{l}_t, \mathbf{l}) = A \cdot r^{-\gamma}, \tag{4}$$

   where A is a scaling constant related to the source signal strength, and γ is the attenuation factor. The perturbation of range sensors are normally due to background noise, hence $\sigma^2(\mathbf{l}_t, \mathbf{l})$ is often modeled as a constant σ^2.

   The observational likelihood function $p(z|\mathbf{l}_t)$ is plotted in Fig. 1a. Pixel locations denote $\mathbf{l}_t$ on a 2-D plane, and graylevels show the likelihood, with bright pixels corresponding to high likelihood values. Given the RSS, the sensor can infer the target location as a ring around itself with some uncertainty.

2. Direction sensors. These sensors measure the direction where the signal comes from, i.e., the bearing of the target. For example, small microphone array can use beamforming techniques to compare the phase of perceived signal at each microphone and determine the incoming angle. For this type of sensors, the measurement in the ideal case is

$$f(\mathbf{l}_t, \mathbf{l}) = atan\left(\frac{\Delta y}{\Delta x}\right). \tag{5}$$

   In practice, the observation of z may deviate from the ideal value (5) due to the presence of noise in the perceived waveform and the imperfection of the beamforming algorithm. The deviation between z and (5) is denoted by n. We model n as Gaussian when its variance is small. Chen et al [5] have characterized the error of estimating direction using microphone arrays. The estimation error grows monotonically with distance. Following this observation, we assume that the variance of n, denoted as $\sigma^2(\mathbf{l}_t, \mathbf{l})$, is an increasing function of distance r. For this type of sensors, the observation likelihood function $p(z|\mathbf{l}_t)$ is shaped like a cone, as plotted in Fig. 1b.

2.2 Fisher Information of Individual Sensors

Given the sensor models, one can quantify the individual sensor contribution based on the FIM. We present the main results without detailed derivations.

Result A. Under the general sensing model (3), the Fisher information matrix takes the form

$$I = \frac{1}{\sigma^2}\begin{bmatrix} \frac{\partial f}{\partial x_t} \\ \frac{\partial f}{\partial y_t} \end{bmatrix} \cdot \begin{bmatrix} \frac{\partial f}{\partial x_t} & \frac{\partial f}{\partial y_t} \end{bmatrix} + \frac{2}{\sigma^2}\begin{bmatrix} \frac{\partial \sigma}{\partial x_t} \\ \frac{\partial \sigma}{\partial y_t} \end{bmatrix} \cdot \begin{bmatrix} \frac{\partial \sigma}{\partial x_t} & \frac{\partial \sigma}{\partial y_t} \end{bmatrix}. \tag{6}$$

The first term is due to the variation of perceived signal with respect to target location. The second term is due to the variation of the noise variance.

This result can be instantiated for more specific sensing models such as the range sensors and the directional sensors.

Result B. For range sensors (4) with constant background noise variance σ^2, the FIM is

$$I = \frac{A^2\gamma^2}{\sigma^2 r^{2\gamma+4}} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} \cdot \begin{bmatrix} \Delta x & \Delta y \end{bmatrix} \tag{7}$$

Result C. For directional sensors (5) with noise standard deviation $\sigma = \sigma(r)$ being an increasing function with distance r, the Fisher information matrix is

$$\begin{aligned} I &= \frac{1}{\sigma^2 r^4}\begin{bmatrix} -\Delta y \\ \Delta x \end{bmatrix} \cdot \begin{bmatrix} -\Delta y & \Delta x \end{bmatrix} \\ &+ \frac{2(\sigma'(r))^2}{r^2\sigma^2}\begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} \cdot \begin{bmatrix} \Delta x & \Delta y \end{bmatrix}, \end{aligned} \tag{8}$$

where $\sigma'(r)$ denotes the derivative of $\sigma(r)$ with respect to r. Note in the results above, Fisher information I is a function of the displacement Δx and Δy. This is because the sensing models are functions of displacements only. Given the sensor location, the contribution can be evaluated straightforwardly.

3 Collective performance of sensors in a region

Now consider a homogeneous sensor network consisting of K identical sensors located at $\mathbf{l}_1, \mathbf{l}_2, \cdots, \mathbf{l}_K$ respectively, with observational model is specified by (3). The localization problem can be formulated as estimating the target location $\mathbf{l}_t$ giving the sensor measurements $\mathbf{z} = \{z_1, z_2, \cdots, z_K\}$. In distributed sensor networks, it is often valid to assume that the measurements across sensors are conditionally independent given the target location, i.e., $p(\mathbf{z}|\mathbf{l_t})) = \prod_i p(z_i|\mathbf{l_t}))$. Under this conditional independence assumption, it is easy to show the Fisher information can be broken down into summation form, i.e.,

$$I = E_{p(\mathbf{z}|\mathbf{l_t})}\left[(\nabla_{\mathbf{l_t}} \log p(\mathbf{z}|\mathbf{l_t})) \cdot (\nabla_{\mathbf{l_t}} \log p(\mathbf{z}|\mathbf{l_t}))^T\right] \quad (9)$$

$$= I_1 + I_2 + \cdots + I_K \quad (10)$$

where I_i is the Fisher information from sensor i.

3.1 FIM of Sensors in a Ring

In collaborative sensing scenarios, it is often of interest to consider the contribution of a region of sensors, rather than individual ones. We have shown that the overall contribution of a sensor region is the summation of all sensors in that region. One may not know the exact sensor locations, but only some general characteristics of how the sensors are located, such as the sensor density. In this case, we may compute the total Fisher information as

$$I_\Omega = \int_{(x,y)\in\Omega} \rho I(x,y) dxdy, \quad (11)$$

where ρ denotes sensor density in Ω. $I(x,y)$ is individual sensor contribution if the sensor is located at (x,y). It is the same as I_i in the previous section. The integral I_Ω reflects the *expected* contribution that may be gained by querying the region Ω if the sensors are uniformly distributed.

We are particularly interested in the total information contribution from the sensors in a circular region or a ring centered at the expected target location, because it will give insight into the efficient planning of sensor tasking without using detailed topology knowledge, and is directly related to a circular incremental inclusion (CII) sensor tasking scheme we will present later. Without loss of generality, consider the case where the target is located at (0,0), and sensors are uniformly distributed around the origin.

Given the sensing models (7) and (8), we can derive the following results of the Fisher information matrix for range sensors and direction sensors in a ring.

3.2 Results for range sensors

Result D. For range sensors (4) with constant background Gaussian noise variance σ^2, the contribution of $\Omega = \{(x,y) : R_0 \le \sqrt{x^2+y^2} \le R_1\}$ has the diagonal form

$$I_\Omega = \alpha \cdot \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad (12)$$

and the scaling coefficient is

$$\alpha = \frac{\pi\rho\gamma A^2}{2\sigma^2}(R_0^{-2\gamma} - R_1^{-2\gamma}). \quad (13)$$

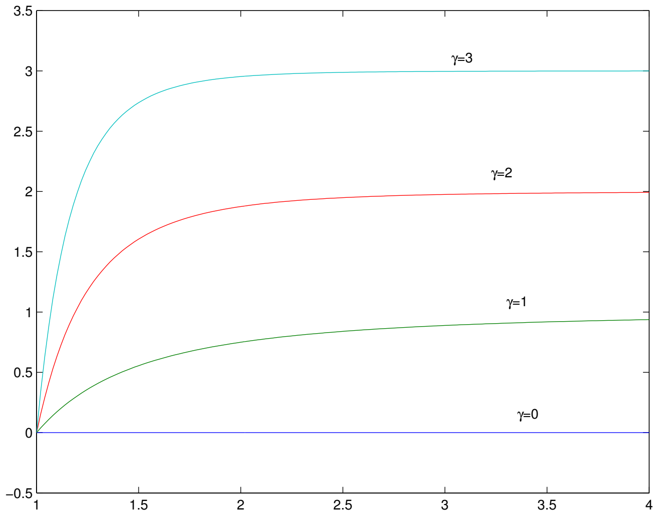

Figure 2: FI (scalar coeff. α) as a function of R_1.

Note that (12) has the nice identity matrix form. It indicates that the errors in x and y direction have the same characteristics. This is due to the symmetry of the sensing model (4) and the sensor region Ω. In localization problems, mean-squared error (MSE) $\hat{x}_t^2 + \hat{y}_t^2$ is often of interest. On average it is the same as the trace of the covariance matrix, hence lower bounded by $2/\alpha$. The higher α is, the better the localization accuracy.

In (13), α is a simple analytic function of R_0 and R_1. Note the following properties:

- The Fisher information is proportional to the sensor density ρ. This implies that higher sensor density will improve localization accuracy monotonically.
- Fixing the inner radius R_0, it is an increasing function of R_1. This implies that with more sensors tasked, the localization accuracy should improve monotonically.
- The rate that α increases with R_1 and the point it saturates depend on the attenuation coefficient γ. Fig. 2 plots the function $\alpha(R_1)$ for different values of γ. The curve increases slowly with small γ and saturates quickly with large γ. The implication on localization is that if the signal does not attenuate too quickly, it pays to include further away sensors. Otherwise, for large γ, the gain is really marginal.

Localization systems based on range sensors have been built and presented in details for example in [6]. Often

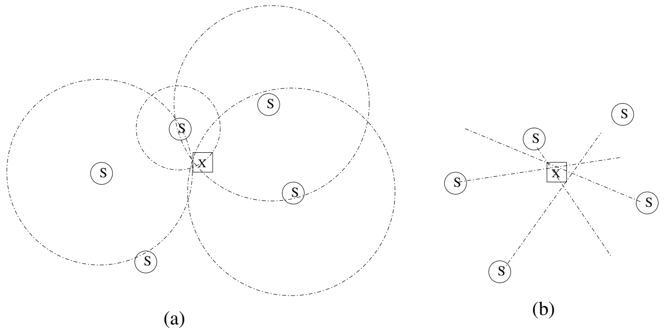

Figure 3: Localization using sensors.

a triangulation type of techniques to infer target location based on their distance measurements, as depicted in Fig. 3a. Our result essentially points out how the ultimate achievable estimation accuracy changes with the system parameters such as the network density, the measurement noise, and the single sensor models, without any artifact of a particular technique.

3.3 Results for direction sensors

A sensor network consisting of direction sensors identifies target location based on line intersection, as illustrated in Fig. 3b. Similar to the range sensor case, we compute the Fisher information over a ring Ω. For the purpose of computing CRLB, any functional form of $\sigma(r)$ can be used. For simplicity of illustration, we assume the noise standard deviation $\sigma(r) = r^{\beta}$, where β is some positive number. The Gaussian assumption of noise n is appropriate in the near field, and less accurate in the far field, since beamforming techniques break down at very low SNR, and the perturbation needs to be modeled differently.

Result E. For direction sensors (5) with noise stand deviation $\sigma(r) = r^{\beta}$, the contribution of $\Omega = \{(x,y) : R_0 \leq \sqrt{x^2+y^2} \leq R_1\}$ has the diagonal form

$$I_{\Omega} = \eta \cdot \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \qquad (14)$$

and the scaling coefficient is

$$\eta = \frac{\pi\rho}{2\beta}(R_0^{-2\beta} - R_1^{-2\beta}) + 2\pi\rho\beta^2 \ln(R_1/R_0). \qquad (15)$$

Fig. 4 plots η as a function of R_1. Note that (15) has two terms: the first term is the contribution due to the variation of signal (5) with respect to target location, and the second term is the contribution due to the variation in $\sigma(r)$. For small R_1, η is dominated by the first term, and increases roughly at the rate of $-R_1^{-2\beta}$. This is similar as in the range sensor case. For large R_1, the second term starts to take effect, and η increases at the rate of $\ln R_1$. This seems to indicate that comparing to the range sensor case, it pays more to tasking direction sensors further away.

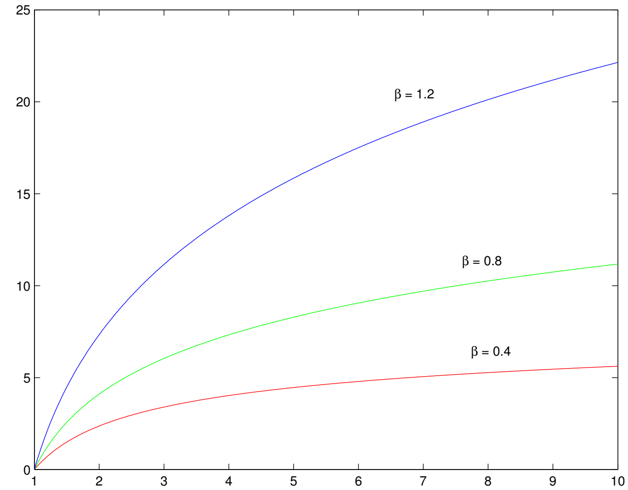

Figure 4: CRLB (scalar coeff. η) as a function of R_1.

3.4 Result for inhomogeneous sensor networks

In practice, heterogeneous sensor networks can be used to better exploit sensing diversity. For example, one may construct a sensor network consisting of both range and direction sensors for localization purpose. To characterize the performance of such inhomogeneous sensor network, CRLB is particularly convenient, since Fisher information can be directly added under the conditional independence assumptions. The results above can be easily combined. For example, for a uniformly distributed sensor network with a certain percentage (p) of range sensors and $1-p$ percentage of direction sensors, the overall contribution of a region Ω is the linear combination of contributions as in (13) and (15), weighted by p and $1-p$ respectively.

---

4 The CII sensor tasking scheme

Results D and E show quantitatively the benefit of collaborative sensing in the context of localization, i.e., how the potential sensing accuracy, represented by FIM (or the CRLB), improves when more and more sensors are included in the sensing group. In this section we discuss how the results help us devise efficient sensor tasking schemes.

4.1 Sensor tasking scheme

To achieve the benefit of collaborative sensing, a sensor group need to be formed and often the distributed sensor data in the sensing group need to be aggregated/pooled to some node in the group for computation. Oftentimes a leader based scheme is used[2]. In such schemes, a leader is either designated or elected to coordinate and compute the collective estimation at each time step, and the improved estimation result can be distributed to interested parties. For simplicity, we do not consider leader election in this paper. But rather, we consider the scenario where a leader has been designated for each data aggregation step and a major task for the leader is to determine and form the sensor group need in an efficient manner, i.e., to determine what other sensors need to contribute their sensed data to the computation and how to include them at each time step. The goal is to achieve certain level of collaborative sensing accuracy with minimum cost.

We adopt an incremental inclusion collaborative sensing strategy. In this strategy, the leader selects one or a few candidate sensor nodes for contributing to the refinement of the sensing result at each step, and stops if the sensing accuracy has reached some prescribed threshold or no eligible sensor is left or time is running out. Fig. 5

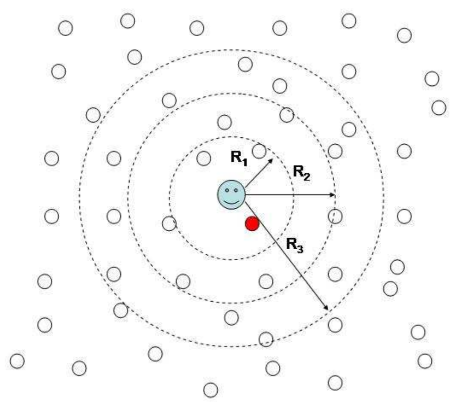

Figure 5: Sensing Group Formation

shows a simple circular incremental inclusion (CII) scheme, where the leader first include the nodes in the circle with radius R_1 in the group. If the leader finds that more inclusion are necessary, it include all nodes within radius R_2 of the estimated target location. This expansion scheme repeats until the stop criteria is met. Note that when all the step size is sufficiently small, the CII scheme becomes a closest distance based one-by-one selection scheme. This CII scheme also admits a more economic communication strategy than the node-by-node inclusion scheme, such as the ones discussed in[7][8]. Note that the scheme to include them one by one requires that the leader know the unicast address and location of each individual node beforehand. This entails bad scalability. In a dynamic environment where the leader role may shift from one node to another[2], the one-by-one inclusion scheme essentially requires everyone knows everyone's network address and location, and requires the network address to be unique as well. On the other hand, the scheme to include the geographically related groups one by one admits a much more scalable solution. The leader can use geographical multicast for requesting the data from sensors in each expansion ring, and does not need to know the sensors' locations and network addresses beforehand. In this sense, the group based inclusion scheme is much more scalable than node-by-node inclusion.

Note that in the CII scheme, each round of expansion entails one round of communication and computation, and means additional cost on both bandwidth and time. Note that the time it takes for an estimation to finish directly affects the accuracy of the estimation if the source location is changing over time. So one would like to reduce the number of rounds to a minimum if possible. In an extreme, one could just one round by including all nodes in the network, but that would entail huge communication cost and defeat the goal of scalability. The "ideal" scenario is that the first joint sensing radius[1] chosen by the leader covers exactly the minimum number of nodes needed for achieving the prescribed accuracy requirement. However, this is a chicken and egg problem, in general it is impossible for the leader to tell this minimum joint sensing radius before the data has been gathered and the computation is done. So instead, we are looking for good heuristics for reducing the number of expansions needed for achieving specific accuracy requirement.

4.2 CRLB as a heuristic for the CII tasking strategy

We propose to use the CRLB results as a heuristic for the leader to find the approximately "right" number of sensors that it initially queries, i.e., to determine the initial joint sensing radius in the CII scheme for forming the initial collaborative sensing group. Note that the CRLB we derived describes how the ultimate expected quality of collaborative localization (of the corresponding sensing models) changes over the joint sensing radius in a uniformly distributed sensor network. Given a sensing precision requirement such as the min variance, the leader can compute the expected joint sensing radius R_1 for achieving the accuracy.

However, as a theoretical lower bound, CRLB only gives an indication of the ultimate achievable performance. In practice, with finite amount of data, the estimators might not able to achieve the CRLB. It is thus of interest to see how well the general trend suggested by the CRLB corresponds with the experimental performance of real estimators. To verify, we use the sequential Bayesian filtering method described in [2]. For the case where the target is stationary, it is the same as a maximum likelihood (ML) estimator. For simulation, we use a sensor field of dimension 300×300 (unit distance)^2, with 150 range sensors uniformly distributed. The target is located at the center of the sensor field. The observational model in this simulation assumes $\gamma = 1$. To localize the target, the sensors are queried incrementally. The query starts at distance $R_0 = 40$, and progressively includes further away sensors. Fig. 6 shows the MSE averaged over 100 runs (the blue curve) as a function of querying distance R_1. The curve drops as R_1 increases and quickly stabilizes. The red smooth curve is the analytical Cramer-Rao bound. The horizontal axis is R_1/R_0. One can see that the trend of

[1]For convenience, henceforth we call the radius of the inclusion circle the "joint sensing group radius" or simply "joint sensing radius".

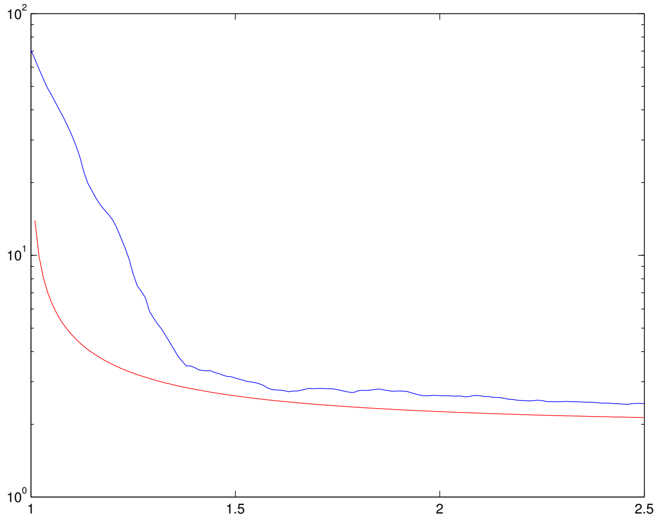

Figure 6: Experimental localization performance.

the MSE improvement agrees well with the CRLB trend, and the experimental value gets closer and closer to theoretical bound as more and more nodes are included. The relatively large difference between 1 and 1.4 is partly due to the limited sampling and the mismatch of the discrete reality of the sensor locations and the continuous nature of the model, as we mentioned earlier. Similar performance curves have been observed in localization problems with various settings, for example see [3].

The experiment result shows that as a theoretical the CRLB can indeed serve as a base for the predication of potential improvement that an expansion can achieve, and helps to choose a good expansion size, in the sense of including the approximately "just right number of" sensors for an accuracy requirement. However, this should be augmented by the dynamic expansion strategy in case of a relatively large discrepancy between the predication and the reality. This is exactly what the design philosophy behind our CII scheme.

5 Tradeoff between CRLB and resource usage

In this section we further analyze the practical implication of the derived FIM and CRLB results.

5.1 Performance limit of collaborative localization

Equation (13) implies that there is a performance limit of collaborative localization by the range sensors. We can see this more clearly by rewriting (13) in the follow form:

$$\alpha = \frac{\rho\pi\gamma A^2}{2\sigma^2 R_0^{2\gamma}}[1-(R_0/R_1)^{2\gamma}] \tag{16}$$

Let R_1 be infinity. α is simplified to

$$\alpha(R_0)_{max} = \frac{\rho\pi\gamma A^2}{2\sigma^2 R_0^{2\gamma}} \tag{17}$$

which is the max information gain possible by including all sensors outsize R_0 and it is finite! Even though this theoretical result indicates that as R_0 goes to 0, the Fisher information goes to infinity, meaning the CRLB can be arbitrarily small, this is partly an artifact of the simple mathematical abstraction. First, the sensors are discretely distributed in the space and not continuous as the density abstraction implies. In reality, the likelihood of a sensor at exactly or infinity close to the target location is very small. Furthermore, the simple sensing model (4) implies the measurement of the sensor can go to infinity as its distance to target goes to zero, this not true in practice either, since all sensors have a finite dynamic range. Without loss of generality, we can address these concern by assuming a minimum non-zero distance R_0 between the sensors and the target. In turn, we see that this implies a bounded collaborative performance gain dictated by (17). One important result is that it gives us a sense of what a given sensor network cannot do, more specifically, what sensing precision it is unlikely to achieve. One way to estimate the R_0 (in turn the α_{max}) is to use the average density knowledge. Let $\frac{1}{\pi R_0^2} \sim \rho$, we have $R_0 \sim 1/\sqrt{\pi\rho}$. This leads to

$$\alpha_{max} \sim \frac{\rho^{\gamma+1}\pi^{\gamma+1}\gamma A^2}{2\sigma^2} \tag{18}$$

In this way, we can see that the performance limit of networked sensing is ultimately determined by the single sensor model parameters (e.g. the noise variance and the signal attenuation factor in this case) and the density of sensor network. We can also clearly see that increasing node density is a more efficient way of improving the overall localization quality for sensing modalities with a higher attenuation factor γ.

5.2 Return on investment

As we pointed out earlier, behind the potential performance gain by collaborative sensing is the cost of communication and computation. In this section we analyze the tradeoff between them. Note that the exact cost of communication and computation depends on the actual estimator and communication protocol being used. In order to achieve certain level of generality in our analysis, we focus on estimating the cost of communication for the class of collaborative sensing strategies using the CII dynamic sensing group formation scheme. The metric we use to characterize the communication cost is the number of successful media access (successful "send") by the sensor nodes in the CII scheme. In general, the cost of each CII ring expansion can be represented by the following quantity:

$$ComCost(R_i, R_{i+1}) = f(R_i) + \rho\pi(R_{i+1}^2 - R_i^2)(1+\beta) \tag{19}$$

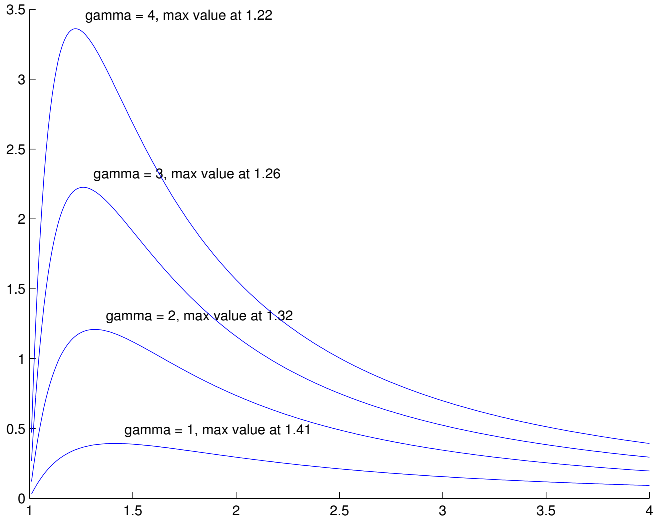

Figure 7: Fisher information gain per communication

The first term represents the cost of routing by the nodes inside the circle of radius R_i for the "join" request of the leader to the ring between R_i and R_{i+1} and for routing the corresponding sensing values back to the leader. The second term represents the cost of flooding the join request in the ring and pushing the sensing values to router nodes inside R_i. β is a parameter characterizing the flooding and routing efficiency in the ring.

In the case of using simple geo-constrained flooding for sending the leader query and the "wait-for-children" type converge-casting for sending the query answers back to the leader, $f(R_i) = 2\rho\pi R_i^2$ (each node has one transmission for query and one transmission for query answer[2]) and $\beta \sim 1$ (each node in the ring forwards the query once), and the total cost for expansion is $ComCost(R_i, R_{i+1}) = 2\rho\pi R_{i+1}^2$. Fig. 7 shows the trend of benefit (in terms of Fisher information gain) per communication (in terms of a successful message transmission) of such schemes. The vertical axis is the average Fisher information, The horizontal axis is R_{i+1}/R_i. Note that the return on investment (ROI) peaked around 20∼40% joint sensing radius expansion in the scenarios studied. When the radius expansion is less than 20%, the cost of routing the request and the replies dominates, so the ROI is low. When more nodes are included in one expansion, the cost of routing is amortized more, thus the ROI increases. However, when the expansion is beyond 40%, the ROI decreases since the information quality of the nodes further away are relatively lower. This ROI peaking phenomena can serve as a base for devising efficient sensing group expansion strategies, especially on how to control each expansion steps in the CII scheme to get maximum benefit with relatively low cost. Note that this observation is also valid when one use more efficient routing mechanism than the simple constrained flooding and converge-cast. The higher routing efficiency is expected to shift the peak location closer to 1.

[2]Note that a more accurate cost model is to take into account of the aggregate data size each node sends. We leave this to future work.

6 Discussion

CRLB analysis is a general tool, applicable to a wide variety of estimation problems. Although in this paper, our illustration is restricted to two sensing modalities (range and direction sensor) with Gaussian noise models, the CRLB analysis is by no means limited to these cases only. Other sensing modalities can be easily incorporated, as long as we know the observation likelihood $p(z|\mathbf{l}_t)$ in differentiable form. This enables the computation of single sensor CRLB. Note that single sensor CRLB itself does not have to be in closed-form. In a localization scenario with multiple sensors, the contribution of each single sensor CRLB can be combined easily (by simple addition or integration) to analyze the overall localization performance, very much like the technique illustrated in this paper. For cases where single sensor CRLB is not in closed-form, numerical integration techniques can be used. Similarly, the noise model is not restricted to Gaussian. The analysis can be easily extended to other distributions. In our examples, we assume sensor noise is independent. If noise is correlated over time, we can modify the single sensor CRLB to take into account the correlation, by treating measurement z as a time-series vector. If the noise is correlated across sensors, the computation will be more complicated, since the conditional independent observation assumption falls apart, and the contribution from individual sensors cannot be directly added or integrated, as is done in this paper. One would need to find a suitable approximation to avoid double-counting, i.e., counting the information redundancy between sensors multiple times.

For simplicity and clarity, we only investigated single target case in this paper. In the case of co-presence of many targets, the analysis is more difficult. For instance, the CII scheme might need to have some additional stopping criteria, e.g., stop at the boundary of Voronoi partition (in the Voronoi diagram of the targets) to minimizes cross-target interference. These are to be explored in our future work.

7 Conclusion

In this paper we studied the fundamental limit of collaborative sensing and the tradeoff between the potential collaborative performance gain and the associated cost of resource consumption. We explored this tradeoff in the target localization context, and analyzed the localization performance gained from collaborative sensing using the Cramer-Rao Lower Bound. We derived closed-form Fisher information matrices and the Cramer-Rao lower bounds for two distance sensitive sensing models and those of multiple sensors uniformly distributed in a ring centered at the target. We also proposed a generic distance-based circular incremental inclusion scheme for dynamic sensor tasking and to use our analytical results as heuristics for improving the efficiency of sensor tasking. Furthermore, we found that there was a performance limit for collaborative sensing ultimately determined by the network density, the measurement noise and the signal attenuation factor. Finally, we also found that the geographical expansion of sensing group size exhibits fast diminishing returns on performance gain in the studied sensor models. This study results in new insights in and provides general guidelines for sensor tasking in practical applications.

REFERENCES

1 Li, D., Wong, K., Hu, Y.H., Sayeed, A.: Detection, classification, and tracking of targets. IEEE Signal Processing Magazine (2002) 17–29

2 Liu, J., Reich, J.E., Zhao, F.: Collaborative in-network processing for target tracking. EURASIP, Journal on Applied Signal Processing **2003** (2003) 378–391

3 Savvides, A., Garber, W., Adlakha, S., Moses, R., Srivastava, M.: On the error characteristics of multihop node localization in ad hoc sensor networks. In: Information Processing in Sensor Networks, Proc. of IPSN 2003. (2003)

4 Poor, V.: An Introduction to Signal Detection and Estimation, *second edition*. New York, NY: Springer-Verlag (1994)

5 Chen, J., Hudson, R., Yao, K.: Maximum-likelihood source localization and unknown sensor location estimation for wideband signals in the near field. IEEE Trans. Signal Processing **50** (2002) 1843–1854

6 Sheng, X., Hu, Y.H.: Energy based acoustic source localization. In: Information Processing in Sensor Networks, Proc. of IPSN 2003. (2003)

7 Wang, H., Yao, K., Pottie, G., Estrin, D.: Entropy-based sensor selection for localization. In: the Proceedings of SenSys'03. (2003)

8 Chu, M., Haussecker, H., Zhao, F.: Scalable information-driven sensor querying and routing for ad hoc heterogeneous sensor networks. International Journal of High-Performance Computing Applications **16** (2002)