It Takes a Network:

Cooperative Geolocation of Wireless Sensors

University of Michigan: Neal Patwari, Alfred O. Hero III
Ohio State University: Josh Ash, Randolph L. Moses
Motorola Labs: Spyros Kyperountas, Neiyer S. Correal

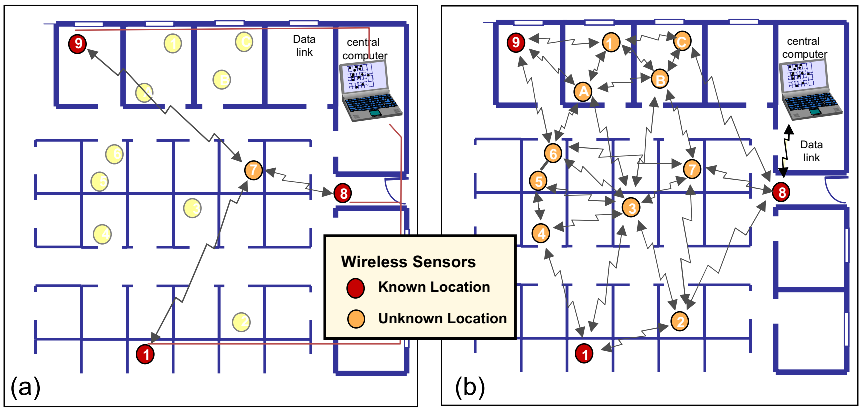

Fig. 1. In cooperative localization (b), measurements made between any pairs of sensors can be used to aid in the location estimate. Traditional multilateration or multi-angulation (a) is a special case in which measurements are made only between an unknown-location sensor and known-location sensors.

Abstract

Accurate and low-cost sensor cooperative localization is a critical requirement for the deployment of wireless sensor networks in a wide variety of applications. Sensors must be inexpensive and energy efficient devices, and sensor networks must be scalable from dozens to millions of devices. Localization methods must conform to these requirements and achieve a desired accuracy (from centimeters to meters). In this article, we describe measurement-based statistical models useful to describe time-of-arrival (TOA), angle-of-arrival (AOA), and received signal strength (RSS) measurements in wireless sensor networks. Wideband & ultra-wideband (UWB) measurements, and RF & acoustic media are discussed. Using the models, we show how to calculate a Cramér-Rao bound (CRB) on the location estimation accuracy possible from a given set of measurements, a useful tool to help system designers and researchers select measurement technologies and evaluate localization algorithms. We also briefly survey a large and growing literature of sensor localization algorithms. This article is meant to emphasize the basic statistical signal processing background necessary to understand the state of the art and to make progress in the new and largely open areas of sensor network localization research.

I. Introduction

Dramatic advances in RF and MEMS IC design have made possible the use of large networks of wireless sensors for a variety of new monitoring and control applications [1], [2], [3], [4]. For example, smart structures will actively respond to earthquakes and make buildings safer; precision agriculture will reduce costs and environmental impact by watering and fertilizing only where necessary, and will improve quality by monitoring storage conditions after harvesting; condition-based maintenance will direct main-

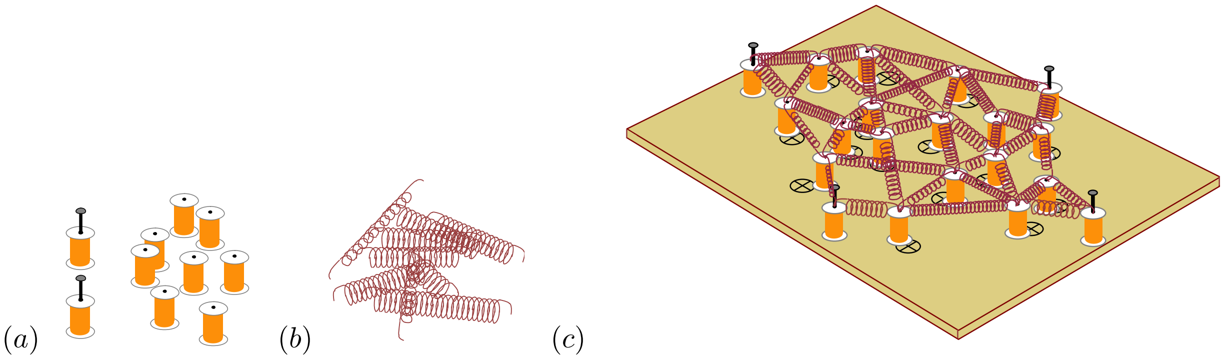

Fig. 2. Cooperative localization is analogous to finding the resting point of (a) masses (spools of thread) connected by a network of (b) springs. First, reference nodes are nailed to their known coordinates on a board. Springs have a natural length equal to measured ranges and can be compressed or stretched. They are connected to the pair of masses whose measured range they represent. After letting go, the equilibrium point (c) of the masses represent a minimum-energy localization estimate; the actual node locations are indicated by $\otimes$.

tenance exactly when and where needed based on data from wireless sensors; traffic monitoring systems will better control stoplights and inform motorists of alternate routes in case of traffic jams; and environmental monitoring networks will sense air, water and soil quality and identify the source of pollutants in real time.

Automatic localization of the sensors in these wireless networks is a key enabling technology. The overwhelming reason is that a sensor's location must be known for its data to be meaningful. As an additional motivation, sensor location information, if it is accurate enough, can be extremely useful for scalable, 'geographic' routing algorithms. Note also that location itself is often the data that needs to be sensed - localization can be the driving need for wireless sensor networks in applications such as warehousing and manufacturing logistics.

To make these applications viable with possibly vast numbers of sensors, device costs will need to be low (from a few dollars to a few cents depending on the application), sensors will need to last for years or even decades without battery replacement, and the network will need to organize without significant human moderation. Traditional localization techniques aren't well suited for these requirements. Including GPS on each device is cost and energy prohibitive for many applications, not sufficiently robust to jamming for military applications, and limited to outdoor applications. Local positioning systems (LPS) [5] rely on high-capability base stations being deployed in each coverage area, an expensive burden for most low-configuration wireless sensor networks.

Instead, we consider the problem in which some small number m of sensors, *reference nodes*, obtain their coordinates - either via GPS, or from a system administrator during startup - and the rest, n *unknown-location nodes*, must determine their own coordinates. If sensors were capable of high-power transmission, they would be able to make measurements to multiple reference nodes, and thus techniques similar to those used in LPS or cellular E-911 could be applied. However, low-capability, energy-conserving devices won't include a power amplifier or the energy necessary for long-range communication. Instead, wireless sensor networks, and thus localization techniques, will be multi-hop (a.k.a. 'cooperative' localization), as shown in Fig. 1. Rather than solving for each sensor's position one at a time, a location solver, analogous to the system of masses connected by springs shown in Fig. 2, will estimate all sensor positions simultaneously.

Network-aided localization systems are an extension of existing location systems such as LPS and cellular E-911. We still allow unknown-location devices to make measurements with known-location references, but in cooperative localization, we additionally allow unknown-location devices to make measurements with other unknown-location devices. The additional information gained from these measurements between pairs of unknown-location devices enhances the accuracy and robustness of the localization system.

A. Motivating Application Example: Animal Tracking

If cooperative localization can be implemented as described above, many compelling new applications can be enabled. For the purposes of biological research, it is very useful to track animals over time and over very wide ranges [6]. Such tracking can answer questions about animal behavior and interactions within their own and with other species. Using current practices, tracking is a very difficult, expensive process, and requires bulky tags that rapidly run out of energy. A typical practice is to attach VHF transmitter collars to animals to be tracked, and then triangulating their location by driving (or flying) to various locations with a directional antenna. Alternatively, GPS-based collars can be used, but are limited by cost concerns, and offer only a short lifetime due to high energy consumption. Using wireless sensor networks can dramatically improve the abilities of biological researchers (as demonstrated by 'ZebraNet' [6]). Using multi-hop routing of location data through the sensor network enables low transmit powers from the animal tags. Furthermore, inter-animal distances, which are of particular interest to animal behaviorists, can be estimated using pair-wise measurements and cooperative localization methods, without resorting to GPS. The end result of the longer battery lifetimes is less frequent re-collaring of the animals being studied.

B. Motivating Application Example: Logistics

As another example, consider deploying a sensor network in an office building, manufacturing floor, and warehouse. Sensors already play a very important role in manufacturing. Monitor and control of machinery has traditionally been wired, but making these sensors wireless reduces the high cost of cabling and makes the manufacturing floor more dynamic. Automatic localization of these sensors further increases automation.

Also, boxes and parts to be warehoused and factory and office equipment are all tagged with sensors when first brought into the facility. These sensors monitor storage conditions (temperature, humidity) and help control the HVAC system. Sensors on mobile equipment report their location when the equipment is lost or needs to be found, and even contact security if it is about to 'walk out' of the building. Knowing where things (parts or equipment) are when they are critically needed reduces the need to have duplicates as back-up, savings which can pay for the system itself.

Radio-frequency identification (RFID) tags, such as those now required by Walmart on pallets and cartons entering in its warehouses [7], represent a first step in warehouse logistics. RFID tags are only located when they pass within a few feet of a reader, thus remain out of access most of their time in the warehouse. Networked wireless sensors, however, can be queried and located as long as they are within range (on the order of 10 m) of the closest other wireless sensors.

The accuracy of cooperative localization increases with the density of sensors, as we show in Section III-E, so having heterogeneous sensors of varied purposes all participate in the same network helps drive localization errors down.

C. Cooperation Requirement: Standardization

One way to ensure that heterogeneous sensors can 'cooperate' to improve localization performance is to pursue standardization of wireless sensor networks. Two major sensor network standards are the IEEE 802.15.4 low data rate physical (PHY) layer and multiple-access (MAC) layer standards and the ZigBee™networking and application layer standards [8]. These standards allow for localization information to be measured between pairs of sensors. In particular, RSS can be measured in the 802.15.4 PHY standard via a physical layer mechanism, the Link Quality Indication (LQI), which reports to higher layers the signal strength associated with a received packet. Finally, we note that these standards enable low power operation required for long-life sensors - the 802.15.4 standard allows duty cycles of less than 0.1%, and when powered on, a transmitter or receiver will consume 40-60 mW.

D. Problem Statement

Before going into detail, it is useful to formally state the cooperative sensor location estimation problem. The 2-D localization problem demonstrated in Fig. 2 is the estimation of $2n$ unknown-location node coordinates $\boldsymbol{\theta} = [\boldsymbol{\theta}_x, \boldsymbol{\theta}_y]$,

$$\boldsymbol{\theta}_x = [x_1, \ldots, x_n], \quad \boldsymbol{\theta}_y = [y_1, \ldots, y_n] \tag{1}$$

given the known reference coordinates $[x_{n+1}, \ldots, x_{n+m}, y_{n+1}, \ldots, y_{n+m}]$, and pair-wise measurements $\{X_{i,j}\}$, where $X_{i,j}$ is a measurement between devices i and j. While we treat the 2-D case here, extension to 3-D appends a vector $\boldsymbol{\theta}_z$ to parameter vector $\boldsymbol{\theta}$ [9]. Measurements $X_{i,j}$ could be any physical reading that indicates distance or relative position, eg. time-of-arrival (TOA), angle-of-arrival (AOA), received signal strength (RSS), or proximity (whether or not two devices can communicate). We do not assume full measurements, so we define the set $H(i)$ to be the set of sensors with which sensor i makes measurements. Clearly, $i \notin H(i)$, and $H(i) \subset \{1, \ldots, n+m\}$. Note that these measurements could be done via different modalities - eg. RF, infrared (IR), acoustics [10], [11], or a combination [12]. Finally, TOA can be measured using different signaling techniques, such as direct-sequence spread-spectrum (DS-SS) [13], [14] or ultrawideband (UWB) [15], [16], [17]. We discuss these measurement methods in Section II-A.

Some research has further assumed that some nodes may have imperfect prior information about their coordinates - for example, that reference coordinates obtained from GPS may be random, but from a known distribution. Also, other localization research has focused on truly 'relative' location, i.e., when no references exist ($m = 0$), and an arbitrary coordinate system can be chosen. These are important directions of research, but to simplify the discussion in this article, we leave these extensions to the references [18], [10], [19].

E. Motivation and Outline

The main goal of this article is to provide an introduction, from a signal processing perspective, to the sensor location estimation problem. We review both theoretical estimation bounds and methods and algorithms being applied to the cooperative localization problem. We believe that signal processing methods will be very useful both for aiding system design decisions and in localization algorithms themselves.

Section II discusses why there are unavoidable limits to localization accuracy, and presents measurement-based statistical models for RSS, TOA, and AOA measurements. Section III uses these models to present lower bounds on sensor location estimation variance. The scope of this article does not include a detailed description of localization algorithms for sensor networks, however, in Section IV, we describe the main categories of approaches and provide references to the growing literature on localization algorithms.

II. WHY ARE MEASUREMENTS RANDOM?

Range and angle measurements used for localization are measured in a physical medium that introduces errors. Generally, these measurements are impacted by both *time-varying errors* and *static, environment-dependent errors*. Time-varying errors (e.g. due to additive noise and interference) can be reduced by averaging multiple measurements over time. Environment-dependent errors are the result of the physical arrangement of objects (e.g. buildings, trees, and furniture) in the particular environment which the sensor network is operating. Since the environment is unpredictable, these errors are unpredictable and we model them as random. However, in a particular environment, objects are predominantly stationary, and thus for a network of mostly stationary sensors, environment-dependent errors will be largely constant over time.

The majority of applications of wireless sensor networks involve mostly stationary sensors, and thus the delay required to make multiple measurements over time is acceptable. To determine how well such a localization system can perform, we must characterize the statistics of measurements *after* time-averaging, that is, measurements primarily in error due to the effects of the environment of deployment.

In Sections II-B to II-E, we present what measurement experiments have indicated about the statistics of RSS, TOA, and AOA measurements in sensor networks. We start in Section II-A by discussing the methodology of these measurement experiments.

A. Measurement Characterization

Ideally, statistical characterization of sensor network measurements would proceed as follows: deploy K wireless sensor networks, each with N sensors deployed with the identical geometry in the same type of environment, but each network in a different place. For example, we might test a sensor network deployed in a grid, in K different office buildings. In each deployment, make many measurements between all possible pairs of devices. Repeat each measurement over a short time period and compute the time-average. Then, the joint distribution (conditional on the particular sensor geometry) of the time-averaged measurements could be characterized. To our knowledge, no such wide-scale measurements have been presented, due to the huge scale of the task. First, a large K would be required to characterize the joint distribution. Secondly, the result would only be valid for that particular N and those particular sensor coordinates - each different geometry would require a different measurement experiment!

Measurements made to date have made simplifying assumptions about the measurement model. Basically, it is assumed that measurements in a network are independent and from the same family of distributions. The independence assumption, which says that observing an error in one link does not provide any information about whether or not errors occur in different links, is a simplifying assumption [20]. Large obstructions may affect a number of similarly-positioned links in a network. Considering correlations between links would make the analysis more difficult, and future research is needed to characterize the effects of link dependencies.

The second simplifying assumption is the choice of a family of distributions. We tend to subtract from each measurement its ensemble mean, and then assume that the error (the difference) is characterized by a particular parameterized distribution, such as Gaussian, log-normal, or mixture distribution. We then use the measurements to estimate the parameters of the distribution, such as the variance. By this method, one set of parameters can be used to characterize the whole set of measurements.

As an online supplement to this article [21], we provide a set of TOA and RSS measured in a 44-node indoor sensor network originally reported in [22] in order to allow researchers to test localization algorithms on measured data. Next, we discuss what those measurements and many other measurements of RSS, TOA, and AOA have indicated about the ensemble averages and the distributions of the error in pair-wise sensor measurements.

B. Received Signal Strength

Received signal strength (RSS) technically refers to the voltage measured by a receiver's received signal strength indicator (RSSI) circuit, but researchers typically also use RSS to refer to measured power. We can consider the RSS of either acoustic or RF signals. Wireless sensors communicate with neighboring sensors, and RSS can be measured by each receiver during normal data communication, without presenting additional bandwidth or energy requirements. Since RSS measurements are relatively inexpensive, they are an important and popular topic of localization research. Yet RSS measurements are notoriously unpredictable. If they are to be useful and part of a robust localization system, their sources of error must be well-understood.

1) Major Sources of Error: In free space, signal power decays proportional to d^{-2}, where d is the distance. In real-world radio channels, multipath signals and shadowing are two major sources of environment-dependence in the measured RSS. Multiple signals with different amplitudes and phases arrive at the receiver, and these signals add constructively or destructively as a function of the center frequency, causing frequency-selective fading. The effect of this type of fading can be diminished by using a spread-spectrum method (eg. direct-sequence or frequency hopping) which averages the received power over a wide range of frequencies. Spread-spectrum receivers are an acceptable solution, since spread-spectrum methods also reduce interference in the unlicensed bands in which wireless sensors typically operate. The measured received power using a wideband method (as the bandwidth $\to \infty$) is equivalent to measuring the sum of the powers of each multipath signal [23].

Assuming that frequency-selective effects are diminished, environment-dependent errors in RSS measurements are caused by shadowing, i.e., the attenuation of a signal due to obstructions (furniture, walls, trees, buildings, and more) that a signal must pass through or diffract around in its path between the transmitter and receiver. As discussed at the start of this section, these shadowing effects are modeled as random - as a function of the environment in which the network is deployed. A RSS model considers the randomness across an ensemble of many deployment environments.

2) Statistical Model: Typically, the ensemble mean received power in a real-world, obstructed channel decays proportional to d^{-n_p}, where n_p is the 'path-loss exponent', typically between 2 and 4. The ensemble mean power at distance d is typically modeled as

$$\bar{P}(d) = P_0 - 10 n_p \log \frac{d}{d_0} \tag{2}$$

where P_0 is the received power (dBm) at a short reference distance d_0.

The difference between a measured received power and its ensemble average, due to the randomness of shadowing, is modeled as log-normal (ie., Gaussian if expressed in dB). The log-normal model is based on a wide variety of measurement results [24], [25], [26], [22] and analytical evidence [27]. The standard deviation of received power (when received power is expressed in dBm), σ_{dB}, has units of (dB) and is relatively constant with distance. Typically, σ_{dB} is as low as 4 and as high as 12 [25]. Thus, the received power (dBm) at sensor i transmitted by j, $P_{i,j}$ is distributed as

$$f\left(P_{i,j} = p|\boldsymbol{\theta}\right) = \mathcal{N}\left(p; \bar{P}(d_{i,j}), \sigma_{dB}^2\right), \tag{3}$$

where $\mathcal{N}(x; y, z)$ is our notation for the value at x of a Gaussian p.d.f. with mean y and variance z, $\boldsymbol{\theta}$ is the coordinate parameter vector from (1), and the actual transmitter-receiver separation distance $d_{i,j}$ is given by

$$d_{i,j} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}. \tag{4}$$

The most important result of the log-normal model is that RSS-based range estimates have variance proportional to their actual range. This is not a contradiction of the earlier statement that σ_{dB} is constant with range. In fact, constant standard deviation in dB means that the multiplicative factors are constant with range, which explains the proportionality. For example, consider a multiplicative factor of 1.5. At an actual range of 100m, we would measure a range of 150m, an error of 50m; at 10m, the measured range would be 15m, an error of 5m, a factor of 10 smaller. This is why RSS errors are referred to as multiplicative, in comparison to the additive TOA errors presented in Section II-C. Clearly, RSS is most valuable in high-density sensor networks.

3) Calibration and Synchronization: In addition to the path loss, measured RSS is also a function of the calibration of both the transmitter and receiver. Depending on the expense of the manufacturing process, RSSI circuits and transmit powers will vary from device to device. Also, transmit powers can change as batteries deplete. Sensors might be designed to measure and report their own calibration data to their neighbors.

Alternatively, each sensor's transmit power can be considered an unknown parameter to be estimated. This means that the unknown vector $\boldsymbol{\theta}$ described in Section I-D is augmented to include the actual transmit power of each sensor along with its coordinates. Or, analogous to time-difference of arrival (TDOA) measurements described in Section II-C, we can consider only the differences between RSS measured at pairs of receivers [28]. The RSS difference between two sensors indicates information about their relative distance from the transmitter, and removes the dependency on the actual transmit power. We leave discussion of localization algorithms until Section IV.

C. Time-of-Arrival

Time-of-Arrival (TOA) is the measured time at which a signal (RF or acoustic) first arrives at a receiver. The measured TOA is the time of transmission plus a propagation-induced time delay. This time delay, $T_{i,j}$, between transmission at sensor i and reception at sensors j, is proportional to the TR separation distance, $d_{i,j}$, by the propagation velocity, v_p. This speed for RF is approximately 10^6 times as fast as the speed of sound – as a rule of thumb, for acoustic propagation, 1 ms translates to 1 ft (0.3 m), while for RF, 1 ns translates to 1 ft.

The cornerstone of time-based techniques is the receiver's ability to accurately estimate the arrival time of the line-of-sight (LOS) signal. This estimation is hampered both by additive noise and multipath signals: multiple time-delayed versions of the transmitted signal.

1) Major Sources of Error: Additive Noise: Even in the absence of multipath, the accuracy of the arrival time is limited by additive noise. Estimation of time-delay in additive noise is a relatively mature field [29]. Typically, the TOA estimate is the time that maximizes the cross-correlation between the received signals and the known transmitted signal. This estimator is known as a simple cross-correlator (SCC). The generalized cross-correlator (GCC) derived by Knapp and Carter [30] (the maximum likelihood estimator (MLE) for the TOA) extends the SCC by applying pre-filters to amplify spectral components of the signal that have little noise and attenuate components with large noise. As such, the GCC requires knowledge (or estimates) of the signal and noise power spectra.

For a given bandwidth and signal-noise ratio (SNR), our time-delay estimate can only achieve a certain accuracy. The Cramér-Rao bound (CRB) provides a lower bound on the variance of the TOA estimate in a multipath-free channel. For a signal with bandwidth B in (Hz), when B is much higher than the center frequency, F_c (Hz), and signal and noise powers are constant over the signal bandwidth [31],

$$\mathrm{var}(TOA) \geq \frac{1}{8\pi^2\, B\, T_s\, F_c^2\, \mathrm{SNR}}, \tag{5}$$

where T_s is the signal duration (s), and SNR is the signal to noise power ratio. By designing the system to achieve sufficiently high SNR, the bound predicted by the CRB (5) can be achieved in multipath-free channels. Thus (5) provides intuition about how signal parameters like duration, bandwidth, and power affect our ability to accurately estimate the TOA. For example, doubling either the transmission power or the bandwidth will cut ranging variance in half. This CRB on TOA variance is complementary to the bound that will be presented in Section III for location variance, because the location variance bound requires, as an input, the variance of the TOA estimates.

2) Major Sources of Error: Multipath: TOA-based range errors in multipath channels can be many times greater than those caused by additive noise alone. Essentially, all late-arriving multipath are self-interference that effectively decrease the SNR of the desired LOS signal. Rather than finding the highest peak of the cross-correlation, in the multipath channel, the receiver must find the first-arriving peak, because there is no guarantee that the LOS signal will be the strongest of the arriving signals. This can be done by measuring the time that the cross-correlation first crosses a threshold. Alternatively, in template-matching, the leading edge of the cross-correlation is matched in a least-squares sense to the leading edge of the auto-correlation (the correlation of the transmitted signal with itself) in order to achieve sub-sampling time resolutions [14]. Generally errors in TOA estimation are caused by two problems:

- *Early-Arriving Multipath*: Many multipath signals arrive very soon after the LOS signal, and their contributions to the cross-correlation obscure the location of the peak from the LOS signal.
- *Attenuated LOS*: The LOS signal can be severely attenuated compared to the late-arriving multipath components, causing it to be 'lost in the noise' and missed completely, causing large positive errors in the TOA estimate.

In dense sensor networks, in which any pair of sensors can measure TOA, we have the distinct advantage of being able to measure TOA between nearby neighbors. As the path length decreases, the LOS signal power (relative to the power in the multipath components) generally increases [32]. Thus the severely attenuated LOS problem is only severe in sparse networks.

While early-arriving multipath components cause smaller errors, they are very difficult to combat. Generally, wider signal bandwidths are necessary for obtaining greater temporal resolution. The peak width of the autocorrelation function is inversely proportional to the signal bandwidth. A narrow autocorrelation peak enhances the ability to pinpoint the arrival time of a signal and helps in separating the LOS signal cross-correlation contribution from the contributions of the early-arriving multipath signals. Wideband direct-sequence spread-spectrum (DS-SS) or ultra-wideband (UWB) signals (see sidebar on UWB) are popular techniques for high-bandwidth TOA measurements. However, wider bandwidths require higher speed signal processing, higher device and possibly higher energy costs. Standards proposals to the IEEE 802.15.3 (see UWB sidebar) quote receiver power consumptions on the order of 200 mW [33]. And, although high-speed circuitry typically means higher energy-consumption, the extra bandwidth can be used to lower the time-average power consumption. Transferring data packets in less time means spending more time in standby mode.

3) Statistical Model: Measurements have shown that for short-range measurements, measured time delay can be roughly modelled as Gaussian,

$$f\left(T_{i,j} = t|\boldsymbol{\theta}\right) = \mathcal{N}\left(t; d_{i,j}/v_p + \mu_T, \sigma_T^2\right), \tag{6}$$

where μ_T and σ_T^2 are the mean and variance of the time delay error, and $\boldsymbol{\theta}$ is defined in (1). Wideband DS-SS measurements reported in [22] supported the Gaussian error model and showed $\mu_T = 10.9$ ns and $\sigma_T = 6.1$ ns. UWB measurements done in on a mostly-empty Motorola factory floor showed $\mu_T = 0.3$ ns and $\sigma_T = 1.9$ ns. This mean error μ_T can be estimated (as a nuisance parameter) by the localization algorithm so that it can be subtracted out.

However, the presence of large errors can complicate the Gaussian model. These errors make the tails of the distribution of measured TOA heavier than Gaussian, and have been modelled using a mixture distribution: with a small probability, the TOA measurement results from a different, higher-variance distribution [34]. Localization systems should be designed to be robust to these large errors, also called non-line-of-sight (NLOS) errors. For TOA measurements made over time in a changing channel, the TOAs which include excess delays can be identified and ignored [34]. Even in static channels, if the number of range measurements to a device are greater than the minimum required, the redundancy can be used to identify likely NLOS errors [35]. Localization algorithm robustness is further addressed in Section IV.

4) Calibration and Synchronization: If wireless sensors have clocks that are accurately synchronized, then the time delay is determined by subtracting from the measured TOA the known transmit time. Sensor network clock synchronization algorithms have reported precisions on the order of 10μs [36]. Because of the difference in propagation speed, such clock accuracies are adequate for acoustic signals [12], but not for RF signals.

For time-of-arrival in asynchronous sensor networks, a common practice is to use *two-way* (or *round-trip*) TOA measurements. In this method, a first sensor transmits a signal to a second sensor, which immediately replies with its own signal. At the first sensor, the measured delay between its transmission and its reception of the reply is twice the propagation delay plus a reply delay internal to the second sensor. This internal delay is either known, or measured and sent to the first sensor to be subtracted. Multiple practical two-way TOA methods have been reported in the literature [37], [38], [13], [16]. Generally each pair of sensors measures round-trip TOA separately in time. But, if the first sensor has the signal processing capability, multiple sensors can reply at the same time, and two-way TOAs can be estimated simultaneously using multi-user interference cancellation [37].

The state of each sensor's clock (its bias compared with absolute time) can also be considered to be an unknown parameter and included in the parameter vector $\boldsymbol{\theta}$. In this case, one-way TOA is measured and input to a localization algorithm which estimates both the sensor coordinates and the biases of each sensor's clock [39]. The difference between the arrival times of the same signal at two sensors is called the time-difference of arrival (TDOA). A TDOA measurement doesn't depend on the clock bias of the transmitting sensor. TDOA methods have been used in source localization for decades for locating asynchronous transmitters, and has application in GPS and cellular localization. Under certain weak conditions, it has been shown that TOA with clock bias (treated as an unknown parameter) is equivalent to TDOA [40].

D. Ultra-Wideband and Localization

[This is intended to be a sidebar] Ultra Wideband (UWB) communication employs narrow pulses of very short (sub-nanosecond) duration that result in radio signals that are broadly spread in the frequency domain. A signal is considered to be UWB if its *fractional bandwidth*, the ratio of its bandwidth to its center frequency, is larger than 0.2. In 2003, the U.S. Federal Communications Commission (FCC) approved the commercialization and operation of UWB devices for public safety and consumer applications.

Among the envisaged applications are wireless networking and localization. Standardization of UWB is underway, including the development of a high bit rate UWB physical layer that supports peer-to-peer ranging, in IEEE task group 802.15.3a [33].

The very high bandwidth of UWB leads to very high temporal resolution, making it ideal for high precision radiolocation applications. Implementations of UWB-based range measurements, reported in [38], [17], [15], [16], have demonstrated RMS range accuracies from 3-5 feet (0.9-1.5 m) to 0.4 ft (12 cm).

E. Angle-of-Arrival

By providing information about the direction to neighboring sensors rather than the distance to neighboring sensors, angle-of-arrival (AOA) measurements provide localization information complementary to the TOA and RSS measurements discussed above.

There are two common ways that sensors measure AOA (as shown in Figure 3). The most common method is to use a *sensor array* and employ so-called *array signal processing* techniques at the sensor nodes. In this case, each sensor node is comprised of two or more individual sensors (*e.g.*, microphones for acoustic signals or antennas for RF signals) whose locations with respect to the node center are known. A four-element Y-shaped microphone array is shown in Figure 3(a). The AOA is estimated from the differences in arrival times for a transmitted signal at each of the sensor array elements. The estimation is similar to time-delay estimation discussed in Section II-C, but generalized to the case of more than two array elements. When the impinging signal is narrowband (that is, its bandwidth is much less than its center frequency), then a time delay τ relates to a phase delay ϕ by $\phi = 2\pi f_c \tau$ where f_c is the center frequency, and AOA estimators are often formulated based on phase delay. See [41], [42], [43] for more detailed discussions on AOA estimation algorithms and their properties.

A second approach to AOA estimation uses the RSS ratio between two (or more) directional antennas located on the sensor (see Figure 3(b). Two directional antennas pointed in different directions such that their main beams overlap can be used to estimate the AOA from the ratio of their individual RSS values.

Either AOA approach requires multiple antenna elements, which can contribute to sensor device cost and size. However, acoustic sensor arrays may already be required in devices for many environmental monitoring and security applications, in which the purpose of the sensor network is to identify and locate acoustic sources [44]. Locating the sensors themselves using acoustics in these applications is a natural extension. RF antenna arrays imply large device size unless center frequencies are very high. However, available bandwidth and decreasing manufacturing costs at millimeter-wave frequencies may make them

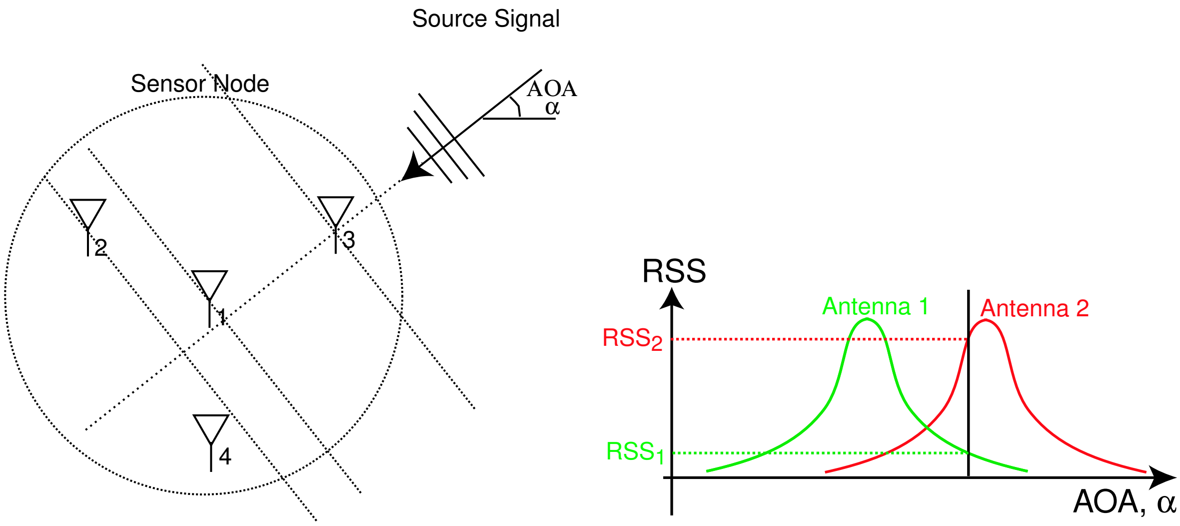

Fig. 3. Angle-of-arrival (AOA) estimation methods. (a) AOA is estimated from the time-of-arrival differences among sensor elements embedded in the node; a 4-element Y-shaped array is shown. (b) AOA can also be estimated from the received signal strength (RSS) ratio RSS_1/RSS_2 between directional antennas.

desirable for sensor network applications. For example, at 60 GHz, higher attenuation due to oxygen absorption helps to mitigate multipath, and accurate indoor AOA measurements have been demonstrated [45].

1) Major Sources of Error and Statistical Model: AOA measurements are impaired by the same sources discussed in the TOA section above - additive noise and multipath. The resulting AOA measurements are typically modeled as Gaussian, with ensemble mean equal to the true angle to the source and standard deviation σ_α. Theoretical results for acoustic-based AOA estimation show standard deviation bounds on the order of $\sigma_\alpha = 2°$ to $\sigma_\alpha = 6°$ depending on range [46]. Estimation errors for RF AOA on the order of $\sigma_\alpha = 3°$ have been reported using the RSS ratio method [47].

2) Calibration and Synchronization: It is not likely that sensors will be placed with known orientation. When sensor nodes have directionality, the network localization problem must be extended to consider each sensor's orientation as an unknown parameter, to be estimated along with position. In this case, the unknown vector $\boldsymbol{\theta}$ (see Section I-D) is augmented to include the orientation of each sensor.

The models presented in Sections II-B-II-E are sufficient to find bounds on localization performance in cooperative localization. These lower bounds are not a function of the particular localization algorithm employed. Thus we present some of these performance limits first, in the following section, before discussing current algorithm research in Section IV.

III. Limits on Localization Covariance

The Cramér-Rao bound provides a means for calculating a lower bound on the covariance of any unbiased location estimator which uses RSS, TOA, or AOA measurements. Such a lower bound provides a useful tool for researchers and system designers. Without testing particular estimation algorithms, a designer can quickly find the 'best-case' possible using particular measurement technologies. Researchers who are testing localization algorithms, like those presented in Section IV, can use the CRB as a comparison for a particular algorithm. If the bound is nearly achieved, then there is little reason to continue working to improve that algorithm's accuracy. Furthermore, the bound's functional dependence on particular parameters helps to provide insight into the behavior of cooperative localization.

The bound on estimator covariance is a function of the following:

1) Number of unknown-location and known-location sensors,
2) Sensor geometry,
3) Whether localization is in two or three dimensions,
4) Measurement type(s) implemented (*i.e.*, RSS, TOA, or AOA),
5) Channel parameters (such as σ_{dB} and n_p in RSS, σ_T in TOA, or σ_α in AOA measurements),
6) Which pairs of sensors make measurements (network connectivity),
7) 'Nuisance' (unknown) parameters which must also be estimated (such as clock bias for TOA or orientation for AOA measurements).

As an online supplement to this article, we provide public access to a multi-featured Matlab-based code and GUI for the calculation of the localization CRB, as shown in Figure 4. This code can consider any combination of RSS, TOA, and AOA measurement methods to be employed. It allows the inclusion of device orientation and clock biases as unknown 'nuisance' parameters. Sensors can be arranged visually using the GUI and the bound can be calculated. For each sensor, the GUI displays the CRB by plotting the lower bound on the 2-σ uncertainty ellipse. The tool also includes the ability to run Monte-Carlo simulations which estimate sensor parameters and coordinates using the Maximum-likelihood estimator that will be discussed in Section IV. The Monte-Carlo coordinate estimates are plotted on screen for comparison with the covariance bound.

In this section, we present some analytical results for the CRB. Our purpose is both to show that it is simple to calculate and as a means to compare the three measurement methods presented in Section II. To keep the formulation short, we make two simplifying assumptions. First, we address 2-D (rather than 3-D) localization. Second, we assume that channel and device parameters (orientation for AOA, transmit powers and n_p for RSS, and clock biases for TOA) are known. Analysis of bounds without these assumptions are left to the references [9], [48], [10], [49], [39] which have presented details of these analytical CRBs, for a variety of different measurement types.

A. Calculating the CRB in Three Steps

The resulting CRB (under these two simplifying assumptions) shows remarkable similarities between the CRBs using RSS, TOA, and AOA measurements. Particular differences can be seen that show how localization performance varies by measurement type. We show how to calculate the CRB for the estimate of $\boldsymbol{\theta}$ as given in (1) in three steps:

1) Calculate Fisher information sub-matrices: First, form three $n \times n$ matrices: $\mathbf{F}_{xx}$, $\mathbf{F}_{xy}$, and $\mathbf{F}_{xx}$. The k, l element of each matrix is calculated as follows:

$$\begin{aligned} \left[\mathbf{F}_{xx}\right]_{k,l} &= \begin{cases} \gamma \sum_{i \in H(k)} (x_k - x_i)^2 / d_{k,i}^s & k = l \\ -\gamma\, I_{H(k)}(l)(x_k - x_l)^2 / d_{k,l}^s & k \neq l \end{cases} \\ \left[\mathbf{F}_{xy}\right]_{k,l} &= \begin{cases} \gamma \sum_{i \in H(k)} (x_k - x_i)(y_k - y_i) / d_{k,i}^s & k = l \\ -\gamma\, I_{H(k)}(l)(x_k - x_l)(y_k - y_l) / \| d_{k,l}^s & k \neq l \end{cases} \\ \left[\mathbf{F}_{yy}\right]_{k,l} &= \begin{cases} \gamma \sum_{i \in H(k)} (y_k - y_i)^2 / d_{k,i}^s & k = l \\ -\gamma\, I_{H(k)}(l)(y_k - y_l)^2 / d_{k,l}^s & k \neq l \end{cases} \end{aligned} \quad (7)$$

Here, γ is a channel constant, and s is an exponent, which are a function of the measurement type and are given in Table I, d_{ij} is the true distance between i and j given in (4), and $I_{H(k)}(l)$ is the indicator function, (which allows us to include the information only if sensor k made a measurement with sensor l), $I_{H(k)}(l) = 1$ if $l \in H(k)$, or 0 if not.

2) Merge sub-matrices to form the FIM: Next, we form the $2n \times 2n$ Fisher information matrix (FIM) $\mathbf{F}$. For TOA or RSS, we select $\mathbf{F} = \mathbf{F}_{\mathbf{TR}}$, while for AOA, we select $\mathbf{F} = \mathbf{F}_{\mathbf{A}}$, where

$$\mathbf{F}_{\mathbf{TR}} = \begin{bmatrix} \mathbf{F}_{xx} & \mathbf{F}_{xy} \\ \mathbf{F}_{xy}^T & \mathbf{F}_{yy} \end{bmatrix}, \quad \mathbf{F}_{\mathbf{A}} = \begin{bmatrix} \mathbf{F}_{yy} & -\mathbf{F}_{xy} \\ -\mathbf{F}_{xy}^T & \mathbf{F}_{xx} \end{bmatrix}, \quad (8)$$

where $\mathbf{F}_{xx}$, $\mathbf{F}_{xy}$, and $\mathbf{F}_{xx}$ are given in (7), and we use the superscript ^{T} to indicate matrix transpose.

3) Invert the FIM to get the CRB: The CRB matrix is equal to $\mathbf{F}^{-1}$, the matrix inverse of the FIM. The diagonal of $\mathbf{F}^{-1}$ contains $2n$ values which are the variance bounds for the $2n$ parameters of $\boldsymbol{\theta}$. To say this more precisely, let an estimator of sensor i's coordinates be $\hat{\mathbf{z}}_i = [\hat{x}_i, \hat{y}_i]^T$. If we define the location variance of the estimator to be σ_i^2,

$$\sigma_i^2 \triangleq \mathrm{tr}\left\{\mathrm{cov}_{\boldsymbol{\theta}}(\hat{\mathbf{z}}_i)\right\} = \mathrm{Var}_{\boldsymbol{\theta}}(\hat{x}_i) + \mathrm{Var}_{\boldsymbol{\theta}}(\hat{y}_i), \quad (9)$$

then the Cramér-Rao bound asserts that,

$$\sigma_i^2 \geq \left(\mathbf{F}^{-1}\right)_{i,i} + \left(\mathbf{F}^{-1}\right)_{i+n,i+n}. \tag{10}$$

TABLE I

DIFFERENCES IN CRB BY MEASUREMENT TYPE.

	Channel Constant γ	Exponent s	FIM $\mathbf{F}$
TOA	$\gamma = 1/(c\sigma_T)^2$	$s = 2$	$\mathbf{F} = \mathbf{F_{TR}}$
RSS	$\gamma = \left(\frac{10n_p}{\sigma_{dB}\log 10}\right)^2$	$s = 4$	$\mathbf{F} = \mathbf{F_{TR}}$
AOA	$\gamma = 1/\sigma_\alpha^2$	$s = 4$	$\mathbf{F} = \mathbf{F_A}$

B. *Results Seen from the CRB*

Even without calculating the CRB for a particular sensor network geometry, we can talk about the scaling characteristics of the variance bound. What happens when the geometry and connectivity of the network is kept constant, but the dimensions of the network are scaled up proportionally?

- *TOA*: TOA bounds will remain constant with a scaling of the dimensions. Note that since $s = 2$ for TOA, the fractions in (7) are unitless - if units of the coordinates were (ft) or even (cm) instead of (m), the ratios would be identical. Instead, the units come from the variance of ranging error, $c\sigma_T$.
- *RSS and AOA*: These bounds on standard deviation are proportional to the size of the system. Since $s = 4$ for RSS and AOA, the geometry ratios in (7) have units of $1/\text{distance}^2$, so the variance bound (the inverse) takes its units of distance^2 directly from this ratio. Note that the channel constant γ is unitless for both RSS and AOA.

Of course, channel parameters will change slowly as the path lengths change (TOA measurements over kilometer links would have higher variance than over 10 m links), but these scaling characteristics are good first order approximations.

Finally, note that the bound on standard deviation of localization error is proportional to $\sqrt{1/\gamma}$. It makes sense that the localization error is proportional to σ_T for TOA and σ_α for AOA. It isn't as obvious, but we also find from the CRB that for RSS, the proportionality is to σ_{dB}/n_p. A RSS-based localization system operating in a high path-loss exponent environment (often found when using ground-level antennas), while requiring higher transmit powers from sensors, also allows more accurate sensor localization.

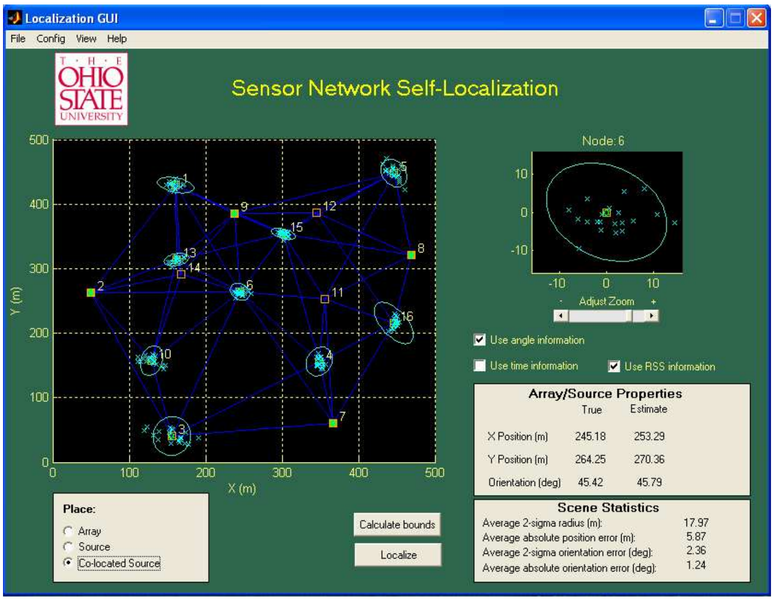

Fig. 4. Lower bounds and Monte-carlo ML estimates can be calculated interactively using this Matlab-based GUI developed by Joshua Ash at Ohio State University and freely available online [21]. Sensors can be placed arbitrarily, and their capabilities and *a priori* location information given. The user may select measurements from AOA, TOA, RSS, and Proximity.

C. *What is the Cramér-Rao Bound?*

[Intended as a sidebar] The Cramér-Rao Bound (CRB) provides a lower bound on the variance achievable by any unbiased location estimator. The bound is useful as a guideline: knowing the best an estimator can possibly do helps us judge the estimators that we implement. Essentially, the CRB is a general uncertainty principle for estimation problems, which we apply in this article to location estimation. It is surprising, to those without *a priori* knowledge of the CRB, that we can calculate the lower bound on estimation variance without ever considering a single estimation method. Without providing a proof [50], we give an intuitive argument on why the CRB does so.

Consider an estimator in which perfect, error-free measurements result in a perfect, error-free position estimate from the localization algorithm. It makes sense that if measurements have a small ϵ error in them, then the location estimate will also have some small error proportional to ϵ. The study of the effect of small errors, i.e. sensitivity analysis, considers the derivatives of the coordinate estimates with respect to changes in the measurements.

The CRB similarly considers derivatives. However, the CRB takes the derivative of the log of the probability density function (pdf) of the measurements (conditioned on the unknown coordinates), a.k.a. the *log-likelihood function*. If the log-likelihood changes rapidly as a function of the unknown coordinates,

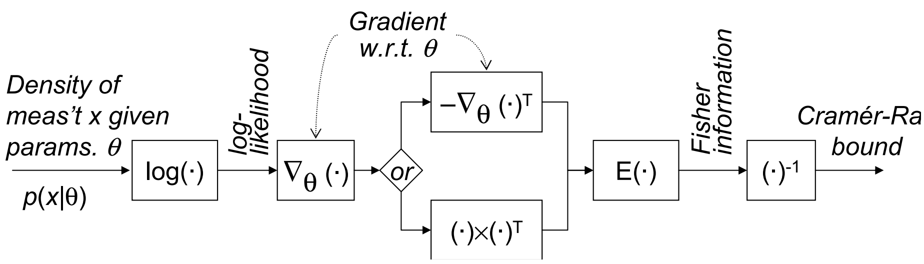

Fig. 5. The CRB can be calculated from the conditional probability density function using a few operators.

then the measurements we observe will provide a great deal of information about the coordinates. We will be able to better guess (with lower errors) the unknown coordinates. Conversely, if the log-likelihood changes slowly with respect to the unknown coordinates, then observing measurements doesn't help much in their estimation. Note that:

- Information is accurately measured by the 'order of magnitude' of a probability. For example, the change from a probability of 0.4 to 0.2 is just as dramatic as a change in probability from 0.2 to 0.1, since in each case the change makes an event half as likely. Taking the logarithm of the probabilities prior to calculating the difference, we get the desired property that the change in log-probability for both cases is the same.
- To bound variance (the expected value of the squared error), the CRB is determined from the expected value of the squared value of the derivatives.
- There are regularity conditions on the density functions - see [50] for details.

The basis for the lower bound is the limitation to unbiased estimators. Such estimators provide coordinate estimates that on average, if we average over enough realizations, are the same as the true coordinates. This unbiasedness must hold regardless of the true value of the coordinates. The limitation serves the same purpose as the requirement given in the above discussion of sensitivity analysis, that when there is no error in the measurements, the estimate is correct. Although unbiased estimation is a very desirable property, there are cases in which we are willing to trade some bias in order to improve upon the variance performance – in such cases, the bound can be adapted [51].

D. What's the difference between the CRB and the GDOP?

[Intended as a sidebar] The 'geometric dilution of precision' (GDOP) has been used to characterize the location estimation accuracy of many TOA and TDOA systems as a function of the geometry of the sensors and the source [52], [53]. The GDOP is defined as the ratio of the standard deviation of location error (achieved by a particular estimator) to the standard deviation of ranging error. Essentially, it provides a multiplicative factor which shows how many times more uncertain our location estimate will be than our range estimate.

The square-root of the CRB bounds the standard deviation of localization error of any unbiased estimator. In some cases, the standard deviation bound is proportional to the standard deviation of ranging error (*proportionality*), and thus provides us with a means to lower bound the GDOP. Further, if the lower bound can be achieved by a particular estimator (*efficiency*), then the GDOP of that location estimator and the CRB will have the same form. Both of these conditions, proportionality and efficiency, are met in the case of Gaussian-distributed TOA measurements and least-squares estimator (LSE) [53], and the lower bound on standard deviation and the GDOP take the same form [22].

E. Numerical Example

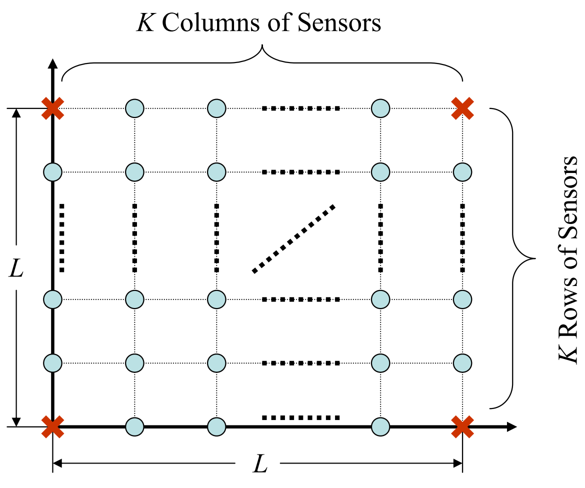

Fig. 6. Diagram showing layout of the K^2 sensors of the example described in Section III-E, with four reference sensors ($\times$) and $K^2 - 4$ unknown-location sensors ($\bullet$) in a $L \times L$ square area.

Consider a sensor network in a 20m by 20m area, with K^2 sensors arranged into K rows and K columns, as shown in Fig. 6. The four sensors in the corners are reference nodes, while the remaining

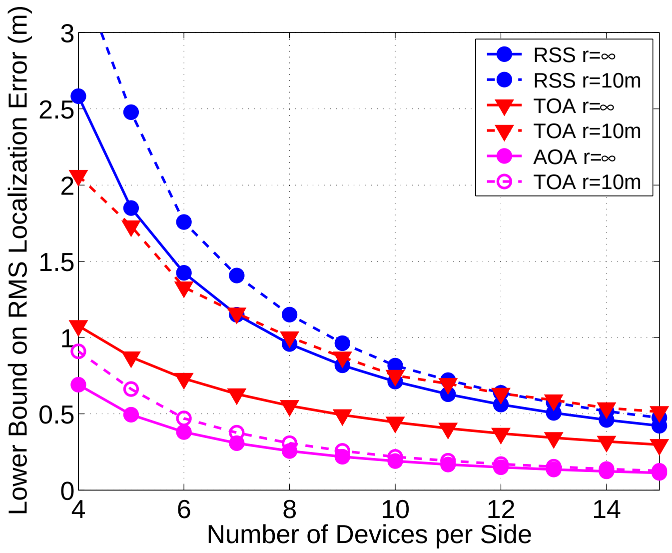

Fig. 7. Lower bounds for localization variance for the example described in Section III-E when measurements are RSS (with $\sigma_{dB}/n_p = 1.7$ [22]), TOA (with $\sigma_{dB}/n_p = 1.7$ [22]), and AOA (with $\sigma_\alpha = 5^o$). Parameter r is the radius of connectivity - only pairs of sensors closer than r make measurements, and for $r = \infty$, all pairs make measurements.

$K - 4$ are unknown-location nodes. Let's consider what happens to the localization variance bound as K increases, for the cases when measurements are:

1) RSS with $\sigma_{dB}/n_p = 1.7$ [22],
2) TOA with $\sigma_T = 6.1$ ns [22],and
3) AOA with $\sigma_\alpha = 5^o$ (see Section II-E.1).

As presented above, the lower bounds for RSS, TOA, and AOA are proportional to these three channel parameters. We start by assuming that each sensor makes measurements with every other sensor in the network. We calculate the RMS value of the localization bound, i.e., $\sqrt{\frac{1}{n}\text{tr}\,\mathbf{F}^{-1}}$, which gives an average of the bound over the entire $K^2 - 4$ unknown-location sensors. The result is shown in the solid lines in Fig. 7 labelled as '$r = \infty$'. Next, we consider the realistic case in which each sensor only makes measurements to those sensors located within $r = 10$ m of itself[1]. In this case, the bound is shown as dotted lines in Fig. 7 and labelled as '$r = 10$ m'.

Comparing performance of the measurement methods for the chosen parameters, AOA outperforms TOA and RSS, while RSS can perform as well as TOA at high sensor densities. Of course, these comparisons are based on the chosen values of the channel parameters and the chosen geometry shown in Fig. 6

[1]Of course, sensors won't really know *exactly* which sensors are within 10 m, but the connectivity implied by the 10 m radius is still a good test.

– these plots are proportional to the channel parameters as described in Section III-B. Note that RSS and AOA bounds decrease more rapidly than TOA as the density increases. Also for RSS and AOA, the difference between the $r = 10$m and $r = \infty$ lines decreases dramatically as density increases. At high densities, the results show that little additional information comes from the distant sensors (> 10 m). For TOA, however, even distant sensors' measurements can provide significant localization information.

IV. Location Estimation Algorithms

So far in this article, we haven't mentioned any particular localization estimators, only a lower bound on their performance. There is a considerable literature in sensor localization algorithms for wireless sensor networks, described alternatively as 'cooperative', 'relative', 'distributed', 'GPS-free', 'multi-hop', or 'network' localization; 'self-localization'; or 'ad-hoc' or 'sensor' positioning.

While positioning and navigation have a long history (as evidenced in this issue), cooperative localization algorithms must extend existing methods by finding ways to use the measurements (range or angle) measured between pairs of unknown-location nodes. The challenge is to allow sensors which aren't in range of any known-location devices to be located, and further, to improve the location estimates of all sensors.

If all sensors were in range of multiple reference nodes, they could directly calculate their own locations. For example, in [54], nodes measure RSS from a dense network of reference nodes and calculate their location to be the mean of the locations of the in-range reference nodes. In most wireless sensor networks, though, to minimize installation expenses, reference nodes will be sparse, and low-energy sensors generally won't be in range of enough references (3 or 4 for 2-D or 3-D localization, respectively).

Several cooperative localization algorithms are reviewed in [19]. Here, we divide methods into centralized algorithms, which collect measurements at a central processor prior to calculation, and distributed algorithms, which require sensors to share information only with their neighbors, but possibly iteratively.

A. Centralized Algorithms

If the data is known to be described well by a particular statistical model (eg. Gaussian or log-normal), then the maximum likelihood estimator can be derived and implemented [22], [10]. One reason that these estimators are used is that their variance asymptotically (as the signal-to-noise ratio goes high) approaches the lower bound given by the CRB (in Section III). As indicated by the name, the maximum of a likelihood function must be found. There are two difficulties with the MLE:

1) *Local Maxima*: Unless we initialize the MLE to a value close to the correct solution, it is possible that our maximization search may not find the global maxima.
2) *Model Dependency*: If measurements deviate from the assumed model (or model parameters), the results are no longer guaranteed to be optimal.

One way to prevent local maxima is to formulate the localization as a convex optimization problem. In [55], convex constraints are presented that can be used to require a sensor's location estimate to be within a radius r and/or angle range $[\alpha_1, \alpha_2]$ from a second sensor. In [39], linear programming using a 'taxi metric' is suggested to provide a quick means to obtain rough localization estimates. More general constraints can be considered if semi-definite programming (SDP) techniques are used [56]. One difficulty which must be overcome in both techniques is their high computational complexity. Towards this end, a distributed SDP-based localization algorithm was presented in [57].

Multi-dimensional scaling (MDS) algorithms (and Isomap [58]) formulate sensor localization from range measurements as a least-squares (LS) problem [59], [60]. In classical MDS, the LS solution is found by eigen-decomposition, which doesn't suffer from local maxima. In order to linearize the localization problem, the classical MDS formulation works with squared distance rather than distance itself, and the end result is very sensitive to range measurement errors. Other MDS-based techniques, not based on eigen-decomposition, can be made more robust by allowing measurements to be weighted according to their perceived accuracy [18].

While MDS and Isomap have complexity $\mathcal{O}(N^3)$, where $N = n + m$ is the total number of sensors, other manifold learning methods, such as local linear embedding (LLE) [61], are also based on eigen-decomposition, but of sparse matrices, and are $\mathcal{O}(N^2)$. Manifold learning performance has been presented for the case when sensor data records are used as location information [62], and will likely play an important role when using other types of measurements. Also adapted from the statistical learning area, 'supervised learning' approaches localization as a series of detection problems [63]. The covered area is split into smaller, overlapping regions, and based on the measurements, each region detects whether or not the sensor is within its boundaries.

B. Distributed Algorithms

There are two big motivations for developing distributed localization algorithms. First, for some applications, no central processor, or none with enough computational power, is available to handle the calculations. Second, for large sensor networks, centralized algorithms cause a communication bottleneck at and near the central processor, due to the forwarding of messages containing the pair-wise measurements.

Distributed algorithms for cooperative localization generally fall into one of two categories:

1) *Network Multilateration*: Each sensor estimates its multi-hop range to the nearest reference nodes. These ranges can be estimated via the shortest path between the sensor and reference nodes, i.e., proportional to the number of hops, or the sum of measured ranges along the shortest path [64][65][66]. Note that finding the shortest path is readily distributed across the network. When each sensor has multiple range estimates to known positions, its coordinates are calculated locally via multilateration [52][67].
2) *Successive Refinement*: These algorithms try to find the optimum of a global cost function, eg., least squares (LS), weighted LS (WLS) [18], or maximum likelihood (ML). Each sensor estimates its location and then transmits that assertion to its neighbors [68][69][70]. Neighbors must then recalculate their location and transmit again, until convergence. A device starting without any coordinates can begin with its own local coordinate system and later merge it with neighboring coordinate systems [71]. Typically, better statistical performance is achieved by successive refinement compared to network trilateration, but convergence issues must be addressed.

Bayesian networks (or factor graphs) provide another distributed successive refinement method to estimate the probability density of sensor network parameters. These methods are particularly promising for sensor localization - each sensor stores a conditional density on its own coordinates, based on its measurements and the conditional density of its neighbors [72]. Alternatively, particle filtering (or Monte-Carlo estimation methods) methods have each sensor store a set of 'particles', *i.e.*, candidate representations of its coordinates, weighted according to their likelihood [73], [74]. These methods have been used to accurately locate and track mobile robots [75], and they will likely find application in future sensor localization and tracking research.

C. Comparison

Both centralized and distributed algorithms must face the high relative costs of communication. The energy required per transmitted bit could be used, depending on the hardware and the range, to execute 1,000 to 30,000 instructions [44]. Centralized algorithms in large networks require each sensor's measurements to be passed over many hops to a central processor, while distributed algorithms have sensors send messages only one hop (but possibly make multiple iterations). When the average number of hops to the central processor exceeds the necessary number of iterations, distributed algorithms will likely save communication energy costs.

There may be hybrid algorithms which combine centralized and distributed features in order to reduce the energy consumption beyond what either one could do alone. For example, if the sensor network is divided into small clusters, an algorithm could select a processor from within each cluster to estimate a map of the cluster's sensors. Then, cluster processors could operate a distributed algorithm to merge and optimize the local estimates. Such algorithms are an open topic for future research.

V. Future Research Needs and Conclusion

Ultimately, actual localization performance will depend on many things, including the localization algorithm used, size and density of the network, the quantity of prior coordinate information, the measurement method chosen, and the accuracies possible from those measurements in the environment of interest (the γ of Table I). However, based on the characteristics of the variance bounds presented in Section III, we can make some broad generalizations. It appears that TOA measurements will be most useful in low-density sensor networks, since it is not as sensitive to increases in inter-device distances as RSS and AOA. Both AOA and TOA are generally able achieve higher accuracy than RSS; however, that accuracy can come with higher device costs. Because of their scaling characteristics, localization based on RSS and AOA measurements can, without sacrificing much accuracy, avoid taking measurements on longer-distance links and focus on those links between nearest neighbors. RSS measurements will allow accurate localization in very dense networks, and will be very attractive to system designers due to their low costs.

Considerable literature has studied cooperative localization with an emphasis on algorithms, less research has placed the emphasis on localization as estimation. Accordingly, bias and variance performance is often a secondary concern. While a notable algorithm comparison is seen in [76], in general, it is difficult to compare the performance of localization algorithms in the literature. Reporting both bias and variance performance along with the Cramér-Rao lower bound will help provide a reference for comparison.

While simulation will be very valuable, the next step in cooperative localization research is to test algorithms using measured data. However, measurements of RSS, TOA, and AOA in wireless sensor networks have only begun to be reported, largely because of the complexity of such measurement campaigns. To conduct measurements in a N sensor network requires $\mathcal{O}(N^2)$ measurements, and multiple sensor networks must be measured in order to develop and test statistical models. Despite the complexity, data from such measurement campaigns will be of key importance to sensor network researchers.

Three other future directions for cooperative localization research are suggested: mobile sensor tracking, the use of connectivity measurements, and routing using virtual coordinates.

1) Mobile Sensor Tracking: This article has not discussed sensor mobility. Mobility creates the problem of locating and tracking moving sensors in real time, and also the opportunity to improve sensor localization. Detecting movement of a sensor in a network of communication or energy-constrained nodes is a distributed detection problem yet to be fully explored. For the problem of passive tracking of sources in the environment, a review is presented in [44], but the problem of tracking active sensors has not been sufficiently addressed as a collaborative signal processing problem. The sensor tracking problem is an important aspect of many applications, including the animal tracking and logistics applications discussed in Sections I-B and I-A.

If a sensor makes multiple measurements to its neighbors as it moves across space, it has the opportunity to reduce environment-dependent errors by averaging over space. The multiple measurements are useful to help improve coordinate estimates for other sensors in the network, not just the mobile node. Researchers have tested schemes which use mobile sources and sensors to achieve cooperative localization [77], [78], and further opportunities to exploit mobility remain to be explored.

2) Connectivity Measurements: Connectivity (a.k.a. proximity) is a binary measurement of whether or not two devices are in communication range of each other. Typical digital receivers have a minimum received power below which it is unlikely that a packet will be correctly demodulated. Connectivity can be considered to be a binary quantization of RSS. As a good approximation, when the RSS is below a power threshold, two devices will not be 'connected', and when the RSS is above the threshold, two devices will be connected . As a binary quantization of RSS, it is clear that connectivity is less informative than RSS and will result in higher localization variance bounds [79]. Research in connectivity-based localization is often called 'range-free' localization. The assumption that connectivity does not suffer from the same fading phenomena as RSS, and instead that radio coverage is a perfect circle around the transmitter, can be a valuable simplification during algorithm development; however, this assumption cannot be used to accurately test the performance of such algorithms. Range-free localization algorithms can be simulated by generating measurements of RSS between each pair of sensors using the log-normal model of (3), and then considering each pair with an RSS measurement above a threshold power to be connected.

3) Routing using Virtual Coordinates: Geographic routing is an application of sensor localization. The use of the coordinates of sensors can reduce routing tables and simplify routing algorithms. Localization errors, however, can adversely impact routing algorithms, leading to longer paths and delivery failures [80]. For the purposes of routing efficiency, actual geographical coordinates may be less useful than ‘virtual’ coordinates [81], *i.e.*, a low-dimensional representation of a sensor’s ‘location’ in the graph of network connectivity. There are often paths in multi-hop wireless networks that consume less power than the shortest, straightest-line path between two nodes, and virtual coordinates may provide a better representation of the network connectivity. The virtual coordinate estimation problem is a dimension reduction problem which inputs each sensor’s connectivity or RSS measurement vector and outputs a virtual coordinate in an arbitrary low dimension, optimized to minimize a communication cost metric. Such research could enable more energy-efficient, scalable routing protocols for very large sensor networks.

A. Conclusion

Cooperative localization research will continue to grow as sensor networks are deployed in larger numbers and as applications become more varied. Localization algorithms must be designed to achieve low bias and as low of variance as possible, and at the same time, be scalable to very large network sizes without dramatically increasing energy requirements.

This article has provided a window into cooperative localization, which has found considerable application in ad-hoc and wireless sensor networks. We have presented measurement-based statistical models of TOA, AOA, and RSS, and used them to generate localization performance bounds. Such bounds are useful, among other design considerations, as design tools to help choose among measurement methods, select neighborhood size, set minimum reference node densities, and compare localization algorithms.

References

[1] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, “Wireless sensor networks: A survey,” *Computer Networks (Elsevier) Journal*, vol. 38, no. 4, pp. 393–422, Mar. 2002.

[2] N. S. Correal and N. Patwari, “Wireless sensor networks: Challenges and opportunities,” in *Proceedings of the 2001 Virginia Tech Symposium on Wireless Personal Communications*, June 2001, pp. 1–9.

[3] M. Perkins, N. S. Correal, and R. J. O’Dea, “Emergent wireless sensor network limitations: a plea for advancement in core technologies,” in *IEEE Sensors*, vol. 2, June 2002, pp. 1505–1509.

[4] B. J. Feder, “Wireless sensor networks spread to new territory,” *New York Times*, July 26 2004. [Online]. Available: http://www.nytimes.com/2004/07/26/business/26sensor.html

[5] J. Werb and C. Lanzl, “Designing a positioning system for finding things and people indoors,” *IEEE Spectrum*, vol. 35, no. 9, pp. 71–78, Sept. 1998.

[6] P. Juang, H. Oki, Y. Wang, M. Martonosi, L.-S. Peh, and D. Rubenstein, “Energy efficient computing for wildlife tracking: Design tradeoffs and early experiences with ZebraNet,” in *Conf. Architectural Support for Programming Lang. and Operating Systems (ASPLOS)*, October 2002.

[7] J. Kumagai and S. Cherry, "Sensors and sensibility," *IEEE Spectrum*, vol. 41, no. 7, pp. 22–28, July 2004.

[8] ZigBee$^{\text{TM}}$Alliance. [Online]. Available: http://www.zigbee.org

[9] W. L. Garber, "Self-localization of sensor networks," Master's thesis, Ohio State University, 2003.

[10] R. L. Moses, D. Krishnamurthy, and R. Patterson, "A self-localization method for wireless sensor networks," *EURASIP Journal on Applied Sig. Proc.*, no. 4, pp. 348–358, Mar. 2003.

[11] J. N. Ash and R. L. Moses, "Acoustic sensor network self-localization: Experimental results," in *Military Sensing Symposia*, 2003.

[12] L. Girod, V. Bychkovskiy, J. Elson, and D. Estrin, "Locating tiny sensors in time and space: a case study," in *IEEE Int. Conf. on Computer Design*, 2002, pp. 214–219.

[13] D. D. McCrady, L. Doyle, H. Forstrom, T. Dempsy, and M. Martorana, "Mobile ranging with low accuracy clocks," *IEEE Trans. on Microwave Theory and Techniques*, vol. 48, no. 6, pp. 951–957, June 2000.

[14] B. B. Peterson, C. Kmiecik, R. Hartnett, P. M. Thompson, J. Mendoza, and H. Nguyen, "Spread spectrum indoor geolocation," *Journal of the Institute of Navigation*, vol. 45, no. 2, pp. 97–102, Summer 1998.

[15] N. S. Correal, S. Kyperountas, Q. Shi, and M. Welborn, "An ultra wideband relative location system," in *IEEE Conf. on Ultra Wideband Systems and Technologies*, Nov. 2003.

[16] R. Fleming and C. Kushner, "Low-power, miniature, distributed position location and communication devices using ultra-wideband, nonsinusoidal communication technology," Aetherwire Inc., Semi-Annual Technical Report, ARPA Contract J-FBI-94-058, Tech. Rep., July 1995.

[17] R. J. Fontana and S. J. Gunderson, "Ultra-wideband precision asset location system," in *IEEE Conf. UWB Syst. & Tech.*, May 2002, pp. 147–150.

[18] J. A. Costa, N. Patwari, and A. O. Hero III, "Distributed multidimensional scaling with adaptive weighting for node localization in sensor networks," *IEEE/ACM Transactions on Sensor Networks*, May 2004, (submitted).

[19] D. Niculescu, "Positioning in ad hoc sensor networks," *IEEE Network*, vol. 18, no. 4, pp. 24–29, July-Aug. 2004.

[20] N. Patwari, Y. Wang, and R. J. O'Dea, "The importance of the multipoint-to-multipoint indoor radio channel in ad hoc networks," in *IEEE Wireless Comm. and Networking Conf.*, March 2002, pp. 608–612.

[21] J. Ash and N. Patwari. Localization CRB Code and Measurement repository. [Online]. Available: http://www.eecs.umich.edu/~hero/locate.html

[22] N. Patwari, A. O. Hero III, M. Perkins, N. Correal, and R. J. O'Dea, "Relative location estimation in wireless sensor networks," *IEEE Trans. Sig. Proc.*, vol. 51, no. 8, pp. 2137–2148, Aug. 2003.

[23] G. D. Durgin, *Space-Time Wireless Channels*. Prentice Hall PTR, 2002.

[24] H. Hashemi, "The indoor radio propagation channel," *Proceedings of the IEEE*, vol. 81, no. 7, pp. 943–968, July 1993.

[25] T. S. Rappaport, *Wireless Communications: Principles and Practice*. New Jersey: Prentice-Hall Inc., 1996.

[26] N. Patwari, R. J. O'Dea, and Y. Wang, "Relative location in wireless networks," in *IEEE VTC*, vol. 2, May 2001, pp. 1149–1153.

[27] A. J. Coulson, A. G. Williamson, and R. G. Vaughan, "A statistical basis for lognormal shadowing effects in multipath fading channels," *IEEE Trans. on Veh. Tech.*, vol. 46, no. 4, pp. 494–502, April 1998.

[28] C. Liu, K. Wu, and T. He, "Sensor localization with ring overlapping based on comparison of received signal strength indicator," in *IEEE Mobile Ad-hoc and Sensor Systems (MASS)*, Oct. 2004.

[29] G. Carter, *Coherence and Time Delay Estimation*. Piscataway, NJ: IEEE Press, 1993.

[30] C. Knapp and G. Carter, “The generalized correlation method for estimation of time delay,” *IEEE Transactions on Acoustics, Speech, and Signal-Processing*, vol. 24, no. 4, pp. 320–327, 1976.

[31] E. Robinson and A. Quazi, “Effect of sound-speed profile on differential time-delay estimation.” *Journal of the Acoustical Society of America*, vol. 77, no. 3, pp. 1086–1090, 1985.

[32] K. Pahlavan, P. Krishnamurthy, and J. Beneat, “Wideband radio propagation modeling for indoor geolocation applications,” *IEEE Comm. Magazine*, pp. 60–65, April 1998.

[33] IEEE Wireless Standards Zone. [Online]. Available: http://standards.ieee.org/wireless

[34] J. Borràs, P. Hatrack, and N. B. Mandayam, “Decision theoretic framework for NLOS identification,” in *IEEE VTC 1998*, vol. 2, May 1998, pp. 1583–1587.

[35] P.-C. Chen, “A non-line-of-sight error mitigation algorithm in location estimation,” in *IEEE Wireless Comm. and Networking Conf.*, Sept. 1999, pp. 316–320.

[36] F. Sivrikaya and B. Yener, “Time synchronization in sensor networks: a survey,” *IEEE Network*, vol. 18, no. 4, pp. 45–50, July-Aug. 2004.

[37] S. Kim, T. Pals, R. A. Iltis, and H. Lee, “CDMA multipath channel estimation using generalized successive interference cancellation algorithm for radiolocation,” in *36th Conf. Information Sciences & Systems*, March 2002.

[38] J.-Y. Lee and R. A. Scholtz, “Ranging in a dense multipath environment using an UWB radio link,” *IEEE J. Sel. Areas Commun.*, vol. 20, no. 9, pp. 1677–1683, Dec. 2002.

[39] E. G. Larsson, “Cramér-Rao bound analysis of distributed positioning in sensor networks,” *IEEE Signal Processing Letters*, vol. 11, no. 3, pp. 334–337, March 2004.

[40] D.-H. Shin and T.-K. Sung, “Comparisons of error characteristics between TOA and TDOA positioning,” *IEEE Trans. Aerospace Electr. Systems*, vol. 38, no. 1, pp. 307–311, Jan. 2002.

[41] B. D. Van Veen and K. M. Buckley, “Beamforming: A versatile approach to spatial filtering,” *IEEE ASSP Magazine*, vol. 5, no. 2, pp. 4–24, Apr. 1988.

[42] P. Stoica and R. L. Moses, *Introduction to Spectral Analysis*. New Jersey: Prentice-Hall, 1997.

[43] B. Ottersten, M. Viberg, P. Stoica, and A. Nehorai, “Exact and large sample ML techniques for parameter estimation and detection in array processing,” in *Radar Array Processing*, S. S. Haykin, J. Litva, and T. Shepherd, Eds. Springer-Verlag, 1993, pp. 99–151.

[44] J. C. Chen, K. Yao, and R. E. Hudson, “Source localization and beamforming,” *IEEE Signal Processing Mag.*, vol. 19, no. 2, pp. 30–39, Mar. 2002.

[45] H. Xu, V. Kukshya, and T. S. Rappaport, “Spatial and temporal characterization of 60 GHz indoor channels,” in *IEEE VTC Fall*, vol. 1, Sept. 2000, pp. 6–13.

[46] S. L. Collier and D. K. Wilson, “Performance bounds for passive sensor arrays operating in a turbulent medium: Plane-wave analysis,” *The Journal of the Acoustical Society of America*, vol. 113, no. 5, pp. 2704–2718, 2003. [Online]. Available: http://link.aip.org/link/?JAS/113/2704/1

[47] J. N. Ash and L. C. Potter, “Sensor network localization via received signal strength measurements with directional antennas,” in *Proceedings of the 2004 Allerton Conference on Communication, Control, and Computing*, 2004.

[48] R. L. Moses, D. Krishnamurthy, and R. Patterson, "An auto-calibration method for unattended ground sensors," in *ICASSP*, vol. 3, May 2002, pp. 2941–2944.

[49] A. Catovic and Z. Sahinoglu, "The Cramér-Rao bounds of hybrid TOA/RSS and TDOA/RSS location estimation schemes," *IEEE Communications Letters*, vol. 8, no. 10, pp. 626–628, Oct. 2004.

[50] H. L. V. Trees, *Detection, Estimation, and Modulation Theory, Part I*. John Wiley & Sons, 1968.

[51] A. O. Hero III, J. A. Fessler, and M. Usman, "Exploring estimator bias-variance tradeoffs using the uniform CR bound," *IEEE Trans. on Sig. Proc.*, vol. 44, no. 8, pp. 2026–2041, Aug. 1996.

[52] D. J. Torrieri, "Statistical theory of passive location systems," *IEEE Trans. on AES*, vol. AES-20, no. 2, pp. 183–198, Mar. 1984.

[53] M. A. Spirito, "On the accuracy of cellular mobile station location estimation," *IEEE Trans. on Veh. Tech.*, vol. 50, no. 3, pp. 674–685, May 2001.

[54] N. Bulusu, J. Heidemann, and D. Estrin, "GPS-less low cost outdoor localization for very small devices," *IEEE Personal Comm.*, vol. 5, no. 5, pp. 28–34, Oct. 2000.

[55] L. Doherty, K. S. J. pister, and L. E. Ghaoui, "Convex position estimation in wireless sensor networks," in *IEEE INFOCOM*, vol. 3, 2001, pp. 1655 –1663.

[56] P. Biswas and Y. Ye, "Semidefinite programming for ad hoc wireless sensor network localization," Dept. of Computer Science, Stanford University, Tech. Rep., Sept 2003.

[57] ——, "A distributed method for solving semidefinite programs arising from ad hoc wireless sensor network localization," Dept. of Computer Science, Stanford University, Tech. Rep., Oct. 2003.

[58] J. B. Tenenbaum, V. de Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, pp. 2319–2323, Dec 2000.

[59] Y. Shang, W. Ruml, Y. Zhang, and M. P. J. Fromherz, "Localization from mere connectivity," in *Mobihoc '03*, June 2003, pp. 201–212.

[60] Y. Shang and W. Ruml, "Improved MDS-based localization," in *IEEE Infocom '04*, March 2004, pp. 2640–2651.

[61] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by local linear embedding," *Science*, vol. 290, pp. 2323–2326, Dec 2000.

[62] N. Patwari and A. O. Hero III, "Manifold learning algorithms for localization in wireless sensor networks," in *IEEE ICASSP*, vol. 3, May 2004, pp. 857–860.

[63] S. N. Simić, "A learning-theory approach to sensor networks," *IEEE Pervasive Computing*, vol. 2, no. 4, pp. 44–49, 2003.

[64] D. Niculescu and B. Nath, "Ad hoc positioning system," in *IEEE Globecom 2001*, vol. 5, April 2001, pp. 2926 –2931.

[65] A. Savvides, H. Park, and M. B. Srivastava, "The bits and flops of the N-hop multilateration primitive for node localization problems," in *Intl. Workshop on Sensor Nets. & Apps.*, Sept. 2002, pp. 112–121.

[66] R. Nagpal, H. Shrobe, and J. Bachrach, "Organizing a global coordinate system from local information on an ad hoc sensor network," in *2nd Intl. Workshop on Inform. Proc. in Sensor Networks*, April 2003.

[67] J. J. Caffery, *Wireless Location in CDMA Cellular Radio Systems*. Kluwer Academic Publishers, 1999.

[68] J. Albowicz, A. Chen, and L. Zhang, "Recursive position estimation in sensor networks," in *IEEE Int. Conf. on Network Protocols*, Nov. 2001, pp. 35–41.

[69] C. Savarese, J. M. Rabaey, and J. Beutel, “Locationing in distributed ad-hoc wireless sensor networks,” in *ICASSP*, May 2001, pp. 2037–2040.

[70] A. Savvides, C.-C. Han, and M. B. Srivastava, “Dynamic fine-grained localization in ad-hoc networks of sensors,” in *IEEE Mobicom*, July 2001.

[71] S. Čapkun, M. Hamdi, and J.-P. Hubaux, “GPS-free positioning in mobile ad-hoc networks,” in 34^{th} *IEEE Hawaii Int. Conf. on System Sciences (HICSS-34)*, Jan. 2001.

[72] A. T. Ihler, J. W. Fisher III, and R. L. Moses, “Nonparametric belief propagation for self-calibration in sensor networks,” in *IEEE ICASSP 2004*, May 2004.

[73] M. Coates, “Distributed particle filters for sensor networks,” in *IEEE Information Processing in Sensor Networks (IPSN)*, April 2004, pp. 99–107.

[74] L. Hu and D. Evans, “Localization for mobile sensor networks,” in *IEEE Mobicom 2004*, Sept. 2004, pp. 45–57.

[75] S. Thrun, “Probabilistic robotics,” *Communications of the ACM*, vol. 45, no. 3, pp. 52–57, March 2002.

[76] K. Langendoen and N. Reijers, “Distributed localization in wireless sensor networks: A quantitative comparison,” *Elsevier Science*, May 2003.

[77] V. Cehver and J. H. McLellan, “Sensor array calibration via tracking with the extended Kalman filter,” in *Fed. Lab. Symp. on Adv. Sensors*, Mar. 2001, pp. 51–56.

[78] A. Galstyan, B. Krishnamachari, K. Lerman, and S. Pattem, “Distributed online localization in sensor networks using a moving target,” in *IEEE Information Processing in Sensor Networks (IPSN)*, April 2004, pp. 61–70.

[79] N. Patwari and A. O. Hero III, “Using proximity and quantized RSS for sensor localization in wireless networks,” in *IEEE/ACM 2nd Workshop on Wireless Sensor Nets. & Applications*, Sept. 2003, pp. 20–29.

[80] T. He, C. Huang, B. M. Blum, J. A. Stankovic, and T. Abdelzaher, “Range-free localization schemes for large scale sensor networks,” in *IEEE Mobicom 2003*, Sept. 2003.

[81] A. Rao, S. Ratnasamy, C. Papadimitriou, S. Shenker, and I. Stoica, “Geographic routing without location information,” in *IEEE Mobicom 2003*, Sept. 2003.