Chapter 1

LOCALIZATION IN SENSOR NETWORKS

Andreas Savvides and Mani Srivastava
Electrical Engineering Department
University of California, Los Angeles
{asavvide,mbs}@ee.ucla.edu

Lewis Girod and Deborah Estrin
Computer Science Department
University of California, Los Angeles
{girod,destrin}@lecs.cs.ucla.edu

Abstract The development of large scale distributed sensor systems is a significant scientific and engineering challenge, but they show great promise for a wide range of applications. The capability to sense and integrate spatial information with other elements of a sensor application is critical to exploring the full potential of these systems. In this article we discuss the range of application requirements, introduce a taxonomy of localization mechanisms, and briefly discuss the current state of the art in ranging and positioning technologies. We then introduce two case studies that illustrate the range of localization applications.

Keywords: Wireless Sensor Network, Localization, Position Estimation, Ranging

Introduction

The development of large scale distributed sensor systems is a significant scientific and engineering challenge, but they show great promise for a wide range of applications. By placing sensors close to the phenomena they sense, these systems can yield increased signal quality, or equivalent signal quality at reduced cost. By reducing deployment overhead, whether in terms of unit cost or installation time, they enable the sensors to be placed in greater numbers.

Relative to other types of distributed systems, distributed sensor systems introduce an interesting new twist: they are coupled to the physical world, and their spatial relationship to other objects in the world is typically an important factor in the task they perform. The term **localization** refers to the collection of techniques and mechanisms that measure these spatial relationships.

When raw sensor data is combined with spatial information, the value of the data and the capability of the system that collects it increases substantially. For example, a collection of temperature readings without location information is at best only useful to compute simple statistics such as the average temperature. At worst, analysis of the data might yield incorrect conclusions if inaccurate assumptions are made about the distribution of physical sampling. By combining the data with location information, the resulting temperature map can be analyzed much more effectively. For instance, statistics can be computed in terms of spatial sampling rather than the count of sensor readings, and the confidence of the results can be assessed more meaningfully. Location also opens up entirely new application possibilities: a model for heat transfer can be applied to filter out noise and pinpoint the location of heat sources.

This simple example is intended to illustrate a more general point. As anyone who has worked with distributed sensor systems is painfully aware, there is a high cost in moving from a centralized, wired application to a large-scale, distributed, wireless application. The application is certain to grow in complexity; new techniques must be developed, new protocols deployed, and the application must be resilient to the whims of nature. In addition, there are the more mundane details of dealing with large numbers of independent parts, each of which needs the right software version and fresh batteries, and each of which can independently fail. But what makes all of this effort worthwhile is the ability to deploy applications that collect data that could never be collected before, and for this we need localization.

In this article, we will first present a range of application requirements in Section 1.1. In Section 1.2, we will introduce a taxonomy of localization mechanisms to satisfy those requirements. In Section 1.3, we will summarize the state of the art in available ranging and positioning technologies. Section 1.4 discusses localization in multihop networks and section 1.5 examines how error behaves with respect to different network parameters. Section 1.6 concludes the chapter with a description of two ad-hoc localization systems.

1. Application Requirements

The field of networked sensor systems encompasses a very broad array of applications, with a broad range of requirements. Often different application requirements can motivate very different systems. While these differences are sometimes tunable parameters, often they are significant structural choices. For example, adding a low-power requirement rules out many possible designs from the start.

To introduce this variety, we first present several points in the application space, and then enumerate a set of requirements axes that characterize the requirements space of localization systems.

Passive Habitat Monitoring

Scientists in numerous disciplines are interested in methods for tracking the movements and population counts of animals in their natural habitat. While there are various techniques currently employed (e.g. rings on birds' claws, "drop buckets" for small animals on the ground), these techniques do not scale well in terms of the time required of experimenters. One of the open challenges in this field is to develop an automated system that can build a record of the passage and habits of a particular species of animal, without disturbing it in its natural habitat.

One possible solution might be built around a passive source localization and species identification system. Such a system would detect and count animals by localizing the sounds they make, then training a camera system on them to aid in counting.

Sensor nodes equipped with microphones would be distributed through the target environment. When an acoustic source is detected by a node, it communicates with nearby nodes to try to estimate the location of the source by comparing the times of arrival of the signals. Analysis techniques such as beam-forming [18] might apply to this application, along with species recognition techniques to filter out acoustic sources not relevant to the task.

From this application we can derive a number of requirements:

- **Outdoor Operation.** The system must be able to operate outdoors, in various weather conditions.
- **Power Efficiency.** Power may be limited, whether by battery lifetime or by the feasibility of providing sufficient solar collectors.
- **Non-cooperative Target, Passive Infrastructure.** The animal does not emit signals designed to be detected by the system (i.e. non-cooperative), and the system does not emit signals to aid in localization.

- **Accuracy.** The system must be accurate enough to be able to produce a reliable count, and to accurately focus a camera on suspected locations.
- **Availability of Infrastructure.** In some cases GPS may be available to localize the sensors themselves. However, in many cases sensors will need to be placed under canopies where GPS signals are unavailable. In these cases, if surveying the sensors is inconvenient, the sensors will need to self-localize. The self-localization system may have a different set of requirements.

Smart Environments

Smart environments are a second class of applications where location awareness is a key component. Smart environments are deeply instrumented systems with very demanding localization requirements. These systems need localization for two different purposes. First, rapid installation and self-configuration of a set of infrastructure "beacons" is required to reduce installation cost and increase flexibility. Second, very fine-grained localization and tracking of the system components is required during normal system operation.

The operation requirements can also be assessed along similar axes:

- **Indoor Operation.** The system must operate indoors. While the weather indoors is generally predictable, there are typically many reflectors that cause multipath interference for both RF and acoustic signals. If the environment is an office environment, acoustic signals should be outside of the range of human hearing. The system will need to operate in the presence of obstacles.
- **Power Requirements.** Some of the infrastructure components may be close to power sources, but a large number of the system components should be untethered and free to move around the space. This implies that the majority of the system components should be as low power as possible.
- **Cooperative Target, Active Infrastructure.** Because the target badge localizes itself relative to the beacons and reports its location, we consider it to be cooperative. The beacon infrastructure is active because it emits periodic signals that the badge receives.
- **Accuracy.** Fine-grained localization of people and objects with 10cm accuracy may be required by many systems.
- **Availability of Infrastructure.** Infrastructure may be present, but infrastructure installation usually becomes a dominant cost

factor. Ideally the infrastructure should be self-configuring in a way that reduces installation cost.

Axes of Application Requirements

These points in the application space demonstrate a broad spectrum of requirements that seem to fall along nine independent axes:

- **Granularity and Scale of Measurements.** What is the smallest and largest measurable distances? For instance, local coordinate systems for a sensor network might scale from centimeters to hundreds of meters, whereas GPS coordinates have a global scale and a granularity on the order of meters.
- **Accuracy and Precision.** How close is the answer to ground truth (accuracy), and how consistent are the answers (precision)?
- **Relation to Established Coordinate System.** Are absolute positions needed, or is a relative coordinate system sufficient? Do locations need to be related to a global coordinate system such as GPS, or an application specific coordinate system such as a forest topography or building floor-plan?
- **Dynamics.** Are the elements being localized fixed in place or mobile? Can a static infrastructure be assumed? What refresh rate is needed? Is motion estimation required?
- **Cost.** Node hardware cost, in terms of both power consumption and monetary cost; Latency of localization mechanism; Cost of installing infrastructure (if needed), in terms of power, money, and labor.
- **Form Factor.** How large can a node be? If the node has multiple sensors separated by a baseline, what kind of baseline is required for sensors to work effectively?
- **Communications Requirements.** What kind of coordination is required among nodes? What assumptions does the system make about being able to send or receive messages at any time? What kind of time synchronization is needed? Does the algorithm rely on the existence of a cluster head or a "microserver"?
- **Environment.** How sensitive is a given technique to environmental influences, and in what range of environments does it work? For instance, indoors (multipath), outdoors (weather variations), underwater, or on Mars?

- **Target Cooperation, System Passivity.** Does the target play a cooperative role in the system? Can the system emit signals into the environment without interfering with the task at hand?

Given such a complex requirement space, it seems that few system designs, techniques or technologies will fit that space uniformly. This complicates efforts to develop reusable designs and components, especially in the early stages of a research program where the canonical applications are not well understood. It also means that the importance of a new technique or system must be evaluated in an appropriate context. In order to equitably compare different mechanisms and systems, we need a taxonomy of mechanisms and system structures to provide context for our comparisons.

2. Taxonomy of Localization Mechanisms

The capsular conclusion of the last section was simple: localization systems will differ not only in details of algorithms and protocols, but also fundamentally in the structure of their system and in the assumptions they make. The challenge of this section is therefore to construct a taxonomy of general system structures that capture the breadth and depth of the solution space. Having done this, hopefully we can better classify localization systems and components for the purposes of comparison and contrast.

Classifying our Examples

Returning to our examples of the previous section, we can motivate a taxonomy by summarizing their important structural features. Perhaps the distinction that stands out the most among our example applications has to do with "cooperative" targets. In the case of an animal localizer, we can't assume that the animal is acting with the intent of being localized. In fact, under some circumstances, targets may intentionally avoid detection. This property is in sharp contrast with the case of locating badges. While a small child might attempt to subvert the localization system, we can assume that the badges themselves will be cooperative in whatever ways will simplify the implementation of the system.

This distinction has numerous repercussions on the overall design of the system. First, while in the non-cooperative case the animal detector must be constantly vigilant, at considerable energy cost, the cooperative smart badge can take advantage of a simple mechanical motion detector (e.g. a jiggle switch) to implement a simple zero-power wakeup mechanism.

Second, while the characteristic sounds of animal calls and motions may be detectable, the detection process is more complex, and ultimately more failure-prone than the detection of a synthetic ranging signal. Not only will the processing be more expensive, both in terms of processing and communications costs, the false positive rate will tend to be higher. In contrast, a cooperative system has the luxury of designing its signals to be detected efficiently and accurately, e.g. low-autocorrelation codes and a simple matched filter detector at each receiver.

Third, a non-cooperative system can only use Time Difference of Arrival (TDoA) techniques to estimate position, because the true time of the signal emission is not known. In contrast, a cooperative system can sometimes use an out-of-band synchronization protocol to establish a consistent timebase, and then provide receivers with the send time so that they can measure Time of Flight (ToF).

These differences result in important structural differences among our example systems. We place the habitat monitoring application in the "Passive Target Localization" category. The smart environment application can be broken into two phases, that fall into different categories: a "bootstrapping" phase in which the infrastructure of beacons self-organizes into a coordinate system, and a "service" phase in which the badges localize themselves with respect to the beacon infrastructure. The bootstrapping phase fits into the "Cooperative Target" category, while the service phase fits into "Cooperative Infrastructure".

A Taxonomy of Localization Systems

We have discussed examples representing three categories of localization system. Extending this process, we can identify six categories of localization system, that appear to cover all the systems of which we are currently aware. These six categories are partitioned into "active" and "passive".

Active Localization. Active localization techniques emit signals into the environment that are used to measure range to the target. These signals may be emitted by infrastructure components or by targets. Within the category of active localization there are three sub-categories:

- **Non-cooperative.** In an active, non-cooperative system, system elements emit ranging signals, which are distorted or reflected in flight by passive elements. The system elements then receive the signals and analyze them to deduce their location relative to passive elements of the environment. Examples include radar systems and reflective sonar systems often used in robotics.

- **Cooperative Target.** In a cooperative target system, the targets emit a signal with known characteristics, and other elements of the system detect the signals and use information about the signal arrivals to deduce the target's location. Often a cooperative target system also involves some synchronization mechanism to readily compute signal ToF. This category includes both infrastructure-less systems and systems that localize with respect to infrastructure receivers. Infrastructure-based systems include the ORL Active Bat and the service phase of the GALORE localization system [7]. Infrastructure-less systems include the bootstrapping phases of both the GALORE and Smart Kindergarten systems.

- **Cooperative Infrastructure.** In a cooperative infrastructure system, elements of the infrastructure emit signals that targets can receive. The infrastructure itself is assumed to be carefully configured and synchronized to simplify the processing done by the target. Another property of this system structure is that receivers can compute their own location passively, without requiring any interaction with the infrastructure. Examples of this type of system include GPS and the MIT Cricket system [14], and the service phase of the Smart Kindergarten system.

Passive Localization. Passive localization techniques differ from active ones in that they discover ranges and locations by passively monitoring existing signals in a particular channel. The term "passive" does not imply that they emit *no* signals, only that the signals they emit are outside the channel that is primarily analysed for time-of-flight measurement. For example, a technique that uses RF signals for synchronization and coordination, but measures range by TDoA of ambient acoustic signals would still be considered passive.

- **Blind Source Localization.** In a blind source localization system, a signal source is localized without any *a priori* knowledge of the type of signal emitted. Typically this is done by "blind beam-forming", which effectively cross-correlates the signals from different receivers. These techniques generally only work so long as the signals being compared are "coherent", which in practice often limits the spacing of receivers because of signal distortion induced by the environment. Coherent combining techniques can generally localize the most prominent source within the convex hull of a sensor laydown, or alternatively can compute a bearing angle to a distant source, but not a range or location. This work is described by Yao et. al.[18].

- **Passive Target Localization.** Similar to blind localization, a passive target localization system is usually based on coherent combination of signals, with the added assumption of some knowledge of the source. By assuming a model for the signals generated by the source, filtering can be applied to improve the performance of the algorithms and to reduce the computational and communications requirements. Examples include our previous example of habitat monitoring, UCLA work on beamforming [19], and some E911 cell phone location proposals.

- **Passive Self-localization.** In passive self-localization, existing beacon signals from known infrastructure elements are used by a target to passively deduce its own location. Most commonly, properties of RF signals from base stations are used to deduce location of a mobile unit. Examples include RADAR [1], which measured RSSI to different 802.11 access points, and the work of Bulusu et. al.[2], which measured RSSI to Ricochet transmitters.

Cross-cutting Issues. As a rule, active and cooperative techniques tend to be more accurate, more efficient, and generally more effective. Because cooperative techniques can design both the receiver and transmitter, the designs can be optimized for performance much more effectively. Cooperative systems can also synchronize explicitly, improving the performance of ranging based on signal propagation time. However, applications such as habitat monitoring can only be addressed using passive techniques. Although passive techniques are attractive because they can leverage existing signaling, they often perform poorly when the signaling is not designed with ranging or localization in mind.

Another aspect of sensor network localization that cuts across these categories is an ability to support ad-hoc deployment and operation. In an ad-hoc setup, there is no guarantee that all the sensor nodes will be in communication and sensing range to each other, nor that the sensing and communications properties will remain constant over time. Thus, regardless of category, systems that can operate in an ad-hoc fashion must collaborate across the sensor nodes, must operate within a multi-hop network, and must react to system dynamics. These issues will be addressed further in Section 1.4.

3. Ranging Technologies

When designing a localization system, an important factor in the design are the mechanisms used to measure physical distances and angles. Typically for cooperative systems this will involve some kind of emitter and detector pair. The selection of these elements has a significant impact on how well the final system will fit the application requirements. In this section, we will discuss the relative merits of three types of ranging mechanism, based on visible light, radio signals, and acoustic signals.

Ranging using RF

RF ranging generally follows one of two approaches: distance measured based on received signal strength, and distance measured based on the ToF of the radio signal.

RF RSS. Received Signal Strength (RSS) is roughly a measure of the amplitude of a detected radio signal at a receiver. If we assume a model for path loss as a function of distance, the received signal strength should generally decrease as a function of distance. The path loss model is highly dependent on environmental factors: in open space the model is $1/R^2$; near the ground, the model is closer to $1/R^4$. Under some conditions (e.g. waveguides, corridors, etc.), path loss can actually be lower than in free space, e.g. $1/R^{1.5}$. Because the path loss model is dependent on details of the environment, automatically choosing a valid model can be difficult.

In practice, the behavior of RSS is dependent on a number of factors, not the least of which is simply whether the RSS estimator is well designed. Some radios, such as the RFM radio used in the Berkeley mote, just sample the baseband voltage to estimate RSS, which is a very crude measurement; other radios have a more capable measurement circuit. Another important factor is the frequency range used by the radio. Multipath fading is a change in RSS caused by the constructive or destructive interference of reflected paths. Multipath fading is dependent on the environment and is in many cases dependent on the frequency of the signals being transmitted. A radio that uses just one frequency will be more susceptible to multipath fades that will cause substantial error in a distance estimate based on RSS. If a variety of frequencies are used and appropriate filtering is applied, the effects of frequency dependent multipath fading may be removed from the RSS estimate, although frequency independent multipath such as ground reflection will still be present. In general, the effectiveness of RSS estimation varies from system to system and cannot be implemented without hooks into the internals of the radio hardware.

Another difficulty with using RSS is that the transmit power at the sender may not be accurately known. In many cases this is a function of specific component values on a given board, or a function of battery voltage. Without knowing the original transmit power it it may not be possible to correctly estimate the path loss.

RF Time of Flight. Measuring the time of flight of radio signals is another possible solution. Because the ToF of a radio signal is not very dependent on the environment, ToF approaches can be much more precise than approaches based on measuring RSS. The two main challenges in implementing an RF ToF scheme are: (1) synchronization must use signals also traveling at the speed of light, and (2) to achieve high precision ToF measurements require high frequency RF signals and fast, accurate clocks.

The timing and synchronization issues are the central problems with RF ToF ranging. The synchonization problem is simplified for infrastructure-based systems where elements of the infrastructure can be synchronized by some out-of-band mechanism, or by taking into account knowledge of their exact locations. However, this does not eliminate the need for accurate clocks, which tend to be expensive in terms of power. For example, GPS satellites carry atomic clocks for timing, but these clocks are continually adjusted to account for relativistic effects as they orbit the earth, and the trajectory of the satellites is carefully measured. In more down-to-earth implementations, in most RF ToF ranging implemenations, the infrstructure elements are connected by carefully measured cables to achieve synchronization.

However, for ad-hoc deployments, the lack of a common timebase means that "round trip" messages must be used in order to compare send and receive times within the same timebase. The time spent "waiting" at the remote transponder must then be measured and subtracted. This requires that the two systems' clocks be running at close to the same rate, or that the turnaround time be a fixed constant. Getting these details to work correctly, given the fact that all the timing must be very precise (30 cm of error per ns), can be quite challenging from a hardware perspective.

Probably the best path toward an ad-hoc RF ToF solution is to leverage the hardware of a sufficiently advanced radio system, such as 802.11 or an ultra-wide band (UWB) receiver. Because these systems operate at high bit rates and must match clock rates in order to inter-operate, there is a better chance that it is possible to exploit their features to implement ranging. Some 802.11 chipsets have ranging support, although at press time we do not have any references to implementations that use them successfully. UWB ranging solutions have been advertised, however because of licensing restrictions and other issues there have been no documented implementations in the sensor networks space.

Bearing estimates for incoming RF signals require similar types of designs. By implementing a radio with an array of antennae, the signals from those independent reception points can be compared to estimate direction of arrival (DoA). While these techniques are commonplace in radar systems and commercial wireless systems, currently this kind of feature is not available on small, low-power platforms. However, similar designs to those that enable ranging might someday also enable DoA estimation.

Ranging using Acoustics

Acoustic ranging is probably the most developed ranging technology in use in sensor networks. There are a number of factors that make acoustics attractive, given currently available COTS components. Acoustic transducers are easy to interface, and simple, inexpensive detector chipsets are available for ultrasound. However the key advantage to using acoustics is that timing and synchronization is much easier to implement. A 32 KHz clock is sufficient to achieve ranging accuracy to 1 cm, and synchronization between sender and receiver can be implemented using most radio modules without modification.

In terms of power, acoustics performs quite well, even near the ground. Whereas RF communication suffers r^4 path loss near the ground because ground reflections are phase-shifted by 180 degrees, this is not the case for acoustic waves. Acoustic path loss near the ground under good conditions is much closer to r^2. Outdoors, acoustics is susceptible to interference from weather conditions, such as wind that causes noise, and convective updrafts that carry signals up and away from the ground.

However, acoustics has a few disadvantages as well. First, acoustic emitters tend to be physically large, especially if they emit low frequencies. The other main disadvantage is that acoustic signals are stopped by solid obstructions. However, for some applications this can be advantageous, such as the case of an asset tracking system which only needs to know which room the asset is in.

When using acoustics, a wide band of frequencies are available for use. Some systems are based on ultrasound frequencies (typically 40 KHz to 1 MHz), while others are based on audible frequencies (100 Hz to 20 MHz). Some systems use tuned piezo emitters at specific frequencies, while others use wide-band acoustic signals. The choice of frequency depends on the application (e.g. is audible sound acceptable), as well as the environment.

Experience with 40 KHz ultrasound systems outdoors indicates a typical range of about 10 meters at a voltage of 3 volts, and about 16 meters at 16 volts. The type of emitter used also has a significant effect on the performance of the system. Many ultrasound emitters are directional, substantially increasing their output in a conical beam. This can be disadvantageous from a packaging perspective, as it may require many emitters and receivers in order to support ad-hoc deployment.

Audible acoustics can be very effective outdoors, because of the wide diversity of wavelengths possible. A wide-band signal will be more robust to environmental interference, because of the process gain in the detection process. A wide-band signal is also less susceptible to narrow-band sources of noise, as well as absorption and scattering of specific frequencies.

Under ideal weather conditions, audible ranging systems have been shown to achieve ranges as large as 100m for power levels of 1/4 Watt. High power emitters such as heavy vehicles are detectable at ranges of 10's of kilometers. Acoustic range is longest at night when the air is still and cool. The worst conditions for acoustics are warm, sunny afternoons, when heated air near the ground rises and deflects signals up and away from other ground-based receivers. Under these conditions, the same acoustic system might achieve only 10m range.

Errors in line-of-sight (LoS) acoustic ranges tend in general to be independent of distance, up to the limit of the signal detector. However, when obstructions or clutter are present, severe attenuation can be observed, as well as radical outliers when the LoS path is completely blocked and a reflected path is detected. When designing positioning algorithms around an acoustic ranging system it is important to take these issues into account.

Bearing estimates for acoustic signals can often be implemented without much difficulty using simple hardware and software solutions. If the baseline between sensors is known with sufficient accuracy, a bearing estimate can be derived from the time difference of arrivals. There are several examples of implemented systems that measure DoA using acoustics, the MIT software compass[20], and beam-forming systems[19].

4. Positioning in Multihop Sensor Networks

The importance and plethora of applications in multihop sensor networks, motivated the development of diverse set of positioning algorithms. The ability to operate in a multihop regime allows nodes with short-range signal transmissions to collaboratively localize themselves across larger areas. These properties make multihop ad-hoc localization an appealing choice in ad-hoc deployed sensor networks, rapidly installable infrastructures and fine-grained localization in indoor settings where the multihop and ad-hoc nature of the system can compensate for the presence of obstacles and many other settings where other infrastructure based technologies such as GPS cannot operate.

Challenges in multihop Ad-Hoc Sensor Networks

Despite the attractiveness of ad-hoc multihop localization, the application requirements need to be carefully reviewed before any design choices are made. Unfortunately the flexibility promised by such localization systems is also coupled with large set of challenges and trade-offs that have so far inhibited their widespread deployment. Some of these challenges are listed here.

Physical Layer Challenges. As described in the previous section, measurements are noisy and can fluctuate with changes in the surrounding environment.

Algorithm Design Challenges. The algorithm designer needs concurrently consider multiple issues when designing such systems.

- **Noisy measurements** call for the use of optimization techniques that minimize the error in position estimates. Despite the well-established body of knowledge in optimization techniques, the use of any optimization algorithm is only as good as the validity of the assumptions on the underlying measurement error distribution in the actual deployment scenario.

- **Computation and communication trade-offs.** Cost and energy limitations force designers to consider the development of lightweight distributed algorithms that can operate on low cost resource constrained nodes, where the computation is performed inside the network.

- **Problem setup.** A large variety of problem setups has appeared in the literature. Some approaches consider the use of a small percentage of location aware anchor nodes spread randomly distributed inside the network. Some other approaches, suggest that one should ensure that enough anchor nodes are placed on the network perimeter, while some others advocate anchor free setups.

Furthermore, the type of measurements used in each case, vary across different solutions, some try to infer locations based on mere connectivity information while others, use angular and/or distance measurements.

- **Error behavior and scalability.** Perhaps the most overlooked aspect of multihop localization in currently proposed solutions is understanding how the network parameters affect the resulting position error behavior and scalability. Network topology and geometry between nodes, network density, ranging accuracy, anchor node concentration and uncertainty in anchor node locations, affect the quality of location estimates; therefore their behavior needs to be formally understood.

System Integration Challenges. All the previously discussed requirements imply a non-trivial system integration effort. Many off-the-shelf measurement technologies are not directly suitable for use in sensor networks, so customized hardware and software often needs to be developed to make a functional system.

Overview of Multihop Localization Methods

Despite the numerous proposals, very few ad-hoc localization systems have been built and evaluated in practice. Furthermore, the side-by-side comparison of different approaches is a non-trivial task due to the differences in problem setup and underlying assumptions. In the remainder of this section we highlight some of the recently proposed approaches by broadly classifying them as connectivity based and measurement based approaches. Later on we also comment on some of the trends associated on position error based Cramér Rao bound analysis.

Radio Connectivity Based Approaches

Connectivity-based approaches try to leverage radio connectivity to infer node locations. Although radio connectivity alone cannot provide fine-grained localization, it can provide a good indication of proximity that is useful in supporting other network level tasks such as geographic routing. The GPS-less low cost localization system described in [2] is an example of a connectivity based system. In this system, a set of predeployed, location aware reference nodes transmit spatially overlapped beacon signals. Other nodes with unknown locations can localize themselves at the centroid of the reference nodes from which they can receive beacon signals. The best results are obtained when the nodes are arranged in a mesh pattern.

The convex position estimation approach proposed by Doherty et. al. in [4] also localizes nodes using radio connectivity. In this case the localization is formulated as a linear or semi definite program that is solved at a central location. This approach also requires a set of nodes with known locations to act as beacons. With careful placement of the beacon nodes on the perimeter of the network the authors have shown that node locations between 0.64 and 0.72R (where R is the radio transmission range) are possible at density of 5.6 neighbors per node.

A more recent proposal based on multidimensional scaling (MDS) can solve for the relative position of the nodes with respect to each other without requiring any beacon nodes [16]. This is done by using a classical MDS formulation that takes node connectivity as inputs and creates a two dimensional relative map of the nodes that preserves the neighborhood information. The connectivity only approach uses hop distances between nodes to initialize a distance matrix. The same MDS formulation can take more accurate inter-node distances to construct more accurate maps.

Measurement Based Approaches

Measurement based approaches build upon a wide range of measurement technologies. While different approaches focus on specific ranging systems, a large source of disparity in measurement-based algorithms stems from different assumptions about measurement error distribution. Some systems assume additive Gaussian noise, while others assume that measurement error is proportional to distance. Furthermore, some algorithms require a set of initial anchor nodes, whereas others perform relative localization and use anchors only at the end of the localization process to translate the derived relative coordinate system to an absolute coordinate system. Because of these reasons, in this section we do not attempt a direct comparison of existing approaches, instead highlighting the key features of each approach. We begin with anchor free approaches, followed by approaches that use anchor nodes.

An example algorithm that does not require anchor nodes is described in [3]. This relative localization system is based on radio ToF measurements and uses geometric relationships to estimate node positions. First, all the nodes compute their locations with respect to their neighbors. The resulting local coordinate systems are then aligned and merged into a global coordinate system using a simple set of geometric relationships. The position estimates acquired by this method are not very accurate due to the noisy ToF measurements and error propagation. Despite this loss of precision, the location estimates are still adequate to help with network level tasks such as geo-routing.

Another notable anchor-free localization method has been developed by [9]. In this work Moses et. al. have shown that sensor node positions and orientations can be estimated using signals from acoustic sources with unknown locations. Each acoustic source generates a known acoustic signal that is detected by the sensor nodes. The sensor nodes in turn measure the ToA and DoA of the signal and propagate this information to a central information-processing center (CIP). The CIP fuses the information using Maximum Likelihood estimation to obtain the location and orientation of the sensor nodes. The authors also consider cases where partial measurements (i.e either ToA or DoA) are available.

The localization system developed as part of the GALORE project at UCLA is another example of an anchor-free ad-hoc positioning system.[7] This system is composed of standard, unmodified iPAQs and Berkeley Motes with acoustic daughter cards. The system operates in two phases: a self configuration phase in which the iPAQs collectively construct a relative 3-D coordinate system, and a service phase in which the iPAQs can localize a mote and report its location back. In this system, iPAQs are acoustic emitters and receivers, and motes are acoustic emitters only. A time synchronization service component maintains time conversions between all adjacent components of the system, and ranges from one node to another are computed by measuring the time of flight of acoustic signals. The positioning algorithm is a centralized algorithm based on relaxation of a spring model in which range measurements map to spring "lengths". A novel element of the spring algorithm is that the spring constants are non-linear: the springs are modeled as easier to compress than to stretch. This has the effect of favoring short ranges over long ranges, which is more consistent with the errors encountered with acoustics, where excess path measurement is more likely than a "short" range.

The Ad-Hoc Positioning System proposed by Nicolescu and Nath in [11] estimates the locations in an ad-hoc network by considering distances to a set of landmarks. This study explores three alternative propagation methods: *DV-hop*, *DV-distance*, and *Euclidean*. In the DV-hop method, landmarks propagate their location information inside the network. Each node forwards the landmark information to its neighbors and maintains a table with the landmark ID, location, and hop distance. When a landmark receives one of the propagated packets with the position of a different landmark, it uses that information to calculate the average hop-distance between the two landmarks. The computed average hop distance is broadcasted back into the network as a correction to previously known hop distances. The nodes that receive this message use the average hop distances to each of the landmarks to estimate their distances to the landmarks. This information is then used to triangulate the node location. The corrections are propagated in the network using controlled flooding. Each node will forward a correction from a certain landmark only once in an effort to ensure nodes will receive only one correction from the closest landmark. This policy tries to account for anisotropies in the network.

The DV-distance approach is similar to DV-hop but uses radio received signal strength measurements to measure distances. Although this approach gives finer level granularity, it is also the most sensitive to measurement error since the received signal strength is greatly influenced by the surrounding environment and therefore not always consistent.

The Euclidean propagation method uses the true distance measurement to a landmark. In this case, nodes that have at least two distance measurements to nodes that have distance estimates to a landmark can use simple trigonometric relationships to estimate their locations. The reported simulation results indicate that the DV-hop propagation method is the most accurate of the three and determines the positions of nodes within one-third of the radio range in dense networks.

Another approach described in [15] uses an algorithm similar to DV-Hop called Hop-TERRAIN in combination with a least squares refinement. The Hop-TERRAIN finds the number of hops to each anchor node and uses the anchor positions to estimate the average hop lengths. The average hop lengths are broadcasted back into the network and are used by nodes with unknown positions to compute rough estimates of their locations. Each node with unknown location that receives a message with the average hop length, estimates its distance to each anchor by multiplying the average hop distance with the number of hops to each anchor. Once a node knows the distance to each anchor, it estimates its location using triangulation. In the refinement phase, each node uses the more accurate distance measurements to its neighbors to obtain a more accurate position estimate using least squares refinement.

The collaborative multilateration approach described in [21] uses a three-phase process to estimate node locations. During the first phase, the nodes compute a set of initial estimates by forming a set of bounding boxes around the nodes. The nodes then organize themselves into over-constrained groups in which their positions are further refined using least squares. The refinement phase is presented in two computation models centralized and distributed. The centralized computation model requires global information over the entire network. The distributed computation model is an approximation of the centralized model in which each node is responsible to compute its own location by communicating with its one-hop neighbors. The key attribute that makes the distributed collaborative multilateration possible is its *in-sequence* execution within an over-constrained set of nodes. In distributed collaborative multilateration, each node executes a multilateration using the initial position estimates of its one-hop neighbors and the corresponding distance measurements. The consistent multilateration sequence helps to form a global gradient that allows each node to compute its own position estimate locally by following a gradient with respect to the global constraints.

In addition to the distance-based approaches, some work has also proposed systems using angular measurements. The Angle-of-Arrival system described in [12] is an example of a system that uses angle measurements in a multihop setup to determine node locations.

5. Network Setup Error Trends

In addition to the error incurred due to noisy measurements, the error in position estimates also depends on network setup parameters such as network size, beacon node concentration and uncertainties in the beacon locations as well error propagation when measurement information is used across multiple hops. In this section we outline the behavior of these effects using results from Cramér Rao bound (CRB) simulations. The Cramér Rao bound is a classical result from statistics that give a lower bound on the error covariance matrix of any unbiased estimator. In our discussion, CRB is used as a tool for analyzing the error behavior in multihop localization systems that use angle and distance measurements with Gaussian measurement error. The details on the actual bound derivation can be found in [6]. A close examination of this error behavior can provide valuable insight for the design and deployment of multihop localization systems.

The first notable trend relates to the behavior of localization error with respect to network density. Intuitively, localization accuracy expected to increase with increasing network connectivity. From the CRB simulations results shown in figures 1.1 and 1.2 one can observe that the rate localization accuracy improves asymptotically with network density. Initially, there is a rapid improvement at densities between 6 and 10 neighbors per node. Later on, as the number of neighbors per node increase, the improvement is much more gradual. Figures 1.1 and 1.2 show the corresponding curves for the cases when distance and angular measurements are used. The y-axis shows the RMS location error normalized by the measurement covariance σ of the measurement technology used.

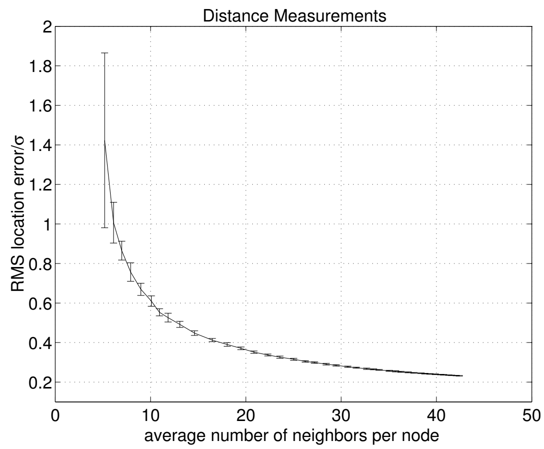

Figure 1.1. Density trend when distance measurements are used

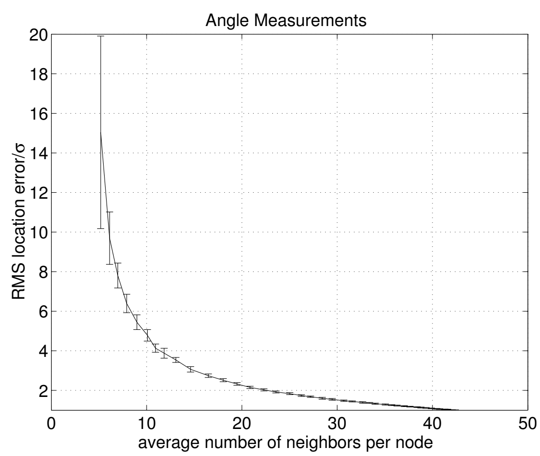

Figure 1.2. Density trend when angle measurements are used

From the figures one can observe that in the case of angle-only measurements, the location error is approximately one order of magnitude more than the error when distance measurements are used. Furthermore, it is important to note that when angle measurements are used, the error increases proportionately with range. In Figure 1.2 the error when angle-only measurements are used decreases faster than the distance measurement case. This is because to increase density, the area of the sensor field was reduced. As a side effect, the distances between nodes have also been reduced, thus reducing the tangential error in the measurement. The opposite effect would take place if the detection range of the nodes were increased, to increase density.

Another important trend to evaluate is error propagation when measurement information is used over multiple hops. Figure 1.3 shows how the error propagates in a hexagonal placement scenario, where all the nodes with unknown locations have exactly six evenly spaced neighbors. The error in both distance and angular measurements has the same trend. Error propagation is sub-linear with the number of hops. Furthermore, error propagates faster when distance measurements are used than when angle measurements are used.

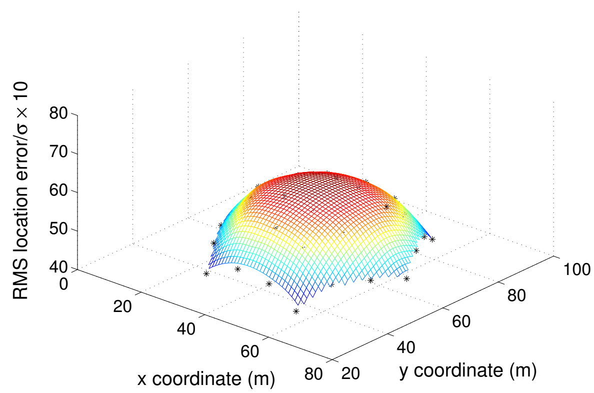

Figure 1.3. Density trend when angle measurements are used

6. Case Studies

Self Configuring Beacons for the Smart Kindergarten

The Smart Kindergarten project at UCLA [17] has developed a deeply instrumented system to study child development in early childhood education. Fine-grained localization is a key component of this system and it is realized using two different sensor node platforms (see figure 1.5. The first one is a wearable tag, called the iBadge [13]. This is attached to a vest worn by the student, or mounted on top of a special cap worn by the student. The localization of these iBadges is made possible by a second type of sensor node, the Medusa MK-2. The Medusa MK-2

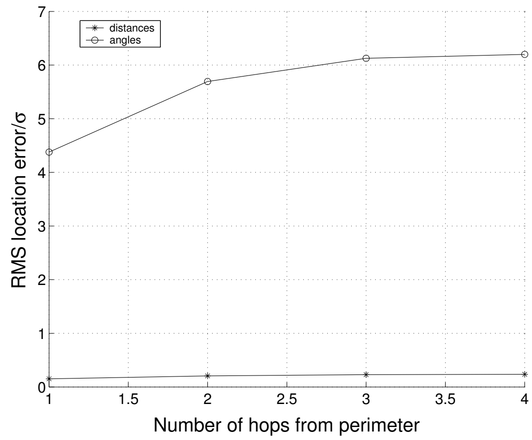

Figure 1.4. Density trend when angle measurements are used

nodes are a set of self-configuring beacon nodes deployed on the ceiling to act as anchors for localizing and tracking the iBadges. These nodes are equipped with four pairs of 40KHz ceramic ultrasonic transceivers (4 transmitters and 4 receivers) capable of omni directional transmission and reception of ultrasonic signals over a hemispherical dome. Once deployed on the classroom ceiling, the self-configuring beacons first execute a bootstrapping phase, before entering their beaconing mode. During the bootstrapping phase, each beacon measures distances to its neighboring beacons using its 40KHz ultrasonic ranging system. The distance measurements are forwarded to a central processing station that computes a relative coordinate system and notifies each beacon node of its coordinates.

Once all the beacon nodes are initialized with their locations, they enter a service mode. When in service mode, the beacons coordinate with each other to broadcast a sequence of radio and ultrasound signals. These signals are concurrently detected by multiple iBadges in the room. The iBadges use the broadcasted signals to measure distances to each other by timing the difference in the time of detection of the radio and ultrasound signals. With this information, the iBadges can compute and track their location using the on-board, DSP processor, or they alternatively they can propagate the raw data back to the central processing station that tracks the iBadges using more powerful tracking algorithms.

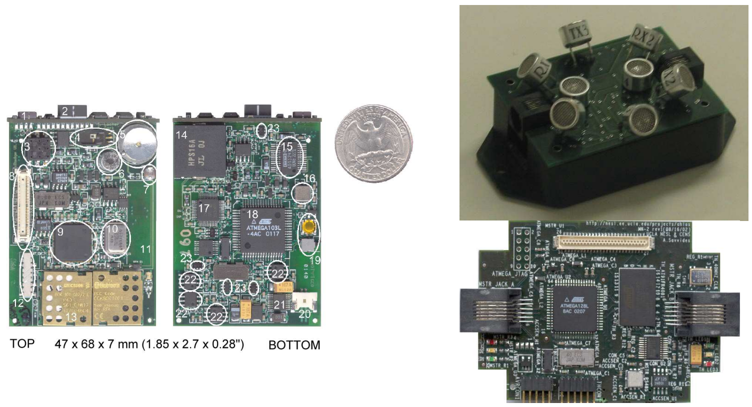

(a) **(b)**

Figure 1.5. Smart Kindergarten Sensor Nodes (a) iBadge - the wearable node, (b) Medusa MK-2 - the self-configuring ceiling beacon node

7. Conclusions

Despite the recent research efforts, and the diversity of proposed solutions, there are still several challenges to be addressed in node localization. In most proposed approaches, the lack of experimental data has so far prevented the evaluation of many localization algorithms under realistic conditions. Furthermore, the multidimensional nature of the problem suggests that schemes using of multiple measurement modalities would improve the robustness of localization systems. Although such schemes have been frequently stated in the literature, the details of fusing information at the measurement level have not been fully explored.

From a theoretical viewpoint, the fundamental characterization of error behavior is still pending. In this chapter we have described a set of initial results characterizing error behavior. Further studies are needed to understand the effects of error propagation under different measurement distributions. Finally we note that node localization is an application specific problem for which a *one size fits all* solution is unlikely to exist for all applications. Since teach application is likely to have its own requirements in terms of accuracy, latency and power consumption, node localization needs to be explored further in the context of each target application.

References

[1] P. Bahl, V. Padmanabhan, "An In-Building RF-based User Location and Tracking System", Proceedings of INFOCOM 2000 Tel Aviv, Israel, March 2000, p775-84, vol 2

[2] N. Bulusu, J. Heidemann and D. Estrin, *GPS-less low cost outdoor localization for very small devices*, IEEE Personal Communications Magazine, Special Issue on Networking the Physical World, August 2000

[3] S. Capkun, M Hamdi and J. P. Hubaux, *GPS-Free Positioning in Mobile Ad-Hoc Networks*, Hawaii International Conference on System Sciences, HICCSS-34 Jan 2001

[4] L. Doherty, L. El Ghaoui and K. S. J. Pister, *Convex Position Estimation in Wireless Sensor Networks*, Proceedings of Infocom 2001, Anchorage, AK, April 2001

[5] J. Elson and D. Estrin, *Time Synchronization for Wireless Sensor Networks*, IPDPS Workshop on Parallel and Distributed Computing Issues in Wireless Networks and Mobile Computing, San Fransisco, April 22, 2001

[6] W. Garber, A. Savvides, R. Moses and M. B. Srivastava *An Analysis of Error Inducing Parameters in Multihop Sensor Node Localization*, under submission to JSAC special Issue on the Fundamentals of Sensor Networks

[7] L. Girod, V. Bychkovskiy, J. Elson, and D. Estrin *Locating tiny sensors in time and space: A case study* In Proceedings of the International Conference on Computer Design (ICCD 2002), Freiburg, Germany. September 16-18 2002. Invited paper.

[8] K. Langendoen and N. Reijers, *Distributed Localization in Wireless Sensor Networks*, Computer Networks (Elsevier), special issue on Wireless Sensor Networks, August 2003

[9] R. L. Moses and R. M. Patterson, *Self-calibration of sensor networks*, in Unattended Ground Sensors Technologies and Applications IV (Proc. SPIE Vol. 4743) (E.M. Carapezza, ed.), pp. 108-119, April 1-4, 2002

[10] A. Nasipuri and K. Li, *A Directionality Based Location Scheme for Wireless Sensor Networks*, Proceedings of First ACM International Workshop on Wireless Sensor Networks and Applications, pp. 105-111, September 28, Atlanta, Georgia 2002

[11] D. Nicolescu and B. Nath, *Ad-Hoc Positioning System*, Proceedings of IEEE GlobeCom, November 2001

[12] D. Nicolescu and B. Nath, *Ad-hoc Positioning System using AoA*, In Proceedings of the IEEE/INFOCOM 2003, San Francisco, CA, April 2003.

[13] S. Park, I. Locher, A. Savvides, M. B. Srivastava, A. Chen, R. Muntz and S. Yuen, *Design of a Wearable Sensor Badge for Smart Kindergarten*, Proceedings of the International Symposium on Wearable Computing, October 2002

[14] N. Priyantha, A. Chakraborty and H. Balakrishnan, *The Cricket Location Support System*, Proceedings of the ACM SIGMOBILE 7th Annual International Conference on Mobile Computing and Networking, Boston, USA, August 2000

[15] C. Savarese, J. Rabay and K. Langendoen, *Robust Positioning Algorithms for Distributed Ad-Hoc Wireless Sensor Networks*, USENIX Technical Annual Conference, June 2002

[16] Y. Shang, W. Ruml, Y. Zhang and M. Fromherz, *Localization from Mere Connectivity*, In proceedings of the Fourth ACM Symposium on Mobile Ad-Hoc Networking and Computing (MobiHoc), Annapolis, MD June 2003

[17] M. Srivastava, R. Muntz and M. Potkonjak, *Samrt Kindergarten: Sensor-based Wireless Networks for Smart Developmental Problem-solving Environments*, Proceedings of the ACM SIGMOBILE 7th Annual International Conference on Mobile Computing and Networking, Rome, Italy, July 2001

[18] K. Yao and R. Hudson and C. Reed and D. Chen and F. Lorenzelli, *Blind Beamforming on a Randomly Distributed Sensor Array System*, Proceedings of IEEE JSAC, Vol 18, No. 8, October 1998

[19] H. Wang and L. Yip and D. Maniezzo and J.C. Chen and R.E. Hudson and J. Elson and and K. Yao, *A Wireless Time-Synchronized COTS Sensor Platform, Part II: Applications to Beamforming*, In Proceedings of the IEEE CAS Workshop on Wireless Communications and Networking, Pasadena, California, September 5-6 2002.

[20] N. Priyantha and A. Miu and H. Balakrishnan and S. Teller, *The Cricket Compass for Context-Aware Mobile Applications*, Proceedings of the ACM SIGMOBILE 8th Annual International Conference on Mobile Computing and Networking, pp. 32-43, August 2001

[21] A. Savvides, H. Park and M. B. Srivastava, *The n-hop Multilateration Primitive for Node Localization Problems*, Proceedings of Mobile Networks and Applications 8, 443-451, 2003, Kluwer Academic Publishers, Netherlands.