

Cooperative Localization in GPS-Limited Urban Environments

Jeffrey Hemmes^{1,*}, Douglas Thain², and Christian Poellabauer²

¹ Air Force Institute of Technology
Wright-Patterson AFB, OH 45433

² University of Notre Dame
Notre Dame, IN 46656

Abstract. Existing localization techniques such as GPS have fundamental limitations which preclude deployment in urban canyons or areas with inconsistent network availability. Augmenting GPS requires specialized infrastructure or tedious calibration tasks which limit general purpose applications. In this paper, we examine the utility of cooperatively sharing location data among connected nodes in order to correct positions with high measurement error in GPS-limited environments. Using simple data sharing and filtering techniques, collaborating users can substantially reduce overall localization error in dead reckoning systems where nodes may have a broad spectrum of location quality. We examine system parameters necessary to fully exploit cooperative localization based on empirical error models and show that mean position error can be reduced by up to 50 percent for given application scenarios. If distance measurement is available, filtering location information based on estimated error and confidence can improve accuracy of pedestrian dead reckoning techniques to approximately that of GPS using trilateration.

1 Introduction

Advances in wireless networking technology and integrated circuit design have opened up vast possibilities for mobile applications and wireless sensor networks [2,12,13,14,15], yet localization still presents significant challenges to designers of location-sensitive mobile applications. Localization of mobile nodes is frequently accomplished via GPS, but GPS positions may be inaccurate, sometimes significantly, under frequently encountered conditions such as in dense urban canyons. When operated in environments characterized by obstructed views of the sky, obtaining a position fix at all may be impossible at times. While techniques have been developed or proposed to facilitate correction of error or lack of GPS availability in specific locations using modified or additional receivers [3], none have proven suitable for all applications.

* The views expressed in this document are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

In our experiences with the TeamTrak mobile testbed [4], an approximation of wearable military computing systems for individual soldiers [11], we have observed that even with sensor hardware of the identical manufacturer and model, error frequently occurs quite independently between devices. GPS receivers placed in the same physical location at the same time and loaded with identical almanacs will not necessarily acquire the same group of satellites with the same geometry, or may acquire different numbers of satellites, resulting in varying degrees of accuracy. The practical effect of this independence of error is that multiple human operators, each carrying portable GPS receivers, may experience much different positioning even when standing beside one another. We have also observed that obtaining an initial GPS fix from a cold start with no almanac can take up to 15 minutes, posing obvious problems for many applications.

Pedestrian dead reckoning techniques, often used to augment GPS, require an accurate initial referencing point [5], which may not be available. They also experience compounding measurement error as a function of the number of steps taken due to both approximation errors and sensor limitations at each step. Although individual personal navigation systems often feature mechanisms for periodic correction [1,6,7], these techniques require preplanned and preinstalled infrastructure, maps, or even manual intervention to correct positioning error, making rapid deployment difficult or placing undue burden on human operators.

To remedy the lack of consistent localization in MANETs deployed in GPS-limited environments, we propose that cooperatively sharing position information among connected users can provide localization of acceptable quality for many applications. Specifically, the contributions of this paper are the following:

- An evaluation of the potential benefit of sharing localization data among mobile nodes in low-cost systems lacking distance measurement tools and in the presence of occasionally significant GPS and dead reckoning error.
- A filtered trilateration algorithm to improve localization in cases where distance measurement is possible but sensor error exists.

In this paper, we examine the utility of sharing location information among nodes capable of localizing to varying degrees of accuracy using either GPS or pedestrian dead reckoning. Through the use of simulation with parameters determined by empirical measurement of sensors, we separately study cases in which spatial separation between nodes can and cannot be measured.

2 Related Work

To date there has been extensive research in the area of localizing mobile wireless nodes. Traditional approaches to distributed localization typically suffer from one of two fundamental limitations. First, techniques such as trilateration or triangulation generally require expensive, specialized hardware for ranging or tedious calibration tasks. Distance measurement using metrics such as received signal strength, frequently described in the literature, can be very unreliable in outdoor environments. Second, many range-free localization techniques rely on

specific topologies or anchor distributions to localize nodes. APIT [8] and DV-HOP [16] are two examples. These techniques, generally best suited for static networks, often suffer in the presence of mobility.

One popular approach to overcoming these limitations is the sequential Monte Carlo method, originally developed for mobile robot localization. A number of different works have employed variants of MCL [17,18,19,20] for robots. Our work differs from such protocols in that for many scenarios, node distribution may be sparse and connectivity sporadic, suggesting that the minimum number of samples might not be consistently obtained. In addition, MCL methods assume only a subset of all nodes are capable of self-localization. This work addresses the question of whether sharing information is worthwhile when all nodes can self-localize to varying extents.

There has been far less work done in the area of pedestrian dead reckoning. The goal of such systems is to reduce overall localization error by providing an alternative to GPS. Like robot localization, pedestrian dead reckoning generally produces positions with decreasing confidence and accuracy as a function of the number of steps taken since correction. Dead reckoning in TeamTrak has many similarities to that of NavMote [9], with one notable difference in that position correction in NavMote occurs largely through manual intervention by the operator as opposed to cooperatively sharing location information.

Finally, this work differs from others by recognizing not only that frequently, significant error can result from dead reckoning due to uncertainties in measuring the noisy motions of a human operator, but that a node designated as an anchor or seed node could have a fairly substantial GPS error as well. The effect of cooperative localization using location information from multiple sources with error that manifests differently has not been widely explored.

3 Cooperative Localization

Throughout this paper, we use a scenario-driven approach to evaluation, with each scenario derived from a single, overarching problem domain. The problem statement can be specified as follows:

Situation: Assume N mobile nodes in the field. A subset of nodes have a functioning GPS receiver. Those that do not rely on dead reckoning for localization. We assume the initial reference point for dead reckoning is determined via GPS, but after establishing the initial point, GPS is no longer available to those nodes. Each delta in position computed by the dead reckoning apparatus is affected by both compass error and stride length estimation error at each step. Additionally, all GPS positions have measurement error, which may be substantial.

Objective: Each node must compute an estimate of its current physical location to display to its human operator. Nodes may combine information collected from multiple remote sources to determine a final result.

Hypothesis: Quality of position estimates from dead reckoning can be improved by sharing information once measurement error grows sufficient

The remainder of this section describes two methods of combining shared location samples based on whether distance measurement is available.

First, to determine whether any position correction should occur, a mobile node \mathcal{N} must determine the quality of its own position. If multiple localization techniques are available, that with the smallest estimated error value e is used. If GPS is available, no further action is taken. Early attempts at correcting erroneous GPS positions by simply computing the average other GPS positions proved counterproductive; doing so eliminated the availability of accurate anchor positions entirely. This quickly propagated large errors among all nodes, especially at longer communication ranges. For cases in which distance measurements are not available, positions of connected peers among n nodes are combined using Algorithm 1, a modified Centroid method we call Simple Averaging.

Algorithm 1. Simple Averaging

```

1:  $(x, y, e) \leftarrow \text{Location-Use-Best}(\mathcal{N})$ 
2: if  $\text{Location-Source}(\mathcal{N}) = \text{GPS}$  then
3:   return  $(x, y, e)$ 
4: end if
5: for  $i = 1$  to  $n$  do
6:   if  $\text{Connected}(\mathcal{N}, i)$  then
7:      $\bar{x} \leftarrow \bar{x} + x_i$ 
8:      $\bar{y} \leftarrow \bar{y} + y_i$ 
9:      $e_{total} \leftarrow e_{total} + error_i$ 
10:     $count \leftarrow count + 1$ 
11:   end if
12: end for
13: if  $count \neq 0$  then
14:    $\bar{x} \leftarrow \bar{x} / count$ 
15:    $\bar{y} \leftarrow \bar{y} / count$ 
16: end if
17: return  $(\bar{x}, \bar{y}, e_{total})$ 

```

Note that Algorithm 1 as stated does not account for error in remote positions. We examine this simple case as well as a modification in which only positions with estimated error less than or equal to that of the node to be localized are considered. We call this modification Selective Averaging.

In cases in which some method for measuring spatial separation between nodes is available, a simple trilateration scheme can be used to correct positions known to have significant estimated error. Unlike static sensor networks, random localization error for a single point cannot be averaged away in mobile networks since the number of available position samples from an anchor that represent a single point may be small. This limitation requires evaluation of the estimated measurement error from each node's location and acceptance of each position before inclusion for trilateration. To accomplish this, we include a filter to a trilateration method that accepts or rejects reported positions from connected peers

based upon the estimated measurement error e from each. Because of the uncertainty associated with reported error, it is helpful to also determine a confidence level for each reported error value for the particular sensor device. In the case of GPS receivers, metrics provided by the device and accompanying each location sample are recorded and correlated a priori with the actual measurement error. Here, confidence is defined as the probability that measurement error is less than or equal to the estimated error metric from the device. We assume that as long as measurement error is likely to be no worse than the metric suggests, that metric can be useful in a filtering scheme to avoid accepting positions with a large measurement error.

To filter positions for trilateration from peers in the presence of both localization error and mobility, let $P_{\mathcal{N}}$ be the (possibly empty) set of all location samples (represented as a 4-tuple) of preferred quality and confidence level at Node \mathcal{N} , the node to be localized. Here, τ is the error tolerance and c_{min} is the minimum acceptable confidence level for each error value as determined for the application. Positions of lower error and higher confidence are accepted:

$$P_{\mathcal{N}} = \{(x, y, e, c) \mid (x \in \mathbb{R}) \wedge (y \in \mathbb{R}) \wedge (e \leq \tau) \wedge (c \geq c_{min})\} \tag{1}$$

and $P'_{\mathcal{N}}$ be the (possibly empty) set of all location samples of arbitrary quality and confidence levels generated by \mathcal{N} as it self-localizes through multiple means, which includes positions of significant error magnitude or low confidence:

$$P'_{\mathcal{N}} = \{(x, y, e, c) \mid (x \in \mathbb{R}) \wedge (y \in \mathbb{R}) \wedge ((e > \tau) \vee (c < c_{min}))\} \tag{2}$$

We take $\ell_{\mathcal{N}}$ to be the locally determined position of highest quality, where $\forall i \forall j (p_i, p_j \in P_{\mathcal{N}}) \wedge (p'_i, p'_j \in P'_{\mathcal{N}})$:

$$\ell_{\mathcal{N}} = \begin{cases} (x_0, y_0, \infty, 0.0) & \text{if } P_{\mathcal{N}} \cup P'_{\mathcal{N}} = \emptyset \\ (x_i, y_i, e_i, c_i) \mid p_i.e = \min\{p_j.e\} & \text{if } P_{\mathcal{N}} \neq \emptyset \\ (x_i, y_i, e_i, c_i) \mid p'_i.e = \min\{p'_j.e\} & \text{if } P_{\mathcal{N}} = \emptyset \end{cases} \tag{3}$$

Let L be the set of location samples from n nodes directly connected to \mathcal{N} :

$$L = \{(x, y, e, c) \mid (x \in \mathbb{R}) \wedge (y \in \mathbb{R}) \wedge (e \in \mathbb{R}) \wedge (c \in [0, 1])\} \tag{4}$$

There are five cases which must be considered depending on the estimated error value e and the associated confidence level c_i of each remote peer node i , where $i = 1..n$.

- Case 1: $e_i \leq \tau$ and $c_i \geq c_{min}$
- Case 2: $e_i \leq \tau$ and $c_i < c_{min}$
- Case 3: $e_i > \tau$ and $c_i < c_{min}$
- Case 4: $e_i > \tau$ and $c_i \geq c_{min}$
- Case 5: $e_{\mathcal{N}} \leq e_i$

Which case represents the minimum threshold for acceptance depends on the application, and different cases can produce much different levels of accuracy. For

each of the above cases, we find the set of remote positions R and corresponding distances D among nodes in R which are directly connected to node \mathcal{N} , whose best position is represented by $\ell_{\mathcal{N}}$.

Therefore, for all positions $p \in L$:

$$R = \begin{cases} p \notin (P_{\mathcal{N}} \cup P'_{\mathcal{N}}) \wedge (p.e \leq \tau) \wedge (p.c \geq c_{min}) & \text{if case 1} \\ p \notin (P_{\mathcal{N}} \cup P'_{\mathcal{N}}) \wedge (p.e \leq \tau) \wedge (p.c < c_{min}) & \text{if case 2} \\ p \notin (P_{\mathcal{N}} \cup P'_{\mathcal{N}}) \wedge (p.e > \tau) \wedge (p.c < c_{min}) & \text{if case 3} \\ p \notin (P_{\mathcal{N}} \cup P'_{\mathcal{N}}) \wedge (p.e > \tau) \wedge (p.c \geq c_{min}) & \text{if case 4} \\ p \notin (P_{\mathcal{N}} \cup P'_{\mathcal{N}}) \wedge q \in (P_{\mathcal{N}} \cup P'_{\mathcal{N}}) \wedge (p.e \leq q.e) & \text{if case 5} \end{cases} \quad (5)$$

Then for all remote positions $r \in R$, the set of distances from each point is:

$$D = \{d \mid d = \sqrt{(r.x - \ell_{\mathcal{N}}.x)^2 + (r.y - \ell_{\mathcal{N}}.y)^2}\} \quad (6)$$

Next, use the points contained in the set R and the pairwise distances to create a set of points $\hat{P}_{\mathcal{N}}$, with each member representing a possible position of node \mathcal{N} . Here, θ_i represents the bearing from \mathcal{N} to node i , $p_i \in R$, and $d_i \in D$.

$$\hat{P}_{\mathcal{N}} = \{((p_i.x - d_i \cos \theta_i), (p_i.y - d_i \sin \theta_i), e_i, c_i)\} \quad (7)$$

Then compute \bar{p} , the mean of all positions $p \in \hat{P}_{\mathcal{N}}$:

$$\bar{p} = \begin{cases} \left(\frac{1}{n} \sum_{p \in \hat{P}_{\mathcal{N}}} p.x, \frac{1}{n} \sum_{p \in \hat{P}_{\mathcal{N}}} p.y \right) & \text{if } \ell_{\mathcal{N}} = (x_0, y_0, \infty, 0.0), \\ \left(\frac{1}{n+1} \left(\ell_{\mathcal{N}}.x + \sum_{p \in \hat{P}_{\mathcal{N}}} p.x \right), \right. & \\ \left. \frac{1}{n+1} \left(\ell_{\mathcal{N}}.y + \sum_{p \in \hat{P}_{\mathcal{N}}} p.y \right) \right) & \text{otherwise.} \end{cases} \quad (8)$$

Thus \bar{p} , the estimated position of \mathcal{N} , is determined with a filtered trilateration procedure shown in Algorithm 2.

4 Sensor Evaluation and Error Models

The TeamTrak mobile ad-hoc network testbed, which is the basis for much of this work, consists of an integrated array of commercially available, low-cost sensor hardware. Components include the Garmin GPS-18 USB and HP iPAQ BT-308 Bluetooth GPS receivers, the OceanServer OS3500 tilt-compensated 3-axis digital compass, and the SparkFun SerAccel 3-axis digital accelerometer, which is employed as a pedometer in tandem with the compass for pedestrian dead reckoning. The remainder of this section discusses measurement of sensor error for these components and development of error models.

In our simulations, both magnitude of GPS error and expected length and direction of drift for anchor nodes are explicitly modeled. Distribution of error in recorded GPS positions, obtained from both the Garmin GPS-18 and HP iPAQ

Algorithm 2. Filtered Trilateration

```

1: if Correction-Required( $\mathcal{N}$ ) then
2:    $(x, y, e, c) \leftarrow$  Location-Use-Best( $\mathcal{N}$ )
3:   for  $i = 1$  to  $n$  do
4:     if  $i \neq \mathcal{N}$  and Connected( $\mathcal{N}, i$ ) then
5:        $L[i] \leftarrow$  Location-Use-Best( $i$ )
6:       if  $L[i].e \geq \tau \wedge L[i].c \geq c_{min}$  then
7:          $R[i] \leftarrow L[i]$ 
8:          $\theta[i] \leftarrow \arctan((L[i].y - y) / (L[i].x - x))$ 
9:          $d[i] \leftarrow \sqrt{(x - L[i].x)^2 + (y - L[i].y)^2} + \epsilon$ 
10:         $R[i].e \leftarrow |R[i].e| + |\epsilon|$ 
11:       end if
12:     end if
13:   end for
14: end if
15: if  $R = \emptyset$  then
16:   return  $(x, y, e, c)$ 
17: else
18:   for  $i = 1$  to  $n$  do
19:     if  $R[i] \neq 0$  then
20:        $R[i].x \leftarrow R[i].x - d[i] \cos \theta[i]$ 
21:        $R[i].y \leftarrow R[i].y - d[i] \sin \theta[i]$ 
22:     end if
23:   end for
24: end if
25:  $\bar{x} \leftarrow \frac{1}{\#R} \sum_{r \in R} r.x$ 
26:  $\bar{y} \leftarrow \frac{1}{\#R} \sum_{r \in R} r.y$ 
27:  $e \leftarrow \frac{1}{\#R} \sqrt{\sum_{r \in R} (r.x - \bar{x})^2}$ 
28:  $c \leftarrow \min \forall r \in R \mid r.c$ 
29: return  $(\bar{x}, \bar{y}, e, c)$ 

```

GPS receivers and shown in Figure 1, was incorporated into a single composite histogram. Each error sample was measured by determining the distance from each position reported by a stationary GPS receiver with a partially obstructed view of the sky to the overall mean of 334,435 position samples collected over a period of four days. Both receivers evaluated have a Rayleigh distribution, but the HP has a mean error of 15.69 meters and $\sigma = 79.73$, while the Garmin has a mean error of 14.27 meters with $\sigma = 23.22$. The model of GPS error takes the actual position of a node at a given time and perturbs it with an offset based upon the distribution. A pseudo-random number is generated and input to an inverse CDF to obtain the magnitude of the error. Similarly, using the GPS data, models of drift length and direction are a function of a specific error value.

Figure 2 shows the confidence in the horizontal error metric, represented as the probability the actual error is no greater than the estimated error, for the Garmin GPS-18 receiver. Such a model is used to determine how reliable a particular error metric might be, which is used in Algorithm 2 for filtering decisions. Note that the confidence level of some values are clearly outliers; for instance, a

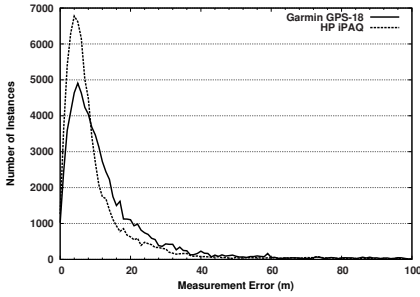


Fig. 1. Distribution of GPS Error

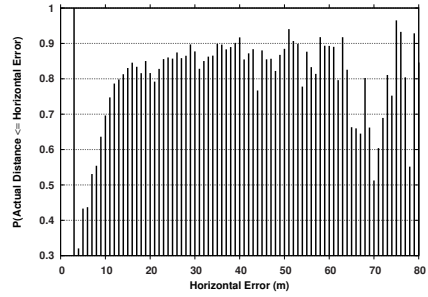


Fig. 2. Error Confidence (Garmin)

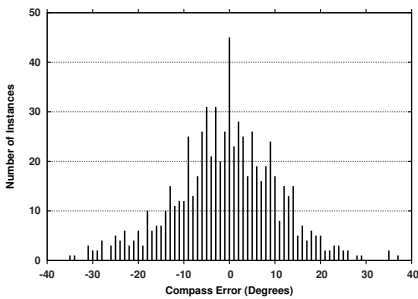


Fig. 3. Distribution of Compass Error




Fig. 4. Dead Reckoning Error in Practice

horizontal error value of 3 uniformly corresponded with a measured error of less than 3 meters. However, the sample size was very small.

A similar approach was used to model dead reckoning error. A series of trials were conducted to measure the distribution of error observed in the OceanServer OS3500 tilt-compensated 3-axis digital compass. Distribution is Gaussian with mean -0.794° and $\sigma = 11.33$, as illustrated in Figure 3. Similarly, stride length error is determined using a model based on that implemented in NavMote [9] and combined with heading to compute the measurement error in the next position. An illustration of localization error as observed in practice using TeamTrak's dead reckoning apparatus is shown in Figure 4. In the figure, navigation begins at the lower right corner of the map. The operator walked along the diagonal path, stopping beneath the trees shown in the upper left corner.

Finally, connectivity is simulated with the anisotropic wireless communication model introduced by He et al. [8]. Using the TeamTrak testbed, we empirically measured the range capability of the Intel PRO/Wireless 2915ABG wireless card, and used the estimated maximum range, illustrated in Figure 5, as an initial simulation parameter. The graph shows the distance in each direction at least one-way connectivity was established between TeamTrak nodes, and is intended to provide an estimate of the maximum communication range in the outdoor environment in which the testbed was deployed.

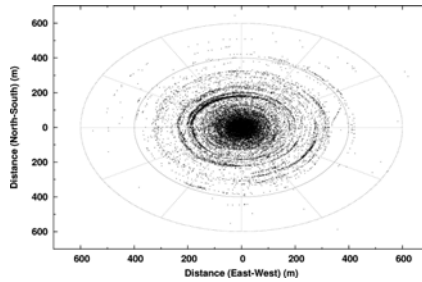


Fig. 5. Wireless Range Measurement

5 Simulation Results

The simulation is intended to model a localized mobile ad-hoc network at steady state. Initial placement of nodes is random with uniform distribution across the simulation field. Mobility is modeled using Random Waypoints [10], a method selected due to its generality, i.e., most real-world scenarios would not involve random patterns of motion and the use of a random model is unlikely to produce overly optimistic results. The speed of each agent is random with a mean of 1.56 m/s, the average walking speed of an adult male, and normal distribution.

Because dead reckoning error itself is generally a function of the number of steps taken, over time this error can significantly eclipse the measurement error of the initial GPS position. Therefore, running simulations for longer periods of time would result in higher average error, but such error changes proportionately with the length of the simulation. In this paper, experiments simulate motion over a period of 20 minutes, over which our experiences with such systems show is sufficient time for uncorrected dead reckoning error to grow quite large.

5.1 Simple Averaging: 300-Meter Maximum Range

In this scenario, any node localized via dead reckoning attempts to correct its position using Algorithm 1, i.e., averaging its own with those of all connected nodes regardless of the source of each neighbor's position, i.e., nodes localized via GPS, dead reckoning, or cooperative averaging may be considered anchors. Each position included in the average is weighted equally. Communication range varies between 200 to 300 meters depending on direction, and is intended to model the anisotropic RF propagation commonly found in wireless radios. With 100 nodes in the system in total, the number of anchors localized with GPS was increased from 1 to 99, with intermediate values in increments of 10.

Figure 6 shows the effect of a simple averaging scheme for dead reckoning nodes with at least one GPS anchor among them. The figure shows the mean position error of all nodes localized via dead reckoning and adjusted using Algorithm 1 over all time steps over 100 independent simulation runs. In this case, increase in the number of anchors in the system effects a reduction in the overall mean position error. However, realizing improvement in mean error suggests the

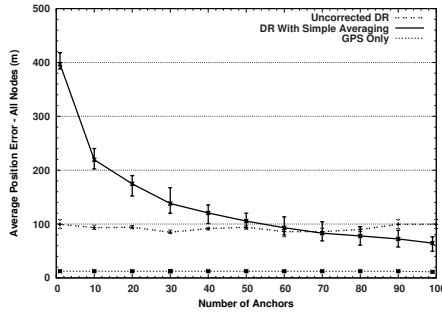


Fig. 6. Simple Averaging: 300m Range

ratio of anchors to other nodes should be greater than 0.80, but realistically that ratio would likely need to be even higher. Above a 0.90 ratio of anchors, mean error can be reduced by as much as approximately 23 percent. However, overall the number of cases in which average error increases is much larger than the number of cases in which it is reduced.

It is important to note that when averaging, the error in the adjusted positions is a function of both the separation between nodes and sensor measurement error. Even when dead reckoning positions are freshly reset, such as at the beginning of each simulation trial, the localization error after adjustment is highly dependent on the arrangement of connected anchors and the spatial separation between nodes. Existing localization techniques such as APIT [8] assume an unlocalized node is surrounded by anchors relatively symmetrically. In cases in which mobility invalidates such an assumption, having very few anchors, poor geometry, or large separation results in an average position whose error can be up to several times worse than in the case of doing nothing.

5.2 Simple Averaging: 100-Meter Range

In this scenario, the communication range is reduced to no more than 100 meters, and again nodes without local GPS information average their positions with all of their connected neighbors using Algorithm 1. The effect of simple averaging at a maximum range of 100 meters is shown in Figure 7. In this configuration, the mean error in the worst case is less than twice that of the uncorrected case, and the variance results in a worst case of no more than three times the mean of the uncorrected case. Increasing the number of anchors in the system results in an improvement in the dead reckoning positions by nearly 50 percent in the cases of networks which are relatively saturated with anchors. The majority of the improvement resulting from averaging is due to the prevention of growth in the dead reckoning error over a large number of steps. With uncorrected dead reckoning, mean error is dominated by the error toward the end of each simulation, i.e., in general, the more steps taken since the last reset, the larger the mean error. It is possible, therefore, for a simple averaging scheme using shared location data to serve as a reset for dead reckoning, even without the availability of reliable ranging. However, accuracy is sensitive to communication range.

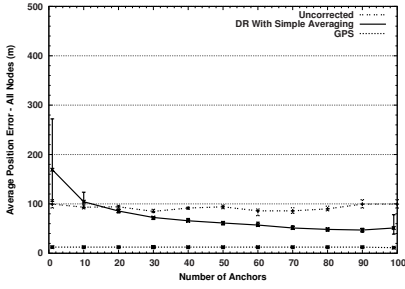


Fig. 7. Simple Averaging: 100m Range

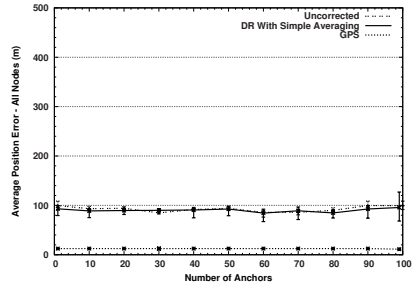


Fig. 8. Simple Averaging: 15m Range

5.3 Simple Averaging: 15-Meter Range

The idea behind averaging positions over lower communication range is to reduce error introduced as a function of the spatial separation between nodes when the geometry of connected anchors is unbalanced. Intuitively, an even shorter range than used in Figure 7 should reduce mean position error further, as each position sample represents locations much closer to the same physical location. However, a further reduction in the communication range to approximately that of Bluetooth radios, i.e., within a minimum of 9 and a maximum of 15 meters, produces no benefit, as shown in Figure 8. In this case, because the communication range is very short, connectivity occurs far less frequently, and therefore fewer opportunities to correct bad positions exist. The mean error is approximately the same as that of the uncorrected case, suggesting that using a very short communication range does not work in the case of general mobility.

It should be noted that such short range could still be used to adjust positions in this manner, since establishing connectivity in such a case implies nodes are collocated. In implementing such a scheme, location information can be shared effectively, but a short-range cooperative sharing scheme requires operators to move within very close proximity for any adjustment to occur. Depending on the application scenario, this may not be suitable.

5.4 Distribution of Error

To further evaluate the effectiveness of using shared location information to reduce position error, the distribution of measurement error, as determined by the Euclidean distance from each position estimated via dead reckoning apparatus or through cooperative averaging, to the actual position for each node at each time step, is shown in Figure 9. The figure shows the effect of two different maximum ranges on the overall distribution of position error: Figure 9(a) shows the distribution of error when nodes are capable of a maximum 300-meter wireless range, while 9(b) shows the result with a 100-meter range. In each case, the distribution of position error with varying numbers of anchors in the system among

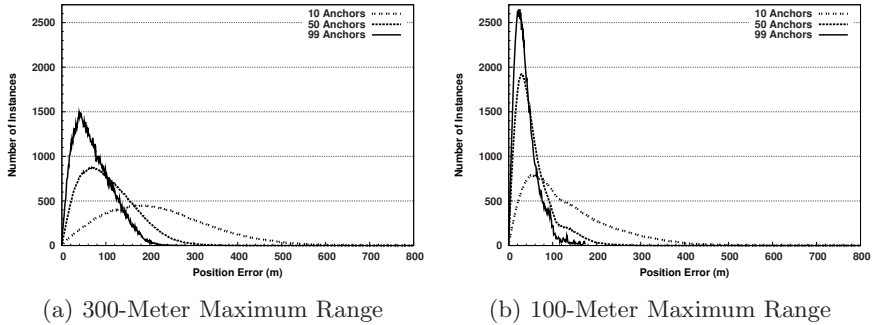


Fig. 9. Distribution of Position Error (Simple Averaging)

100 nodes in total is shown, and each histogram is normalized to account for the differing number of nodes whose positions are adjusted.

At longer range, a significant number of samples exhibit a very large magnitude of error with a low density of anchors. In the case of a 10 anchors among 100 nodes, the error has a Rayleigh distribution, but is very flat, suggesting a high variance of error with fewer occurrences of each. Increasing the number of GPS-enabled anchors among in the network, reduces both variance and mean error. However, as shown previously in Figure 6, the resultant mean error in this case is still unacceptably large, particularly with low anchor density, due primarily to the wide spatial separation between nodes.

Conversely, at a shorter communication range the histograms for different anchor ratios are grouped much more closely. This suggests that when using a simple range-free combination technique, localization error is less sensitive to changes in the number of anchors than it is to changes in the maximum communication range, provided a sufficient deployment scale.

5.5 Selective Averaging

While evaluating whether distributing location information is beneficial for systems encompassing a range of complementary localization techniques, we have varied the maximum communication range, but have not considered measurement error or estimated error of the shared location information. Algorithm 1 will dutifully include positions known to be bad. To address this limitation, we introduce a modification, dubbed Selective Averaging, in which a remote position determined to have a higher measurement error than the position to be adjusted is ignored, even if it is from the only available anchor. Such a method requires that a node to be localized consider the estimated error of all positions, including its own, which could easily be reported along with the location itself. If a dead reckoning node has a more accurate position than its connected neighbors based on its own estimate of its localization error, no adjustment occurs.

Figure 10 shows the effect of such selective averaging in a scenario with 100-meter maximum communication range. Note that while the mean error in the

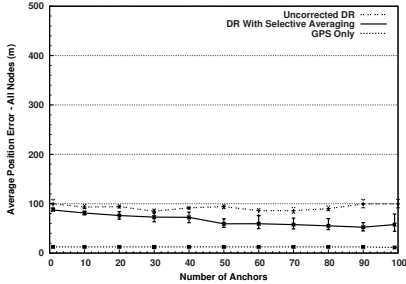


Fig. 10. Selective Averaging

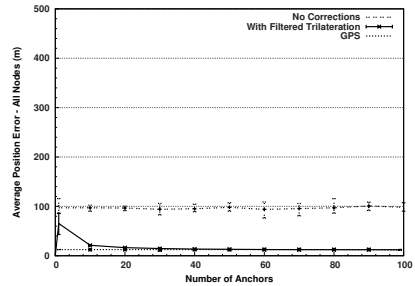


Fig. 11. Filtered Trilateration

cases of relatively high anchor density is similar to that of the case of simple averaging, i.e., a reduction in mean position error of approximately 50 percent, the worst cases are roughly the same as for uncorrected dead reckoning, so this method has the advantage of not making positions worse on the average. By accounting for estimated error, substantial increases in mean error which occur whenever a stale dead reckoning position is combined with either another stale position or an anchor whose GPS position is drifting substantially can be reduced. In addition, because examining error suggests that corrections occur less frequently than in the naive approach, error attributable to significant spatial separation is less of a factor. This is not to say that worsening of localization does not occur, but the overall net effect is a reduction in dead reckoning error. A selective averaging technique such as this could most likely be used to correct positions in urban environments in which systems are deployed with plentiful GPS-enabled anchors, but in which there exist pockets lacking GPS availability.

5.6 Filtered Trilateration

Having examined the effects of simple variations of range-free combination techniques for shared location information, focus now shifts to the use of filtered trilateration, specified in Algorithm 2. Ordinarily, trilateration requires at least three anchors, but with a very stale dead reckoning-measured position, even a single anchor can provide improvement regardless of localization source. Because ranging for outdoor mobile ad-hoc networks currently is not sufficiently mature and reliable, for this paper we assume that such a method exists.

Simulation parameters are the same as in previous cases. In all simulation trials discussed in this paper, GPS localization has a mean measurement error of approximately 12 meters based on empirical measurement, and dead reckoning error is effectively unbounded. Error in distance measurement follows a Gaussian distribution with σ equal to half of the maximum communication range. As shown in Figure 11, even with a relatively low density of anchors, the mean error of pedestrian dead reckoning that begins with a GPS-measured initial reference point suggests that filtered trilateration can achieve overall localization

Table 1. Location Adjustments (Filtered Trilateration)

Anchor Density (%)	Positions Improved (%)	Mean Error Reduction (m)	Positions Worsened (%)	Mean Error Increase (m)
WiFi Range				
100	5.10	6.37	6.60	7.94
67	6.75	6.28	3.68	3.36
33	4.66	6.42	6.58	1.40
0	0.16	1.04	14.3	0.95
Bluetooth Range				
100	0.33	4.34	0.27	4.49
67	0.19	6.13	0.08	1.93
33	0.16	6.26	0.66	2.26
0	0.00	1.50	0.05	0.35

accuracy for dead reckoning comparable to that of GPS. As previously mentioned, this error is only the mean; it is certainly possible for localization accuracy for individual positions to be reduced due to the uncertainty of available metrics.

To compare net reduction versus net increase in error, Table 1 shows the percentages of positions adjusted, with the effects either positive or negative, for various densities of GPS anchors. In this experiment, the large percentage of positions left unadjusted is due to the maximum error tolerance of 15 meters, the minimum confidence level of 70 percent, and the small incremental error in dead reckoning at each step illustrated in Figure 4. Note that while for cases in which the percentage of positions improved or worsened is approximately the same, particularly for WiFi range, the mean of the improvement is roughly twice that of the increased error introduced. The net change overall results in improvement.

Similarly, in the case of 100% anchor density, filtered trilateration does result in a net worsening of positions. However, if all nodes in the system have GPS available, there is no need to use cooperative localization, so the case is irrelevant. Any worsening is particularly apparent in cases with very low anchor density, due to the inherent challenges in estimating error in pedestrian dead reckoning. The modest reduction in mean localization error achieved for networks with relatively few anchors, coupled with both the lack of robust distance measurement and the uncertainty in dead reckoning error, also suggests that any cooperative localization scheme would be employed best as a corrective measure for a relatively small number of nodes in a larger overall deployment.

Regardless of which algorithm is used for combining shared location information, filtering based on estimated error has real benefits in terms of reduction of mean error. Those cases in which location error is reduced or maintained dominate those in which positions are made worse either in terms of either frequency of occurrence or magnitude of the change.

6 Future Work

The work in this paper represents the first steps in using cooperative computing in connected navigation systems to reduce localization error. Additional work must be done to further refine the error model for pedestrian dead reckoning. A more robust error model can provide superior metrics in terms of both error magnitude and confidence level, which can be used more effectively for filtering. Secondly, a more robust model for estimating localization error with combination techniques when larger distances are involved should be developed and integrated in this system, particularly for range-free approaches. Finally, efforts to integrate devices such as laser range finders into mobile testbeds, which would permit correction of drifting GPS positions, are currently being explored.

7 Conclusion

This paper presents an evaluation of the potential benefit arising from cooperatively sharing location information among users of MANETs deployed to urban settings or other environments in which GPS capability is limited. When GPS is augmented with other localization techniques such as dead reckoning, error manifests much differently and is affected by entirely different factors. Cooperatively sharing location information can reduce overall localization error even without the ability to accurately measure spatial separation between nodes.

Because using a simple combination method to average position introduces error as a function of such separation, cooperative localization is most effective if the communication range is relatively short, but not excessively so. An overly long range introduces error from wide separation, while an excessively short range does not facilitate connectivity without requiring human intervention. The results show that with the error experienced in the commodity off-the-shelf sensor hardware such as that found in TeamTrak, a range of approximately 100 meters can reduce mean position error by roughly half in anchor-rich environments, while at the same time minimizing the compounding error resulting from a network with a low anchor density. Furthermore, filtering based on error metrics can reduce the rapid increase in error in low density deployments.

Traditional trilateration can be improved to account for measurement error in localization by simply filtering positions using available metrics with each position. Excluding positions with large measurement error or low confidence levels reduces the likelihood of worsening positions, while simultaneously leveraging anchors with accurate localization to reset dead reckoning. With filtering, trilateration could improve dead reckoning localization to approximately that of GPS. While adjusting GPS positions using other remote GPS positions with a simple averaging technique may increase overall localization error, the growth in error observed in dead reckoning can be large and varied across nodes. This suggests that some dead reckoning nodes could serve as anchors for others, and in cases where ranging is not available, even a simple averaging technique such as discussed in this paper may be beneficial in a significant number of instances.

References

1. Leonard, J., Durrant-Whyte, H.: Mobile Robot Localization by Tracking Geometric Beacons. *IEEE Transactions on Robotics and Automation* 7(3) (1991)
2. Committee on Networked Systems of Embedded Computers: Embedded Everywhere: A Research Agenda for Networked Systems of Embedded Computers (2001), http://books.nap.edu/html/embedded_everywhere
3. Corvallis Technology, Inc.: Introduction to the Global Positioning System for GIS and TRAVERSE. CMT Inc. (1996)
4. Hemmes, J., Thain, D., Poellabauer, C., Moretti, C., Snowberger, P., McNutt, B.: Lessons Learned Building TeamTrak: An Urban/Outdoor Mobile Ad Hoc Network Testbed. In: *Proc. of WASA 2007* (2007)
5. Mezentsev, O., Lachapelle, G., Collin, J.: Pedestrian Dead Reckoning - A Solution to Navigation in GPS Signal Degraded Areas. *Geomatica* 59(2) (2005)
6. Liu, H., Pang, G.: Accelerometer for Mobile Robot Positioning. In: *Proc. of IEEE Industry Applications Conf.* (1999)
7. Pang, G., Liu, H.: Evaluation of a Low-cost MEMS Accelerometer for Distance Measurement. *Journal of Intelligent and Robotic Systems* 30(3) (2001)
8. He, T., Huang, C., Blum, B., Stankovic, J.: Range-Free Localization Schemes for Large Scale Sensor Networks. In: *Proc of Mobicom 2003* (2003)
9. Fang, L., Antsaklis, P., Montestruque, L., McMickell, M., Lemmon, M., Sun, Y., Fang, Y., Koutroulis, I., Haenggi, M., Xie, M., Xie, X.: Design of a Wireless Assisted Pedestrian Dead Reckoning System - The NavMote Experience. *IEEE Transactions on Instrumentation and Measurement* 54(6) (2005)
10. Johnson, D., Maltz, D.: Dynamic Source Routing in Ad Hoc Wireless Networks. *Mobile Computing* 53, 153–181 (1996)
11. Zieniewicz, M., Johnson, D., Wong, D., Flatt, J.: The Evolution of Army Wearable Computers. *IEEE Pervasive Computing* 1(4), 30–40 (2002)
12. Akyildiz, I., Su, W., Sankarasubramaniam, Y., Cayirici, E.: Wireless Sensor Networks: A Survey. *Computer Networks (Elsevier) Journal* 38(4), 393–422 (2002)
13. Correal, N., Patwari, N.: Wireless Sensor Networks: Challenges and Opportunities. In: *Proc. of the 2001 Virginia Tech. Symposium on Wireless Personal Communications* (2001)
14. Feder, B.: Wireless Sensor Networks Spread to New Territory. *New York Times*, July 26 (2004), <http://www.nytimes.com/2004/07/26/business/26sensor.html>
15. Perkins, M., Correal, N., O’Dea, R.: Emergent Wireless Sensor Networks Limitations: A Plea for Advancement in Core Technologies. In: *Proc. of IEEE Sensors* (2002)
16. Niculescu, D.: Ad-hoc Positioning System. In: *Proc of IEEE Global Telecommunications Conference, GLOBECOM* (2001)
17. Dellaert, F., Fox, D., Burgard, W., Thrun, S.: Monte Carlo Localization for Mobile Robots. In: *IEEE Intl. Conference on Robotics and Automation, ICRA* (1999)
18. Thrun, S., Fox, D., Burgard, W., Dellaert, F.: Robust Monte Carlo Localization for Mobile Robots. *Artificial Intelligence* 128(2), 99–141 (2000)
19. Thrun, S., Fox, D., Burgard, W.: Monte Carlo Localization With Mixture Proposal Distribution. In: *AAAI Natl. Conference on Artificial Intelligence* (2000)
20. Köse, H., Akin, H.: The Reverse Monte Carlo Localization Algorithm. *Robotics and Autonomous Systems* 55(6), 480–489 (2001)