A Practical Approach to Cooperative Localization in GPS-Limited Urban Environments

Technical Report 2008-14

Jeffrey Hemmes[1]
Department of Electrical and Computer Engineering
Air Force Institute of Technology
Wright–Patterson AFB, OH 45433
Email: jhemmes@afit.edu

Douglas Thain and Christian Poellabauer
Department of Computer Science and Engineering
University of Notre Dame
Notre Dame, IN 46656
Email: dthain, cpoellab@cse.nd.edu

***Abstract*—Existing localization techniques have fundamental limitations which make deployment under commonly encountered conditions difficult. Techniques used to augment GPS generally require specialized infrastructure, tedious calibration tasks, or other requirements that make general purpose application difficult, yet may still be unreliable in practice. In this paper, we examine the utility of cooperatively sharing location data among connected nodes for the purpose of correcting positions with high measurement error, while doing so in a simple manner that can be deployed anywhere without onerous setup tasks or highly specialized hardware. Using a simple data sharing and filtering technique, collaborating users can substantially reduce overall localization error in dead reckoning systems deployed to urban areas where nodes may have a broad spectrum of location quality. We evaluate the effectiveness of this method and examine the system parameters necessary to fully exploit such cooperative localization. We show that mean position error can be reduced by nearly 50 percent for given application scenarios. If distance measurement is available, filtering positions based on estimated error can improve localization accuracy of pedestrian dead reckoning techniques to approximately that of GPS.**

I. Introduction

Advances in wireless networking technology and integrated circuit design have opened up vast possibilities for mobile applications and wireless sensor networks [1]–[5], 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 [6], none have proven suitable for all applications.

Pedestrian dead reckoning techniques, frequently used to augment GPS, require an accurate initial reference point [7], which may not always be available, and also experience compounding measurement error as a function of the the number of steps taken due to both approximation errors and sensor limitations at each step. Although individual personal navigation systems often feature some mechanism for periodic correction [8]–[10], these techniques require preplanned and preinstalled infrastructure, maps, or even manual intervention to correct positioning error, which makes rapid deployment difficult or places undue burden on human operators.

In our experiences with the TeamTrak mobile testbed [11], we have observed that even with sensor hardware equipment of the identical manufacturer and model, error frequently occurs independently between devices. GPS receivers, for instance, 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. Additionally, we have observed that obtaining an initial GPS fix can take as long as 15 minutes, posing obvious problems in many applications.

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 most applications. Corrective localization schemes like this become challenging when spatial separation cannot be reliably determined, error metrics are not guaranteed, reliably estimating the error in pedestrian dead reckoning techniques is difficult, and wireless connectivity in ad-hoc networks are at best haphazard and sporadic. However, like the stepwise refinement used to localize static sensor networks, our goal is to improve lower quality positions using higher quality positions, but without necessarily relying on measuring distance between widely separated objects, which has proven notoriously tricky to implement reliably in practice. In this paper, we examine the utility of cooperatively sharing location information among MANET nodes to provide localization of acceptable quality, but doing so in a manner that is easy to implement, does not require

[1]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.

Fig. 1. The TeamTrak Mobile Testbed

preinstalled infrastructure, and can be accomplished using low-cost, off-the-shelf commodity hardware components.

We examine this problem in two distinct cases: those where estimates of spatial separation are not available, and those where distance can be measured, albeit with measurement error, and thus a simple trilateration scheme can be employed. The effectiveness of each is evaluated using discrete event simulation. The simulation environment is calibrated by incorporating models of sensor error, which are constructed from empirical measurements using outputs from an array of commercially available hardware. The sensor devices evaluated are fully integrated into the TeamTrak platform, one configuration of which is shown in Figure 1, intended to model wearable military computing systems for individual soldiers [12]. 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 purposes of pedestrian dead reckoning. Our simulations demonstrate that mean overall position error can be reduced by as much as 49 percent when the maximum wireless range is limited to approximately 100 meters. Furthermore, mean error can be reduced by as much as approximately 23 percent without any modification to the existing hardware.

II. Cooperative Localization

We first discuss methods for using shared location data. To determine whether any position correction should occur, a mobile node $\mathcal{N}$ first must determine the source and quality of its own position. If multiple localization techniques are available, that with the smallest estimated error value e is used. In TeamTrak, availability of sensors is easily detected and the estimated measurement error is determined either directly from the device, as in the case of GPS, or from a previously calibrated error model in the case of dead reckoning. If GPS localization is available, no further action is taken. Previous experiences attempting to correct erroneous GPS positions by averaging other GPS positions proved counterproductive; doing so essentially eliminated the availability of accurate anchor positions entirely. This quickly propagated large errors among all nodes, especially at longer communication ranges. Therefore, if a GPS fix is available, simply using it uncorrected for localization is probably the better option compared to averaging remote positions. Because of the potentially very large measurement error arising from dead reckoning, adjusting a GPS position by averaging with a dead reckoning position tends to cause a relatively accurate position to be dramatically offset by another with a high estimated error, and this worsening quickly poisons the locations of all connected nodes in the system. Therefore, no attempt is made to correct one GPS position with another. Otherwise, in cases in which distance measurements are not available, positions of connected peers among n nodes are combined using Algorithm 1.

Algorithm 1 Location Average

$(x, y, e) \Leftarrow$ Location-Use-Best($\mathcal{N}$)
if Location-Source($\mathcal{N}$) = GPS **then**
  return (x, y, e)
end if
for $i = 1$ to n **do**
  if Connected($\mathcal{N}$, i) **then**
    $\bar{x} \Leftarrow \bar{x} + x_i$
    $\bar{y} \Leftarrow \bar{y} + y_i$
    $e_{total} \Leftarrow e_{total} + error_i$
    $count \Leftarrow count + 1$
  end if
end for
if $count \neq 0$ **then**
  $\bar{x} \Leftarrow \bar{x} \ / \ count$
  $\bar{y} \Leftarrow \bar{y} \ / \ count$
end if
return $(\bar{x}, \bar{y}, e_{total})$

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 cannot be averaged away in mobile networks since the number of position samples from an anchor representing a single point may be small. This limitation requires an evaluation of the estimated measurement error from each node and acceptance of each position before inclusion in the set of points used 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, which is a random variable, 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 such as horizontal error or horizontal dilution of precision are recorded and correlated a priori with the actual measurement error. Here, confidence is defined as the probability that the actual 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.

Once a position has been accepted based on its error metrics, bearing is determined based on the estimated positions of the node pair. If direction of arrival of the signal could be determined from the network card, that might be an alternative approach, but the low cost, consumer-grade hardware currently employed in the testbed does not provide such capability.

To employ a filtered trilateration scheme in the presence of both localization error and mobility, consider the following:

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.

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

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, which includes positions of significant error magnitude or low confidence:

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

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

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

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

$$L = \big\{(x, y, e, c) \bigm| (x \in \mathbb{R}) \wedge (y \in \mathbb{R}) \wedge (e \in \mathbb{R}) \wedge (c \in [0, 1])\big\}$$

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 $\ell_{\mathcal{N}}$. For all positions $p \in L$:

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

Then for all remote positions $r \in R$:

$$D = \big\{d \bigm| d = \sqrt{(r.x - \ell_{\mathcal{N}}.x)^2 + (r.y - \ell_{\mathcal{N}}.y)^2}\big\}$$

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 $d_i \in D$.

$$\hat{P_{\mathcal{N}}} = \big\{\big((p_i.x - d_i \cos\theta_i), (p_i.y - d_i \sin\theta_i), e_i, c_i\big)\big\}$$

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

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

Thus, $\bar{p}$, the estimated position of $\mathcal{N}$, is determined with a filtered trilateration algorithm as follows:

Algorithm 2 Filtered Trilateration

if Correction-Required($\mathcal{N}$) **then**
  $(x, y, e, c) \Leftarrow$ Location-Use-Best($\mathcal{N}$)
  for $i = 1$ to n **do**
    if $i \neq \mathcal{N}$ and Connected($\mathcal{N}, i$) **then**
      $L[i] \Leftarrow$ Location-Use-Best(i)
      if $L[i].e \geq \tau \wedge L[i].c \geq c_{min}$ **then**
        $R[i] \Leftarrow L[i]$
        $\theta[i] \Leftarrow \arctan\left((L[i].y - y) \ / \ (L[i].x - x)\right)$
        $d[i] \Leftarrow \sqrt{(x - L[i].x)^2 + (y - L[i].y)^2} + \epsilon$
        $R[i].e \Leftarrow |R[i].e| + |\epsilon|$
      end if
    end if
  end for
end if
if $R = \emptyset$ **then**
  return (x, y, e, c)
else
  for $i = 1$ to n **do**
    if $R[i] \neq 0$ **then**
      $R[i].x \Leftarrow R[i].x - d[i]\cos\theta[i]$
      $R[i].y \Leftarrow R[i].y - d[i]\sin\theta[i]$
    end if
  end for
end if
$\bar{x} \Leftarrow \frac{1}{\#R}\sum_{r \in R} r.x$
$\bar{y} = \frac{1}{\#R}\sum_{r \in R} r.y$
$e \Leftarrow \frac{1}{\#R}\sqrt{\sum_{r \in R}(r.x - \bar{x})^2}$
$c \Leftarrow \min \forall r \in R \mid r.c$
return $(\bar{x}, \bar{y}, e, c)$

We evaluate this approach through discrete event simulation, in which the error models for sensor data are based on empirical measurements of physical hardware determined via outdoor experiments using TeamTrak. The following section discusses the experimental results used to calibrate the simulation.

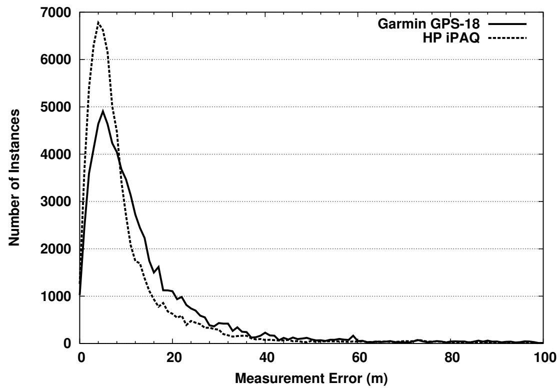

Fig. 2. Distribution of GPS Error

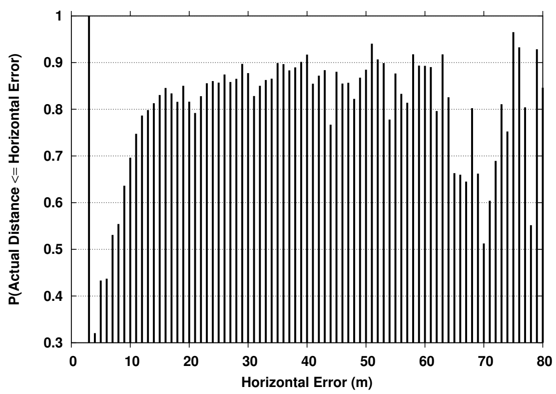

Fig. 3. Confidence of GPS Horizontal Error (Garmin)

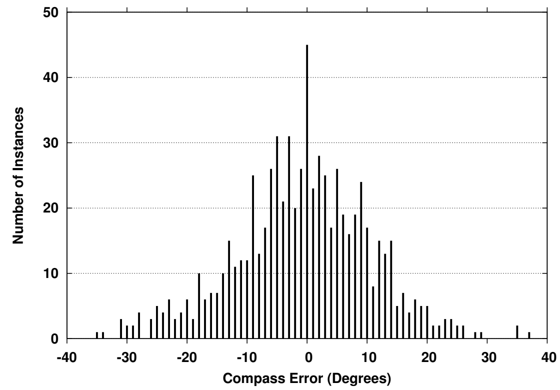

Fig. 4. Distribution of Compass Error

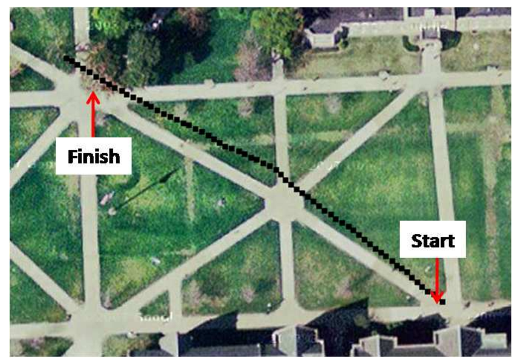

Fig. 5. Dead Reckoning Error in Practice

III. Sensor Evaluation and Error Models

In this paper, we examine the effectiveness of cooperative localization using both the simple averaging in Algorithm 1 and filtered trilateration in Algorithm 2. Due to the inherent difficulty evaluating ideas at scale in mobile networks using physical implementations, we evaluate these approaches via discrete-event simulation using empirically determined parameters. This section discusses measurement of sensor error and development of error models.

Simulation parameters were determined experimentally based on data collected using the testbed. Both the magnitude of GPS error and the expected length and direction of drift are explicitly modeled. The distribution of errors in recorded GPS positions, shown in Figure 2, was incorporated into a single composite histogram. Here, the Garmin GPS-18 and the HP iPAQ GPS receivers, both part of the baseline platform, were used to collect error data. Error was measured by determining the distance from each position sample of 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, we construct models of GPS drift length and direction as a function of a specific error value.

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

A similar approach is taken for 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 4. Similarly, stride length error is determined using a model based on that implemented in NavMote [13] 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 5. 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.

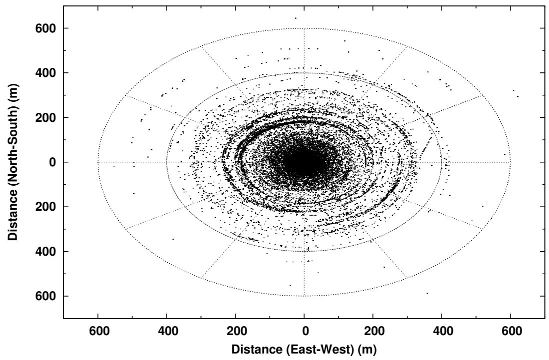

Fig. 6. Wireless Range Measurement

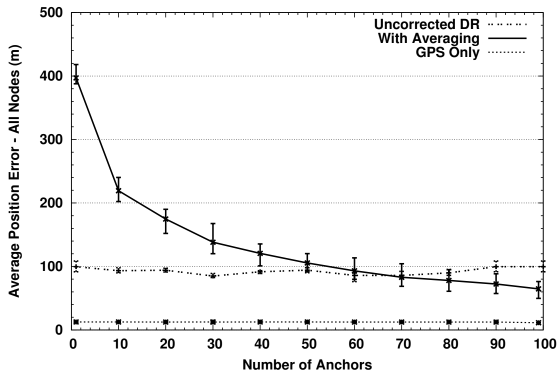

Fig. 7. Averaging (300-Meter Maximum Range)

Connectivity is simulated with the anisotropic wireless communication model introduced by He et al. [14]. 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 6, as an initial simulation parameter. The graphs 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.

IV. 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 [15], 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 in the simulation is based on the average walking speed of an adult male.

A. Simple Averaging: 300-Meter Maximum Range

In this scenario, any node whose location is determined 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. A communication range variable from 200 to 300 meters depending on direction was used, intended to model 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 in increments of 10.

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

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 [14] assume an unlocalized node is surrounded by anchors relatively symmetrically. In a case where mobility invalidates such an assumption, having very few anchors results in an average position whose error is several times worse than in the case of doing nothing.

Figure 7 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 over all time steps in the simulation over 100 independent simulation runs. In this case, more anchors in the system effects a reduction in the overall mean position error. However, realizing improvement in mean error suggests the ratio of anchors to other nodes should be more than 80 percent, but realistically that ratio would likely need to be even higher. Above a 90 percent ratio of anchors, in this scenario, mean error can be reduced by approximately 23 percent.

B. 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. 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. The mean error with uncorrected dead reckoning is dominated by the error towards the end of each

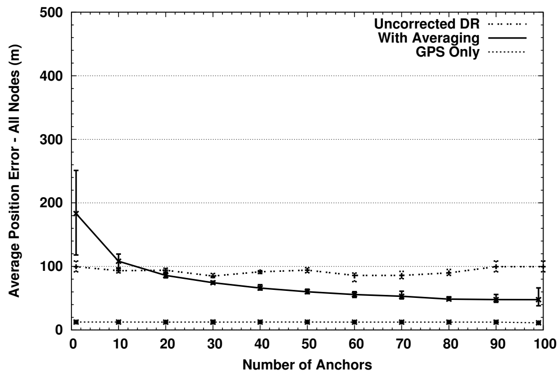

Fig. 8. Averaging (100-Meter Maximum Range)

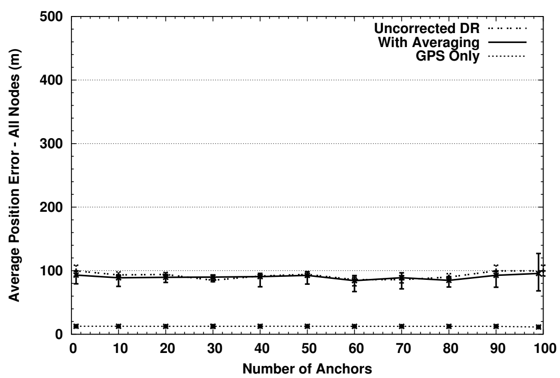

Fig. 9. Averaging (15-Meter Maximum Range)

simulation, i.e., in general, the more steps taken since the last reset, the larger the mean error. Averaging shared location data therefore serves as a reset for dead reckoning, but without the availability of reliable ranging techniques, accuracy is sensitive to communication range. The effect of simple averaging at a maximum range of 100 meters is shown in Figure 8.

C. Simple Averaging: 15-Meter Range

The idea behind averaging positions over lower communication range is to reduce the 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 8 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 9. 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, simply because of the few opportunities for adjusting positions which occur with highly sporadic connectivity.

It should be noted that such short range can be used to adjust positions, since establishing connectivity implies nodes are collocated. In such a case, location information can be

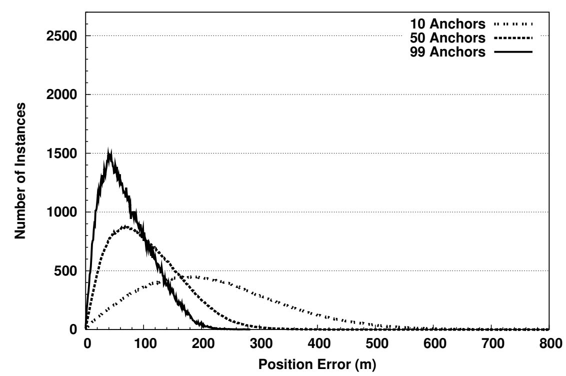

(a) 300-Meter Maximum Range

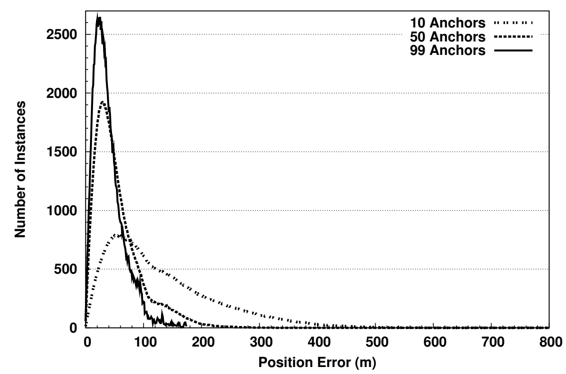

(b) 100-Meter Maximum Range

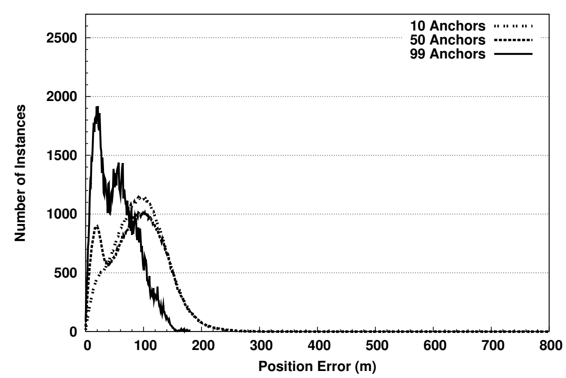

(c) 30-Meter Maximum Range

Fig. 10. Distribution of Position Error (Averaging)

shared effectively, but requires operators to move within close proximity for any adjustment to occur. This method does not work in the case of general mobility.

D. Distribution of Error

To further evaluate the effectiveness of cooperatively sharing location information and using such to reduce position error, the distribution of measurement error, as determined by the 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 10. The figure shows the effect of three different maximum ranges on the overall distribution of position error: Figure 10(a) shows the distribution of errors when nodes are capable of a maximum 300-meter wireless range, while 10(b) shows the result with a 100-meter range, and 10(c), a 30-meter range. In each case, the distribution of position error with varying numbers of anchors in the system is shown, and each histogram is normalized to account for the differing number of nodes whose positions have been adjusted.

In Figure 10(a), 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 7, in this case the resultant mean error is unacceptably large, particularly in the cases with low anchor density, due primarily to the wide spatial separation between nodes, and is thus reflected in the error histogram.

Similarly, with a very short communication range, shown in Figure 10(c), there is relatively little difference in the histograms among all of the anchor ratios due to highly infrequent connections. In the case of 100-meter range (Figure 10(b)), the histograms for different anchor ratios are grouped much more closely, suggesting that when using a simple averaging technique, provided a sufficient deployment scale, localization error is less sensitive to changes in the number of anchors than it is to changes in the maximum communication range.

E. Selective Averaging

While evaluating the effectiveness of cooperative localization, we have varied the maximum communication range, but have not considered measurement error or estimated error in the averages. In the cases of simple averaging, a position which might be known to be bad is included. In this case, dubbed selective averaging, a position that is determined to have a higher measurement error than the position to be adjusted is not considered and is ignored, even if it is the only available anchor. Such a case is the logical next step from the simple case, and 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 in any data sharing or possibly routing protocol, and only include those whose estimated error is no greater than its own. 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 11 shows the effect of such selective averaging in a scenario with 100-meter maximum communication range. Note that while the mean error in the 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

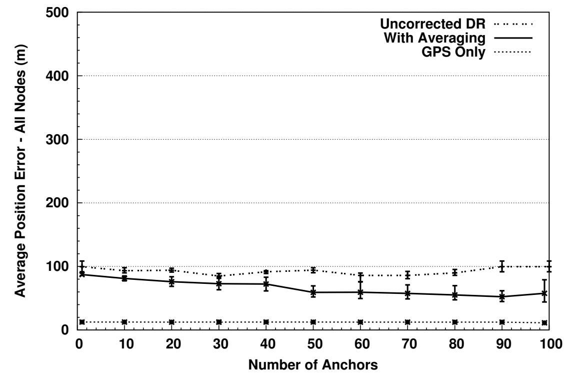

Fig. 11. Selective Averaging (100-Meter Maximum Range)

mean error which occur whenever one node with a very stale position averages its position with either another or an anchor whose GPS position is drifting substantially can be reduced. A selective averaging technique would most likely be used to correct positions in urban environments where GPS-enabled anchors exist but have GPS-limited areas in which fixes are poor or impossible to obtain, so an approach such as this could be effective in the absence of ranging techniques, particularly since this method offers significant improvement without increasing mean overall position error. While with range-free localization techniques such as this some positions may be worsened, the overall net effect is a reduction in dead reckoning error.

F. Filtered Trilateration

Having examined the effects of variations of range-free averaging techniques to reduce overall error in localization from different sources with significant differentials in location quality between nodes, focus now shifts to the use of filtered trilateration, which requires an estimate of both the distance and bearing between a node and one or more anchors. Ordinarily, at least three anchors are required for trilateration, but with a very stale dead reckoning-measured position, even a single anchor can provide improvement. Naturally, since no distance measurement technique for outdoor mobile ad-hoc networks is sufficiently mature and reliable at the present, this work is based largely on an assumption that such a method exists. We are currently exploring techniques to include ranging devices such as laser range finders, and such measurements would permit correction of drifting GPS positions, but such efforts are left to future work as of the time of this paper.

Simulation setup is the same as the previous cases, except instead of averaging the positions of neighbors, the filtered trilateration method specified in Algorithm 2 is implemented. Here, because of the error present in GPS positions as well as in distance measurement, averaging is used to smooth out the random error present in both. Nodes localized with GPS have a mean measurement error of approximately 12 meters. Distance measurement error follows a Gaussian distribution with σ equal to half of the maximum communication range. As shown in Figure 12, by using the filtered trilateration

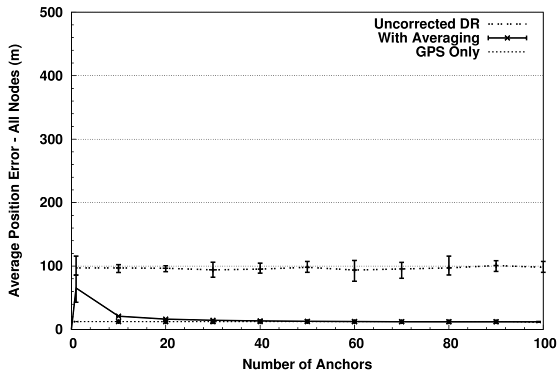

Fig. 12. Filtered Trilateration (100-Meter Maximum Range)

algorithm, localization accuracy for dead reckoning similar to that of GPS can be obtained, but these results depend on the existence of a fairly reliable distance determination. It is important to note that in cases with very low anchor density, positions can still be made worse by attempting to correct using this method. This suggests that any cooperative localization scheme is best employed as a corrective measure for relatively small numbers of nodes in a larger deployment when localization methods such as GPS fail or are unavailable.

V. Conclusion

This paper presents an evaluation of the potential benefit arising from cooperatively sharing location information among users in MANETs when deployed to urban settings or other environments in which GPS capability is limited. Due to the enormous challenges inherent to building a robust ranging scheme for all but the most carefully controlled environments, methods which rely on distance measurement cannot be assumed to be available for any system under development.

With arbitrary distances of inter-node separation, averaging positions of connected neighbors can reduce error in dead reckoning localization, provided that such localization is relatively stale. Because averaging introduces error as a function of the separation between connected peers, simple combination schemes using cooperative localization are most effective if the communication range is neither too long or too short. An overly long range introduces error from wide spatial separation, while an excessively short range does not facilitate connectivity without requiring the operators to manually intervene to correct their position. 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.

Even better results can be achieved by simply filtering positions using estimated measurement error associated with each position. Excluding positions with large measurement error or reduces the likelihood of worsening positions, while simultaneously using anchors with accurate localization to reset dead reckoning. Additionally, filtering positions of high error or low confidence applied to trilateration can improve dead reckoning localization to approximately that of GPS. The implication might seem to be that adjusting positions estimated via similar means, e.g., dead reckoning versus dead reckoning or GPS versus GPS, would be ineffective, but that is not generally the case. While adjusting GPS positions using other remote GPS positions with a simple averaging technique frequently worsens overall localization, the growth in error observed in dead reckoning frequently can be sufficiently large and varied across nodes that some dead reckoning nodes can serve as anchors for others, and in cases where distance measurement 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] Committee on Networked Systems of Embedded Computers, "Embedded Everywhere: A Research Agenda for Networked Systems of Embedded Computers," National Academy Press, 2001, report of the National Research Council Committee on Networked Systems of Embedded Computers, 2001. [Online]. Available: http://books.nap.edu/html/embedded_everywhere

[2] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirici, "Wireless Sensor Networks: A Survey," *Computer Networks (Elsevier) Journal*, vol. 38, no. 4, pp. 393–422, March 2002.

[3] N. Correal and N. Patwari, "Wireless Sensor Networks: Challenges and Opportunities," in *Proceedings of the 2001 Virginia Tech Symposium on Wireless Personal Communications*, June 2001, pp. 1–9.

[4] B. Feder, "Wireless Sensor Networks Spread to New Territory," *New York Times*, July 2004. [Online]. Available: http://www.nytimes.com/2004/07/26/business/26sensor.html

[5] M. Perkins, N. Correal, and R. O'Dea, "Emergent Wireless Sensor Networks Limitations: A Plea for Advancement in Core Technologies," in *Proceedings of IEEE Sensors*, vol. 2, June 2002, pp. 1505–1509.

[6] Corvallis Technology, Inc., *Introduction to the Global Positioning System for GIS and TRAVERSE*. CMTinc.com, 1996. [Online]. Available: http://www.cmtinc.com/gpsbook/"

[7] O. Mezentsev, G. Lachapelle, and J. Collin, "Pedestrian Dead Reckoning A Solution to Navigation in GPS Signal Degraded Areas," *Geomatica*, vol. 59, no. 2, pp. 175–182, 2005.

[8] J. Leonard and H. Durrant-Whyte, "Mobile Robot Localization by Tracking Geometric Beacons," *IEEE Transactions on Robotics and Automation*, vol. 7, no. 3, 1991.

[9] H. Liu and G. Pang, "Accelerometer for Mobile Robot Positioning," in *IEEE Industry Applications Conference*, July 1999.

[10] G. Pang and H. Liu, "Evaluation of a Low-cost MEMS Accelerometer for Distance Measurement," *Journal of Intelligent and Robotic Systems*, vol. 30, no. 3, pp. 249–265, March 2001.

[11] J. Hemmes, D. Thain, C. Poellabauer, C. Moretti, P. Snowberger, and B. McNutt, "Lessons Learned Building TeamTrak: An Urban/Outdoor Mobile Ad Hoc Network Testbed," in *WASA '07: International Conference on Wireless Algorithms, Systems, and Applications*, August 2007.

[12] M. Zieniewicz, D. Johnson, D. Wong, and J. Flatt, "The Evolution of Army Wearable Computers," *IEEE Pervasive Computing*, vol. 1, no. 4, pp. 30–40, 2002.

[13] L. Fang, P. Antsaklis, L. Montestruque, M. McMickell, M. Lemmon, Y. Sun, H. Fang, I. Koutroulis, M. Haenggi, M. Xie, and X. Xie, "Design of a Wireless Assisted Pedestrian Dead Reckoning System - The NavMote Experience," *IEEE Transactions on Instrumentation and Measurement*, vol. 54, no. 6, pp. 2342–2358, 2005.

[14] T. He, C. Huang, B. Blum, and J. Stankovic, "Range-Free Localization Schemes for Large Scale Sensor Networks," in *Proceedings of Mobicom*, September 2003.

[15] D. Johnson and D. Maltz, "Dynamic Source Routing in Ad Hoc Wireless Networks," *Mobile Computing*, vol. 353, pp. 153–181, 1996.