LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

UCRL-TR-201930

Position Estimation of Access Points in 802.11 Wireless Networks

C. A. Kent, P. K. Atwal, W. J. Lennon, F. U. Dowla

January 21, 2004

This work was performed under the auspices of the U.S. Department of Energy by University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

Position Estimation of Access Points in 802.11 Wireless Networks

Claudia A. Kent, Paul K. Atwal, William J. Lennon, Farid U. Dowla
Lawrence Livermore National Laboratory
Livermore, CA 94550 USA
[kent12, atwal1, lennon4, dowla1]@llnl.gov

***Abstract*—We developed a technique to locate wireless network nodes using multiple time-of-flight range measurements in a position estimate. When used with communication methods that allow propagation through walls, such as Ultra-Wideband and 802.11, we can locate network nodes in buildings and in caves where GPS is unavailable. This paper details the implementation on an 802.11a network where we demonstrated the ability to locate a network access point to within 20 feet.**

I. INTRODUCTION

In many 802.11 wireless networks it is desirable to know the location of the network nodes. We address the need for network security where an access point may be providing connectivity to unapproved users, transmitting unwanted data, or otherwise acting in a non-compliant manner, and we propose a method to locate its position. Experimental limitations at this stage require us to address only the inadvertent violator scenario. In a real world application, we would refine the communications to use system-level transactions allowing utility in a more hostile environment. Assuming all nodes communicate with each other via an access point, and the 802.11 signals propagate through walls, a range measurement between a node and an access point is proportional to their distance. The transaction we choose is the PING. The version distributed by the Microsoft Corp. website measures time-of-flight in milliseconds; instead of this, we use a version where trip delay is given in microseconds, hrPING distributed by cFOS Corp. in Denmark.

Given multiple range measurements from several separate locations, we employ the position estimation technique developed in [3] that combines multiple round-trip time-of-flight measurements between a network transmitter and receiver in a closed-form position estimate. In [3], we analyzed system characteristics such as the relationship between the accuracy of a range measurement and the accuracy of the position estimate, whether a "located" node can be used to find another node, the number of independent range measurements necessary for an accurate position estimate, and the degree of improvement with additional measurements. We found that a quality position estimate could be calculated with as few as four noisy range measurements, and this paper details a hardware demonstration of this using an 802.11 network of laptops using PING to find the position of their network access point.

There are several reasons PING is sub-optimal for an 802.11 range measurement. First, PING is a high-level protocol and a low-priority in the CPU stack. The microseconds spent doing "other things" can dramatically reduce the accuracy of the time of flight measurement. Second, a PING requires full cooperation from the receiver, potentially nullifying an obvious application of this technique, which is to locate an "out of compliance" network node. If a node were maliciously out of compliance, one can assume it would not respond to a PING request. We therefore assume that a non-compliant node is acting unintentionally. The solution to both of these problems is to replace PING with a communication protocol on the physical-layer, or MAC level. This would solve CPU stack-priority issues and could potentially allow communication in a non-cooperative environment.

II. SIMULATION SOFTWARE

We developed a MATLAB communications and simulation environment to achieve two goals: to simulate networks of virtual transmitters and receivers, and to act as an interface on a real transmitter or receiver. Both goals require a ranging mechanism, a data sharing communications infrastructure, and position estimation algorithms. A screenshot of the interface is shown in Fig. 2. In the simulation environment we sought answers to network questions such as number of transmitters needed for a position estimate and importance of network geometry, and this analysis is detailed in [3]. Here we implement the software on an actual wireless 802.11 network to test the capability of ranging and positioning.

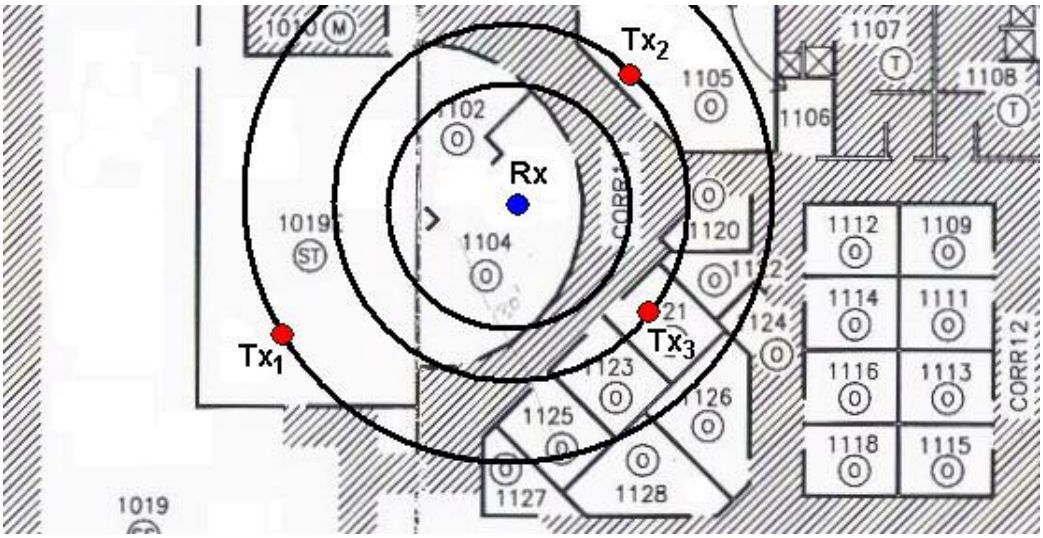

Figure 2. In the MATLAB GUI-based software the transmitters measure the range to a receiver and incorporate multiple range measurements to estimate position.

A. Range Measurement Error

For every PING issued by a transmitter, it receives a batch of replies verifying a connection and noting the elapsed time. As each batch arrives we send it through two stages of filtering to extract the real PING time, since signal multi-path and unknown computer processing time can impart substantial variation within a batch and between batches. In the first filter, we distribute the data in a histogram of 100 microsecond width bins. The data in Fig. 3 (a) (i) is shown in a Histogram in Fig. 3 (b), where the primary subset, or "first hump" is filtered, and the results are shown in Fig. 3 (a) (ii). This stage removes the disproportionately large spikes in the data of Fig. 3 (a) (i), leaving the data within a range of approximately 100-300 microseconds, as opposed to the original 5ms range.

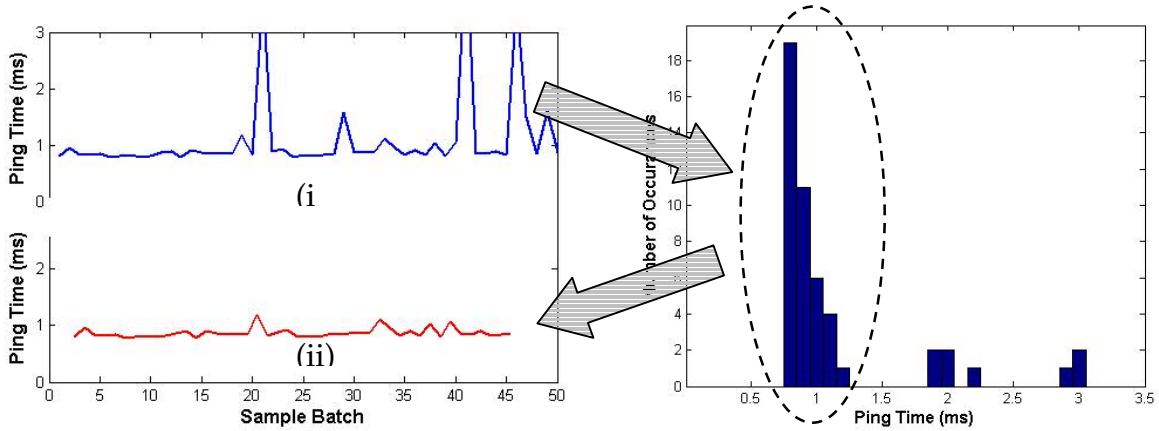

(a) (b)

Fig. 3. The histogram filtering takes the noisy data in (a)(i) and removes the outliers, keeping only the first subset of data in (b). The results are in (a)(ii).

The second filter stage is a recursive weighted least-squares estimator, chosen for it's ability to predict a the true value of a variable given sequential batches of "noisy" variable measurements over time. The filter works recursively by *updating* the least-squares solution after every new batch of data arrives. For the PING issued at the k^{th} sampling interval, we receive a batch of m new measurements $\mathbf{z}_k$, and we estimate the PING time at the next interval z_{k+1}, and call it $\hat{z}_{k+1}$. To achieve this, we assume $\mathbf{z}_k$ takes the form

$$\mathbf{z}_k = \mathbf{H}\mathbf{x}_k + \mathbf{n}_k \text{ where } \mathbf{H} = \begin{bmatrix} 1 & t_1 \\ \vdots & \vdots \\ 1 & t_m \end{bmatrix}. \quad (1)$$

The matrix $\mathbf{H}$ defines the behavior of the system, we assume a first-order system of constant velocity, the vector $\mathbf{n}_k$ is the residual measurement error, and $\Delta t = t_n - t_{n-1}$ is the sampling time. If we knew the value of $\mathbf{x}_k$, we could simply solve for $\hat{z}_{k+1}$, the estimate of PING time at the next measurement. The WLS solution to (1) is

$$\hat{\mathbf{x}}_k = \hat{\mathbf{x}}_{k-1} + \mathbf{K}_k(\mathbf{z}_k - \mathbf{H}\hat{\mathbf{x}}_{k-1}) \quad (2)$$

which is the estimate of $\mathbf{x}_k$ that minimizes a quadratic cost function of residual error. A thorough derivation of (2) is found in [5]. The solution consists of the previous estimate plus the residual error scaled by a gain matrix. The gain matrix is

$$\mathbf{K}_k = \mathbf{P}_{k-1}\mathbf{H}^T(\mathbf{H}\mathbf{P}_{k-1}\mathbf{H}^T + \mathbf{R}_k)^{-1} \tag{3}$$

where $\mathbf{P}_k$ is the error covariance matrix representing the error after the kth estimate.

$$\mathbf{P}_k = (\mathbf{P}_{k-1}{}^{-1} + \mathbf{H}^T\mathbf{R}_k{}^{-1}\mathbf{H})^{-1} \tag{4}$$

Finally, we presume some of our measurements are better than others, and we define a "weighting matrix" $\mathbf{R}_k$ proportional to each new measurement's variation from the previous estimate, or

$$\mathbf{N}_k = \mathbf{I} \otimes \mathbf{z}_k - \mathbf{I}\hat{z}_{k-1} \tag{5}$$

$$\mathbf{R}_k^{-1} = \left(\mathbf{N}_k^T\mathbf{N}_k\right)^{-1} = \begin{bmatrix} \frac{1}{(z_{k1} - \hat{z}_{k-1})^2} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \frac{1}{(z_{k1} - \hat{z}_{k-1})^2} \end{bmatrix} \tag{6}$$

where the operator $\otimes$ is the element-by-element product of the measurement vector $\mathbf{z}_k$ with the identity matrix, resulting in a diagonal matrix of measurement values. The weights "reward" the points that are more closely equal to the previous estimate in a feedback sense.

An example batch of data whose outliers have been removed by the first filter stage is shown in Fig. 4(a), and the results of the second stage WLS filtering is shown in Fig. 4(b), where the final point is the most recent update.

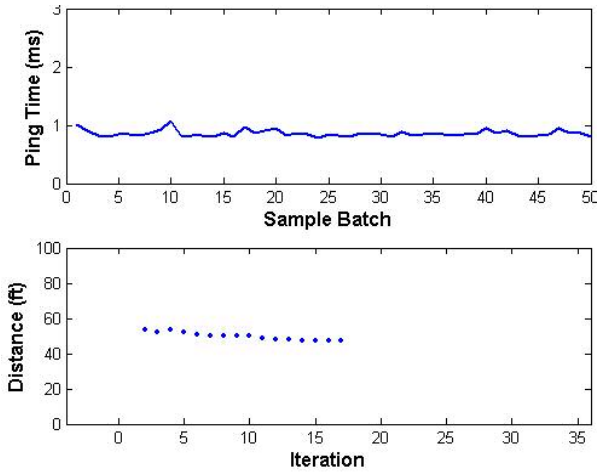

Fig. 4. The second filter stage takes the data from the first stage through a Recursive Weighted Least-Squares filter, where data is weighted within a batch as well as between batches. Finally, a single scalar distance measurement is calculated using Eq.'s (1)-(6).

B. Generating Position Estimates

To share range estimates between transmitters, we employ the communications infrastructure developed in [3] where multiple transmitters maintain information on range measurements between themselves and all receivers in the network, and they share the filtered range measurements with all the other transmitters in the network. Once a transmitter has range measurements between a receiver and three separate transmitters, it can calculate the receiver's position estimate. This calculation was developed in [4], tested in [3], and implemented here. A graphical representation of the technique is shown in Fig. 5, where the Pythagorean Theorem requires two range measurements, R^*_1 and R^*_2 for a target position estimate, and a third range measurement to eliminate the ambiguity between the target position and its "alternate image."

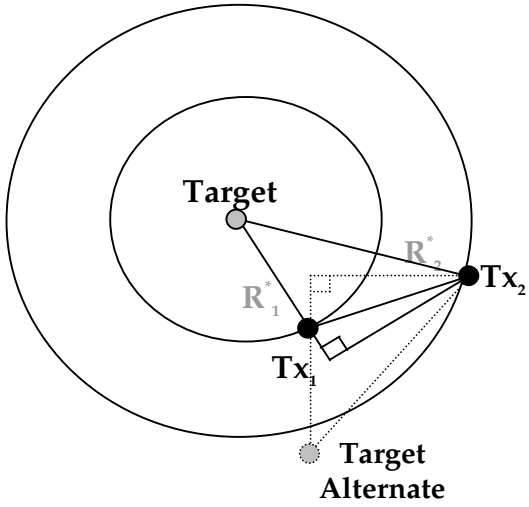

Figure 5. In this graphical representation of the closed-form least squares position estimation method developed in [1], the range measurements from multiple transmitters are combined using the Pythagorean Theorem for an estimate of position.

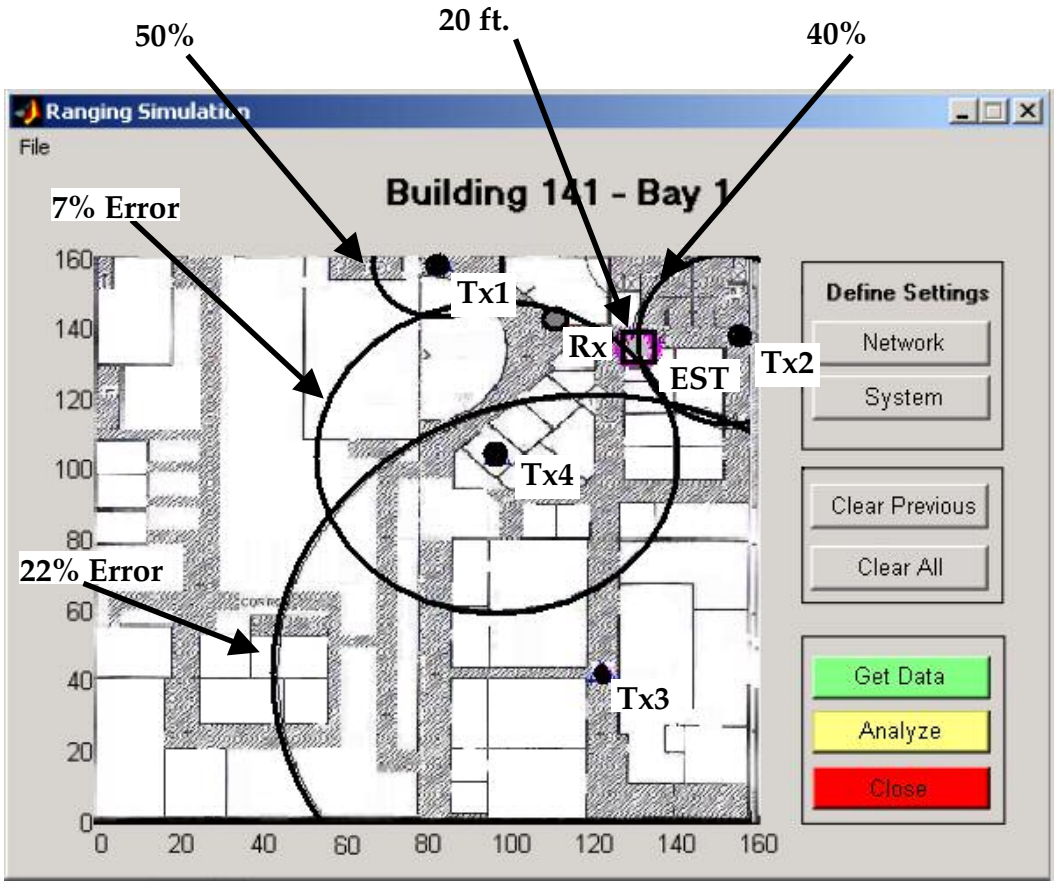

Figure 5. The results from the implementation show that a meaningful position estimate can still be calculated with 50% range measurement error.

III. Conclusion

A representative example of the results from our 802.11 implementation is shown in Figure 6. All data was collected here at LLNL in building 141 where walls and metal filing cabinets create plenty of signal reverberation. Using 802.11b in this environment gave too little variation in our microsecond measurement resolution to be useful. 802.11a however provided large error, but with enough variation between range measurements to be usefully incorporated into a position estimate. Range measurement error using 802.11a varied up to 60% of the total distance, yet a position estimate could still be provided which was within 20 feet of the real position. An example of this is shown in Figure 6. The ability to predict position with such a high range measurement error is due to signal filtering in combination with the powerful position estimation algorithm developed in [4], and tested extensively in [3]. The algorithm can handle large measurement errors as long as additional measurements are introduced.

AUSPICES

This work was performed under the auspices of the U. S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

REFERENCES

[1] N. Bulusu, J.Heidemann, and D. Estrin, "GPS-less low cost outdoor localization for very small devices," IEEE Personal Communications Magazine, vol. 7, no. 5, pp. 28-34.

[2] D. Estrin, L.Gilrod, G. Potie, and M. Srivastava, "Instrumenting the world with wirelesss sensor networks," in Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing, vol. 4, pp. 2033-2036, Salt Lake City, Utah, USA May 2001.

[3] C. Kent, F. Dowla, "Position Estimation of Transceivers in Communication Networks," Technical Report UCRL-JC-?, Lawrence Livermore National Laboratory, October 2003.

[4] J. Smith and J. Abel, "Closed-Form Least-Squares Source Location Estimation from Range-Difference Measurements," *IEEE Transactions on acoustics and Speech*, vol. ASSP-35, no. 12, pp. 1661-1669, 1987.

[5] R.F. Stengel, *Optimal Control and Estimation* (Dover Publications, New York, 1994).