

An Experimental Study of Localization Using Wireless Ethernet

Andrew Howard Sajid Siddiqi Gaurav S. Sukhatme
USC Robotics Research Laboratory, Computer Science Department
University of Southern California, Los Angeles, California, U.S.A
ahoward@usc.edu siddiqi@usc.edu gaurav@usc.edu

Abstract

This paper studies the use of wireless Ethernet (Wi-Fi) as a localization sensor for mobile robots. Wi-Fi-based localization relies on the existence of one or more Wi-Fi devices in the environment to act as beacons, and uses signal strength information from those beacons to localize the robot. Through the experiments described in this paper, we explore the general properties of Wi-Fi in indoor environments, and assess both the accuracy and utility of Wi-Fi-based localization.

1 Introduction

This paper presents an experimental study exploring the use of wireless Ethernet (Wi-Fi) as a localization sensor. Wi-Fi-based localization relies on the existence of one or more Wi-Fi devices in the environment to act as beacons, and uses signal strength information from those beacons to localize the robot. Compared with traditional localization sensors, such as cameras and laser range-finders, Wi-Fi devices are cheap, light-weight and have relatively low power consumption. Moreover, an increasing number of environments come pre-equipped with suitable beacons in the form of Wi-Fi access points. For robots that are too small or inexpensive to carry a laser range-finder or camera, Wi-Fi-based localization offers a viable alternative.

The basic method for Wi-Fi-based localization is as follows. First, a number of Wi-Fi devices placed in the environment to act as beacons; pre-existing Wi-Fi access points, embedded devices, or other robots may serve in this role. Second, one or more robots is used to build a Wi-Fi signal strength map of the environment; this map specifies the expected signal strength for each beacon at every location in the environment. We assume that, during the mapping phase, robots are localized using some other technique. Finally, armed only with a signal strength map, a Wi-Fi adapter and odometry, robots may localize themselves using a variant of the standard Monte Carlo Localization algorithm [3, 9]. Note that this approach is inspired by the work of a number of authors [2, 8] on the subject

of Wi-Fi-based localization. Our key contributions are the embedding of the problem within the context of Monte-Carlo Localization (MCL), the development of appropriate Wi-Fi signal strength maps, and the presentation of comprehensive experimental results.

The experiments described in this paper address four key questions:

1. How does Wi-Fi signal strength vary over time, and to what extent is it affected by day-to-day activity in human environments?
2. How does signal strength vary as a function of robot pose, and is it possible to construct signal strength maps capturing this variation?
3. Are signal strength measurements consistent across robots, such that the signal strength map acquired by one robot can be used to localize another?
4. What level of accuracy is achievable with Wi-Fi-based localization?

In answering these questions, we aim to determine both the accuracy and the practical utility of Wi-Fi-based localization.

2 On Monte-Carlo Localization

For the sake of the discussion that follows, we will briefly sketch the basic theory underlying MCL (see [3] and [9] for a more complete presentation). MCL is a form of Bayes filtering; in the context of localization, the Bayes filter maintains a probability distribution $p(x_t)$ over all possible robot poses x at time t . We interpret the probability associated with each pose as our degree of *belief* that the robot is pose x at time t and denote this $Bel(x_t)$. The belief distribution is updated in response to two events: the robot performs some action, or the robot records a new sensor observation. The filter update rules have the following general form:

$$Bel(x_t) \stackrel{a_t}{\leftarrow} \int p(x_t|x_{t'}, a_t) Bel(x_{t'}) dx_{t'} \quad (1)$$

$$Bel(x_t) \stackrel{s_t}{\leftarrow} p(s_t|x_t) Bel(x_t) \quad (2)$$




Figure 1: Building floor-plans of the SAL2 environment, showing the location of the four wireless beacons A, B, C and D. The occupancy grid used for contact sensing and ground truth pose determination is also shown (free space is shown in white, occupied space in black and unknown space in gray).

where $a_{t'}$ is an action performed at time $t' < t$ and s_t is a subsequent sensor reading. Normalization factors have been omitted for the sake of clarity. The terms $p(s_t|x_t)$ and $p(x_t|x_{t-1}, a_{t-1})$ are known as the sensor and action models, respectively, and must be provided a priori. In this paper, we develop a sensor model for Wi-Fi signal strength (Section 5) and evaluate the utility of this model for robot localization (Section 6).

While conceptually simple, the Bayes filter can be difficult to implement. The sensor and action models tend to be non-parametric, and the pose distribution $Bel(x_t)$ is often multi-modal. MCL seeks to address this difficulty through the use of *particle filters*. Particle filters approximate the true distribution by maintaining a large set of weighted samples. Roughly speaking, each sample in the particle set represents a possible robot pose, and the filter update rules are modified such that the action update step 1 modifies the sample poses, while the sensor update step 2 modifies their weights. *Re-sampling* is used to weed out unlikely samples and focus computation on the more likely parts of the distribution. See [1] for a good tutorial on particle filters and re-sampling techniques.

3 Experimental Setup

The experiments described in this paper were conducted in a typical office environment consisting of rooms and corridors (Figure 1). Four wireless devices were placed at the indicated locations to act as beacons (two iPaq's, a laptop and an Intel Stayton unit). The devices were used in ad-hoc mode, and the iwspy utility (Linux) was used to collect signal strength information. The environment was pre-mapped with a scanning laser range-finder to produce the

occupancy grid shown in Figure 1. During subsequent experiments, the occupancy grid was used in conjunction with a laser-based localization algorithm to generate “ground truth” pose estimates. These estimates were used for both Wi-Fi map generation and the evaluation of Wi-Fi-based localization.

Two Pioneer2DX robots (Fly and Bug) were used in these experiments. Each robot was equipped with odometry, a SICK LMS200 scanning laser range-finder, an Orinoco Silver 802.11b PCMCIA card, and a range-extender antenna. The antenna was placed in an unobstructed location on the top of each robot in order to minimize any correlation between signal strength and robot orientation. The robots use the Player robot device server [5, 4], which includes drivers for measuring Wi-Fi signal strength and algorithms for laser-based (and now Wi-Fi-based) MCL.

In the remainder of this paper, we investigate the general properties wireless signal strength (in both space and time), the construction of signal strength maps, and the use of those maps for localization.

4 Properties of Wi-Fi Signal Strength

Figure 2(a) shows a plot of the signal strength recorded by one of the robots over a 48 hour period. The robot was located adjacent to beacon A in Figure 1, and recorded the signal strength for beacon B. The period captured includes two full working days, with people moving about in the corridors and offices, opening and closing doors, and so on. Some of the variation in the signal strength plot is likely to be a result of such changes in the environment. More importantly, however, all of the variation is confined to a relative narrow band of around 10 dB.




Figure 2: Signal strength measurements for wireless beacon B. (a) Signal strength plotted as a function of time over a 48 hour period. (b) Signal strength as a function of beacon range (log scale). (c) Signal strength as a function of robot orientation.




Figure 3: (a) Signal strength recorded by robot Fly over two complete circuits of the environment. (b) Signal strength recorded by robot Bug over a similar circuit.

Figure 2(b) shows a plot of signal strength as a function of range from one the beacons. This data was gathered by one of the robots over two complete laps around the environment; the robot was localized using the laser-based method described above, allowing the range to the beacon to be accurately determined. In free-space, we expect signal strength (a logarithmic measure) to vary as $-2\log r$, as indicated by the lines in the figure. In indoor environments, however, radio is known to have complex propagation characteristics, with reflections, refraction and multipath effects [6]. Hence it is not surprising that while Figure 2(b) follows the correct general trend for free-space propagation, it also shows significant local departures from this trend. Figure 3(a) shows the same set of data plotted as a function of robot position. In this plot, we note that there is clear variation in signal strength across the environment, and that the local signal strength values remain consistent over multiple passes (within about 5 dB). Somewhat to our surprise, the signal strength values are also consistent when

measured by different robots. Figure 3(b) plots the results generated by a second robot for a similar circuit of the environment; the signal strength measurements are indistinguishable from those acquired by the first robot. Finally, Figure 2(c) plots signal strength as a function of robot orientation. While there does appear to be some correlation between signal strength and orientation, this correlation is weak; the variance over the full range of orientations is at most twice that seen in the static time series plot.

Three important implications can be drawn from the data presented in Figures 2 and 3. (1) There is less variance associated with the position plot than the range plot; hence we expect that a signal-strength *map* should yield better localization results than a simple parametric model. (2) Raw signal strength measurements are consistent across different robots having identical hardware. This implies that a signal strength map acquired by one robot can be used to localize another, greatly increasing the practical utility of this approach. (3) Signal strength is largely invariant with




Figure 4: Interpolated signal strength maps generated using the filters K_1 and $K_1 \cdot K_2$ (described in the text). The maps were generated using the sample set shown in Figure 2(a).

respect to robot orientation, at least for the hardware configuration used in these experiments. This result greatly simplifies the construction of Wi-Fi signal strength maps.

5 Mapping Wi-Fi Signal Strength

While it is clearly impractical to probe the signal strength at every point in the environment, it is relatively easy to collect a representative set of samples and construct an interpolated map. We make two important assumptions: (1) during the sampling process, the robot’s pose is known (in our case, this pose is provided by a laser-based system), and (2) signal strength is invariant with respect to orientation, reducing map making to a two-dimensional problem.

For simplicity, we encode the signal strength map using a regular grid. Each grid cell records the interpolated signal strength value at a particular location, and separate grids are used for each beacon. The grid is generated from raw signal strength data using a low-pass filter, as follows. Let $\Psi = \{(x_0, \psi_0), (x_1, \psi_1), \dots\}$ denote a set of samples such that each sample i has a position x_i and signal strength ψ_i . Let $\Phi = \{(x_0, \phi_0), (x_1, \phi_1), \dots\}$ denote the set of grid cells, where the interpolated signal strength ϕ_i at position x_i is given by:

$$\phi_i = \frac{\sum_j K(|x_j - x_i|) \psi_j}{\sum_j K(|x_j - x_i|)}. \quad (3)$$

$K(s)$ is a weight function whose value depends on the distance $s = |x_j - x_i|$ between the sample at x_j and the cell at x_i . There are many possible choices for the weight function $K(s)$; to date, we have achieved our best results using a combination of two filters. The first filter uses the weight function:

$$K_1(s) = \begin{cases} 1 & \text{if } s < d \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

This filter considers only those samples that lie within dis-

tance d of a cell, and generates an unweighted local average. Figure 4(a) shows the map generated by this filter ($d = 1$ m) when applied to the sample data shown in Figure 3(a). Note that the filter generates good values in those regions visited by the robot, but leaves large ‘holes’ in the unvisited portion of the map. To fill these holes, we apply a second filter with weight function:

$$K_2(s) = s^{-m} \quad (5)$$

where m is generally a low integer value. This is a fairly typical interpolation filter that considers all samples, but assigns higher importance to those closer to the cell. Figure 4(b) shows the final interpolated map.

To use the signal strength map for localization, we must augment it with an appropriate sensor model 2. If we assume that the sensor noise is normally distributed, we can write down the sensor model $p(\phi|x)$ for Wi-Fi signal strength:

$$p(\phi|x) = \exp \frac{-(\phi - \phi_i)^2}{2\sigma^2} \quad (6)$$

where ϕ_i is the interpolated signal strength for the cell i containing pose x , and σ^2 is the expected variance in the signal strength. Based upon time-series plots such as the one shown in Figure 2(a), we typically choose σ to be 10 dB. Note that it is not our intention to suggest that this particular sensor model (or the interpolation procedure describe above) is the the *best* possible model for Wi-Fi-based localization. Rather, we propose that this is a *sufficient* model, and seek to determine its utility empirically.

6 Localization

In order to assess the comparative utility of Wi-Fi-based localization, we compare the localization results achieved using three different combinations of sensors: Wi-Fi, contact, and Wi-Fi plus contact sensing. In all three cases we




Figure 5: Localization results using different combinations of Wi-Fi and contact sensing. The plots on the left show the estimated robot trajectory (the true trajectory is indicated by the ‘+’ symbols); the plots on the right show the error in the pose estimate as a function of the distance travelled by the robot. (Top) Wi-Fi sensing only. (Middle) Contact sensing only. (Bottom) Both Wi-Fi and contact sensing.

assume that odometry is also available. Note that the ‘contact’ sensor in our case is logical rather than physical, and simply asserts that the robot cannot be co-located with another object. Thus, given an occupancy map of the environment, the ‘contact’ sensor rejects those poses that lie in occupied space.

Our basic experimental methodology is as follows. Data was collected from a robot performing a series of circuits of the environment, and processed off-line using different combinations of the recorded sensor data. In all cases, the initial pose of the robot was entirely unknown. Furthermore, *different robots* were used for the map acquisition and localization phases, and these two robots executed their circuits in *opposite* directions.

The combined localization results are presented in Fig-

ure 5. The top row shows the localization results using Wi-Fi sensing only: the left hand figure plots the estimated robot trajectory, while right hand figure plots the error in the pose estimate as a function of the distance travelled by the robot (i.e., the distance between the true pose and the estimated pose). The key feature to note is that the pose estimate converges very quickly (within about 10 m of robot travel) to a steady state error of 0.40 ± 0.09 m. The second row in Figure 5 shows the results using contact sensing only. Here, the estimate takes much longer to converge: the robot has travel around 80 m before it can gather enough data to make an unambiguous determination of the robot’s pose. The steady-state error, however, is only 0.26 ± 0.03 m; better than that obtained using the Wi-Fi sensor alone. The third and final row in Figure 5 shows the results of combin-

ing Wi-Fi and contact sensing. Here, convergence is rapid, and the steady-state error is only 0.25 ± 0.02 m. It would appear that these two sensors complement each other extremely well: Wi-Fi ensures rapid convergence when the initial pose is unknown, while contact sensing improves the subsequent accuracy of the estimate. One can, of course, further improve the estimate by adding data from additional sensors; with sonar or laser range data, we expect to achieve steady-state errors of less than 0.10 m with convergence distances of a few meters.

The results described above were generated using all four of the Wi-Fi beacons placed in the environment. Fewer beacons can be used, with a consequent decrease in localization accuracy. The steady-state errors for different combinations of the beacon A, B, C and D shown in Figure 1 are as follows.

Beacons	Error (m)	Beacons	Error (m)
A	0.87 ± 0.34	A,B,C	0.45 ± 0.14
A,B	0.55 ± 0.12	A,B,C,D	0.40 ± 0.09

Given the basic geometry of trilateration in two dimensions, it comes as no surprise that best results are achieved using two or more beacons. It should also be noted that for this particular set of experiments, the location of the beacons was selected based on the likelihood that it would yield good localization accuracy. Other configurations may yield lower accuracy for the same number of beacons.

7 Conclusion

Four major conclusions can be drawn from the results presented in this paper.

1. Signal strength values are stable over time, and relatively unaffected by environmental changes induced by day-to-day activity.
2. Signal strength values vary relatively smoothly with increasing range from the beacon. As a result, it is possible to produce interpolated signal strength maps from a relatively sparse sampling of the environment.
3. Signal strength measurements are consistent across robots using identical Wi-Fi hardware, and thus maps generated by one robot can be used to localize another.
4. Given a sufficient number of beacons and a signal strength map, robots can be localized to within 0.50 m. Accuracy can be increased by adding additional beacons and/or other forms of sensing.

Clearly, much experimental work remains to be done; this paper does not, for example, consider the effect of different environments, different beacon configurations, or heterogeneous hardware. Nevertheless, the results presented here indicate that Wi-Fi is a very effective localization sensor.

Resources

Wi-Fi-based localization has been incorporated into the Player robot device server [5], which can be downloaded from the Player/Stage web-site [4]. The data-sets used in this paper are also available on the Radish (Robotics Data Set Repository) web-site [7].

Acknowledgments

This work is sponsored in part by DARPA grants DABT63-99-1-0015 and 5-39509-A (via UPenn) under the Mobile Autonomous Robot Software (MARS) program.

References

- [1] S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for on-line non-linear/non-gaussian bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, Feb. 2002.
- [2] P. Bahl and V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *INFOCOM (2)*, pages 775–784, 2000.
- [3] F. Dellaert, D. Fox, W. Burgard, and S. Thrun. Monte carlo localization for mobile robots. In *IEEE International Conference on Robotics and Automation (ICRA99)*, May 1999.
- [4] B. Gerkey, R. Vaughan, and A. Howard. Player/Stage homepage. <http://playerstage.sourceforge.net>, September 2001.
- [5] B. P. Gerkey, R. T. Vaughan, and A. Howard. The player/stage project: Tools for multi-robot and distributed sensor systems. In *Proceedings of the International Conference on Advanced Robotics (ICAR 2003)*, Coimbra, Portugal, June–July 2003. To appear.
- [6] M. Hassan-Ali and K. Pahlavan. A new statistical model for site-specific indoor radio propagation prediction based on geometric optics and geometric probability. *IEEE Transactions on Wireless Communication*, 1(1):112–124, Jan 2002.
- [7] A. Howard and N. Roy. Radish: the robotics data set repository. <http://radish.sourceforge.net>, May 2003.
- [8] A. M. Ladd, K. E. Bekris, G. Marceau, A. Rudys, D. S. Wallach, and L. E. Kavraki. Using wireless Ethernet for localization. In *Proceedings of the 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Lausanne, Switzerland, Sept 2002.
- [9] S. Thrun, D. Fox, W. Burgard, and F. Dellaert. Robust Monte Carlo localization for mobile robots. *Artificial Intelligence Journal*, 128(1–2):99–141, 2001.