Practical WiFi Localization for Autonomous Industrial Vehicles

Frederic Moster[1,2] **and Ashley Tews**[1]
[1]Autonomous Systems Laboratory, CSIRO ICT Centre
PO Box 883, Kenmore, Qld 4069, Australia
Email: `Firstname.Lastname@csiro.au`
[2]Centrale Nantes
1 rue de la Noe BP 92101, 44321 Nantes Cedex 3, France

Abstract

Localization for industrial vehicles is a key issue in the development and use of autonomous systems on industrial worksites. The WiFi technology being widely used in such environments can be used as a base for an easy to set up system requiring little hardware. In this paper we propose a system estimating the location of industrial vehicles, based only on WiFi signals, that can operate both in and out of buildings and sheds, and that can deal with the kidnapped robot problem. It is based on the use of a signal strength intensity map and a particle filter, and has been demonstrated in a typical industrial worksite.

Figure 1: Vehicle working on industrial worksite.

1 Introduction

Real−time accurate knowledge of the location of mobile vehicles in industrial worksites is a key issue in the use of autonomous systems, for both safety and efficiency reasons. Various types of sensors and methods can be used to achieve this goal (e.g. odometry, lasers, GPS, WiFi), each one having his own advantages and disadvantages (good results are usually obtained by combining different types of complementary sensors).

We want to build a reliable region−based localization system that can be used on mobile industrial vehicles (e.g. Figure 1). It needs to:

- be easy to set up
- be robust to varying environmental conditions
- be able to deal with the kidnapped robot problem (presented in [Thrun *et al.*, 2005])
- be able to work inside and outside of buildings
- require little hardware other than what can usually be found on an industrial worksite

The WiFi technology frequently used on industrial worksites consists of access points that form the pre−existing wireless networks. These networks can be considered as the base for a localization system. The only additional hardware requirement for such a system is a WiFi card (the receptor), which can be coupled with an external antenna for better signal reception. Moreover, depending on the coverage of the distribution of the access points, a localization system based on WiFi signals is theoretically able to give an estimation of the location instantly, anywhere on the site. However several problems make the design and use of such a system difficult, particularly the WiFi signal characteristics: problems of interference, multi-path and reflections are more prolific in outdoor environments than in indoor environments, and can make the recorded signal fluctuate.

Our approach is as follows: we acquire the WiFi signals recorded from the different 802.11b access points distributed over an industrial worksite, and calculate the means of the signal strength over time for each access point. The means are used by the sensor model of the particle filter, compared with an *a priori* intensity map. By using an intensity map, the system does not need other knowledge such as the position, number and type of access points.

The localization system is divided into three sub−systems: signal processing, the particle filter, and the comparison model used by the particle filter. We analyze each sub−system in order to optimize it to our needs. The remainder of this paper is organized as follows: in Section 2 we present related work in WiFi localization systems. In Section 3 we discuss the choices made concerning the signal processing. In Section 4 we present the comparison model (an *a priori* intensity map of the site) used by the particle filter. In Section 5 we present the particle filter used. Section 6 contains the experimental study and the results obtained. And finally our conclusions are presented in the Section 7, as well as a discussion on future possible work on the subject.

2 Related Work

Location−aware computing is an active field of research. Many systems have been built using various technologies designed for various environments and using various methods. Many of the technologies used for location−aware computing are presented in [Hightower and Borriello, 2001].

Many WiFi localization systems have been developed for indoor environments, where the WiFi signal is more predictable and reliable than outdoors. In indoor environments it is possible to use a signal strength intensity map as well as probabilistic models to predict the signal strength at a given place − both techniques have been demonstrated to be relatively reliable. Good accuracy can be achieved when the WiFi technology is associated with a complementary sensor such as odometry: [Howard *et al.*, 2003] obtained an accuracy of less than 1 metre, [Serrano *et al.*, 2004] obtained an accuracy of less than 3 metres.

Systems using only WiFi have also demonstrated good results. Various techniques have been tried to either improve the accuracy or to decrease the calibration effort required for setting the system's parameters. Localization systems using Bayesian localization frameworks [Ladd *et al.*, 2004] or Monte Carlo methods [Dellaert *et al.*, 1999] have been successful, as well as using the signal strength information to triangulate a user's coordinates [Bahl and Padmanabhan, 1999]. The decision to use an intensity map [Haeberlen *et al.*, 2004] or probabilistic models to predict the signal strength at a given location, or a combination of both methods [Roos *et al.*, 2002], have also been much studied. For example, RADAR is a system based only on WiFi signals, using the existing infrastructure of an indoor RF wireless LAN to identify the location of a user's device using previously collected WiFi 'fingerprints' at known locations [Bahl *et al.*, 2000]. The median resolution of this system was in the range of 2 to 3 metres.

LOCADIO [Krumm and Horvitz, 2004] is another indoor localization system, that uses the signal strength from existing access points to infer the motion and location of a client, using Hidden Markov Models, but as with most of the systems described in this section, the calibration effort is important.

Although the indoor systems have achieved good performances, extending them for use outdoors often requires precise calibration and extensive installation efforts. Systems such as Place Lab [LaMarca *et al.*, 2005] have been built especially for large environments (metropolitan areas). Place Lab uses pre−existing hardware to reduce the cost and time effort for configuring the system. An accuracy of 13−20 metres in dense urban areas, and 40 metres in suburban neighborhoods can be achieved [Cheng *et al.*, 2005], using GSM beacons and fixed Bluetooth devices as well as WiFi access points.

Another system has been built, working both in indoor and outdoor environments [Letchner *et al.*, 2005]. It requires minimum calibration, is able to bootstrap from sparse training data, and has minimal hardware requirements. It uses a hierarchical Bayesian Sensor Model (using a particle filter) and a graph−based location estimation system. Using only WiFi this system achieves good accuracy (12−26 metres) in outdoor environments.

Our research differs from previous research in that we focus on the needs and specificities of localizing vehicles in industrial worksites which requires low installation, high reliability and the capability to be accurate inside and outside of large buildings.

3 Signal Processing

One of the biggest issues concerning a localization system based only on WiFi is the unreliability and instability of WiFi signals. In this section we present an overview of the main factors of perturbation of signals, and explain the choice of the metric used by our localizer considering these issues.

3.1 WiFi Signal Characteristics

There are three main phenomena that perturb the WiFi signal and cause multi-path effects: diffraction, reflection and scattering.

Diffraction occurs when an electromagnetic wave encounters a surface with irregular edges, and travels along a path other than the line of sight.

Reflection occurs when an electromagnetic wave encounters an obstacle whose size is bigger than the wavelength of the signal. The propagating wave then loses power while passing through the obstacle, and reflected waves can be created and propagate along different paths.

Scattering occurs when an electromagnetic wave encounters an obstacle whose size is smaller than the wavelength of the signal.

These effects can cause a signal to follow multiple paths, not only the line of sight. At the intersection of these paths, constructive and destructive interference between the intersecting waves can cause the signal to be increased or decreased. As a result, recording the WiFi signal from one access point at a given position leads to a signal spread over a large band of dB (this can be seen on Figure 2).

Moreover, we noticed during experiments that the signal strength at a given position varies with the orientation of the antenna (confirmed by [Bahl and Padmanabhan, 1999]) and with the difference between the height of the access point's antenna and the height of the receiver's antenna (confirmed

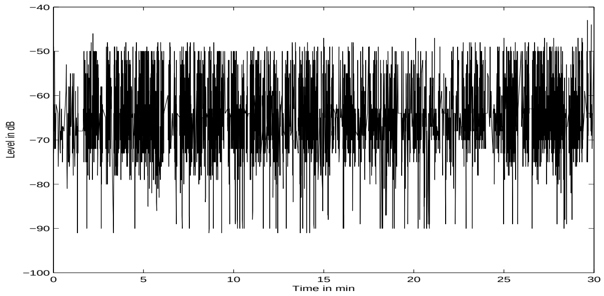

(a) Signal Strength.

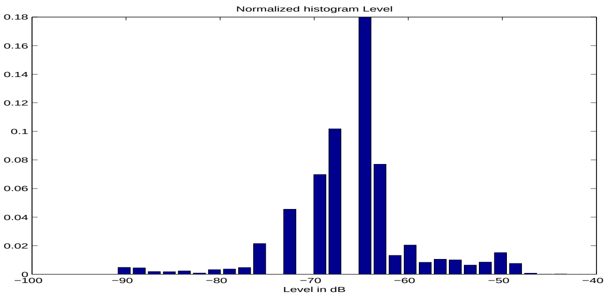

(b) Histogram.

Figure 2: WiFi signal recorded at a given location during 30 minutes.

by [Kotz *et al.*, 2003]). Finally, we have not analyzed the influence of atmospheric conditions such as temperature and humidity on the signal strength, but it is common knowledge that these conditions impact on the propagation characteristics of electro−magnetic signals.

3.2 Signal Preprocessing for Localization

WiFi localization ideally requires a method that stabilizes the received signals at a location, and uniquely identifies each location in the environment. As stated in the previous section, the raw signal strength is not suitable for accurate and reliable positioning since the signal at a given location is not stable enough to give usable information.

An alternative method for estimating a user's location is the response rate metric (defined in [Cheng *et al.*, 2005]). The response rate is the percentage of time that a given access point was heard in all of the WiFi scans at a specific location. The idea being that moving away from one access point towards others reduces the rate of packets received from the initial access point, which indicates an increasing distance from it. However this is not true in areas where the receiver detects only one access point or the receiver moves away at the same rate (therefore the same distances) from all detected access points. As we want a localization system that is flexible and adaptive to different environments and access points distributions, we do not consider this method.

If the raw signal strength is not usable, it is however possible to process it to into a more useful and stable form. Our method uses the mean of the signal strength over a certain amount of time. The time period is the largest possible considering the estimated maximum velocity: V m/s of the target platform and the required accuracy: A m. The time T required then becomes: T $<=$ A/V seconds.

This mean absorbs variations in the signal, which can then be used as a reliable method of characterizing the signal at a given location in real−time. To confirm this we record the mean of the strength of the signal emitted by a single access point as a function of the distance between this access point and the receiver. The mean is calculated over 10 sec, the receiver going away from the access point in straight line with a speed of 0.5 m/s. Doing this experiment several times, we compare the shapes of the curves obtained (examples can be seen on Figure 3). Those shapes present a similar profile, confirming that this method stabilizes the signal and enables us to use it to approximately identify the region where the mean of the signal strength has been recorded.

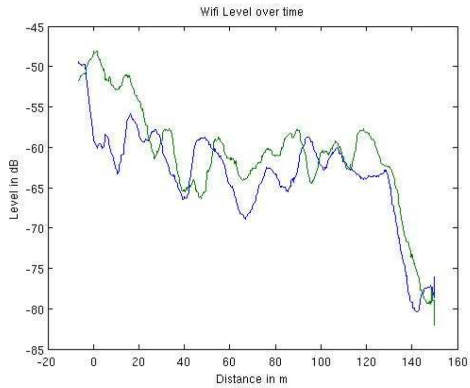

Figure 3: Comparison of the curves obtained by calculating the mean of the signal strength recorded over time and by drawing it as a function of distance between the access point and the receptor.

4 Comparison Model: Intensity Map

The localizer has to be able to compare the data acquired in real−time to some comparison model, which can be either a propagation model for each access point (method used in [Serrano *et al.*, 2004], [Bahl and Padmanabhan, 1999] in indoor environments) or an *a priori* intensity map. A propagation model does not suit the requirements of our localizer, because of the lack of flexibility it creates and the fact that in an outdoor environments, the signal is not predictable and stable enough with too many parameters influencing and perturbing it. In this section, we present the structure of the intensity map used by our localizer and the method used to build it.

4.1 Structure of the Map

The intensity map is a discretization of the industrial site to be covered by the localizer in regions, each region being characterized by its position (coordinates x and y) and by the fingerprints of all the access points in range at this position, that is to say by the signal strengths of each access point recorded in the region.

All the regions have the same size, that has to be small enough to achieve a good accuracy (the size of the regions directly determining the maximum accuracy of the localizer), but big enough to fit with the speed of the vehicle and the time of calculation of the means of the signal strengths that is used.

For example if the vehicle's nominal speed is V m/s and the means of the signal strengths is calculated over T sec, than the size of the regions must be at least V∗T m.

4.2 Building a Map

A map is built by calculating for each region the mean of all the signal strengths received in this region. To do this we drive around the site with another localizer (presented in [Duff *et al.*, 2006]) providing us ground truth, to determine which region we are in when we acquire a WiFi packet.

Of course the more packets recorded in each region, the more reliable the value of signal strength for this region is, so building the map requires a low speed. Moreover all the regions that we want to be represented in the map have to be covered during this phase.

Figure 4 is the visualization of the environment discretized into 5m by 5m regions. The white squares are the locations where a signal has been recorded, and correlate to the drivable areas in the environment. The access points of the site are represented by the blue squares.

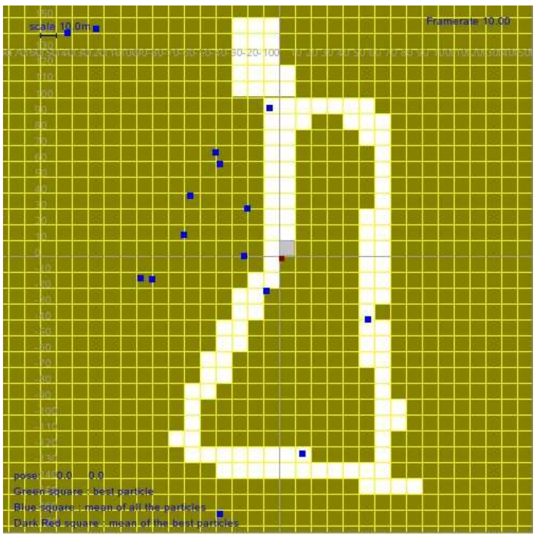

Figure 4: Example of the environment discretized into 5m by 5m regions.

5 The Particle Filter

Particle filters are sample−based probabilistic approximations whose main objective is to track a variable of interest as it evolves over time (see [Thrun *et al.*, 2005] for a good study of particle filters). In this Section we present the general principles of particle filters, before explaining the adaptation for our WiFi localization system.

5.1 General Principle

Particle filters have been successfully used by the robotics community for tracking robot positions and location estimation ([Haeberlen *et al.*, 2004], [Gustafsson *et al.*, 2002]). Many studies have been made to measure their efficiency. For example [Hightower and Borriello, 2004] is a case study of applying particle filters to location estimation using a multi−sensor location system, and shows that particle filters are a good choice to implement location estimation for ubiquitous computing. This is due to their accuracy as good deterministic algorithms, their practicality and their flexibility.

[Arulampalam *et al.*, 2002] presents a review of Bayesian algorithms for non−linear/non−Gaussian tracking problems with a focus on particle filters and their different variants. It demonstrates that the critical parameter to take into account when designing a particle filter for particular applications is the choice of importance density. Another review of Bayes filter implementations can be found in [Fox *et al.*, 2003].

A particle filter is a sequential Monte Carlo method, implementing a Bayes filter. It represents the probability distribution for the location estimate as a set of weighted samples, with each sample being a discrete hypothesis about the location of the object. For robot localization, the standard approach used for implementing a particle filter is based on the Sequential Importance Sample with Resampling (SISR) procedure:

- a motion model is used to predict each sample's motion;
- a sensor model is used to weight all the samples considering the current observation (the weights being normalized so that their sum equals to 1);
- a resampling is performed using the importance factors of the samples (particles with infinitesimally small weights are eliminated);

The particle filter used by our localizer is based on this procedure, and has been adapted to fit with the localizer's requirements. These adaptations are presented in the next section.

5.2 Specific characteristics

Our localizer works only with information provided by the WiFi packets recorded. The action model of the particle filter is therefore reduced to the noise model used, that perturbs each particle up to two times the size of a region at each step. But we have to introduce an additional step of resampling after the motion model, that reinitializes particles with infinitesimally small weights. During this phase, every particle that has an importance below a fixed threshold is moved to a random location on the map, enabling a cloud of particles to jump from one location on the map to another, which is necessary if the localizer has been temporarily lost (if the vehicle drives in a zone where there is no signal for example) and has to quickly catch up with the vehicle.

This characteristic of our particle filter influences the choice of the method used to determine the estimated position (E) of the mobile vehicle. We use three methods:

- the mean of all the particles weighted by the importance of the particles which gives the position M1
- the mean of the top 10 percent of the particles weighted by their importance which gives the position M2
- the best particle which gives the position B

Considering the relative positions of these three points, we choose the one being the best adapted to the current state of the distribution of the particles ($|\mathrm{M1} - \mathrm{B}| < 2$ meaning that M1 and B are situated in the same region or in adjacent regions):

if $|\mathrm{M1} - \mathrm{B}| < 2$ **then** $\mathrm{E} = \mathrm{M1}$
else if $|\mathrm{M2} - \mathrm{B}| < 2$ **then** $\mathrm{E} = \mathrm{M2}$
else $\mathrm{E} = \mathrm{B}$

Indeed several clouds of particles can coexist at the same time in different parts of the map (if two regions have exactly the same fingerprint for example, which can happen in zones where very few access points are available), and then M1 and M2 does not represent the estimated position of the vehicle anymore, but a location between the several clouds.

Concerning the evaluation model of the particle filter, it simply compares the means of the signal strengths recorded in real−time by the localizer to the signal strengths expected in the region where the current particle is. The mean of the error for each sensed access point is used as parameter of a Gaussian that gives the particle its importance value. The fact that only sensed access points are taken into account (and not missing access points) prevents the localization system from crashing when an access point is broken. It will lose accuracy because regions covered by this access point will be less differentiated, but positioning is still available in these regions even if an expected access point is not seen.

The particle filter running can be seen on Figure 5: the accessible locations of the site are in white, the access points of the site are the blue (undetected access points) and red (detected access points) squares, the particles are the grey points, the true location of the vehicle being the light green square.

6 Experiments and Results

In this Section we present the experiments made in an industrial worksite and the results obtained by our localizer.

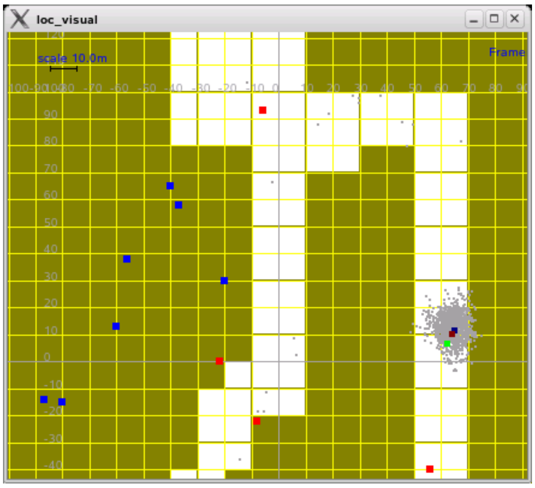

Figure 5: The particle filter running in the worksite.

6.1 Experimental Testbed

To acquire WiFi packets we use an 802.11b PCMCIA Orinoco Gold WiFi card with a 3dB external antenna plugged in to a laptop situated on a Toro ride−on tractor (Figure 6 shows the tractor with the external antenna mounted on a pole). The external antenna has been mounted on a pole to limit the influence of the difference of height between it and the worksite's access points' antennas that are mostly situated several metres above the ground, and to simulate the height of an industrial vehicle like a forklift.

The software architecture is composed of several parts. A packet sniffer called Kismet [Kismet, 2006], enables us to acquire WiFi packets from every access point on every channel. Kismet identifies networks by passively (without sending any loggable packets) collecting packets and detecting standard named networks, detecting hidden networks, and inferring the presence of non-beaconing networks via data traffic. To find as many networks as possible, Kismet also uses channel-hopping.

The communication between the sniffer and the localizer is carried out by Dynamic Data eXchange (DDX) (see [Corke *et al.*, 2004]). DDX is a distributed software architecture that allows programs to share data through an efficient shared memory mechanism. The WiFi packets are processed by the localizer, that also runs the particle filter (Figure 7 represents the architecture of the system and the links between its components).

The environment for the experiments is an industrial worksite of approximately 400m by 250m. The localizer does not have to cover all this area, only the places where the vehicles can actually go, which represents 168 square regions of 10 metres on 10 metres. The site is classified into zones of regions (Figure 8 showing the locations of the different zones on the site on which we are testing the localizer), each with different characteristics which allows us to test the localizer

Figure 6: Toro Tractor used for the experiments.

in different conditions:

- a large open area (called the Bay Compound) of about 50 metres on 40 metres (about 7 access points available in this zone);
- a road (Road 1) surrounded with buildings which is well covered by access points (about 8 access points available in this zone);
- a road (Road 2) in an open area surrounded by bushland and forest (about 6 access points available in this zone);
- a road (Road 3) in an open area surrounded by bushland and forest which is not well covered by access points (less than 4 access points available in this zone);

In these experiments we evaluate the accuracy of our localization system by manually driving the tractor around the site. We use different sized regions for the intensity map, and evaluate the system using the error generated between the WiFi nominated region and actual vehicle location provided by an external system (the same used to build the intensity maps). The global results are presented followed by how the results can be improved considering the type of zone to be covered. Since the industrial vehicles targeted by our localizer having a maximum length of $3-5$ metres, we do not aim at an accuracy less than 5 metres. Moreover, the time of mapping increases quickly when the size of the regions decreases, so we use intensity maps with regions of size 5m by 5m and 10m by 10m. In region sizes larger than 10m by 10m, the signal varies significantly. We built and use two maps of each type

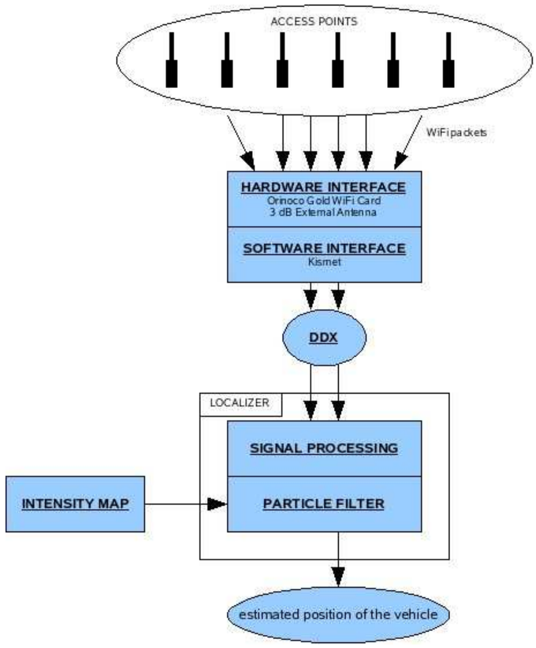

Figure 7: Architecture of the system.

Table 1: Results obtained on the whole site, using different maps (2 maps 5*5, 2 maps 10*10) and three experiments per map.

	$5*5$ Day 1	$5*5$ Day 2	$10*10$ Day 1	$10*10$ Day 2
Error (m)	25.95	27.5	21.3	18.9

(created on different days) to verify their reliability over time. All experiments use 1000 particles for the particle filter.

6.2 Global Results

Table 1 shows the mean results obtained with 10m by 10m and 5m by 5m regions. Two different intensity maps created on two different days are used for each region size to verify the reliability of the map. Each map has been tested three times (on three different days). In Table 1, the global error of the positioning is in the range of 15 meters to 30 meters, regardless of the map and region size. However this error is the mean of the errors recorded at each position. The design of the localizer to be able to localize anywhere and be able to deal with the kidnapped robot problem, influence the behavior of the particle filter. It is possible for the cloud of particles to jump from one region to another that can be far away, knowing that different regions can have quite similar signal strength fingerprints. This can provoke high error values over a short period (the time for the vehicle to move into more differentiated regions), which are included in the calculation of the global mean of the errors in Table 1.

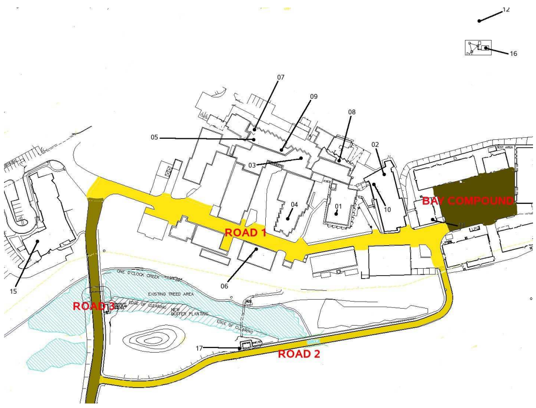

Figure 8: Map of the site (approximately 400m x 250m), with the different zones highlighted and the access points.

Figure 9 shows the distribution of the errors for the different region sizes. The X−axis represents the number of regions of error of the localizer (for x=0 the localizer is correct). The Y−axis represents the percentage of time that the localizer is wrong by a given number of regions of error. Once again for each size of region we use two different intensity maps built on two different days. Most of the time the error is less than three regions, and big jumps due to the specific characteristics of the localizer are rare (they can occur when two regions have similar fingerprints of signal strengths, which is more probable when fewer access points are detected).

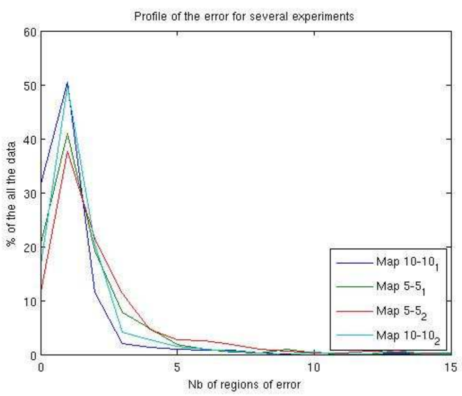

Figure 9: Distribution of the error for four different experiments

To show the reliability of an intensity map over time, we compare the results obtained by driving around the site on three different days using the same map. Table 2 presents those results for two 10m by 10m maps.

The influence of the size of the regions composing the intensity map is complex. As stated in the previous section, using regions smaller than 5m by 5m is not feasible due to the time to calculate the mean of the signal strength. Using regions bigger than 10m by 10m induces the problem of the signal strength varying too much in one region for the value stored in the map to be reliable. Although it appears the accuracy is better for the 10m by 10m versus the 5m by 5m regions, the difference is sufficient to highlight a real advantage of using such a size. However the accuracy recorded during those experiments is only a global mean of the error recorded at each position, and does not give any information on its distribution over the site (the localizer being more accurate in some zones than in others). It is an indication of what range of accuracy can be expected in a typical industrial worksite. In the next section we analyze the evolution of this accuracy over the different types of environment (zones) that compose the site.

Table 2: Results for three different runs made using the same intensity map.

	Run 1	Run 2	Run3
Map 1 Error (m)	22.9	16.9	24.2
Map 2 Error (m)	14.7	16.8	25.3

6.3 Analysis of the Different Zones

In the experiments, the site was considered of comprising of four distinct zones, each having specific physical characteristics that can affect the behavior of WiFi signals (those four zones are represented on Figure 8). For each zone we record the number of regions of error of the localizer while driving around with the tractor. The experiment is conducted using two different maps built on two different days, to verify the reliability of the maps over time. Figure 10 shows the distribution of the error for each zone for regions' size of 10m by 10m. It can be seen that for both experiments the results have the same profile, which validate the fact that the influence of the type of zone on the accuracy of the localizer is reliable.

Table 3 shows the results for each zone. Three different runs have been made for each map, the results being the mean of the error for each experiment.

There is a link between the accuracy of the localizer and the type of zone which the vehicle is in. Several conclusions can be drawn from these results:

- the best results are obtained in the road surrounded with buildings (Road 1), which shows that the buildings and other static obstacles do not influence the behavior of the localizer, because the intensity map was built in the same environment and so implicitly takes their existence into account.
- the worst results are obtained in the Road 3, in open area not well covered by access points (less than 4 access points available over this road). Therefore the number of access points in range has an influence on the accu-

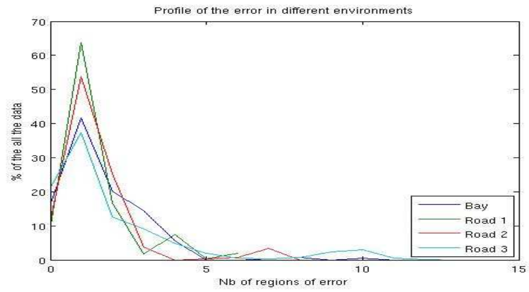

(a) Map 1.

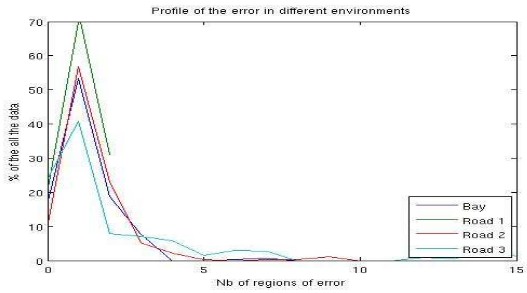

(b) Map 2.

Figure 10: Distribution of the error using 2 different maps 10m by 10m

racy of the localizer: the more access points that are detected in a region, the more complex the fingerprint in this region can be, so the region is differentiated from the others.

- the results for the two maps do not refute each other, which means that the maps are reliable enough to be used over different days.

7 Conclusion and Future Work

We have built a localization system based only on WiFi, requiring little in the way of hardware, installation and calibration effort. It has been demonstrated in a typical industrial worksite, using the existing WiFi access points of the site. The accuracy achieved is typically 25 meters, depending on the environment, the number of access points detected and their distribution. This accuracy is also closely linked to the speed of the vehicles using the localizer: the slower the vehicle is travelling, the more information is received from the access points. Even if such an accuracy is not good enough for typical mobile robotic application, our system can deal with the kidnapped robot problem, and provide a valuable reference for bootstrapping another localizer using different types of sensors.

This localizer presents a good preliminary system but can be enhanced by using higher level signal preprocessing, trajectory prediction, or other methods to improve its accuracy and performance. Among these methods we believe that some work can be done to improve the use of the intensity map, for example by using auto−generated regions of variable size depending on the local variations of the signal strength, instead of regions of constant size.

Table 3: Error obtained in the different zones of the site, using two different maps.

	10m∗10m Day 1	10m∗10m Day 2
Error (m) Bay Compound	22.2	19.1
Error (m) Road 1	15.6	12.0
Error (m) Road 2	22.5	26.0
Error (m) Road 3	32.6	24.1

Acknowledgments

This work was funded by the CSIRO ICT Centre under the ROVER and Dependable Field Robotics projects. The authors would like to thank the Autonomous Systems Laboratory team for their support of this work. Special thanks go to Polly Alexander, Stephen Brosnan, Peter Corke, Leslie Overs, Cedric Pradalier, Jonathan Roberts, Kane Usher, John Whitham and Graeme Winstanley who all contributed to the development of our experimental autonomous tractor or provide advice and assistance.

References

[Arulampalam *et al.*, 2002] S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, February 2002.

[Bahl and Padmanabhan, 1999] P. Bahl and V. N. Padmanabhan. User location and tracking in an in-building radio network. Technical report msr-tr-99-12, Microsoft Research, 1999.

[Bahl *et al.*, 2000] P. Bahl, V. Padmanabham, A. Balachandran, Y. Wang, and W. Russell. Radar: Location-aware services in an in-building environment. *INFOCOM*, pages 775–784, 2000.

[Cheng *et al.*, 2005] Y. Cheng, Y. Chawathe, A. LaMarca, and J. Krumm. Accuracy characterization for metropolitan-scale wi-fi localization. In *Proceedings of MobiSys '05*, Seattle, June 2005.

[Corke *et al.*, 2004] P. Corke, P. Sikka, J. Roberts, and E. Duff. Ddx: A distributed software architecture for robotic systems. In *Proc. Australian Conf. Robotics and Automation*, Canberra, December 2004.

[Dellaert *et al.*, 1999] 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.

[Duff *et al.*, 2006] E. Duff, K. Usher, and P. Ridley. Swing loader traffic control. Internal CSIRO Technical Report 06/192, CSIRO, Jan 2006.

[Fox *et al.*, 2003] D. Fox, J. Hightower, L. Liao, D. Schultz, and G. Borriello. Bayesian filters for location estimation. *IEEE Pervasive Computing Magazine 2(3)*, pages 24–33, July-September 2003.

[Gustafsson *et al.*, 2002] F. Gustafsson, F. Gunnarsson, N. Bergman, U. Forssell, J. Jansson, R. Karlsson, and P. Nordlund. Particle filters for positioning, navigation and tracking. *IEEE Transactions on Signal Processing*, pages 50:425–435, February 2002.

[Haeberlen *et al.*, 2004] A. Haeberlen, E. Flannery, A. M. Ladd, A. Rudys, D. S. Wallach, and L. E. Kavraki. Practical robust localization over large-scale 802.11 wireless networks. In *Proceedings of ACM MobiCom*, Philadelphia, September 2004.

[Hightower and Borriello, 2001] J. Hightower and G. Borriello. Location systems for ubiquitous computing. *IEEE Computer*, pages 34(8):57–66, August 2001.

[Hightower and Borriello, 2004] J. Hightower and G. Borriello. Particle filters for location estimation in ubiquitous computing: A case study. In *In Proceedings of International Conference on Ubiquitous Computing (Ubicomp)*, 2004.

[Howard *et al.*, 2003] A. Howard, S. Siddiqi, and G. S. Sukhatme. An experimental study of localization using wireless ethernet. In *The 4th International Conference on Field and Service Robotics*, July 2003.

[Kismet, 2006] Kismet. www.kismetwireless.net, April 2006.

[Kotz *et al.*, 2003] D. Kotz, C. Newport, and C. Elliott. The mistaken axioms of wireless-networks research. Technical Report TR2003-467, Dept. of Computer Science, Dartmouth College, July 2003.

[Krumm and Horvitz, 2004] J. Krumm and E. Horvitz. Locadio: Inferring motion and location from wi-fi signal strengths. In *In Proceedings of International Conference on Mobile and Ubiquitous Systems: Networking and Services (MobiQuitous'04)*, 2004.

[Ladd *et al.*, 2004] A. M. Ladd, K. E. Bekris, A. P. Rudys, D. S. Wallach, and L. E. Kavraki. Using wireless ethernet for indoor localization. *IEEE Transactions on Robotics and Automation*, June 2004. 555-559.

[LaMarca *et al.*, 2005] A. LaMarca, Y. Chawathe, S. Consolvo, J. Hightower, I. Smith, J. Scott, T. Sohn, J. Howard, J. Hugues, F. Potter, J. Tabert, P. Powledge, G. Borriello, and B. Schilit. Place lab : Device positioning using radio beacons in the wild. In *International Conference on Pervasive Computing*, 2005.

[Letchner *et al.*, 2005] J. Letchner, D. Fox, and A. LaMarca. Large-scale localization from wireless signal strength. In *Proceedings of the National Conference on Artificial Intelligence*, 2005.

[Roos *et al.*, 2002] T. Roos, P. Myllymki, P. Misikangas H. Tirri, and J. Sievnen. A probabilistic approach to wlan user location estimation. *International Journal of Wireless Information Networks*, 9(3), July 2002.

[Serrano *et al.*, 2004] O. Serrano, J. M. Canas, V. Matell, and L. Rodero. Robot localization using wifi signal without intensity map. In *WAF*, March 2004.

[Thrun *et al.*, 2005] S. Thrun, W. Burgard, and D. Fox. Probabilistic robotics. 2005.