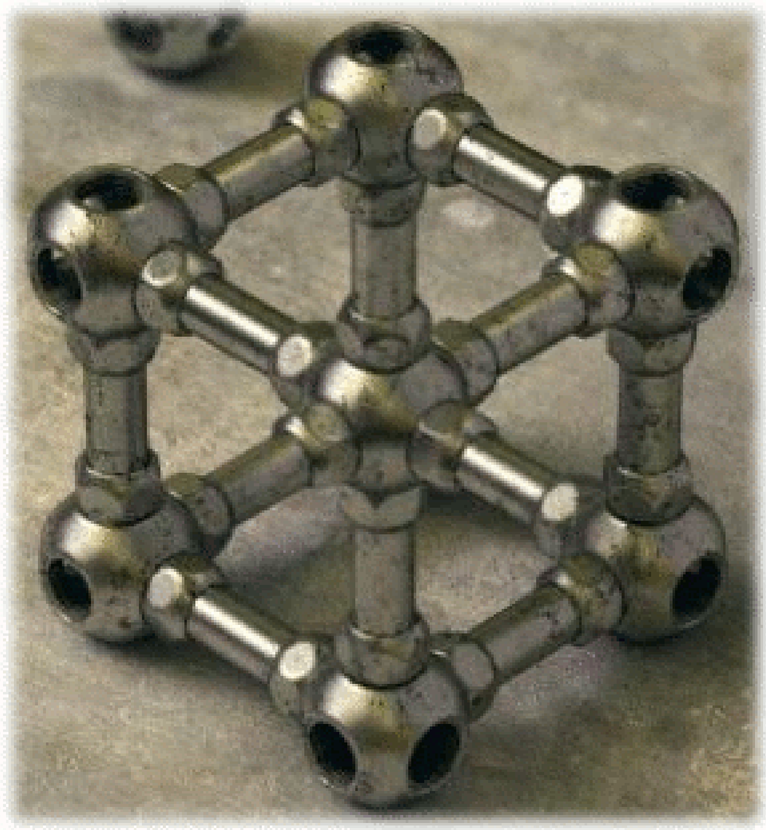

WBLS: a Signal Presence-based Wi-Fi Localisation System for Mobile Devices in Smart Environments [1]

Author(s):

André Iasi Moura
Carlos Henrique Costa Ribeiro
Anna Helena Reali Costa

[1]This work was supported by Fapesp Project LogProb, grant 2008/03995-5, São Paulo, Brazil.

Submitted for a Special Edition

Intelligent Agents and Services for Smart Environments
Special Edition Guest Editors: F. S. Corra da Silva and S. Bandini

WBLS: a Signal Presence-based Wi-Fi Localisation System for Mobile Devices in Smart Environments

André Iasi Moura
Universidade de São Paulo, PCS/EPUSP
Av. Prof. Luciano Gualberto, trav.3, n.158
05508-900 São Paulo SP Brazil
andre.moura@poli.usp.br

Carlos Henrique Costa Ribeiro
Instituto Tecnológico de Aeronáutica
Divisão de Ciência da Computação
São José dos Campos - SP, Brazil
carlos@ita.br

Anna Helena Reali Costa*
Department of Computer Engineering, Universidade de São Paulo
Av. Prof. Luciano Gualberto, trav.3, n.158
05508-900 São Paulo SP Brazil
anna.reali@poli.usp.br

Abstract

The proliferation of mobile computing devices and wireless networks has set the stage for the development of smart environments rich in computing and communication capabilities, yet gracefully integrated with human users. This paradigm has fostered a growing interest in localisation-based systems and services for portable devices, especially in indoor environments. However, designing indoor localisation systems with increasing estimation capabilities and decreasing cost installation is a challenge. An interesting approach to reach such requirements consists in using the wireless local area network (WLan) infrastructure that is already installed in many places. Most reported WLan localisation approaches use a map of received signal strength and signal presence frequency collected from multiple channels at different physical localisations in the environment, which can be very noisy. This work proposes a new localisation system, WBLS (Wireless Based Localisation System), that considers the unreliability of information on signal presence frequency in the estimation process, in an attempt to eliminate its associated noise. Experiments considering mobile agents carrying devices and moving at human walk speeds show that the most important feature of WBLS is a robustness to access points shutdowns that may happen without any warning in an environment where there is little control over the infrastructure.

1 Introduction

In the last years, the increase in the use of high performance portable devices such as laptop computers and personal digital assistants (PDAs) has been remarkable. Such devices are now

*Corresponding author: phone +55 11 3091 5397, fax +55 11 3091 5294

part of our daily lives, a new condition that brings new technological possibilities and challenges, such as the development of context-dependent applications which consider situational aspects of users and devices, their actions and past states.

Context is in fact a critical aspect in many ubiquitous computing applications. Such applications may infer the activities and goals pursued by the user, and when executed in a mobile device, a potentially important information for the inference processes is the localisation of the device in the environment. Context-oriented applications that use information about device localisation have been presented in the literature. [23] and [1] report on school and museum portable device guides which can provide users with information about near objects, suggestions of places to visit depending on the region the user currently is, and even trajectory plans for reaching new parts of the environment. Other interesting examples are the projects for smart houses [7]. In such projects, all the residential devices operate coordinately, possibly via a central management system that analyses the behaviour and actions of house users, automatically turning lamps on and off, setting air conditioning, operating access restriction policies and providing residents with information that can help them in using appliances with maximum efficiency and comfort. Device localisation information can be used in those cases to provide data about user habits and context information for the applications when they communicate with the central management system. Context-oriented applications thus demand reliable localisation processes, but it must be noticed that many applications for routing, search, assistance and resource allocalisation – even when not context-based – may be benefited by or even require localisation capabilities.

As far as localisation systems are concerned, the most popular approaches are those based on GPS (*Global Positioning System*). GPS is a georeferenced localisation system that offers an estimation error of no more than 5 meters in open areas, and with a technological improvement based on a reference state (DGPS – Differential GPS) it can produce estimation errors of no more than a few centimetres [24]. However, it is well known that both GPS and DGPS, for being based on triangulation of signals sent by satellites, are prone to failure in indoor or heavily occupied tall building spaces. This is an important pitfall, as many ubiquitous computing applications such as those presented in the last paragraph are actually for indoor use.

Cell phone companies use localisation systems based on the signals received in the antennas from the cell phones. Such systems in fact can be used to locate the devices, but – as for the case of GPS systems – indoor localisation can be critically affected, mainly due to reflections and other physical interactions with obstacles and objects in the environment [19]. Additionally, there is the common problem of too few antennas at a sufficiently close range from the device, a situation that is caused by signal power control policies for avoiding signal interference among cell phones located in the same cell [4].

Some high precision indoor localisation systems have been developed. Among these, we can mention systems based on ultrasound [12] and infrared [15] signal detection, both with estimation errors of a few centimetres. Unfortunately, a ubiquitous localisation system based on such technologies would require an extremely large number of corresponding sensors, a requirement that is still too expensive.

Despite the large number of papers on Wi-Fi localisation, only a few ones do consider performance analysis for moving devices. Such situation – certainly very common in various applications [1, 23] – induce particular characteristics that imply much harder conditions for localisation systems. In fact, comparative results of localisation systems for stationary and moving devices have shown considerably worse performance in the latter [3, 17], making it clear that mobility situations are critical and do require a deeper analysis.

In this paper, we are interested in inferring the localisation of agents moving in indoor environments, over an existing WLan infrastructure over which the localisation system has very restricted control. Furthermore, we assume that the WLan is based on the IEEE 802.11b standard, commonly called Wi-Fi. Current localisation solutions based on existing indoor Wi-Fi networks generally use the Received Signal Strength Indication (RSSI) received from Wi-Fi access points (APs) to localise a user carrying a portable device. Both Wi-Fi APs and mobile Wi-Fi clients are becoming more ubiquitous. Many buildings such as shopping malls, schools and museums, usually have a good number of APs, so that we can connect to the Wi-Fi network and use this existing infrastructure to infer the agent localisation. The advantage of this approach is then to reduce, or even eliminate, the need for installing or altering existing devices for producing localisation estimates. However, a critical problem for a localisation system in such environments is the fact that operation must be carried out on extremely noisy signal conditions. Wi-Fi signals in indoor, non-structured environments are prone to interference and reflections from multiple objects and obstacles, and even from other mobile devices and agents. Characteristically, large variations and unpredictability in RSSIs are expected, and in fact, experimental evidence largely reported in the literature show that some APs can produce signals intermittently. As many localisation systems use the signal presence frequency as important information for localisation estimation, this is certainly an important source of error.

The localisation system we propose in this paper is called WBLS (*Wireless Based Location System*). It is based on a probabilistic modelling using Hidden Markov Models, but it does not consider signal presence frequency for localisation estimation. Moreover, it should be stressed that a localisation system operating on non-structured indoor environments must be robust to unexpected structural changes such as transmission failure, AP shutting-off and so forth. For not considering signal presence frequency for producing estimates, WBLS is naturally resilient to such problems.

The paper is organized as follows. Section 2 describes the Wi-Fi localisation problem and highlights the main challenges associated to mobile device localisation based on signal detection. Section 3 then presents a Markov localisation system, which will serve as a basis for the localisation system WBLS (*Wireless Based Location System*) proposed in Section 4. WBLS is shown to be derived from an analysis of system noise characteristics, together with some performance-motivated modifications on the model presented in Section 3. Section 5 presents the experimental results for WBLS. Finally, Section 6 presents the conclusions of the paper and points to some future works that are worth pursuing.

2 Localisation in Wi-Fi networks

Wi-Fi localisation is based on radio signal detection and analysis of attributes such signals hold: radio signal strength, difference in time-of-arrival, and angle of arrival. It is acknowledged that the more precise the information from these attributes is received, the more accurately localisation is estimated.

For indoor environments, ideal conditions for measuring such attributes almost never hold. In fact, such environments induce radio wave propagation through multiple paths, almost always as NLoS (*Non Line of Sight*) propagation, therefore usually not including direct pair wise antenna transmission and reception [20, 13]. NLoS reflections, refractions, diffractions, and absorptions produce both cross-interference and weakening of signal power, making it very difficult to reliably interpret reading of signal information from the devices.

Moreover, other agents and objects (including interfering radio signal transmitting devices such as microwave ovens [25]), and non-stationary conditions (mobility of agents and transmission devices, variations in temperature and humidity) usually found in non-structured environments contribute to make Wi-Fi localisation even more difficult.

An important aspect to be taken into account when designing Wi-Fi localisation systems is that usually there is already an infra-structure for Wi-Fi localisation in the environment. But at the same time that this brings opportunities for implementation cost reduction, also sets forth the problem of having to produce accurate localisation estimates using devices over which there is no control whatsoever. As the devices are operated by different and independent agents, phenomena such as unexpected moves, turning off and even permanent disconnection are common and may affect localisation processes. In addition, users can usually install new APs quite easily (increasing the probability of co-channel interference, unless there is some common policy supervised by a central manager), and therefore even the assumedly fixed structure of the network is not necessarily stable. Naturally, the lack of control can also involve the physical structure of the environment: installation of new walls, removal or positioning of objects can radically affect the radio propagation characteristics in the environment.

The maximum transmission range of APs is another relevant characteristic of Wi-Fi systems. Usually, the output power of standard APs (one inch distance from the device) is around 20 dBm, which implies -10 dBm at one meter distance. The signal can be detected if its power is not less than -100 dBm (this may depend on the device, but it is usually a good rule of thumb, as described in [18]). Thus, the signal varies 90 dBm from a one meter distance from the emitter antenna to the furthest point where it can still be detected by a receiver antenna. One of the basic models for characterizing the relationship between signal power (in dBm) and emitter-receiver distance in indoor environments suggests a linear relationship between those parameters [11]. Data collected in tests conducted in our environment estimated a distance power loss coefficient of 2.4 dBm/m, in reasonable agreement with results by [8] (2.0 dBm/m). As a result, one can argue that the range for an AP must be between 40m and 60m. As the relationship between power and distance is extremely dependent on the environment configuration, one can expect that the distance power loss coefficient can suffer huge variation among different environments. As a matter of fact, for obstacle-free environments the range of an AP can be as large as 10km [10]. In practice, however, Wi-Fi networks with ranges larger than 100m are not common.

Since most successful Wi-Fi localisation systems are based on RSSI, it is important to highlight some of its characteristics. Many environmental influences can cause the received signal intensity to vary over time, and such time-varying effects can have severe implications on the localisation accuracy. In order to illustrate some of this effects, we collected a series of measurements of the radio signal transmitted from an AP and measured by a laptop in a fixed position but in four different directions (see figure 1-a) and also at a fixed localisation in a building over a period of approximately two hours (see figure 1-b). As we can see in figure 1-a, variations of up to 10 dBm occur in the signal intensities averaged over each direction, affecting *circa* 3 to 5 meters the agent localisation estimates. In figure 1-b there are signal variations of up to 15 dBm, which can cause variations of up to 7 meters in the localisation estimates. Not only many changes occur in the environment which affect the observed signal intensity, but also agent moves in the environment further complicates the task of maintaining an accurate position estimate.

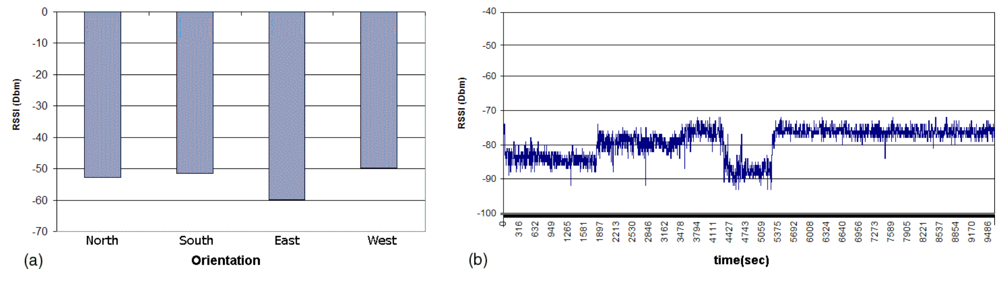

Figure 1: Examples of signal intensity variations for an AP signal measured from a laptop in a building: (a) RSSI values averaged over four different directions at a fixed position, and (b) RSSI values over a 2-hour period in a fixed localisation.

Most Wi-Fi localisation systems are based on discrete environmental models that use patterns of RSSI variation as reference for the localisation system. Such systems, as reported in the literature, are compounded by two phases: training and run-time [3, 8, 9, 14, 16, 22, 26]. During the training phase, a database is built measuring several values of RSSI and associated signal presence frequencies received from different APs in each reference point (RP) in the environment. This database is called the RSSI map of the environment. In the run-time phase, the RSSI map is then used as a model to infer the localisation based on the comparison of the RSSI value observed with the model value. The way this is done varies in different works, but the most successful localisation systems use probabilistic methods in the localisation inference.

3 A basic system for Wi-Fi localisation

The localisation system we propose in this paper is based on a probabilistic modelling using Hidden Markov Models (HMM). The localisation problem consists in determining the agent state (or localisation) given one or more observations, and its solution is also based on a training phase and a run-time phase.

3.1 Training phase

In the training phase, we define a discrete grid composed by a set of area cells that form a partition of the locomotion surface of the agent in the environment, with each cell having a point associated with it, called measurement point (MP). The number of cells must be large enough to cover all regions of interest in the environment, and the area of each cell should respect a size trade-off to satisfy both the desired and possible localisation estimation accuracies. The discrete grid can be either uniformly distributed (defining a set of regularly spaced MPs), or could follow topological aspects, associating differently sized cells to rooms of interest.

Let $A = \{p_a | a = 1, 2, ..., \mathcal{A}\}$ be the set of APs, where $\mathcal{A}$ represents the total number of APs in the environment. Observations collected at MPs are vectors $O_t = [o_{a,t} | a = 1, 2, ..., \mathcal{A}]$, where $o_{a,t}$ is the RSSI value received from $p_a \in A$ in time t. When an AP signal intensity is bellow the minimal receiver sensitivity $\mathcal{R}min$, then $o_{a,t}$ receives a value meaning that AP p_a could not be detected (here, we use $o_{a,t} = \mathcal{R}min - 1$).

Measurements can be done in three different ways: (i) collecting RSSI values over some time period in a fixed localisation (normally, for all MPs, a scan is done at each position for each orientation)[3, 16]; (ii) collecting RSSI values over some time period in a fixed position but varying slowly the height and the orientation of the receiver [14]; and (iii) walking around slowly for some time in each cell, in order to cover the entire cell [9]. The first method has the disadvantage of having a RSSI map four times bigger than the second method, and the third is the one that probably best represents the signal characteristics in the entire cell. The time period spent in the measurements can be either a pre-defined fixed value [9, 14] or an adaptive time period according to some procedure or heuristic [16].

The data sampled is then stored in a RSSI map. The effort needed to build the RSSI map depends on the number of active channels present in the environment and on the number of MPs (cells) considered. This means that a lot of effort is often requested by the training phase when the environment is made up by many APs and small localisation errors are required.

Formally, a RSSI map M is composed by a set of $\mathcal{I}$ records, $M = \{R_i | i = 1, 2, ..., \mathcal{I}\}$, where $\mathcal{I}$ is the number of MPs. Each record is given by pairs $R_i = (e_i, C_i)$, where e_i is the identification of the i-th MP – usually given by the MP localisation coordinates (x_i, y_i) in the environment metric map – and C_i is an n-dimensional vector, $C_i = [Rssi_{i,a} | a = 1, 2, ..., \mathcal{A}]$, where $Rssi_{i,a}$ contains the RSSI values sampled at the position MP i for AP a. $Rssi_{i,a}$ should represent a distribution of signal strengths for each AP a, and we can use either histograms or Gaussian-fit curves to store these distributions. In the first, for each AP a and each record i in the RSSI map, $Rssi_{i,a} = [b_{i,a,n} | n = 1, 2, ..., \mathcal{N}]$, where $b_{i,a,n}$ represents the frequency count of value n for AP a at MP i, and $\mathcal{N}$ is the discrete values number considered in the histogram. The use of histograms has the advantage of representing any signal strength distribution but it demands more storage space. On the other hand, fitting the data to a Gaussian only requires storing two numbers for each distribution: its mean and standard deviation [9, 26], so that $Rssi_{i,a} = (\mu_{i,a}, \sigma_{i,a}, f_{i,a})$, where $\mu_{i,a}$ is the mean and $\sigma_{i,a}$ is the standard deviation, both calculated only when the signal of AP a is present in the site, and $f_{i,a}$ represents the frequency with which AP a was not detected in cell i.

3.2 Run-time phase

The basic system we describe is based on the work described in [9]. We model a moving agent trying to track its localisation as a Hidden Markov Model (HMM).

Formally, an HMM is:

- a set of states $E = \{e_i = (x_i, y_i) | i = 1, 2..., \mathcal{I}\}$, where each state e_i is given by the coordinates (x_i, y_i) of the corresponding MP i that represents the environment cell;
- a set of observations $O_t = [o_{a,t} | a = 1, 2, ..., \mathcal{A}]$, where $o_{a,t}$ is the signal strength captured from AP a in time t, discretized in an interval ranging from $\mathcal{R}min - 1$ (when the AP is not detected) to $\mathcal{R}max$ (the highest sensitivity value of the detector);
- the probability of observing $\mathbf{o}_a$ (with signal strength $rssi$ captured from AP a) while at state $\mathbf{q} = e_i$, i.e.,

$$P(O_t | \mathbf{q}_t = e_i) = \prod_{a=1}^{\mathcal{A}} P(\mathbf{o}_a = o_{a,t} | \mathbf{q} = e_i), \tag{1}$$

  where $\mathbf{q}_t$ is a random variable for the localisation state at time t;
- the probability of transiting from state e_j in t to e_i in $t + 1$, given by

$$P_{trans}(e_j, e_i) = P(\mathbf{q}_{t+1} = e_i | \mathbf{q}_t = e_j) \tag{2}$$

  where

$$\sum_{i=1}^{\mathcal{I}} P_{trans}(e_j, e_i) = 1; \tag{3}$$

- an initial probability distribution over states $P(\mathbf{q}_0)$, where $P(\mathbf{q}_0)$ is an uniform distribution, when no previous estimation of the agent initial localisation (state) is known; or $P(\mathbf{q}_0 = e_k) = 1$ and $P(\mathbf{q}_0 = e_i) = 0$, $i \neq k$ and $i = 1, 2, ..., k-1, k+1, ..., \mathcal{I}$, when there is a perfect knowledge about the agent initial localisation; or any distribution between these two limits.

When the Gaussian-fit method is used to store $Rssi_{i,a}$ in the training phase, we can compute the observation probability as follows:

$$P(\mathbf{o}_a = rssi | \mathbf{q} = e_i) = \eta_{a,i}[G_{a,i}(rssi) + \lambda], \tag{4}$$

where

$$G_{a,i}(rssi) = \int_{rssi-0.5}^{rssi+0.5} \frac{exp[-(x - \mu_{a,i})^2 / 2\sigma_{a,i}^2]}{\sigma_{a,i}\sqrt{2\pi}} dx. \tag{5}$$

and $G_{a,i}$ represents a discretization of the Gaussian curve with mean $\mu_{a,i}$ and standard deviation $\sigma_{a,i}$, $\eta_{a,i}$ is a normaliser that ensures

$$\sum_{rssi=\mathcal{R}min-1}^{\mathcal{R}max} P(\mathbf{o}_a = rssi | \mathbf{q} = e_i) = 1, \tag{6}$$

and λ is a value that ensures $P(\mathbf{o}_a = rssi | \mathbf{q} = e_i) > 0$. The probability of not detecting signal from AP a, i.e., $P(\mathbf{o}_{a=\mathcal{R}min-1} | \mathbf{q} = e_i)$, can be computed from $f_{a,i}$.

The transition probability $P_{trans}(e_j, e_i)$ can be computed as a function proportional to the inverse of the (direct) Euclidean distance between every pair $(e_j, e_i), i \neq j$, and by defining a fixed value for the reflexive transition ($i = j$). An alternative is to compute the minimal path from e_j to e_i (distance in the environment) and use this measurement to define the value of $P_{trans}(e_j, e_i)$. A topological map can reflect the physical restrictions on the environment, and it is very useful in this respect since it can give support to set $P_{trans}(e_j, e_i) = 0$ if there is no path from e_j to e_i in the environment.

Given the HMM model, we now present the algorithm *forward procedure* [21] used to update the state estimate between each set of observations, considering constraints on how the agent can move from state to state, agent walking speed and topological constraints. Suppose that at time t the state estimate is $P(\mathbf{q}_t = e_j|O_t)$ [1]. Between time t and $t+1$ the agent moves to an unknown state e_i. At time $t+1$ observations $o_1, ..., o_k$ are received. Then, the transition module computes

$$\sum_{j=1}^{\mathcal{I}} P(\mathbf{q}_t = e_j|O_t) P_{trans}(e_j, e_i), \tag{7}$$

which represents the probability sum of transiting from each state e_j in t to the state e_i in $t+1$, and the observation module computes $P(O_{t+1}|\mathbf{q}_{t+1} = e_i)$ based on information stored in the RSSI map and in the observation vector O_t (see Eq.1, 4, 5 and 6). Figure 2 illustrates the localisation algorithm.

Results of both modules are multiplied and normalized so that

$$P(\mathbf{q}_{t+1} = e_i|O_{t+1}) = \eta\, P(O_{t+1}|\mathbf{q}_{t+1} = e_i) \sum_{j=1}^{\mathcal{I}} P(\mathbf{q}_t = e_j|O_t) P_{trans}(e_j, e_i), \tag{8}$$

and for each HMM state we then have $\mathbf{P}(\mathbf{q}_{t+1}|O_{t+1})$.

Finally, for generating a localisation information for the user, an estimation module can receive as input the calculated distribution and then generate a localisation estimation, which can be either the state with the largest probability or (more commonly) the expected value of state [5].

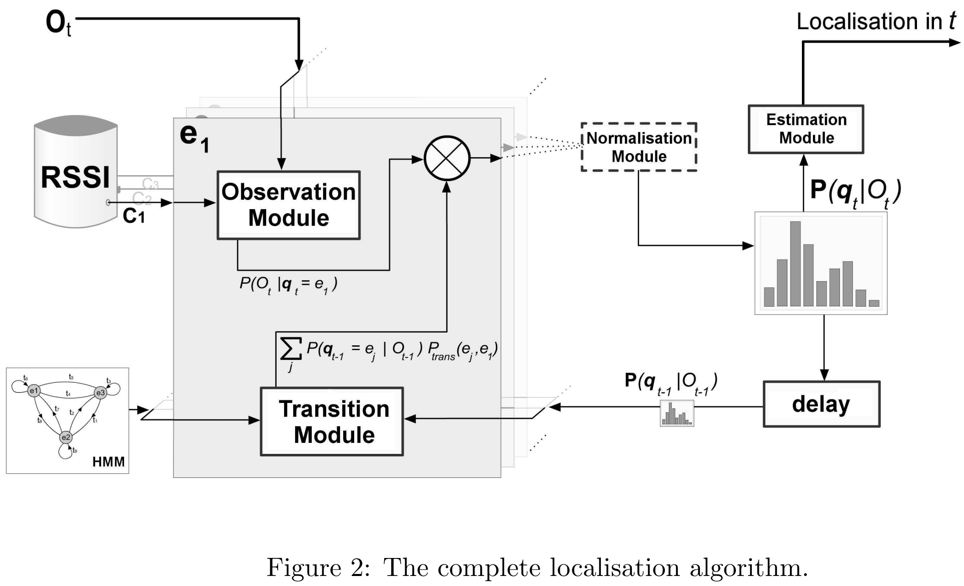

Figure 2: The complete localisation algorithm.

[1] In $t = 0$, $P(\mathbf{q}_t = e_j|O_t) = P(\mathbf{q}_0)$

3.3 Specification of the basic Wi-Fi localisation system

Once we defined the localisation system, we must specify a) how the measurements for building up the RSSI map will be carried out; b) how to represent the RSSI distribution; c) how to calculate the transition probabilities; and d) how to define the initial probability distribution.

Measurements of RSSI were carried out in each cell, varying both height and orientation of the measurement device. For representing the RSSI distribution, two comparative studies can be found in the literature [9, 22], with controversial reports on which representation produces better results. We chose a Gaussian-based representation both for the basic model and for the proposed one presented in the next section, in a similar fashion to [9], which served as a basis for our proposal.

Regarding the calculation of the transition probabilities, an approach based on topological maps and expected behaviour of the mobile agents was adopted. For the sake of illustration, consider that the environment is composed by square cells with 1.5m side length, and that observations were taken every 0.2s [2] In the best case, a person would have to walk at a minimum 6m/s to traverse two contiguous cells (3m distance) between two consecutive observations. It is reasonable to assume that a human agent carrying a portable device would be running at speeds that are much lower than this. Thus, for the transition probability model we defined a zero probability for the transitions between non-neighbouring vertices (each pair of neighbouring vertices corresponds to a topological cell neighbourhood), and a nonzero and single probability value for every transition from each node to all its neighbours[3].

We must also define values for the reflexive transition probabilities, *i.e.*, probabilities of standing in the same cell in consecutive observations. This was defined empirically through observations of the customary behaviour of a large number of agents walking in the environment. It was realized that, on average, each agent remained 15 iterations in the same cell before moving to a new one, indicating a 0.95 probability of a reflexive transition. We could have used the agent customary behaviour related in each specific cell in the environment, so that more representative reflexive transition probabilities could be produced, but for the sake of simplicity we decided to use an average value for all cells.

Finally, we assumed an uniform distribution for the initial state, representing the uncertainty about the initial localisation of the device.

To summarize, the main considerations and specifications both for the basic model and for the one proposed in the next section are: a) it is based on a discrete model for cells with a maximum dimension of 1.5m; b) it adopts a probabilistic approach for estimating localisation; c) it builds up a RSSI map during a training phase, with variations of the capture in the cells (height and orientation of the agent) and storage of information as Gaussian distributions and frequency of signal absence; d) it considers a mobile device that is conducted by a walking agent (human), and uses a topological map of the environment and physical limitations of this agent to define state transition probabilities; e) it defines (as a function of experiments whose results depend on the grid cell size and agent customary behaviour in the environment) a reflexive transition probability as 0.95; f) it considers an uniform probability distribution for the initial state.

3.4 Analysis of the basic Wi-Fi localisation system

Our aim in the analysis of the basic Wi-Fi localisation system was to grasp a better understanding of some problems faced by Wi-Fi localisation systems in general, mainly due to the noisy nature of the detected signals. We considered three main aspects, related respectively to the estimation, transition and observation modules.

[2] Empirical results show that a 0.2s sampling rate is actually a conservative rule-of-thumb for the proposed algorithms and hardware used in the experimental results reported herein.

[3] Naturally, the topological maps and the transition probability model must be modified for the case in which the mobile agents move at higher speeds (*e.g.*, cars).

Regarding the estimation module, there are indications that a measure of the expected value of the device localisation (*i.e.*, a sum of the possible discrete cell localisations weighted by the respective estimated probabilities) produces better results than simply choosing the cell localisation that holds the largest associated probability of occurrence [26]. However, our experimental results show that it is not always the case that highest probability peaks are concentrated in small regions (see Figure 3-a), and therefore simply obtaining the expected localisation by a weighted sum over the whole cell space can produce strange results, such as unlikely localisations resulting from a weighted average of two distant and almost equally probable localisations. We thus adopted the following heuristics, which aims at combining the reportedly adequate expected value approach and a mechanism for avoiding unlikely localisations: calculating the expected value using only a cell neighbourhood around the localisation with highest probability.

The transition module can be improved by using knowledge acquired from experience, reflecting in the transition probabilities not only topological and device dynamical restrictions, but also more specific patterns such as preferred trajectories performed by the agents [6, 2].

As the signals received from the APs are very noisy, the observation module calculates conditional probabilities of observations given states that can vary significantly in different time steps, resulting in situations such as the one depicted in Figure 3, which shows the distribution $P(O_t|\mathbf{q}_t = e_i)$ on the state space for two consecutive time steps. It can be noticed that the most likely observation in (a) has a very low probability in (b), and in fact, the distance between the states that produce the maximum values of $P(O_t|\mathbf{q}_t = e_i)$ in (a) and (b) is nearly 10m.

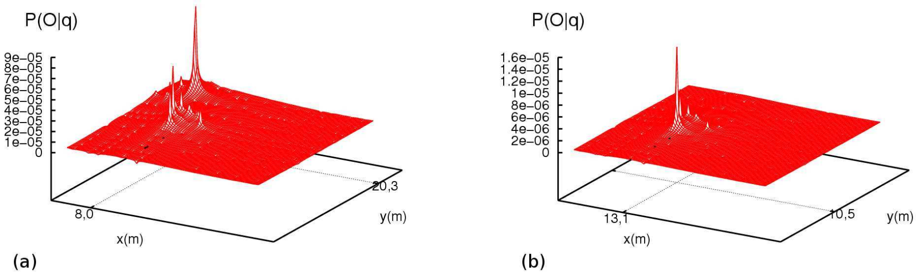

Figure 3: Distribution of observation probabilities in two consecutive time steps. Vertical axis represents $P(O|\mathbf{q})$ and the horizontal axes define environment coordinates.

The large variations on the conditional probabilities induced by noise led us to analyse more closely the relationship between signal power and frequency of detected signal presence. Initially, we considered the hypothesis that low power signals are difficult to detect. Experiments showed that, for signals with power in the range [-96 dBm, -90 dBm], the observed frequency of detected signal presence was less than 15%, corroborating this hypothesis[4]. This was actually expected, but surprisingly, the dual condition for high power signals was not as obvious. Repeating the experiments for observations with signal strength higher than 60 dBm led to 35% of situations where the frequency of detected signal presence was below 80%, and 17% of situations with a frequency below 50%. Table 1 summarizes the results regarding the frequency of signal presence of an AP in six points located in a room, with measurements taken at intervals of less than 10 minutes. Notice that the RSSI values do not vary much (min 63 dBm, max 70 dBm) for points that are spatially close, yet the presence frequency does have a large variation (min 55%, max 90%).

Such results indicate that information on frequency of signal presence is very noisy and probably unreliable, suggesting a possibly harmful effect from the use of this information for

[4]The device sensitivity was -96 dBm

Table 1: Signal presence frequency and average RSSI in a single room.

MPs coordinates		Capture information	
X (m)	Y (m)	Freq. (%)	RSSI mean (dBm)
12,3	15,3	71	67
10,0	15,3	55	65
12,3	18,3	90	67
10,0	18,3	82	63
12,3	22,8	67	70
10,0	22,8	89	70

the purpose of localisation estimation. The next section presents a proposal that takes this analysis into account.

4 The proposed system: WBLS

WBLS is a localisation system that considers the unreliability of information on frequency of signal presence in the measurements for the construction of a RSSI map. It basically defines a new form of calculation of the conditional observation probabilities. We detail WBLS in the following subsections.

4.1 WBLS changes in the basic localisation system

Fundamentally, WBLS considers only the probabilities corresponding to the present signals. Thus, for each AP the probability must be conditioned also on the presence of the signal. In what follows, we present the corresponding modifications for the model that considers gaussian distributions for RSSI, but the concept is also valid for histogram-based (or any other) representation scheme.

In WBLS, the conditional probability of $\mathbf{o}_a = rssi$ is calculated as follows:

$$P(\mathbf{o}_a = rssi|\mathbf{q} = e_i, \mathbf{S}_a) = \eta_{a,i}[G_{a,i}(rssi) + \lambda] \tag{9}$$

where $rssi$ is an observed RSSI value, $G_{a,i}$ is the discretization of the gaussian distribution of AP a in state i according to the RSSI map, $\eta_{a,i}$ is a normalization factor and λ is a regularization factor to prevent observations with near zero probabilities. Notice that this equation is nearly identical to equation 4, the only difference being the presence of the condition $\mathbf{S}_a$, which means the presence of the AP signal, *i.e.*, $\mathbf{o}_a \geq \mathcal{R}min$. There is however an additional change regarding the normalization factor, which is calculated considering the restriction:

$$\sum_{rssi=\mathcal{R}min}^{\mathcal{R}max} P(\mathbf{o}_a = rssi|\mathbf{q} = e_i, \mathbf{S}_a) = 1. \tag{10}$$

Here, the RSSI interval $[\mathcal{R}min, \mathcal{R}max]$ differs from the one for the basic system (Equation 6), which also considers $\mathcal{R}min - 1$ (representing the absence of signal). In this way, RSSI calculation is based solely on the observations actually made by the device that is being localised.

4.2 A formal analysis of WBLS

Consider equation 1. We can represent the probabilities that compose the product for determining the conditional probability of an observation [14] as:

$$P(\mathbf{o}_a = o_{a,t}|\mathbf{q} = e_i) = \begin{cases} P(\mathbf{o}_a = o_{a,t}|\mathbf{q} = e_i, \mathbf{S}_a)P(\mathbf{S}_a|\mathbf{q} = e_i) & \text{if } o_{a,t} \geq \mathcal{R}min \\ P(\overline{\mathbf{S}_a}|\mathbf{q} = e_i) & \text{otherwise} \end{cases} \tag{11}$$

where $P(\mathbf{S}_a|\mathbf{q}=e_i)$ is the probability that the signal from AP a is present at state e_i and $P(\overline{\mathbf{S}_a}|\mathbf{q}=e_i)$ is the probability that the signal from AP a is not present at state e_i.

The probability of being in state e_i given observation O_t (Equation 8) can be written as:

$$P(\mathbf{q}_t = e_i|O_t) = \eta P(O_t|\mathbf{q}_t = e_i)T(e_i), \tag{12}$$

where

$$T(e_i) = \sum_{j=1}^{\mathcal{I}} P(\mathbf{q}_{t-1} = e_j|O_{t-1})P_{trans}(e_j, e_i). \tag{13}$$

Equation 12 can be written as

$$P(\mathbf{q}_t = e_i|O_t) = \eta \prod_{a=1}^{\mathcal{A}} P(\mathbf{o}_a = o_{a,t}|\mathbf{q} = e_i)\; T(e_i), \tag{14}$$

where $P(\mathbf{o}_a = o_{a,t}|\mathbf{q}=e_i)$ is determined as in equation 11.

WBLS then calculates $P(\mathbf{q}_t = e_i|O_t)$ in the following way:

$$P(\mathbf{q}_t = e_i|O_t) = \eta' \prod_{a=1}^{\mathcal{A}} \left\{ \begin{array}{ll} P(\mathbf{o}_a = o_{a,t}|\mathbf{q} = e_i, \mathbf{S}_a) & \text{if } o_{a,t} \geq \mathcal{R}min \\ 1 & \text{otherwise} \end{array} \right\} T(e_i), \tag{15}$$

where $\eta' = \eta F(e_i, O_t)$ and

$$F(e_i, O_t) = \prod_{a=1}^{\mathcal{A}} \left\{ \begin{array}{ll} P(\mathbf{S}_a|\mathbf{q} = e_i) & \text{if } o_a \geq \mathcal{R}min \\ P(\overline{\mathbf{S}_a}|\mathbf{q} = e_i) & \text{otherwise} \end{array} \right\}. \tag{16}$$

The reasoning above can be carried out for any state model, and therefore:

$$F(e_i, O_t) = F(e_j, O_t)\ \forall i, j. \tag{17}$$

Equation 17 must be valid for any observation. Consider two possible observations O' and O'', such that all APs signals present in the first are present in the second, and all APs signal absent in the first are absent in the second, with the exception of the AP signal indexed by 1, present in observation O' but absent in O''. We determine $F'(e_i)$ as:

$$F'(e_i, O) = \prod_{a=2}^{\mathcal{A}} \left\{ \begin{array}{ll} P(\mathbf{S}_a|\mathbf{q} = e_i) & \text{if } o_a \geq \mathcal{R}min \\ P(\overline{\mathbf{S}_a}|\mathbf{q} = e_i) & \text{otherwise} \end{array} \right\}, \tag{18}$$

in such a way that $F'(e_i, O') = F'(e_i, O'')\forall i$. It is then possible to write Equation 17 for each of the possibilities as:

$$P(\mathbf{S}_1|\mathbf{q} = e_i)F'(e_i, O') = P(\mathbf{S}_1|\mathbf{q} = e_j)F'(e_j, O'), \tag{19}$$

$$P(\overline{\mathbf{S}_1}|\mathbf{q} = e_i)F'(e_i, O'') = P(\overline{\mathbf{S}_1}|\mathbf{q} = e_j)F'(e_j, O''). \tag{20}$$

Replacing $P(\overline{\mathbf{S}_1}|\mathbf{q} = e_i)$ by $1 - P(\mathbf{S}_1|\mathbf{q} = e_i)$ in Equation 20, it follows that

$$[1 - P(\mathbf{S}_1|\mathbf{q} = e_i)]F'(e_i, O'') = [1 - P(\mathbf{S}_1|\mathbf{q} = e_j)]F'(e_j, O''). \tag{21}$$

Dividing the terms in Equation 21 by the terms in Equation 19, we get

$$\frac{[1 - P(\mathbf{S}_1|\mathbf{q} = e_i)]}{P(\mathbf{S}_1|\mathbf{q} = e_i)} = \frac{[1 - P(\mathbf{S}_1|\mathbf{q} = e_j)]}{P(\mathbf{S}_1|\mathbf{q} = e_j)}, \tag{22}$$

which implies that

$$P(\mathbf{S}_1|\mathbf{q} = e_i) = P(\mathbf{S}_1|\mathbf{q} = e_j). \tag{23}$$

The analysis above can be carried out to any AP in the environment. It means that — according to the WBLS model — the probability of presence of signal from an AP is the same for every state of the environment.

An advantageous aspect brought by this condition is that the model represents situations where the signal is affected in every direction, *e.g.* when there is signal interference close to the AP, and adapts itself to a signal presence frequency that can change along time. On the other hand, the model does not explicitly considers the coverage area of an AP, unlikely the basic system.

In other words, WBLS reduces Wi-Fi noise by eliminating information regarding signal presence. This brings forward an interesting issue that has been largely ignored in the literature: Wi-Fi localisation is almost always seem as based on the use of an already installed infrastructure, but this advantage carries together the problem of a lack of control over the same infrastructure. Harmful effects — from the point of view of localisation — can be produced as consequence of uncontrollable events such as maintenance, device failures, energy blackouts and so forth. WBLS, for considering only signals that are present in the observations, is much more robust to this kind of situation. In the next section, we produce empirical evidence that supports this claim.

5 Experiments

We conducted a series of experiments and compared the results achieved from the basic localisation system and from WBLS in the same physical environment with the same data.

5.1 Testing environment

We tested our system in a building at Escola Politécnica (University of São Paulo, Brazil), in a 45m × 25m area. The plant map is shown in figure 4, where we can see the 181 MPs used in the RSSI mapping, spaced around 2m from each other. Signals from 18 APs were detected, 4 of which (represented as diamonds) located in the mapped area. It was possible to detect an average of 10.9 APs in each MP. The graph corresponding to the topological map of the environment has MPs represented as a set of vertices v_i, with average vertex degree $\delta(v_i)$ =2.87 and 260 edges, which represent navigable paths in the mapped area.

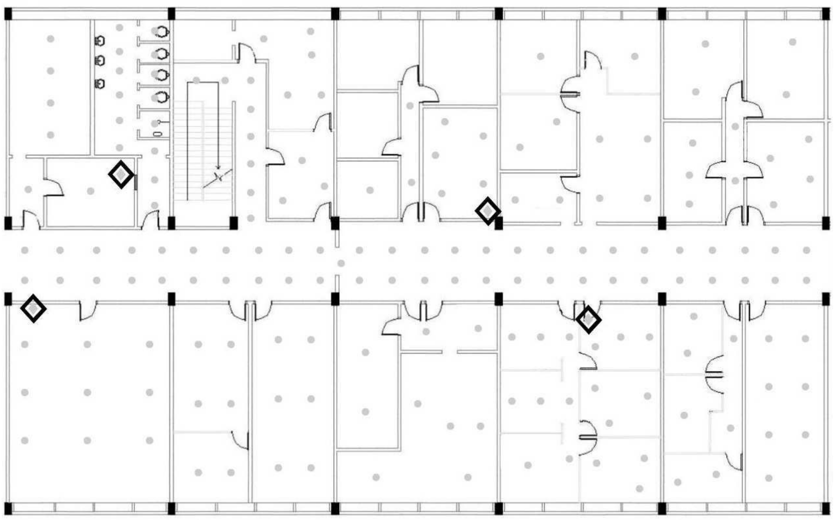

Figure 4: Map of the tested plant: MPs are represented as points and APs as diamonds.

Experiments were performed on a Toshiba Satellite 6500 notebook running Linux 2.4.25

kernel and WLan PCMCIA cards with the Atheros RX5004G chipset. 150 measurements in each of the 181 MPs were collected. A person doing the training spent approximately 30s in each cell (represented by an MP), and walked around slowly in order to cover the entire cell. Data collection took two four-hour periods in two days overall.

5.2 Capturing observation sequences in the walking paths

The procedure we used to capture observation sequences in our building is based on the methodology described in [17] and [14]. We first manually drew a set of 5 feasible walking routes through the building floor, as shown in figure 5. Each route was traversed in both directions (clockwise and anticlockwise), resulting in 10 paths. Once all the lines had been drawn, each line end point was marked in the environment map. The agent is placed on the path beginning (a line end point) and notifies the observation capture program when the moviment starts. The program then records both localisation and signal strengths for each time step. Localisation is estimated by considering both the agent click given on each line end point and a constant velocity along each line segment of the walking path. Each path was traversed seven times, resulting in 70 stored observation sequences.

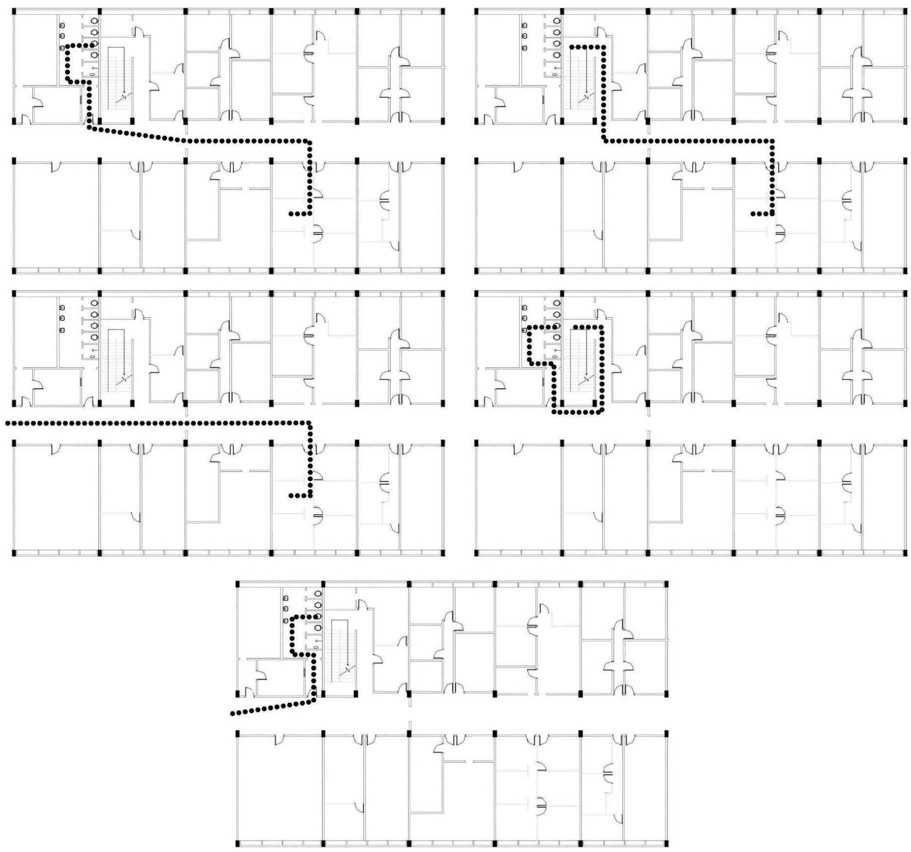

Figure 5: Walking routes for the experiments on localisation with moving devices.

5.3 Performance analysis

We conducted four experiments in order to analyse the performance of WBLS, and used in all of them the RSSI map acquired in the training phase, without any modification. We inspected all observed signals in the training phase and identified both the less and the most frequently observed AP.

In the first experiment, the signal of the less frequently observed AP was eliminated from each observation of the 70 observation sequences stored during he walking paths, in order to simulate its removal from the WLan, and then we ran WBLS and the basic system using this new observation sequence to localise the agent during the walking. In the second experiment we removed the signal from the most frequently observed AP. Finally, in the last experiment we removed both APs: the most and the less frequently observed ones.

In each experiment, we tested a series of 100 different AP combinations on both systems — WBLS and the basic system. For each series we varied the number of APs used in the observation sequence, ranging from 7 to 16 observed APs (not including the AP removed in the step described before). As we used the RSSI maps acquired in the training phase, APs that were not included in the observation sequence actually simulate failures. Localisation estimate error was calculated as the average of the estimated error over all observed APs for each one of the 100 combinations. Results are shown in figure 6a-c, where we see that WBLS presents better results than the basic system in all but the first experiment, where similar performances were achieved. This means that failures in less frequently observed APs do not affect the system performance, as expected.

Finally, we conducted a fourth experiment where no AP failure occurred, i.e., the same APs modelled in the RSSI map are observed in the localisation process executed during walking paths. Results are shown in figure 6d. Performances are similar; however, for reduced AP numbers, the basic system shows a slightly superior performance.

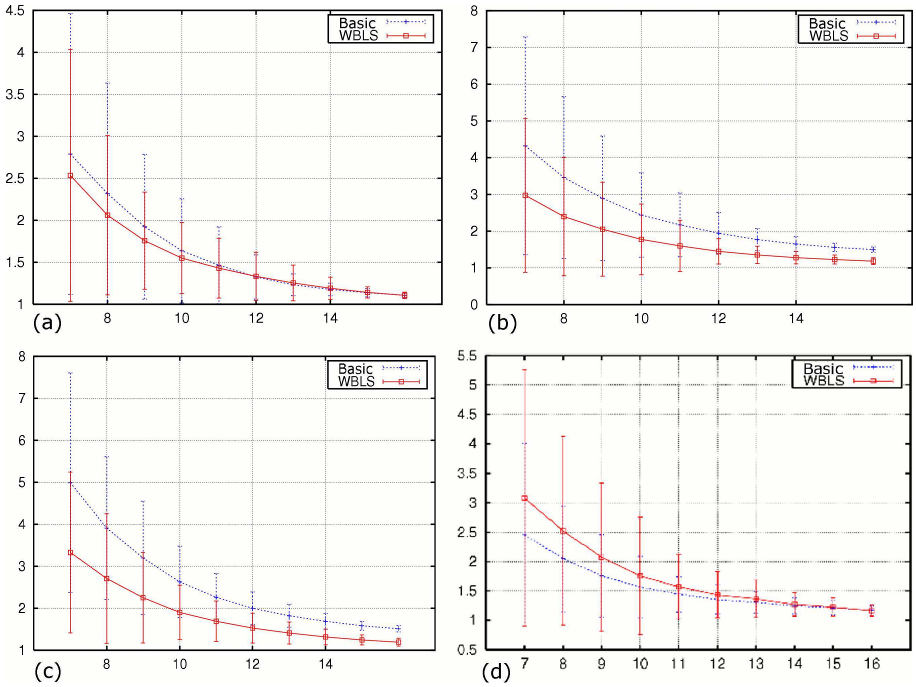

Figure 6: Performance comparison (number of APs $\times$ average localization error (in meters) between WBLS and the basic system when APs are removed during the localisation process: (a) removal of the less frequently observed AP; (b) removal of the most frequently observed AP; (c) removal of both the less and the most frequently observed AP in the training phase; (d) performance comparison between WBLS and the basic system when all APs are observed.

Results based on the Student T test assess that there is a statistically significant difference of performance between WBLS and the basic system in the second, third and fourth experiments. Table 2 shows the t-values for the four experiments, where a t-value>1.96 indicates a confidence level for the performance difference larger than 0.95. For the fourth experiment, results show that there are small differences between both systems only for less than 12 APs.

It is worth noticing that differences between WBLS performance and basic system performance decrease with the increase of the number of observed APs. This means that the more information is available, the less sensitive is the localisation system to AP removals. Moreover, standard deviations for the basic system performance are larger than those depicted by WBLS, indicating that WBLS might be more stable against channel deactivation than the basic localisation system.

Table 2: T test for the four experiments.

	Number of APs									
	7	8	9	10	11	12	13	14	15	16
T Exp. 1	1.51	2.12	2.14	1.54	0.81	0.24	1.16	1.22	0.71	0.06
T Exp. 2	4.96	5.18	5.29	5.91	6.85	9.87	14.51	18.94	26.69	35.31
T Exp. 3	6.84	6.98	7.31	9.07	10.23	11.94	14.42	18.65	28.25	35.52
T Exp. 4	3.62	3.05	2.43	2.78	2.27	2.25	1.54	0.72	0.68	1.93

6 Conclusions

Computing context from WLan infrastructure is attractive because many buildings already have Wi-Fi access points, and more and more mobile devices have been equipped with Wi-Fi hardware. As a consequence, a wide spectrum of new context-based applications based on pervasive computing concepts could be designed to address the needs of a variety of residential, commercial and industrial service users. Many of these applications depend on reliable information about the mobile device localisation in the building in order to successfully accomplish their goals.

This paper presented a new proposal of a localisation system – WBLS – for mobile devices. WBLS exploits the fact that Wi-Fi signal strength varies with localisation, and uses an HMM on a graph of localisation nodes whose transition probabilities are a function of the building floor plan, probabilistic signal strength pattern, and expected speeds of a pedestrian agent holding a portable device. WBLS eliminates signal presence frequency information due to its associated noise, allowing for an increase in the accuracy of the localisation estimation, despite the information about the signal presence that is discarded. Experiments in real environment show that WBLS presents a high degree of robustness regarding failures of access points, a common condition in WLan infrastructures. Moreover, in ideal conditions WBLS was shown to have performance that is similar to the best proposals related in the literature.

Our future work will concentrate on reducing calibration efforts while improving the localization accuracy to make systems of this type more attractive. We are also interested in investigating ways to better reflect the agent customary moving habits and the building uses in the probabilistic model, so that the localization system can not only improve its performance with experience, but also automatically adapt to new situations.

References

[1] Abowd, G. D., Atkeson, C. G., Hong, J., Long, S., Kooper, R., and Pinkerton, M. Cyberguide: a mobile context-aware tour guide. *Wireless networks 3*, 5 (oct 1997), 421–433.

[2] Ashbrook, D., and Starner, T. Learning significant locations and predicting user movement with GPS. In *Proceedings...* (Piscataway, NJ, US, 2002), International Symposium on Wearable Computers, 6th, 2002, IEEE, pp. 101–108.

[3] Bahl, P., and Padmanbham, V. N. Radar: an in-building RF-based user location and tracking system. In *Proceedings...* (Piscataway, NJ, US, mar 2000), INFOCOM - the Joint Conference of the IEEE Computer and Communications Societies, 9th, 2000, IEEE, pp. 775–784.

[4] Balbach, O. UMTS - competing navigation system and supplemental communication system to GNSS. In *Proceedings...* (Fairfax, VA, sep 2000), ION GPS Meeting, 2000, Salt Lake City, ION, pp. 519–527.

[5] Bozic, S. M. *Digital and Kalman filtering.* Edward Arnold Ltd., London, 1979.

[6] Bruce, A., and Gordon, G. Better motion prediction for people-tracking. In *Proceedings...* (Piscataway, NJ, US, 2004), International Conference on Robotics and Automation - ICRA, 2004, Barcelona, IEEE.

[7] Brumitt, B., Meyers, B., Krumm, J., Kern, A., and Shafer, S. Easyliving: technologies for intelligent environments. In *Proceedings...* (Berlin/Heidelberger, sep 2000), international symposium on Handheld and Ubiquitous Computing, 2nd, 2000, Springer, pp. 12–29.

[8] Brunato, M., and Kall, C. K. Transparent location fingerprinting for wireless services. In *Proceedings...* (Laxenburg, AUT, sep 2002), Mediterranean Ad Hoc Networking Workshop - Med-Hoc-Net, 1st, 2002, Sardegna, IFIP.

[9] Haeberlen, A., Flannery, E., Ladd, A. M., Rudys, A., Wallach, D. S., and Kavraki, L. E. Practical robust localization over large-scale 802.11 wireless networks. In *Proceedings...* (New York, sep 2004), international conference on Mobile computing and networking - mobicom, 10th, 2004, Philadelphia, ACM press, pp. 70–84.

[10] Halliday, D., Resnick, R., and Walker, J. *Fundamentals of physics*, 6 ed. John Wiley and Sons, Hoboken, NJ, US, jun 2000.

[11] Hashemi, H. The indoor radio propagation channel. *Proceedings of the IEEE 81* (jul 1993), 943–968.

[12] Hazas, M. A novel broadband ultrasonic location system. In *Proceedings...* (Berlin/Heidelberg, sep 2002), international conference on Ubiquitous Computing - Ubicomp, 4th, 2002, Gteborg, SWE, Springer, pp. 264–280.

[13] Hightower, J., and Borriello, G. A survey and taxonomy of location sensing systems for ubiquitous computing. Tech. rep., Washington: Department of Computer Science and Engineering, University of Washington, 2001. (Technical Report UW-CSE 01-08-03).

[14] Krumm, J., and Horvitz, E. LOCADIO: Inferring motion and location from Wi-Fi signal strengths. In *Proceedings...* (Piscataway, NJ, US, aug 2004), International Conference on Mobile and Ubiquitous Systems - MOBIQUITOUS, 1st, 2004, Boston, IEEE, pp. 4–13.

[15] Krumm, J., Williams, L., and Smith, G. Smartmovex on a graph: an inexpensive active badge tracker. In *Proceedings...* (Berlin/Heidelberg, sep 2002), international conference on Ubiquitous Computing - Ubicomp, 4th, 2002, Gteborg, SWE, Springer-Verlag, pp. 299–307.

[16] Ladd, A. M., Bekris, K. E., Marceau, G., Rudys, A., Wallach, D. S., and Kavraki, L. E. Using wireless ethernet for localization. In *Proceedings...* (Piscataway, NJ, US, sep 2002), International Conference on Intelligent Robots and Systems - iros, 2002, Lausanne, SWI, IEEE, pp. 402–408.

[17] Ladd, A. M., Berkis, K. E., Rudys, A., Kavraki, L. E., and Wallach, D. S. Robotics-based location sensing using wireless ethernet. *Wireless networks 11*, 1-2 (jan 2005), 189–204.

[18] Langfeldt, N. A survey of linux and wi-fi, 2008. http://users.linpro.no/janl/hardware/wifi.html.

[19] Neskovic, A., Neskovic, N., and Paunovic, G. Modern approaches in modeling of mobile radio systems propagation environment. *IEEE communications surveys and tutorials 3* (2000), 2–12.

[20] Pahlavan, K., Li, X., and Makela, J.-P. Indoor geolocation science and technology. *IEEE communications magazine 40* (feb 2002), 112–118.

[21] Rabiner, L. R. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE 77* (feb 1989), 257–286.

[22] Roos, T., Myllymaki, P., Tirri, H., Misikangas, P., and Sievanen, J. A probabilistic approach to WLAN user location estimation. *International journal of wireless information networks 9*, 3 (jul 2002), 155–164.

[23] Sparacino, F. Sto(ry)chastics: a bayesian network architecture for user modeling and computational storytelling for interactive spaces. In *Proceedings...* (Berlin/Heidelberg, oct 2003), international conference on Ubiquitous Computing - Ubicomp, 5th, 2003, Tokyo, Springer-Verlag.

[24] Wbbena, G., Bagge, A., Seeber, G., Bder, V., and Hankemeier, P. Reducing distance dependent errors for real-time precise DGPS applications by establishing reference station networks. In *Proceedings...* (Fairfax, VA, sep 1996), ION GPS Meeting, 9th, 1996, Kansas City, ION, pp. 1845–1852.

[25] Wisocki, T. A., and Zepernick, H.-J. Characterization of the indoor radio propagation channel at 2.4 GHz. *Journal of telecommunications and information technology* (mar 2000), 84–90.

[26] Youssef, M., and Agrawala, A. The Horus WLAN location determination system. In *Proceedings...* (New York, jun 2005), conference on Mobile systems, applications, and services - MobiSys, 3rd, 2005, Seattle, ACM press, pp. 205–218.