

Probabilistic Indoor Tracking of Mobile Wireless Nodes Relative to Landmarks *

Keyong Li,[‡] Dong Guo,[§] Yingwei Lin,[¶] and Ioannis Ch. Paschalidis[†]

Abstract

The profile-based approach is known to be advantageous when it comes to inferring positions of mobile wireless devices in complex indoor environments. The past decade has seen a significant body of work that explores different implementations of this approach, with varying degrees of success. Here, we cast the profile-based approach in a probabilistic framework. Launching from the theoretical basis that this framework provides, we provide a suite of carefully designed methods that make use of sophisticated computations in pursuit of high localization accuracy with low hardware investment and moderate set-up cost. More specifically, we use full distributional information on signal measurements at a set of discrete locations, termed *landmarks*. Positioning of a mobile node is done relative to the resulting landmark graph and the node can be found near a landmark or in the area between two landmarks. Key elements of our approach include profiling the signal measurement distributions over the coverage area using a special interpolation technique; a two-tier statistical positioning scheme that improves efficiency by adding movement detection; and joint clusterhead placement optimization for both localization and movement detection. The proposed system is practical and has been implemented using standard wireless sensor network hardware. Experimentally, our system achieved an accuracy equivalent to less than 5 meters with a 95% success probability and less than 3 meters with an 87% success probability. This performance is superior to well-known contemporary systems that use similar low-cost hardware.

* Research partially supported by the NSF under grants EFRI-0735974, CNS-1239021, and IIS-1237022, by the ARO under grants W911NF-11-1-0227 and W911NF-12-1-0390, and by the ONR under grant N00014-10-1-0952.

‡ Center for Information & Systems Eng., Boston University, e-mail: likeyong@ieee.org.

§ Center for Information & Systems Eng., Boston University, e-mail: dong.dongguo@gmail.com.

¶ Center for Information & Systems Eng., Boston University, e-mail: yingwei@bu.edu.

† Corresponding author. Dept. of Electrical & Computer Eng., Division of Systems Eng., and Center for Information & Systems Eng., Boston University, 8 St. Mary's St., Boston, MA 02215, e-mail: yannisp@bu.edu, url: <http://ionia.bu.edu/>.

Index Terms

Wireless sensor networks, localization, probabilistic profiling, hypothesis testing, optimal deployment.

I. INTRODUCTION

Demand for reliable indoor positioning of mobile wireless devices is rapidly emerging. Using such a functionality, businesses can better manage their equipment and personnel; museums can provide automatically guided tours and enhance visitor experiences; hospitals can track their patients, personnel, and valuable mobile equipment; rescue workers can navigate through a disaster site more easily; malls can guide shoppers to specific stores they seek; large warehouses can track their fleet of forklifts [1]; and security agencies can strengthen the protection of critical assets such as nuclear and biochemical materials.

In contrast to indoor positioning, technologies for outdoor positioning are relatively mature. These include the GPS technology that is widely used today, but also technologies using the cellular network (see [2], [3]). Translating, however, these successes to indoor environments is far from straightforward. First, the GPS technology is hardly operational indoors due to heavy signal attenuation by the building structures. Moreover, the cellular-based technologies for outdoor use cannot produce satisfactory positioning accuracy when confronted with the rich effects of the indoor environment on the signals. The indoor environment can be very complex, and also dynamic due to, for example, people moving and doors opening and closing. The triangulation or trilateration techniques, on which GPS or cellular technologies are based, can be inaccurate and ineffective under such conditions.

The main objective of this chapter is to describe a new probabilistic approach to indoor localization of mobile wireless devices. As we will see, the technique is quite general and can handle any type of information from wireless signals that is correlated with location; from very basic information related to signal strength and available with almost any hardware, to more sophisticated information such as angle and time of arrival. The approach we propose is built on rigorous decision theory and that rigor allows us to provide performance guarantees and tackle associated problems such as tracking and optimal system deployment.

A. Related work

Given the interest in indoor positioning, a wide range of indoor positioning solutions have been proposed, with varying degrees of success. [4] implemented an infrared-based positioning system (*Active Badge*) for low-accuracy applications. [5] proposed an ultrasound-based system (*Cricket*) that achieved high accuracy, but the system requires the installation of a dense network of ultrasound beacons. Acoustic signals are also used by the system proposed in [6] but without the need of infrastructure nodes. In the wireless positioning camp, many types of signal measurements can be useful, ranging from the most basic — *Received Signal Strength Indication (RSSI)* — to more sophisticated ones including signal phase, time-of-arrival (TOA), angle-of-arrival (AOA), and multipath components (MPC) [7], [8].

Among the RSSI-based approaches, the methods in [9], [10], [11], [12] compare mean RSSI measurements to a pre-computed signal-strength map. In particular, [9] tested two methods. The first one uses a map of mean signal strength that is profiled offline, while the second assumes a signal propagation model taking into account how many walls are in the path. The former was shown to be superior, achieving an accuracy of 75% error < 5 meters and 50% < 3 meters. This system and others succeeded in demonstrating the feasibility of meaningful positioning services using wireless sensor networks and injected enthusiasm into the field. Their performance, though, leaves room for improvement. Many other works followed. [13] improved upon [9] by taking the probabilistic nature of the problem into account through the use of Bayesian network techniques. Another probabilistic method appeared in [14] where both probabilities of RSSI measurements and a hidden Markov model of device movements are being estimated. [10] proposed a method based on mean RSSI profiling, but used power-level diversity to achieve an accuracy of 80% < 3 meters. However, further improvement of the performance within this framework (e.g., by increasing the density of “reference signatures”) seems unlikely. Another related idea that combines information from GPS at some locations to reduce the need of profiling was proposed in [15] and achieved accuracies between 2 and 7 meters. Yet another class of systems such as in [16], [17] use stochastic triangulation techniques but rely on a path loss model, thus, introducing a modeling error. In addition to the evidence provided by [9], our earlier related work based on signal strength profiling [18], [19] has also been shown to reduce the mean error distance (by a factor of about 3.5) compared to stochastic triangulation. References to many other systems can

be found in [20] and [21].

B. Key contributions

Despite the extensive literature, several fundamental questions remain. On one hand, there have been some efforts to understand the theoretical limit of wireless-signal-based device positioning. For example, [22] obtained a Cramer-Rao bound of wireless-based positioning, but the result was built upon the presumption that positioning is done using distance and angle of arrival measurements, and these measurements follow Gaussian distributions. A wider range of techniques were considered by [23], but similar assumptions (especially the Gaussian assumption) were put in place. Such assumptions might be overly simplistic in the indoor positioning context. On the other hand, most positioning systems proposed to date have shunned away from sophisticated decision rules for the fear of implementation difficulties on mobile devices. To a certain extent, our work suggests that by careful algorithm design that enables distributed processing, it is possible to implement complex decision rules to achieve robust and accurate positioning using existing low-cost hardware.

Due to the rather diverse applications of indoor positioning and the rapid evolution of the hardware capabilities found in wireless devices, the positioning problem should not be considered in just a single context. Instead, it is meaningful to explore a spectrum of techniques that reflect different constraints and trade-offs of hardware investment, computational complexity, set-up cost, and positioning accuracy. In particular, the present work strives to achieve high accuracy with low hardware investment and moderate set-up cost, at the cost of fairly sophisticated computations. More specifically, our approach does not require advanced signal measurements such as angle and time of arrival (AOA and TOA), and the amount of computation needed is kept under a threshold such that contemporary hardware suffices. One can argue that we treat computational power as a constraint and seek to minimize set-up cost and maximize accuracy.

Even though our primary motivating application is to locate/track nodes of a wireless sensor network, other settings can be served by the exact same techniques we develop. In particular, the nodes can correspond to a wireless phone with Wi-Fi access, or some other Wi-Fi device that we want to locate. Many of the practical applications we outlined earlier can easily be implemented in a smartphone “app,” thus, locating and tracking the phone in an indoor environment becomes important.

Our contribution goes beyond a successful positioning/tracking system with attractive experimental results. The approach differs from what has been considered in the literature and is based on a set of formal techniques that result in analytical performance guarantees. As we will see, positioning is done relative to a *landmark graph*. The nodes of this graph are a chosen set of landmarks, or places of interest; and the edge defined between any two nodes corresponds to the contiguous geographical area between the landmarks. Which points in the coverage area are defined to be the landmarks depends on the specifics of the application. What is important is that the device’s position is resolved either to a node of the landmark graph, if the device is in its vicinity, or to an edge if the device is in the area between two landmarks. The following topics are then covered:

- A. We construct appropriate probabilistic descriptors associated with a device’s position from a limited amount of RSSI measurements. This process is commonly known as *profiling*. The descriptors that we construct go beyond mean values and variances to record the shapes of the measurement distributions. We also associate a parametrized family of probability density functions (pdfs) to each location, which introduces some analytical challenges, but proved to produce more robust performance. Our experimental results show that such information is of significant value to performance.

By “a limited amount” of measurements, we mainly mean that one can directly measure the RSSI distributions only for a finite number of positions during the profiling phase. Without much loss of generality, we assume that these positions coincide with the landmarks. In order to construct the probabilistic descriptors of a device’s position, we adopt a pdf interpolation technique which originated in statistical physics simulations [24]. To the best of our knowledge, this approach is novel for localization problems. Practices that are equivalent to some form of interpolation are not foreign (e.g., [14]), but these have been carried out rather implicitly without formal design and evaluation. Thus, the present chapter makes a contribution in introducing a formal technique to the field, which is also validated by experimental data.

Earlier versions of probabilistic-descriptors have been explored in [25], [18], [19]. However, the descriptors in [25] are single pdfs rather than pdf families. Although [18] used pdf-family descriptors, the way that they were constructed lacks formalism. A related question

to the profiling discussion is: “In what length-scale is pdf interpolation meaningful for RSSI signatures?” Clearly, we would like to minimize measurements but interpolating the pdfs of two very distant locations would not make sense. This question is investigated experimentally, and interestingly, the answer confirms the intuition expressed in earlier work.

- B. We develop a two-tier device tracking system that relies on RSSI measurements made by a set of *clusterheads* positioned at some of the landmarks. Clusterheads are simply static nodes positioned at these landmarks with the exact same capabilities as other devices, potentially though with line power to accommodate their heavier use. The motivation for a two-tier system is that we would like to exploit the fact that most mobile agents in indoor environments are on the move only occasionally. Fortunately, we found it possible to detect whether a mobile node has moved from its previously known position based on observations from a single clusterhead. We call this tier *movement detection* (the lower tier). When movement of a device has been detected, the upper tier is invoked, which detects the new position of the device using multiple clusterheads. We call this tier *localization*, as it is how localization is commonly construed. The decisions of both tiers are formulated as composite hypothesis testing problems. We develop the requisite theory and characterize the probability of detection error. Movement detection has a lower run-time cost, and in many applications the device being tracked can remain at one location for a long period of time. The two-tier design thus results in significant savings. We also consider and address the problem of optimally placing clusterheads in order to minimize the probability of making incorrect decisions. Similar results have been largely lacking in the literature except for [25], [18], [19], which considered a single tier. The present chapter extends the optimization to jointly considering localization and movement detection, and establishes that this is computationally feasible.
- C. We present a working system that demonstrates the practicality of our approach. Our system achieved an accuracy equivalent to $95\% < 5$ meters and $87\% < 3$ meters, which should be considered of high-quality compared with well-known contemporary systems. We also examine the accuracy of our formal pdf interpolation and find that interpolations with two end points 9 meters apart may replace empirically measured pdfs with very good precision. This is significant for improving the efficiency of profiling and reinforces the

findings in [10].

Notation: We use bold lower case letters for vectors, and bold upper case letters for matrices. Our discussions will involve both probability density functions (pdfs) and probability mass functions (pmfs). With a slight abuse of terminology, we will use the term pdf throughout.

Organization: The remainder of this chapter is organized as follows. Section II formulates the problem and identifies the key components of our approach, namely:

- profiling,
- sensor placement,
- localization, and
- movement detection.

Profiling is handled in Section III and the key technique presented is that of interpolating pdfs. The main theoretical underpinnings of our system and the associated algorithms are presented in Section IV. More specifically, Subsection IV-A reviews some basic facts from information theory which are important in the mathematical development of our algorithms. Subsection IV-B describes how (binary) localization and movement detection decisions can be taken by a single clusterhead. Subsection IV-D describes how multiple clusterheads can collaborate (in a distributed fashion) to arrive at an accurate localization decision. Subsection IV-C presents our approach to optimal clusterhead placement. Section V contains all our experimental results and Section VI gathers some concluding remarks.

II. TRACKING PROBLEM FORMULATION

An overview of the tracking problem formulation is shown in Figure 1. Consider the problem of tracking a wireless sensor network node in a contiguous space \mathcal{X} , which typically corresponds to some indoor environment, e.g., a hospital building, a convention center, or a warehouse. As a way of discretizing this space, we consider a given set of *landmarks* and construct what we call a *landmark graph* as follows. The node set of this undirected graph is the set of landmarks $\mathcal{V} = \{V_i \mid i = 1, \dots, M\}$. A landmark can be a room, a reception area, a cubical, a storage area or an intersection of aisles. We draw an edge between landmarks that are neighbors (in some geographic sense); i.e., the edge set is $\mathcal{E} = \{E_{ij} \mid i = 1, \dots, M, j > i, V_j \in \mathcal{N}_i\}$, where \mathcal{N}_i is the set of neighboring landmarks to V_i . In reality, such an edge may represent a section of a corridor or pathway. There are many different ways to formally define a neighborhood and this is left to

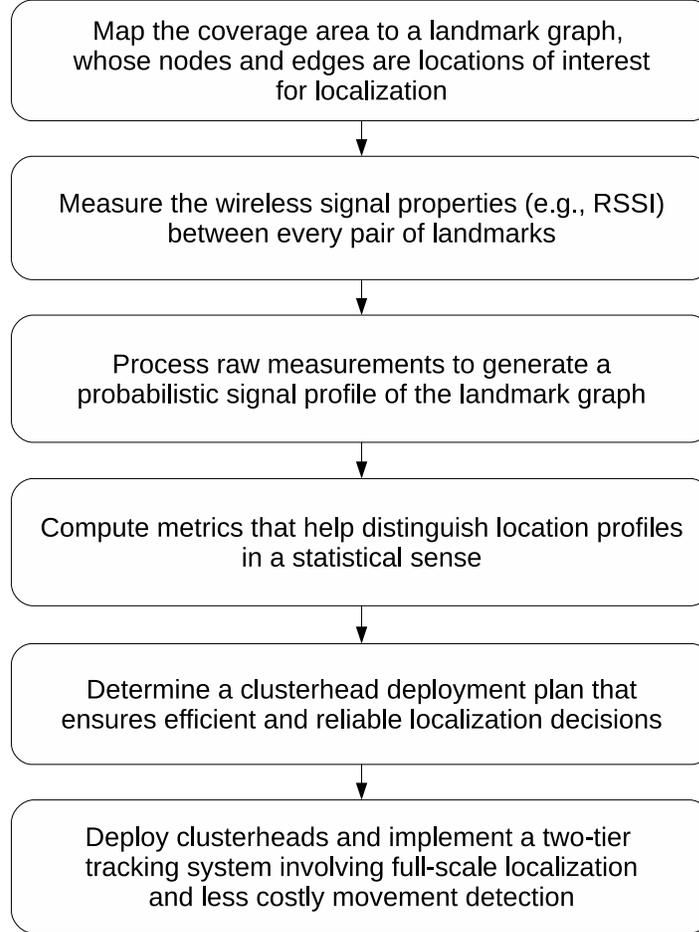


Fig. 1. Overall tracking problem formulation.

the user; one approach could be to set a radius and consider all landmarks within that radius of a node to be its neighbors. With a slight abuse of notation, we sometimes also write: (1) $j \in \mathcal{N}_i$ if $V_j \in \mathcal{N}_i$, and (2) $(i, j) \in \mathcal{E}$ when $E_{ij} \in \mathcal{E}$. In what follows, a *location* refers to either a node or an edge. The set of all locations will be denoted by $\mathcal{L} = \{L_l \mid l = 1 \dots, N\}$, where $N = M + |\mathcal{E}|$.

The next step is profiling, i.e., to associate to various locations appropriate probabilistic descriptors of some features of the wireless signal. Here we use the RSSI, which is measured between all pairs of landmarks. (Additional RF features may also be used if available.) Let $Y^{(k)} \in \{\eta_1, \dots, \eta_R\}$ be the RSSI received at landmark k , which takes values from an R -

dimensional discrete set. We then have a collection of empirical distributions:

$$q_i^{(k)}(y) = \text{Freq}(Y^{(k)} = y|V_i), \quad y = \eta_1, \dots, \eta_R, \quad i, k = 1, \dots, M, \quad (1)$$

where k is the index of the receiving landmark and i is the index of the transmitting landmark. In (1), $\text{Freq}(Y^{(k)} = y|V_i)$ simply denotes the fraction of measurements for which $Y^{(k)} = y$. Using these empirical distributions, we build the probabilistic descriptors of all locations using methods introduced in the sections that follow. As the result of profiling, we obtain pdfs of RSSI that characterize the signals transmitted from each location and received at each landmark. In fact, for improved robustness we associate with each node and edge of the landmark graph a family of pdfs parametrized by vectors θ_i and θ_{ij} , respectively. These are the location descriptors, or profiles:

$$\begin{aligned} p_i^{(k)}(\cdot|\theta_i), \quad i = 1, \dots, M, \quad k = 1, \dots, M; \\ p_{ij}^{(k)}(\cdot|\theta_{ij}), \quad (i, j) \in \mathcal{E}, \quad k = 1, \dots, M. \end{aligned} \quad (2)$$

We may also list the pdf families in terms of the locations, using the notation

$$p_{Y^{(k)}|\theta_l}(\cdot), \quad l = 1, \dots, N, \quad k = 1, \dots, M. \quad (3)$$

The former notation will be used when we discuss profiling, while the latter will be used while introducing the decision rules for positioning.

While the wireless signal transmissions are profiled for every pair of landmarks, for the actual operation that follows we do not suggest placing a device at every landmark to listen to the signals transmitted by the mobile nodes; that would be wasteful. Instead, devices for that purpose, which we call “clusterheads,” shall be placed only at the “best” $K \leq M$ landmarks where the RSSI measurements carry the greatest amount of location information. The clusterhead placement decisions are naturally based on the results of profiling, and we will cast this problem as a *Mixed Integer Linear Programming (MILP)* problem.

Finally, movement detection and localization are done by “comparing” the clusterheads’ RSSI measurements with the location profiles. Formulated as hypothesis testing problems, these decisions are made in statistically meaningful ways.

III. PROFILING

This section focuses on how to generate the location profiles, i.e., (2) and (3), using the empirical RSSI distributions, i.e., (1). The key technique is the interpolation of RSSI distributions, or more generally, pdfs.

A. Interpolation of pdfs

A naive way of interpolating pdfs is to calculate a simple weighted average. However, one may quickly find that the naive way can produce unnatural results. For example, given two Gaussian pdfs with different means, their naive interpolation always has two peaks.

A more sophisticated approach has appeared in a work on statistical physics simulation [24], which we adopt with several generalizations. Given K pdfs, $p_1(x), p_2(x), \dots, p_K(x)$, let $\mu_1, \mu_2, \dots, \mu_K$ and $\sigma_1^2, \sigma_2^2, \dots, \sigma_K^2$ be their means and variances, respectively. Let $\boldsymbol{\rho} \in \mathbb{R}^K$ have elements $\rho_1, \rho_2, \dots, \rho_K \in [0, 1]$ satisfying $\sum_{i=1}^K \rho_i = 1$. These are the weights that we assign to the K source pdfs. We are now seeking an interpolated pdf $p_{\boldsymbol{\rho}}(x)$, whose mean and variance are

$$\mu_{\boldsymbol{\rho}} = \sum_{i=1}^K \rho_i \mu_i \quad \text{and} \quad \sigma_{\boldsymbol{\rho}}^2 = \sum_{i=1}^K \rho_i \sigma_i^2. \quad (4)$$

Let

$$\xi_i(x) = \frac{\sigma_i}{\sigma_{\boldsymbol{\rho}}} (x - \mu_{\boldsymbol{\rho}}) + \mu_i, \quad i = 1, \dots, K.$$

Then define

$$p_{\boldsymbol{\rho}}(x) \triangleq \sum_{i=1}^K \rho_i \frac{\sigma_i}{\sigma_{\boldsymbol{\rho}}} p_i(\xi_i(x)). \quad (5)$$

Intuitively, this operation can be described as: take a copy of each source pdf, stretch and shift it such that the variance and mean equals $\sigma_{\boldsymbol{\rho}}^2$ and $\mu_{\boldsymbol{\rho}}$ respectively, and finally sum these copies together with the weights adjusted in proportion to the standard deviations of the source pdfs.

We next prove that the mean and variance of $p_{\boldsymbol{\rho}}(x)$ are indeed given by (4). Our proof is different from the one given in [24], but is inspired by the discussions in that paper. We believe our proof is more complete, concise, and intuitive for the present context.

Proposition III.1 *The pdf given in (5) satisfies*

$$\begin{aligned} (a) \quad & \int p_{\boldsymbol{\rho}}(x) dx = 1, \\ (b) \quad & \int x p_{\boldsymbol{\rho}}(x) dx = \mu_{\boldsymbol{\rho}}, \\ (c) \quad & \int x^2 p_{\boldsymbol{\rho}}(x) dx = \sigma_{\boldsymbol{\rho}}^2 + \mu_{\boldsymbol{\rho}}^2. \end{aligned} \quad (6)$$

Proof: In general, the α -th moment of the pdf in (5) can be written as

$$\begin{aligned} \int x^\alpha p_\rho(x) dx &= \sum_{i=1}^K \rho_i \int x^\alpha \frac{\sigma_i}{\sigma_\rho} p_i(\xi_i(x)) dx \\ &= \sum_{i=1}^K \rho_i \int \left[(\xi_i - \mu_i) \frac{\sigma_\rho}{\sigma_i} + \mu_\rho \right]^\alpha p_i(\xi_i) d\xi_i. \end{aligned}$$

The last equality holds because $d\xi_i/dx = \sigma_i/\sigma_\rho$. Denote the integral in the above by

$$I_i(\alpha) \triangleq \int \left[(\xi_i - \mu_i) \frac{\sigma_\rho}{\sigma_i} + \mu_\rho \right]^\alpha p_i(\xi_i) d\xi_i.$$

We will show that $I_i(\alpha)$ equals 1, μ_ρ , and $\sigma_\rho^2 + \mu_\rho^2$, respectively for $\alpha = 0, 1$, and 2. Note that if these were true, then $I_i(\alpha)$ is constant with respect to i , thus the weighted sum of $I_i(\alpha)$ over i simply equals $I_i(\alpha)$, and the proposition is proved.

The case of $\alpha = 0$ is trivial.

For $\alpha = 1$,

$$\begin{aligned} I_i(1) &= \int \left[(\xi_i - \mu_i) \frac{\sigma_\rho}{\sigma_i} + \mu_\rho \right] p_i(\xi_i) d\xi_i \\ &= \left[\int \xi_i p_i(\xi_i) d\xi_i - \mu_i \right] \frac{\sigma_\rho}{\sigma_i} + \mu_\rho \\ &= (\mu_i - \mu_i) \frac{\sigma_\rho}{\sigma_i} + \mu_\rho \\ &= \mu_\rho. \end{aligned}$$

For $\alpha = 2$,

$$\begin{aligned} I_i(2) &= \int \left[(\xi_i - \mu_i) \frac{\sigma_\rho}{\sigma_i} + \mu_\rho \right]^2 p_i(\xi_i) d\xi_i \\ &= \int \left[(\xi_i - \mu_i)^2 \frac{\sigma_\rho^2}{\sigma_i^2} + 2(\xi_i - \mu_i) \frac{\sigma_\rho}{\sigma_i} \mu_\rho + \mu_\rho^2 \right] p_i(\xi_i) d\xi_i \\ &= \sigma_i^2 \frac{\sigma_\rho^2}{\sigma_i^2} + 2(\mu_i - \mu_i) \frac{\sigma_\rho}{\sigma_i} \mu_\rho + \mu_\rho^2 \\ &= \sigma_\rho^2 + \mu_\rho^2. \end{aligned}$$

■

In addition to the nice feature of interpolating the mean and variance, formula (5) also interpolates the shapes of the pdfs in some sense. For example, when the original pdfs are Gaussian, the interpolation is also Gaussian.

When the random variable takes discrete values, an issue is that the transformation $\xi(x)$ may produce a value for which probability mass is not defined. An approximate formula that

solves this issue is also provided by [24] for the case of interpolating two pdfs. We generalize the method to more than two pdfs. Assume that the probabilities are defined for values $-\infty, \dots, -1, 0, 1, \dots, \infty$. For integers j and l , and for $i = 1, \dots, K$, let

$$\gamma_{ijl} = \max \left\{ 0, \min \left\{ \begin{array}{l} \xi_i(j+0.5) \\ l+0.5 \end{array} \right\} - \max \left\{ \begin{array}{l} \xi_i(j-0.5) \\ l-0.5 \end{array} \right\} \right\}, \quad (7)$$

where σ_ρ and ξ_i are defined as before. To understand the quantity γ_{ijl} , let $C_{ij,\text{tran}}$ denote the transformed cell $[\xi_i(j-0.5), \xi_i(j+0.5)]$, and $C_{l,\text{pmf}}$ denote $[l-0.5, l+0.5]$, the cell for which the probabilities are defined. Then γ_{ijl} is the proportion of $C_{l,\text{pmf}}$ that overlaps with $C_{ij,\text{tran}}$; see Fig. 2. Note that the length of $C_{l,\text{pmf}}$ is taken to be 1; otherwise, γ_{ijl} should be normalized by the length of $C_{l,\text{pmf}}$. The approximate formula is then

$$p_\rho(j) = \sum_{i=1}^K \rho_i \sum_l \gamma_{ijl} \cdot p_i(l). \quad (8)$$

Formula (7) is chosen such that

$$\sum_l \gamma_{ijl} = \frac{\sigma_i}{\sigma_\rho}, \quad \forall i, j, \quad \text{and} \quad \sum_j \gamma_{ijl} = 1, \quad \forall i, l.$$

Formula (8), to which we refer as *linear interpolation*, is what we use in our experiments. Hereafter, we denote the linear interpolation of K pdfs with the coefficient vector $\rho \in \mathbb{R}^K$ by $\text{Interpol}(\rho, p_1, p_2, \dots, p_K)$.

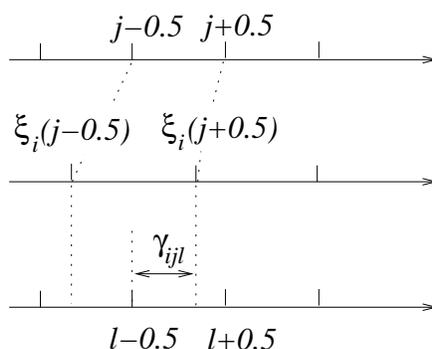


Fig. 2. Interpretation of γ_{ijl} .

B. Associating pdf families to locations

It suffices to consider the RSSI profile of all locations observed by a clusterhead placed at one of the landmarks. The index of the clusterhead is thus suppressed in all formulae of this subsection.

First, we “regularize” the empirical pdfs to eliminate zero elements. This is necessary because the number of our sample measurements during profiling is finite. As a result, some RSSI value η_r that is possible but rare for a location L_i might not be observed during profiling, leaving the r th element of the empirical pdf equal to zero. If we use the empirical pdf directly as the probabilistic descriptor of the location, then when η_r appears, we would rule out L_i immediately, regardless of how many total observations are made and how the rest of the observations resemble the profile of location L_i . This is clearly undesirable. To mitigate this problem, we mix the empirical pdf with a discretized Gaussian-like pdf of the same mean and variance. Namely, let q be an empirical pdf with mean μ and variance σ^2 . Let $\phi(\mu, \sigma^2)$ be a Gaussian-like pdf whose domain is discretized to the set $\{\eta_1, \dots, \eta_R\}$. Let $\gamma \in (0, 1)$ be a chosen mixing factor — typically we set γ to a small value such as 0.1 or 0.2. Then the pdf after regularization is

$$\tilde{q} = \text{Interpol} \left(\begin{bmatrix} 1 - \gamma \\ \gamma \end{bmatrix}, q, \phi(\mu, \sigma^2) \right).$$

Second, consider the areas represented by the landmarks. Here, we use the pdf-family framework to achieve robustness. Specifically, suppose V_i has I neighbors — $\mathcal{N}_i = \{V_{j_1}, \dots, V_{j_I}\}$. Let $\boldsymbol{\theta}_i = (\theta_i^{(1)}, \dots, \theta_i^{(I)}) \in \mathbb{R}^I$, $\boldsymbol{\theta}_i \geq \mathbf{0}$ elementwise, and $\sum_{j=1}^I \theta_i^{(j)} < 1$. Let

$$\boldsymbol{\rho}_{\boldsymbol{\theta}_i} = \begin{pmatrix} 1 - \sum_{j=1}^I \theta_i^{(j)} \\ \theta_i^{(1)} \\ \vdots \\ \theta_i^{(I)} \end{pmatrix}.$$

Then the pdf family associated with V_i can be defined as an interpolation of $I + 1$ empirical pdfs:

$$p_i(\cdot | \boldsymbol{\theta}_i) \triangleq \text{Interpol} \left(\boldsymbol{\rho}_{\boldsymbol{\theta}_i}, \tilde{q}_i(\cdot), \tilde{q}_{j_1}(\cdot), \dots, \tilde{q}_{j_I}(\cdot) \right).$$

This is motivated by the fact that when localization decisions are made, the mobile node we seek to locate will never be exactly on top of a landmark and the pdf of the RSSI measurements

may not match well with the pdf associated with that landmark. By interpolating among all pdfs corresponding to neighboring landmarks, we create a pdf family which is very likely to have a member that matches well our measurements.

Last, consider the edges of the landmark graph. For the same reasons we outlined above, we associate with the edge (i, j) another pdf family defined as a set of parameterized interpolations of two pdfs, where the two source pdfs are drawn from the pdf families representing landmarks i and j . Specifically, let $\vartheta_{ij} \in (0, 1)$ and $\boldsymbol{\theta}_{ij}$ be a vector concatenating $\boldsymbol{\theta}_i$, $\boldsymbol{\theta}_j$, and ϑ_{ij} . The pdf family associated with edge (i, j) is

$$p_{ij}(\cdot|\boldsymbol{\theta}_{ij}) \triangleq \text{Interpol} \left(\begin{bmatrix} \vartheta_{ij} \\ 1 - \vartheta_{ij} \end{bmatrix}, p_i(\cdot|\boldsymbol{\theta}_i), p_j(\cdot|\boldsymbol{\theta}_j) \right).$$

The collection of $p_i(\cdot|\boldsymbol{\theta}_i)$'s and $p_j(\cdot|\boldsymbol{\theta}_j)$'s corresponds to the pdfs defined in (2), which are then rearranged as in (3).

C. An alternative interpolation technique

One may propose another theoretically attractive interpolation technique based on geodesics in the space of probability distributions. We only discuss the two-pdf case for simplicity. A well-known geodesic that connects two pdfs $p_0(x)$ and $p_1(x)$ is a ρ -parameterized curve

$$p_\rho(x) = \frac{p_0^{1-\rho}(x)p_1^\rho(x)}{\int p_0^{1-\rho}(x)p_1^\rho(x)dx}, \quad (9)$$

where $\rho \in [0, 1]$. We will refer to this as the *geometric interpolation*. Like in the linear interpolation case, one can verify that the geometric interpolation of two Gaussian pdfs remains Gaussian. However, our experiments show that the performance of this interpolation is inferior to that of linear interpolation. The reason is that it emphasizes too much the common support of both source pdfs.

D. An alternative Gaussian model

In the above, we associate a family of generally shaped pdfs to each location. If a Gaussian model of the RSSI is used instead, this task can be greatly simplified. One may then ask whether using generally shaped pdfs is worth the effort. The answer to this question may depend on circumstances. However, our experiments show that significant information regarding the signals

transmitted from a location is captured by our approach, but would be neglected if we assume the Gaussian model. This will be discussed at length later.

Now, having the location profiles in the forms of pdf families, we are in a good position to describe our tracking system.

IV. TWO-TIER TRACKING SYSTEM

In positioning a mobile wireless node, one typically needs to draw on measurements from multiple (likely more than 3) clusterheads. However, a somewhat surprising observation is: if we already know the previous location of a wireless device and ask the question of whether it has moved to another location, the RSSI signature observed by a *single* clusterhead provides sufficient information. Furthermore, that clusterhead can be selected for each location based on the profiling results. We will refer to such a test that determines whether a device has moved as *movement detection*. Thus, if many mobile nodes in the system actually stay at some locations for prolonged periods of time (e.g., as is the case for office workers), then switching into the movement detection mode can significantly reduce the workload of the system. This motivates a two-tier positioning scheme of the following generic form with parameters T_L , H , and T_M .

Localization Tier: As will be explained, we will make sure that for every pair of locations, there is at least one clusterhead that can distinguish them. It ensures that our system is capable of statistically localizing a mobile device at any location without ambiguity. When a mobile device enters the coverage area, we have an initial set a candidate locations based on the set of clusterheads that can detect the signal from this device. Suppose there are n locations. There is no need to conduct all $n \times (n-1)$ hypothesis tests of these pairs. Instead, we carry out the process in a greedy, single-elimination manner that guarantees the most reliable $n - 1$ decision are employed given any deployment. This process is repeated every T_L seconds. If the device's position remains the same for H rounds of localization, then switch to the movement detection tier.

Movement Detection Tier: Every T_M seconds, use one clusterhead (selected during the initial setup for the mobile device's last known position) to test the hypothesis that the device's position is unchanged. If the hypothesis is rejected, then return to the localization tier. It is worth noting that the single clusterhead, denoted by C , responsible for detecting whether the device remains at a landmark, denoted by L , is selected such

that all neighboring landmarks of L can be distinguished from L by the signal profile viewed from C . Most likely, C is at a location that is neither too far nor too close to L .

Next we address the issues of how to make the localization and movement detection decisions and how to ensure that the deployed clusterheads can perform the tasks of both tiers satisfactorily for every location. We will present a generalized framework of statistical decisions that covers both tiers. Furthermore, we optimize both tiers jointly in the clusterhead placement phase.

A. Preliminaries

To present the decision rules of the two tiers in a unified framework, we begin by recalling the Kullback-Leibler (KL) distance [26]. For two distribution functions (assuming discrete random variables) \mathbf{p} and \mathbf{q} , the KL distance of \mathbf{q} from \mathbf{p} is defined as

$$D(\mathbf{q}||\mathbf{p}) \triangleq \sum_y q(y) \log \frac{q(y)}{p(y)} \quad (\text{assuming } 0 \cdot \log \frac{0}{a} \equiv 0, \forall a \geq 0). \quad (10)$$

Intuitively, the KL distance reflects the difference between two distributions in a statistically meaningful way (see [27]). It is not a true distance metric (lacks symmetry and does not satisfy the triangle inequality) but it is non-negative, where zero is achieved if and only if $\mathbf{p} = \mathbf{q}$. Typically, \mathbf{q} is some sample distribution and \mathbf{p} is a model distribution. For example, if a certain event is impossible according to the model distribution (\mathbf{p} has a zero element somewhere) but the event occurred in the sample (the corresponding element of \mathbf{q} is positive), then the KL distance equals ∞ , and we immediately know that the sample is not drawn from that model distribution. This example also illustrates that sometimes we may want to avoid proposing a model distribution with a zero element, which is why we regularize the empirical pdfs in Section III-B.

Another important quantity that we draw from information theory is the entropy of the distribution \mathbf{p} . Assume a discrete random variable taking values in $\{\eta_1, \dots, \eta_R\}$. The entropy of the distribution is defined as:

$$H(\mathbf{p}) = - \sum_y p(y) \log p(y). \quad (11)$$

Consider an i.i.d. sequence $\mathbf{y}^n = (y_1, \dots, y_n)$ drawn from \mathbf{p} and let $\mathbf{q}_{\mathbf{y}^n}$ be its sample (empirical) distribution defined as in (1). It is easy to verify that the probability of \mathbf{y}^n can be written as

$$p(\mathbf{y}^n) = \prod_{i=1}^R p(\eta_i)^{nq_{\mathbf{y}^n}(\eta_i)} = e^{-n[H(\mathbf{q}_{\mathbf{y}^n}) + D(\mathbf{q}_{\mathbf{y}^n}||\mathbf{p})]}. \quad (12)$$

Recall now from Equation (3) that the pdf family associated with the RSSI from location l to clusterhead k is $p_{Y^{(k)}|\theta_l}(\cdot)$, $l = 1, \dots, N$, $k = 1, \dots, M$, as constructed in the profiling phase. For either localization or movement detection, suppose the clusterhead makes n i.i.d. observations $\mathbf{y}^{(k),n} = (y_1^{(k)}, \dots, y_n^{(k)})$, with a corresponding sample distribution $\mathbf{q}_{\mathbf{y}^n}^{(k)}$, upon which the decision rules apply. The efficiency of the decision rules to be presented will be quantified using the *error exponent*, defined as

$$d \triangleq - \lim_{n \rightarrow \infty} \sup \frac{1}{n} \log \mathbf{P}(\text{error}). \quad (13)$$

This represents the exponentially decay rate at which the probability of error converges to zero.

B. Statistical decisions contributed by a single clusterhead

Next, consider the decisions that can be made by one clusterhead (hence the clusterhead index is dropped). For localization, the core task is to distinguish between two candidate locations of the mobile device. It was shown in [18], [19] that the question boils down to a binary composite hypothesis testing problem (composite because the distributions involved have unknown parameters — the θ 's), and one can use the well-known *Generalized Likelihood Ratio Test (GLRT)*. In particular, the GLRT decides location i over j if

$$\frac{1}{n} \log \frac{\max_{\theta_j} p_{Y|\theta_j}(\mathbf{y}^n)}{\max_{\theta_i} p_{Y|\theta_i}(\mathbf{y}^n)} < \lambda_{ij},$$

for some appropriate threshold λ_{ij} . We have

$$\begin{aligned} \lambda_{ij} &> \frac{1}{n} \log \max_{\theta_j} p_{Y|\theta_j}(\mathbf{y}^n) - \frac{1}{n} \log \max_{\theta_i} p_{Y|\theta_i}(\mathbf{y}^n) \\ &= \max_{\theta_j} [-H(\mathbf{q}_{\mathbf{y}^n}) - D(\mathbf{q}_{\mathbf{y}^n} \| p_{Y|\theta_j})] \\ &\quad - \max_{\theta_i} [-H(\mathbf{q}_{\mathbf{y}^n}) - D(\mathbf{q}_{\mathbf{y}^n} \| p_{Y|\theta_i})] \\ &= \min_{\theta_i} D(\mathbf{q}_{\mathbf{y}^n} \| p_{Y|\theta_i}) - \min_{\theta_j} D(\mathbf{q}_{\mathbf{y}^n} \| p_{Y|\theta_j}). \end{aligned}$$

Thus, in terms of the KL distance, the GLRT rule can be expressed as:

$$\begin{aligned} &\text{Decide location } i \text{ relative to } j \text{ if and only if} \\ &\min_{\theta_i} D(\mathbf{q}_{\mathbf{y}^n} \| p_{Y|\theta_i}) - \min_{\theta_j} D(\mathbf{q}_{\mathbf{y}^n} \| p_{Y|\theta_j}) < \lambda_{ij}. \end{aligned} \quad (14)$$

A slightly conservative estimate of the corresponding error exponent, denoted by d_{ij} , was also derived in [19] (including an explanation of why λ_{ij} is not always zero and how it can be

selected). Moreover, [19] describes how a decision among many potential locations can be taken by using a sequence of binary tests; interestingly enough, such a decision can be taken in a distributed manner by appropriate collaboration between the clusterheads.

Let us now turn our attention to the movement detection tier. Suppose that the last known location of a mobile device is location j and we would like to determine if the device remains at j based on the n i.i.d. observations in \mathbf{y}^n by a single clusterhead. The following *Generalized Hoeffding Test (GHT)* [28], expressed analogously to (14) using the KL distance, is applicable:

$$\begin{aligned} &\text{Report "no movement" if and only if} \\ &\min_{\theta_j} D(\mathbf{q}_{\mathbf{y}^n} \| p_{Y|\theta_j}) < \lambda_j. \end{aligned} \tag{15}$$

Let α_j and β_j be the error probabilities of false alarm and missed detection, respectively. The following proposition shows that the GHT is optimal in a *Generalized Neyman-Pearson* sense.

Lemma IV.1 *The test in (15) is optimal in a generalized Neyman-Pearson sense, that is,*

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \alpha_j < -\lambda_j, \quad \forall \theta_j, \tag{16}$$

and $-\limsup_{n \rightarrow \infty} \frac{1}{n} \log \beta_j$ is maximized among all tests satisfying (16) uniformly for all θ_i characterizing some alternative pdf.

Proof: Let $\mathbf{P}_j(\cdot)$ denote a probability conditional on the mobile node having not moved from location j . Let also $\mathcal{Q}_n = \{\nu \mid \nu = \mathbf{q}_{\mathbf{y}^n} \text{ for some } \mathbf{y}^n\}$ denote the set of all empirical measures that can be obtained from an n -length observation sequence and $T_n(\nu) = \{\mathbf{y}^n \mid \mathbf{q}_{\mathbf{y}^n} = \nu\}$ the set of n -length observation sequences with an empirical measure equal to ν .

First, letting \mathbf{P}_ν denote a probability under the measure ν and $|\cdot|$ the cardinality, we have

$$\begin{aligned} \mathbf{P}_\nu[T_n(\nu)] &= \sum_{\{\mathbf{y}^n \mid \mathbf{q}_{\mathbf{y}^n} = \nu\}} \mathbf{P}_\nu[\mathbf{y}^n] \\ &= \sum_{\{\mathbf{y}^n \mid \mathbf{q}_{\mathbf{y}^n} = \nu\}} \nu(\eta_1)^{n\nu(\eta_1)} \dots \nu(\eta_R)^{n\nu(\eta_R)} \\ &= |T_n(\nu)| e^{-nH(\nu)}, \end{aligned}$$

which implies

$$|T_n(\nu)| \leq e^{nH(\nu)}. \tag{17}$$

Now, for all θ_j the false alarm probability is given by

$$\begin{aligned}
\alpha_j &= \mathbf{P}_j[\{\mathbf{y}^n \mid \min_{\theta_j} D(\mathbf{q}_{\mathbf{y}^n} \| p_{Y|\theta_j}) \geq \lambda_j\}] \\
&= \sum_{\{\mathbf{q}_{\mathbf{y}^n} \mid \min_{\theta_j} D(\mathbf{q}_{\mathbf{y}^n} \| p_{Y|\theta_j}) \geq \lambda_j\}} |T_n(\mathbf{q}_{\mathbf{y}^n})| p_{Y|\theta_j}(\mathbf{y}^n) \\
&\leq \sum_{\{\mathbf{q}_{\mathbf{y}^n} \mid \min_{\theta_j} D(\mathbf{q}_{\mathbf{y}^n} \| p_{Y|\theta_j}) \geq \lambda_j\}} e^{nH(\mathbf{q}_{\mathbf{y}^n})} e^{-n[H(\mathbf{q}_{\mathbf{y}^n}) + D(\mathbf{q}_{\mathbf{y}^n} \| p_{Y|\theta_j})]} \\
&= \sum_{\{\mathbf{q}_{\mathbf{y}^n} \mid \min_{\theta_j} D(\mathbf{q}_{\mathbf{y}^n} \| p_{Y|\theta_j}) \geq \lambda_j\}} e^{-nD(\mathbf{q}_{\mathbf{y}^n} \| p_{Y|\theta_j})} \\
&\leq (n+1)^R e^{-n\lambda_j}.
\end{aligned}$$

For the first inequality above we have used (12) and (17). In the last inequality above we used the fact that the set of all possible empirical measures, \mathcal{Q}_n , has cardinality upper bounded by $(n+1)^R$ (a symbol of length R with each element taking values from $\{\frac{0}{n}, \dots, \frac{n}{n}\}$). This establishes (16).

Let now \mathcal{S}_n be some other decision rule satisfying (16). It is well known that the empirical measure is a sufficient statistic so any rule will depend only on that. Let $\alpha_{\mathcal{S}_n}$ and $\beta_{\mathcal{S}_n}$ denote the corresponding false alarm and missed detection probabilities. For all $\epsilon > 0$ and all large enough n we have

$$\alpha_{\mathcal{S}_n} \leq e^{-n(\lambda_j + \epsilon)}. \quad (18)$$

Meanwhile for all $\epsilon > 0$, all large enough n , and any \mathbf{y}^n such that \mathcal{S}_n declares ‘‘movement’’ it holds

$$\begin{aligned}
\alpha_{\mathcal{S}_n} &= \sum_{\{\mathbf{q}_{\mathbf{y}^n} \mid T_n(\mathbf{q}_{\mathbf{y}^n}) \text{ implies movement}\}} |T_n(\mathbf{q}_{\mathbf{y}^n})| p_{Y|\theta_j}(\mathbf{y}^n) \\
&\geq \sum_{\{\mathbf{q}_{\mathbf{y}^n} \mid T_n(\mathbf{q}_{\mathbf{y}^n}) \text{ implies movement}\}} (n+1)^{-R} e^{-nD(\mathbf{q}_{\mathbf{y}^n} \| p_{Y|\theta_j})} \\
&\geq e^{-n[D(\mathbf{q}_{\mathbf{y}^n} \| p_{Y|\theta_j}) + \epsilon]},
\end{aligned}$$

where the first inequality above uses [29, Lemma 2.1.8]. Comparing the above with (18) it follows that if \mathbf{y}^n implies ‘‘movement’’ then for all θ_j it should hold $D(\mathbf{q}_{\mathbf{y}^n} \| p_{Y|\theta_j}) \geq \lambda_j$. Therefore, the GHT in (15) should declare movement as well, which implies that $\beta_{\mathcal{S}_n} \geq \beta_j$ for all θ_i characterizing an alternative pdf, where $i \neq j$. The latter establishes that the GHT maximizes the exponent of the missed detection probability. \blacksquare

Eq. (16) provides a bound on the exponent of the false alarm probability. To bound the exponent of the missed detection probability we can use Sanov's theorem [29, Chap. 2]. Specifically,

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \beta_j \leq - \min_{\mathbf{q} \in \mathcal{H}_j} D(\mathbf{q} \| p_{Y|\theta_i}), \quad \forall \theta_i, \quad (19)$$

where $\mathcal{H}_j = \{\mathbf{q} | \min_{\theta_j} D(\mathbf{q} \| p_{Y|\theta_j}) < \lambda_j\}$.

Defining

$$\begin{aligned} Z_{j,\theta_i}(\lambda_j) &= \min_{\mathbf{q}} D(\mathbf{q} \| p_{Y|\theta_i}) \\ \text{s.t. } & \min_{\theta_j} D(\mathbf{q} \| p_{Y|\theta_j}) < \lambda_j, \end{aligned} \quad (20)$$

and

$$Z_j(\lambda_j) = \min_{i \neq j} \min_{\theta_i} Z_{j,\theta_i}(\lambda_j), \quad (21)$$

we can write

$$- \limsup_{n \rightarrow \infty} \frac{1}{n} \log \beta_j \geq Z_j(\lambda_j).$$

The error exponent of the whole test equals the lesser between the exponents of α_j and β_j . On the other hand, it is straightforward to show that $Z_j(\lambda_j)$ is monotonically decreasing in λ_j . If we could compute $Z_j(\lambda_j)$, then the solution of $Z_j(\lambda_j) = \lambda_j$ gives the optimal value of λ_j as well as the best achievable error exponent. However, the computational cost of finding the exact solution is significant. This is mostly because the constraint in (20) is non-convex and also high-dimensional. To side-step this problem, we replace $Z_{j,\theta_i}(\lambda_j)$ with an estimate based on dual relaxation:

$$\begin{aligned} \tilde{Z}_{j,\theta_i}(\lambda_j) &\triangleq \max_{\mu \geq 0} \min_{\theta_j} \min_{\mathbf{q}} D(\mathbf{q} \| p_{Y|\theta_i}) + \mu(D(\mathbf{q} \| p_{Y|\theta_j}) - \lambda_j) \\ \text{s.t. } & \sum_y q(y) = 1, \\ & q(y) \geq 0, \quad \forall y. \end{aligned} \quad (22)$$

As a relaxation, it holds that $\tilde{Z}_{j,\theta_i}(\lambda_j) \leq Z_{j,\theta_i}(\lambda_j)$. The computational gain comes from the fact that the minimization over \mathbf{q} is now convex. In fact, the optimal solution can be found in closed-form:

$$q^*(y) = \frac{p_{Y|\theta_i}^\rho(y) p_{Y|\theta_j}^{1-\rho}(y)}{\sum_{\eta} p_{Y|\theta_i}^\rho(\eta) p_{Y|\theta_j}^{1-\rho}(\eta)}, \quad \text{where } \rho = \frac{1}{1 + \mu}. \quad (23)$$

Then $\tilde{Z}_j(\lambda_j) = \min_{i \neq j} \min_{\theta_i} \tilde{Z}_{j,\theta_i}(\lambda_j)$, which is also monotonically decreasing in λ_j . The solution of $\tilde{Z}_j(\lambda_j) = \lambda_j$, denoted by h_j , gives a near-optimal threshold and the error exponent at the same time.

C. Optimal clusterhead placement

From the above, for each landmark (also candidate clusterhead location) k , we have a collection of localization error exponents d_{ijk} , and movement-detection error exponents h_{jk} , all of which are optimized (note that the index of the clusterhead location is now added back).

In the profiling phase, every location is a potential place of clusterhead installation. In the actual operation of the tracking system, however, we may not want to place a clusterhead at every location, because if we do that, some of the clusterheads may not be in a good position to provide useful information. The budget of the system may further reduce the number of clusterheads that we are allowed to deploy. For these reasons, we need to place the clusterheads carefully. In the case of the dense deployment mentioned above, placement amounts to selecting which nodes from the ones already deployed will play the role of a clusterhead.

In formulating the clusterhead placement problem, our goal is to ensure:

- 1) For the localization tier, there should be at least one clusterhead k for each pair of locations i, j such that L_i, L_j can be clearly distinguished using observations of clusterhead k under the GLRT.
- 2) For the movement detection tier, there should be at least one clusterhead k for each location j such that L_j can be distinguished from alternative locations using clusterhead k under the GHT.

More precisely, for a given number K of clusterheads to deploy, we maximize the error exponent level ϵ that is met (or exceeded) for every location in the movement detection tier, and for every pair of locations in the localization tier. Furthermore, the error exponent for each location or each pair of locations is given by the clusterhead that can best resolve the case (with greatest error exponent). Formally, let \mathcal{P} denote the set of landmarks where clusterheads are placed. The problem can be written as

$$\begin{aligned}
 & \max_{\mathcal{P} \subset \mathcal{V}} \quad \epsilon \\
 & \text{s.t.} \quad |\mathcal{P}| = K, \\
 & \quad \min_{i,j} \max_k d_{ijk} \geq \epsilon, \\
 & \quad \min_j \max_k h_{jk} \geq \epsilon.
 \end{aligned} \tag{24}$$

It is important to note that the optimal value of this problem provides an upper bound on the probability of error for both localization and movement detection decisions. That is, if ϵ^* is an

optimal solution, the maximum probability of error of our system satisfies

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{P}[\text{error}] \leq -\epsilon^*. \quad (25)$$

Problem (24) can be reformulated as the *Mixed Integer Linear Programming (MILP)* problem shown in Fig. 3.

$$\begin{aligned}
 & \max \quad \epsilon \\
 & \text{s.t.} \quad \sum_{k=1}^M x_k = K \\
 & \quad \sum_{k=1}^M y_{ijk} = 1, \quad \forall i, j = 1, \dots, N, i < j, \\
 & \quad \sum_{k=1}^M z_{jk} = 1, \quad \forall j = 1, \dots, N, \\
 & \quad y_{ijk} \leq x_k, \quad \forall i, j, i < j, k = 1, \dots, M, \\
 & \quad z_{jk} \leq x_k, \quad \forall j, k = 1, \dots, M, \\
 & \quad \epsilon \leq \sum_{k=1}^M d_{ijk} y_{ijk}, \quad \forall i, j, i < j, \\
 & \quad \epsilon \leq \sum_{k=1}^M h_{jk} z_{jk}, \quad \forall j, \\
 & \quad y_{ijk} \geq 0, \quad \forall i, j, i < j, \forall k, \\
 & \quad z_{jk} \geq 0, \quad \forall j, \forall k, \\
 & \quad x_k \in \{0, 1\}, \quad \forall k.
 \end{aligned}$$

Fig. 3. Clusterhead placement MILP formulation.

Software exists for solving generic MILPs. Furthermore, we have designed an algorithm that solves the MILP in Fig. 3 much faster than commercial general purpose MILP-solvers. Our approach is to solve this problem by an iterative *feasibility* algorithm along the lines proposed in [30]. In particular, we use a modified version of a *two-phase* algorithm proposed in [31]. Its computational advantage lies in the fact that we solve a feasibility problem in each iteration that contains only $O(M)$ variables and $O(N^2)$ constraints, instead of the $O(N^2M)$ variables and $O(N^2M)$ constraints that appear in the formulation.

D. Putting everything together: the decision procedure involving multiple deployed clusterheads

While movement-detection requires one clusterhead at a time, localization generally requires more than one clusterhead. No matter how good the decision rule is, a single clusterhead is normally not enough for distinguishing every pair of locations in a wide coverage area, because given any single clusterhead, there may be pairs of locations whose RSSI profiles are similar. So, we use the following procedure to “chain” together the decisions of multiple clusterheads.

When a wireless device is detected, we can have an initial estimate of the candidate locations based on the set of clusterheads that received signals from that device. Suppose there are n candidate location. Instead of comparing all $n \times (n - 1)$ pairs of them, we apply a greedy selection process.

To understand this procedure, first recall that the profiling and clusterhead placement phases assigned one clusterhead for each pair of locations that need to distinguished. Furthermore, the statistical reliability with which this pair of locations can be distinguished (based on the clusterhead of our choosing) is associated with an error exponent, the greater value of which indicates higher reliability.

Now, we sort the pairs of candidates in descending order of the associated error exponent. Using GLRT, we first make a judgment between the first pair of candidates in this ordered list. That is, we start with the most reliable decision that can be made. The candidate that loses the judgment is eliminated, so are the pairs that include the eliminated candidate. Then we move our focus to the remaining pair that is associated with the next greatest error exponent. In the end, $n - 1$ judgments are made and a single winner is left.

V. EXPERIMENTS

Our testbed is set up on the first floor of a Boston University building (see Fig. 4). The wireless device used is the MPR2400 (MICAz) “mote” from Crossbow Technology Inc. (now Memsic Inc.).

A. Testing the complete system

Our two-tier tracking system covers 10 rooms and the corridors, which are mapped to 30 landmarks, marked by either a green circle or a red square on the floor plan (Fig. 4), where the latter marks the clusterhead positions obtained from solving the optimal placement problem.

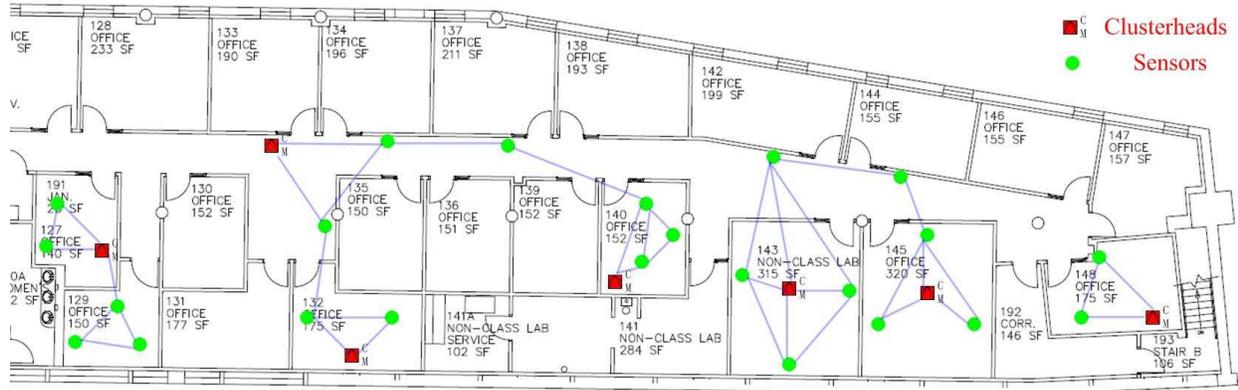


Fig. 4. Floor plan with the landmarks for the testbed.

The landmark graph is then constructed resulting in 39 edges. Adding the landmarks and edges together, we have a total of 69 locations. Hence, $N = 69$ and $M = 30$ in this experiment. A mote is placed at each landmark location, but only some of them will serve as clusterheads. All 30 motes are connected to a base MICAz through a mesh network. The base mote is docked on a Stargate node which forwards the messages back to the server.

The experimental validation of our localization approach can be divided into the following six phases, which have a one-to-one correspondence to those in Fig. 1:

Phase 1: Map the coverage area to a landmark graph.

Phase 2: We obtained the empirical pdfs for the landmarks corresponding to Eq. (1). The 30 motes placed at the landmark locations took turns to broadcast packets, specifically, when one was transmitting the others were listening and recording the RSSI. A total of 200 packets were transmitted by each mote. The data collection was repeated for the combinations of two frequencies and two power levels; details will be given below.

Phase 3: We used the methods in Section III to construct the pdf families corresponding to Eq. (2), which are the descriptors of all 69 locations. Note that the interpolation technique allowed us to construct high quality descriptors without densely covering the area with landmarks.

Phase 4: We obtained d_{ijk} and h_{jk} as described in Sec. IV.

Phase 5: We solved the MILP to optimize clusterhead placement and simultaneously obtained the performance guarantee of Eq. (25). In the MILP formulation, we needed to input K ,

the total number of clusterheads. By varying K from 1 to 30, we discovered that the performance guarantee reached a satisfactory level after $K = 7$, and somewhat flattens afterward. Thus, we placed clusterheads at 7 locations (again, marked by red squares in Fig 4).

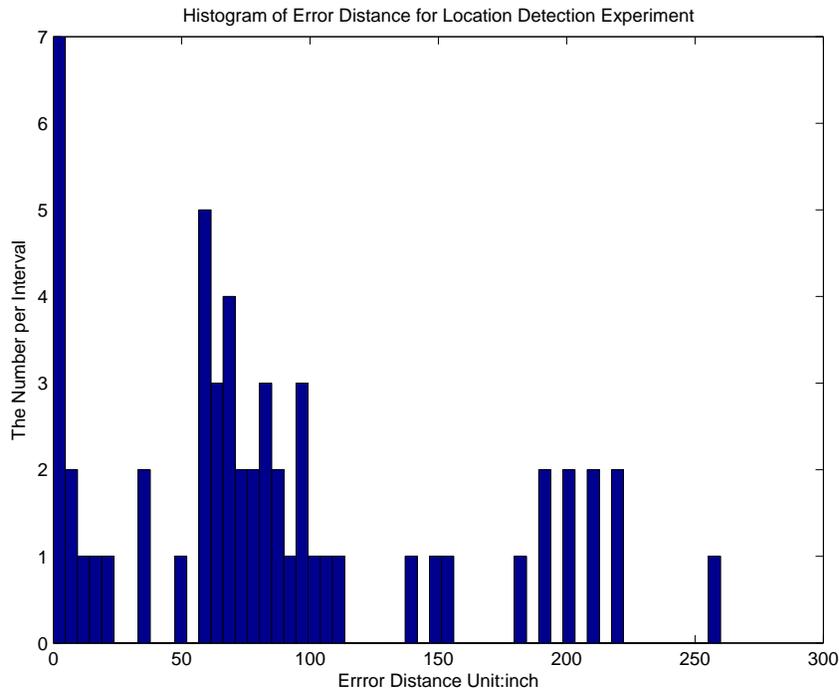
Phase 6: We introduced mobile motes in the coverage area and let the system make localization and movement detection decisions.

For Phase 1, the coverage area is mapped to a landmark graph as shown in Fig. 4.

We let Phase 2 (a completely automated procedure) stretch over 24 hours to acquire data under diverse conditions of the surrounding environment. The objective is to capture the indoor environment in all possible “modes” and configurations so that the pdfs we generate can model any one of these conditions. Phase 3 takes virtually no time. Phase 4 takes another 24 hours on our computer, although further optimization of our code may reduce the computation time significantly. Phase 5 only takes about half an hour. Note that all these steps are performed only once, after which the amount of real-time computation needed for each localization decision is very small, such that the resources on the clusterheads (typically motes plugged into a wall outlet or a large enough battery) are sufficient.

We know from previous experiences [19] and the literature that frequency and power diversity provide better performance. The mote to be located broadcasted 20 packets over the combination of 2 frequencies (2.410 GHz and 2.460 GHz) and 2 power levels (0 dBm and -10 dBm), with 5 packets corresponding to each combination. We achieved a mean error distance of 87.32 inches, which is better than our earlier result of 96.08 inches [19] based on techniques that do not use a formal method of pdf interpolation. The percentile of errors < 3 meters (118 inches) also improved from 80% to 87%. One may also count from Fig. 5 that the percentile of errors < 5 meters (197 inches) is 95%.

The total coverage area (we have excluded the rooms that are in the floor plan but to which we did not have access) was 1827 feet², that is, about 61 feet² per landmark. With a mean error distance of $\bar{D}_e = 7.3$ feet the mean area of “confusion” was $7.3^2 = 53$ feet². It is evident that we were able to achieve accuracy on the same order of magnitude as the area “covered” by a landmark; this is the best possible outcome with a “discretized” system such as ours. That is, the system was identifying the correct location or a neighboring location most of the time. We used a clusterhead density of 1 clusterhead per $1827/7 = 261$ feet². Note that our system is *not* based on the “proximity” to a clusterhead; the ratio of locations to clusterheads is $69/7$, or



$$\bar{D}_e = 87.32 \text{ inches}$$

Fig. 5. Localization result.

about 10.

We also obtained results for the movement detection tier. The mote to be monitored now transmits 20 packets at a single frequency and power level depending on its *a priori* location. The use of a single frequency/power level was selected because the diversity only benefits decisions involving multiple clusterheads. The resulting error probability was 8%. Again, we emphasize that we are exploiting only the most basic RF measurements in obtaining these results. Yet, the approach is easily generalizable to include additional information if available (in that case, instead of scalar RSSI measurements we will be dealing with an observation vector).

B. A closer look at pdf interpolation

We have proposed a rather sophisticated interpolation technique for generating location profiles. One concern is: if the interpolated pdfs were merely low-quality approximations of the actual pdfs, then we might be better-off using a Gaussian approximation, which is computationally

ally cheaper. In our experiments however, the interpolated pdfs did a very good job preserving the shape information of the empirical pdfs which were not close to a Gaussian. As will be shown, the decision accuracy using the interpolated pdfs dominates that of the Gaussian approximation by a significant margin. Another question that we attempt to answer is: At what length-scale does pdf interpolation make sense? It turns out that the interpolation is very meaningful when the two end points are about 30 feet (or 9 meters) apart, but not when they are 60 feet apart.

The first experiment is conducted in a roughly straight corridor of about 60 feet long, mapped to 5 locations roughly 15 feet apart. Labeling the locations consecutively as location 1, 2, . . . , 5, we place the clusterhead at location 1 (which is at one end of the corridor). To measure the signals transmitted from each location, one of the coauthors stood at that location holding a transmitting mote, which sends a packet every 5 seconds. We chose to have a person hold the mote because this is close to an actual application scenario. The clusterhead received the packets and recorded the RSSI values. During the experiment, a total of 150 packets were sent from each location. Due to packet loss, the number of actual samples taken by the clusterhead is less, but we still obtained more than 100 samples for each location. Then, we mix a Gaussian component into each of the six empirical distributions as described earlier with a mixing factor of 0.2, i.e., regularized empirical distribution = 0.8 measured + 0.2 Gaussian. The empirical distributions for the six locations after regularization are denoted by q_1, q_2, \dots, q_5 .

We compare three interpolation methods. First, in what is labeled “linear short interpolation,” the interpolated pdf of location i is generated using q_{i+1} and q_{i-1} :

$$p_{i,\text{short}} = \text{Interpol} \left(\begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}, q_{i-1}, q_{i+1} \right), \quad i = 2, 3, 4.$$

Second, in what is labeled “linear long interpolation,” the interpolated pdfs are generated using q_1 and q_5 :

$$p_{i,\text{long}} = \text{Interpol} \left(\begin{bmatrix} \frac{5-i}{4} \\ \frac{i-1}{4} \end{bmatrix}, q_1, q_5 \right), \quad i = 2, 3, 4.$$

Third, we adopt the Gaussian model instead and interpolate the pdf of each location with adjacent

locations:

$$\begin{aligned}
 p_{i,\text{gaussian}} &= \text{Interpol} \left(\begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}, \phi(\mu_{i-1}, \sigma_{i-1}^2), \phi(\mu_{i+1}, \sigma_{i+1}^2) \right) \\
 &= \phi \left(\frac{\mu_{i-1} + \mu_{i+1}}{2}, \frac{\sigma_{i-1}^2 + \sigma_{i+1}^2}{2} \right), \quad i = 2, 3, 4.
 \end{aligned}$$

Last, we also include the geometric interpolation (cf. Sec. III-C) for comparison, denoted by $p_{i,\text{geo}}$, where location $i \in \{2, 3, 4\}$ is interpolated by location $i - 1$ and $i + 1$ as in the linear short case.

1) *Qualitative study*: We visually compare the various interpolation results for location 2 as an example (Fig. 6). The short interpolation seems to capture some shape information of the

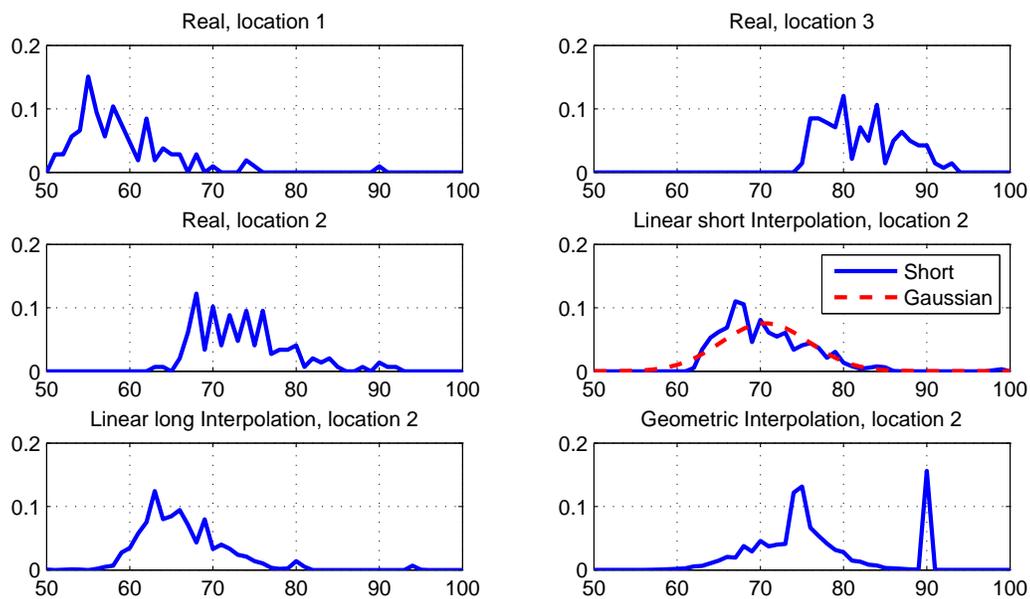


Fig. 6. Visual comparison of interpolated pdfs for location 2. The horizontal axis represents signal strength reading of the notes.

actual pdf that is missed by the Gaussian model. For example, the empirical pdf is skewed to the left. The interpolated pdf also exhibits the skewness, while the Gaussian pdf is always symmetrical. One may also notice that the linear long and the geometric interpolations appear very different from the actual pdf.

2) *Quantitative observation*: First, it is of interest to compare the qualities of the different interpolations using the KL distance (10) as a metric of distance between pdfs. As we have seen, the KL distance appears in both localization and movement detection error exponents.

The comparison is plotted in Fig. 7. It is very interesting to see that *the quality of linear short interpolation dominates that of the Gaussian model in the KL sense*. For example, the KL distance of short-interpolation-to-empirical for location 4 is only a little over one third of that of the Gaussian model. For locations 2 and 3, the difference is roughly a factor of 1.5, which is still significant. Similarly, the linear short interpolation appears superior to the geometric (also short) interpolation. The long interpolation on the other hand clearly departs from the actual distribution.

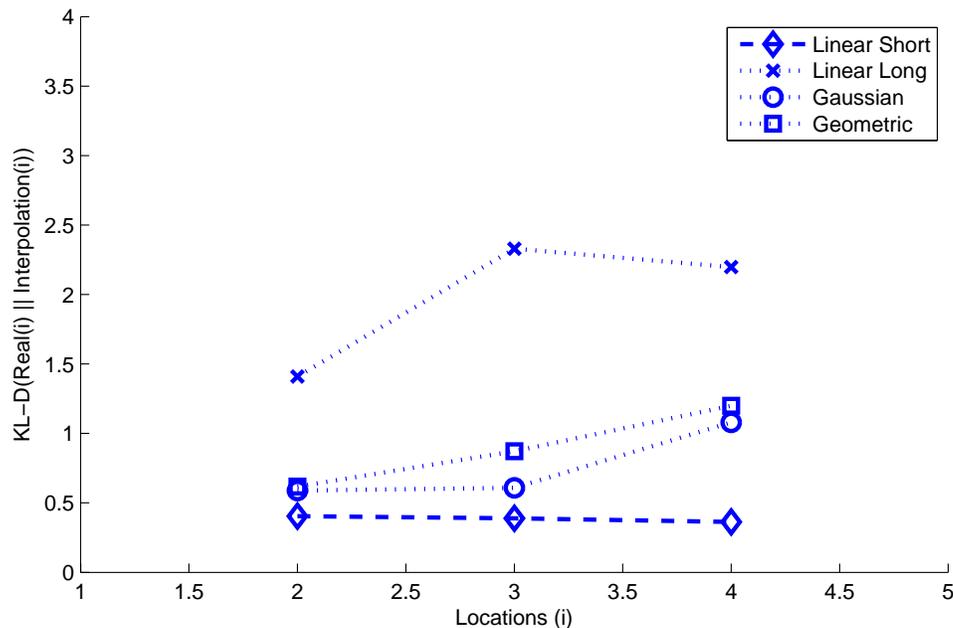


Fig. 7. Kullback-Leibler distance comparison.

Next, we attempted location distinction using these pdfs. We omit the geometric and the linear long cases as, based on the discussion above, they are not suitable for our purposes. We will hence change the label “linear short interpolation” to simply “linear interpolation”. For each location $j = 2, 3$, and 4, and each sample size $n = 5, 10, 15, \dots, 30$, we tested the hypothesis that “the wireless node is at location j ” using n RSSI measurements drawn randomly from a

large pre-compiled data set. Each sample contained measurements associated with one single location i , where $i \in \{1, \dots, 6\}$. The GHT was used to make the decision, in which the threshold λ was optimized for each j and n . For each pair of j and n , we repeated the trial 4000 times and calculated the empirical error probability, which is a weighted sum of all types of errors.

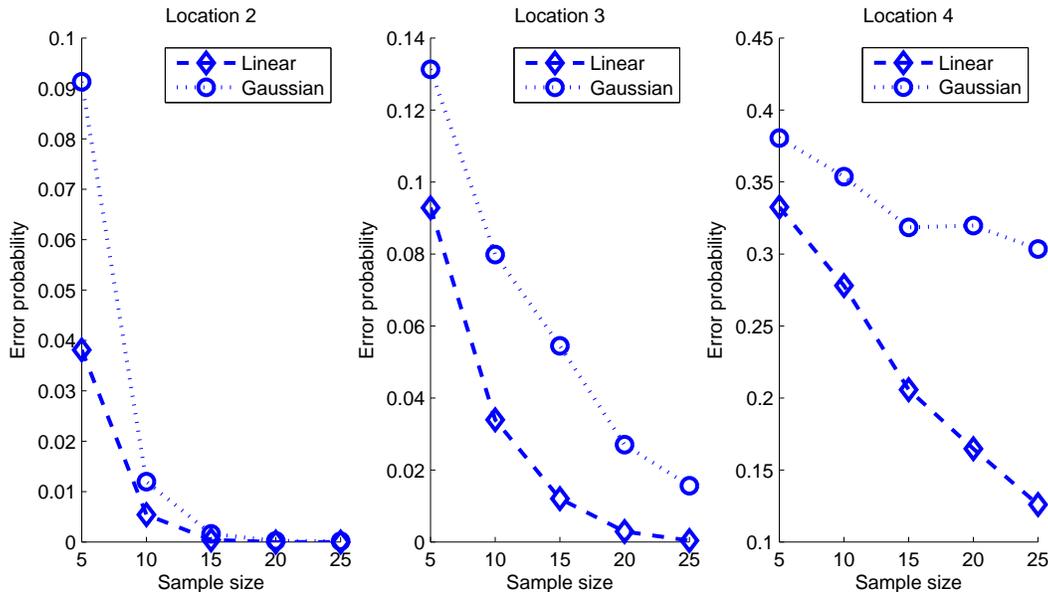


Fig. 8. Comparison of linear interpolation and the Gaussian model in terms of the error probability of GHT using the interpolated pdf.

Several observations are in order. First, very low error probability is achieved using the linearly interpolated pdfs. Second, and maybe more interestingly, the error probabilities using linearly interpolated pdfs are lower than those using the Gaussian model in all instances, and the difference is quite large in most cases. Further, as the sample size increases, the former decreases faster than the latter. These demonstrate that the approximation qualities of the interpolated pdfs are fairly notable, and the added computational effort (compared to the Gaussian model) may indeed be well justified.

Yet another interesting observation emerges in the comparison with [10], which has shown that when the spacing of “reference signatures” goes below roughly 10 meters, the improvement in performance diminishes. The spacing of the “reference signatures” there is analogous to the distance between the two end-point locations in our pdf interpolation. The two end points in

the working version (the “short”) of our pdf interpolation happen to be a little more than 9 meters (30 feet) apart. This result reinforces that of [10], as both indicate that taking empirical measurements at a spacial density of less than 9 or 10 meters apart, or roughly 1 per 25 sq. meters, carries diminishing benefit.

The interpolations above were done along a line. We also tested interpolations in a triangle, as shown in Figure 9. The clusterhead (the listener) was placed at location 1 and the profile of location 5 was interpolated using those of location 2,3, and 6. Figure 10(a) shows the KL distance from the interpolated profile of location 5 to the empirical profile of itself and other locations. The figure shows that under linear interpolation, the distance to self is clearly lower than those to others — the plot has a clear minimum at location 5 — while the distinction under Gaussian interpolation is less pronounced. Finally, Figure 10(b) shows the error probability of GHT using different profiles. Consistent with the KL distance, the linear interpolated profile out-performs the Gaussian one significantly.

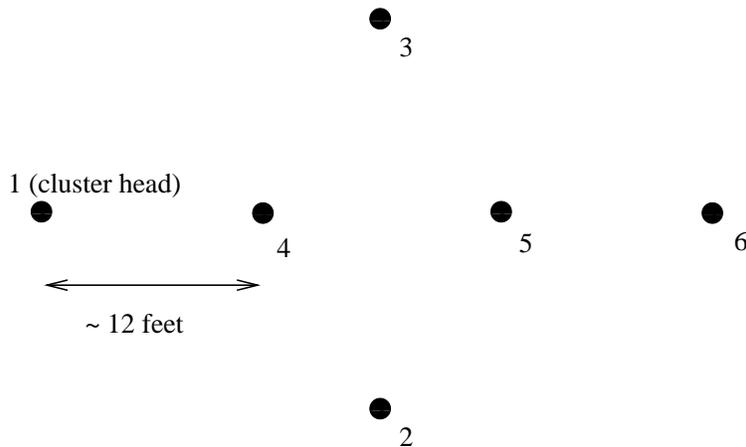
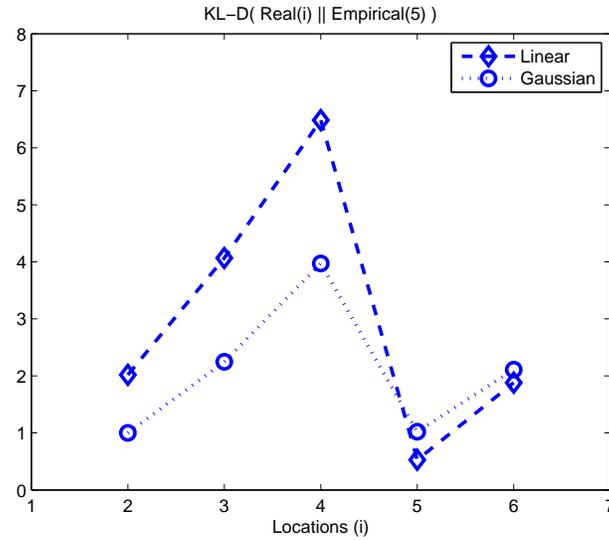


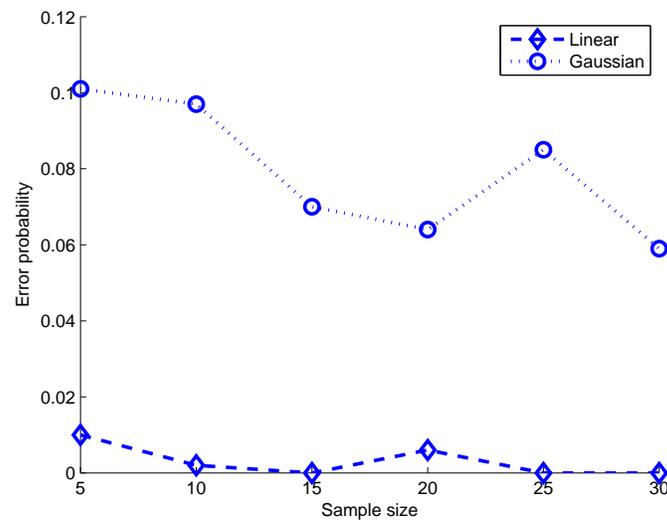
Fig. 9. Layout of the experiment. Node 1 is the clusterhead (listener).

C. Change of signal profiles over time

Consideration of short-term changes such as the effect of people walking around has already been built into our model. With that taken care of, we can consider the signal profiles to be relatively stable. In the long run, some re-profiling would be inevitable. Our experiences indicate



(a) KL distances of the empirical pdfs to the interpolated pdf of location 5 (a clear minimum at 5 is desired).



(b) Error rate of movement detection for location 5.

Fig. 10. Comparison of linear interpolation and the Gaussian model in terms of (a) the KL distance, (b) the error rate in GHT.

that the profile experiences virtually no change for the time scale of one month, but re-profiling may be needed for a period beyond 4 months.

First, to assess the effect of an outdated profile on localization accuracy, we performed a three-part experiment that used a current profile, an one-month-old profile, and a four-month-

TABLE I
KL DISTANCE OF SIGNAL PROFILE DRIFTING

	KL-D Early Nov vs. Late Nov	KL-D July vs. Early Nov
Location 1	0.54	0.57
Location 2	0.68	2.15
Location 3	0.58	1.12
Location 4	0.40	3.19

old profile, respectively. When we used the one-month-old profile the probability of accurately reporting the correct landmark deteriorated by 2.27% compared to the one obtained by using the current profile. On the other hand, when we used the four-month-old profile, the probability of accurately reporting the correct landmark deteriorated by 6.5% compared to the one obtained by using the current profile. Real systems are typically overprovisioned so such a performance deterioration may be tolerable. Of course, the percentage drift in the corresponding probabilities of error was substantially higher but in any case we work with small probabilities of error. Furthermore, even when we make an error, the error is small because the mistaken location is a neighboring landmark in virtually all of the cases.

To gain more insight, we used again the KL distance to measure the drift of signal profiles. Signal profiles of four locations that resemble a neighborhood in our landmark graph were measured in 2010, first in July, then in the beginning of November (i.e., after 4 months), and then at the end of November. Table II shows that the KL distances of signal profiles between pairs of locations are on the scale of 3 to 8. On the other hand, Table I shows that, the signal profiles changed by a KL distance of around 0.5 from the beginning to the end of November, which is negligible compared with the pairwise distances. However, the maximum drift in KL distances from July to November is around 3, which is large enough (compared to the pairwise distances) to affect the localization results.

VI. CONCLUSION

This chapter reviews a set of carefully designed rules and algorithms that underpins a successful two-tier wireless RSSI-based positioning system. Both theoretical and experimental justifications are provided. In addition to working out the details of the design, the experimental

TABLE II
KL DISTANCE BETWEEN LOCATIONS

KL-Distance	Location 1	Location 2	Location 3	Location 4
Location 1	-	4.42	6.44	6.49
Location 2	3.29	-	5.17	6.96
Location 3	6.86	5.48	-	3.41
Location 4	8.01	7.87	3.45	-

validation of the formal pdf interpolation and that of the single-clusterhead-based movement detection are valuable pieces of information to practitioners.

Our approach to the localization problem is formal. The coverage area is modeled as a landmark graph, whose nodes and edges are locations of interest. The signal profiles between locations are processed using a rigorous pdf interpolation technique. The decision rules are GLRT and GHT, which provide error probability estimates. And these error probability estimates are used in a MILP formulation to come up with optimal system deployment.

Furthermore, these formal approaches are validated by experiments. We showed that the overall accuracy of our system compared favorably against other state-of-the-art methods using the same low-cost hardware (although we recognize that the conditions of experimentation are not strictly comparable). The idea of combining localization and movement detection to improve system efficiency is also shown to be practical, i.e., an efficient deployment solution that satisfies the needs of both localization and movement detection does exist. In addition to validating the entire system, we also zoomed in to the particularly interesting pdf interpolation technique. In these more focused experiments, we not only showed that this theoretically appealing interpolation technique actually works, but also showed that the reason why we expected it to work is also valid, i.e., the interpolated pdf carries more information (as measured by the KL distance) for distinguishing between locations.

It is worth noting that the use of advanced decision theory does not make our approach more difficult to use; quite the contrary. We do record full pdf information and use formal optimization to decide clusterhead placement, which makes the initialization of the system a little more complex than other methods. However, after implementing these algorithms, deploying the

system is very easy exactly because the approach is “formal” at every step. There is little need for trial-and-error type of adjustments. In short, it is our hope to contribute to the rational decision making process in constructing a localization system.

REFERENCES

- [1] R. M. Estanjini, Y. Lin, K. Li, D. Guo, and I. C. Paschalidis, “Optimizing warehouse forklift dispatching using a sensor network and stochastic learning,” *IEEE Transactions on Industrial Informatics*, vol. 7, no. 3, pp. 476–486, 2011.
- [2] J. Caffery and G. Stuber, “Subscriber location in CDMA cellular networks,” *IEEE Transactions on Vehicular Technology*, vol. 47, no. 2, pp. 406–416, 1998.
- [3] A. Weiss, “On the accuracy of a cellular location system based on RSS measurements,” *IEEE Transactions on Vehicular Technology*, vol. 52, no. 6, pp. 1508–1518, 2003.
- [4] R. Want, A. Hopper, V. Falcao, and J. Gibbons, “The active badge location system,” *ACM Transactions on Information Systems*, vol. 10, pp. 91–102, January 1992.
- [5] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan, “The cricket location-support system,” in *Mobile Computing and Networking*, 2000, pp. 32–43. [Online]. Available: “<http://citeseer.nj.nec.com/priyantha00cricket.html>”
- [6] S. Tarzia, P. Dinda, R. Dick, and G. Memik, “Indoor localization without infrastructure using the acoustic background spectrum,” in *Proceedings of the 9th international conference on Mobile systems, applications, and services (MobiSys)*, 2011, pp. 155–168.
- [7] I. Guvenc, C.-C. Chong, and F. Watanabe, “NLOS identification and mitigation for UWB localization systems,” in *Wireless Communications and Networking Conference (WCNC 2007)*, March 2007, pp. 1571–1576.
- [8] N. Patwari and S. Kasera, “Robust location distinction using temporal link signatures,” in *Proceedings of the 13th annual ACM international conference on Mobile computing and networking (MobiCom '07)*, 2007, pp. 111–122.
- [9] P. Bahl and V. Padmanabhan, “RADAR: An in-building RF-based user location and tracking system,” in *Proceedings of the IEEE INFOCOM Conference*, Tel-Aviv, Israel, March 2000.
- [10] K. Lorincz and M. Welsh, “Motetrack: A robust, decentralized approach to RF-based location tracking,” *Springer Personal and Ubiquitous Computing, Special Issue on Location and Context-Awareness*, pp. 1617–4909, 2006.
- [11] K. Kaemarungsi and P. Krishnamurthy, “Modeling of indoor positioning systems based on location fingerprinting,” in *Proceedings of the IEEE INFOCOM Conference*, 2004.
- [12] J. Hightower, R. Want, and G. Borriello, “SpotON: An indoor 3d location sensing technology based on RF signal strength,” University of Washington, Department of Computer Science and Engineering, Seattle, WA, UW CSE 00-02-02, February 2000.
- [13] P. Castro, P. Chiu, T. Kremenek, and R. Muntz, “A probabilistic location service for wireless network environments,” in *Proceedings of Ubicomp*. Atlanta, GA: ACM, September 2001.
- [14] J. Krumm and E. Horvitz, “Locadio: Inferring motion and location from Wi-Fi signal strengths,” in *First Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services (MobiQ 2004)*, 2004, pp. 4–13.
- [15] K. Chintalapudi, A. Padmanabha Iyer, and V. Padmanabhan, “Indoor localization without the pain,” in *Proceedings of the sixteenth annual international conference on Mobile computing and networking*. ACM, 2010, pp. 173–184.

- [16] N. Patwari, A. O. Hero, M. Perkins, N. S. Correal, and R. J. O’Dea, “Relative location estimation in wireless sensor networks,” *IEEE Transactions on Signal Processing*, vol. 51, no. 8, pp. 2137–2148, 2003.
- [17] K. Yedavalli, B. Krishnamachari, S. Ravula, and B. Srinivasan, “Ecolocation: A sequence based technique for RF-only localization in wireless sensor networks,” in *The Fourth International Conference on Information Processing in Sensor Networks*, Los Angeles, CA, April 2005.
- [18] I. C. Paschalidis and D. Guo, “Robust and distributed localization in sensor networks,” in *Proceedings of the 46th IEEE Conference on Decision and Control*, New Orleans, Louisiana, December 2007, pp. 933–938.
- [19] —, “Robust and distributed stochastic localization in sensor networks: Theory and experimental results,” *ACM Trans. Sensor Networks*, vol. 5, no. 4, pp. 34:1–34:22, 2009.
- [20] M. A. Youssef, “Collection about location determination papers available online,” 2008, http://www.cs.umd.edu/~moustafa/location_papers.htm.
- [21] M. Kjærsgaard, “A taxonomy for radio location fingerprinting,” *Lecture Notes in Computer Science: Location-and Context-Awareness*, vol. 4718, pp. 139–156, 2007.
- [22] C. Chang and A. Sahai, “Cramer-Rao-type bounds for localization,” *Journal on Applied Signal Processing*, vol. 2006, p. 113, 2006.
- [23] S. Gezici, “A survey on wireless position estimation,” *Journal Wireless Personal Communications*, vol. 44, no. 3, pp. 263–282, 2008.
- [24] F. H. Bursal, “On interpolating between probability distributions,” *Applied Mathematics and Computation*, vol. 77, pp. 213–244, 1996.
- [25] S. Ray, W. Lai, and I. C. Paschalidis, “Statistical location detection with sensor networks,” *Joint special issue IEEE/ACM Trans. Networking and IEEE Trans. Information Theory*, vol. 52, no. 6, pp. 2670–2683, 2006.
- [26] T. Cover and J. Thomas, *Elements of Information Theory*. Wiley, New York, 1991.
- [27] T. Cover and J.A.Thomas, *Elements of Information Theory*. John Wiley & Sons, Inc, 1991.
- [28] W. Hoeffding, “Asymptotically optimal tests for multinomial distributions,” *Ann. Math. Statist.*, vol. 36, pp. 369–401, 1965.
- [29] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*, 2nd ed. NY: Springer-Verlag, 1998.
- [30] M. Daskin, *Network and Discrete Location*. Wiley, New York, 1995.
- [31] F. Özsoy and M. Pinar, “An exact algorithm for the capacitated vertex p-center problem,” *Computers and Operations Research*, vol. 33, no. 5, pp. 1420–1436, 2006.