A Probabilistic Clustering-Based Indoor Location Determination System

Moustafa A. Youssef, Ashok Agrawala, A. Udaya Shankar, Sam H. Noh[1]
Department of Computer Science and UMIACS
University of Maryland,
College Park, MD 20742
{moustafa,agrawala, shankar, noh}@cs.umd.edu

CS-TR-4350 and UMIACS-TR-2002-30
March 22, 2002

ABSTRACT

We present an indoor location determination system based on signal strength probability distributions for tackling the noisy wireless channel and clustering to reduce computation requirements. We provide two implementation techniques, namely, Joint Clustering and Incremental Triangulation and describe their tradeoffs in terms of location determination accuracy and computation requirement. Both techniques have been incorporated in two implemented context-aware systems: User Positioning System and the Rover System, both running on Compaq iPAQ Pocket PC's with Familiar distribution of Linux for PDA's. The results obtained show that both techniques give the user location with over 90% accuracy to within 7 feet with very low computation requirements, hence enabling a set of context-aware applications.

1 Introduction

As ubiquitous computing becomes more popular, the need for context-aware applications increases. One of the most important contextual information is the user location, with which the system can provide location-specific information and services. There have been many systems that provide context-aware services to the users based on their locations [1] including automatic call forwarding to the user based on his current location, helping shoppers through the stores based on their location, providing information to the tourist about his current location and office assistant that interacts with visitors and manages the office owner's schedules.

Many systems over the years have tackled the problem of determining and tracking user position. Examples include GPS [2], wide-area cellular-based systems [3], infrared-based systems [4][5], magnetic tracking systems [6], various computer vision systems [7], physical contact systems [8], and radio frequency (RF) based systems [9]-[14]. Of these, the class of RF-based systems that use an underlying wireless data network [12]-[14], such as 802.11, to estimate user location has gained attention recently, especially for indoor application. Unlike infrared-based systems, which are limited in range, RF-based techniques provide more ubiquitous coverage and do not require additional hardware for user location determination, thereby enhancing the value of the wireless data network.

We present an RF-based location determination system that achieves better positioning accuracy than existing systems with low computation overhead. Given an indoor region covered by multiple access points, the system collects access point signal strengths at various locations and constructs a histogram-based **radio map**. Then given a new signal strength reading from an arbitrary location, the system estimates the closest map location corresponding to the arbitrary location. The estimation procedure has two key features:

- It uses the histogram distributions (rather than just the mean) to enhance accuracy and tackle the noisy nature of the wireless channels.
- It uses clustering of map locations to reduce the computation requirements. We present two techniques: The Joint Clustering (JC) technique that uses explicit clustering and the Incremental Triangulation (IT) technique that features implicit clustering.

We have evaluated the system in an indoor space of corridors spanning a 20,000 square foot floor of a building. Results obtained show that using the signal strength values collected from the access points, both the Joint Clustering and Incremental Triangulation techniques give the user location with over 90% accuracy to within 7 feet with very low computation requirements.

The closest related work to ours in the area of indoor location determination are the RADAR system [12], the CMU system proposed in [13], and the Nibble system from UCLA [14]. Our approach differs from RADAR and the CMU approach in that we use probabilistic ranking and clustering to better handle the noisy wireless channel and to reduce the search space. While our approach and Nibble are similar in some ways, there are significant differences: (a) we store only the marginal distribution of each access point, rather than the joint distribution of all the random variables of the system, thereby reducing the computational cost and significantly enhancing system scalability; (b) we use the received signal strength instead of the signal to noise ratio (SNR) because the former is a stronger function of location [12]; (c) we have a much finer quantization of the received signal strength, thereby achieving better accuracy; (d) we use clustering to control the computational cost. A detailed comparison of our approach with these approaches and other

[1] Also with the School of Information and Computer Engineering, Hongik University, Mapo-Gu Sangsoo Dong 72-1, Seoul, Korea

approaches for location determination is presented in Section 7.

The rest of the paper is organized as follows. Section 2 presents our general architecture for location determination systems. Section 3 presents the details of radio map construction and location estimation with the Joint Clustering technique. Section 4 presents the details of location estimation with the Incremental Triangulation technique. In Section 5, we describe the evaluation of the techniques in the indoor space and the obtained results. Section 6 describes two applications implementing the general architecture for location determination, incorporating the JC and IT techniques as their location determination algorithms. Section 7 surveys related work and compares the new techniques with previous RF-based location determination approaches that do not require additional hardware. Finally Section 8 concludes the paper.

2 Location Determination System Architecture

Figure 1 shows our location determination system architecture. The hardware layer covers mobile devices, such as laptops and handhelds, and fixed devices that need location information (e.g., for automatic configuration). All these devices are equipped with wireless cards. The operating system layer includes the operating systems running on the devices. The device driver interacts with the wireless card to collect the signal strength values from the access points in range. The Location Determination System layer runs the location determination algorithm, e.g. the JC algorithm that uses the signal strength values to estimate the user location. A wireless API provides, in a device driver-independent way, the Location Determination System layer with a method to get the required information from the driver, such as the access point MAC addresses and received signal strengths. In the same way, a Location API provides the user application with the device's current position in a way independent of the location determination algorithm.

In Section 6, we present 2 examples on implementing this architecture. In the next two sections, we describe the JC technique and the IT technique, respectively, which are part of the Location Determination System layer.

3 The Joint Clustering Technique

The Joint Clustering technique is based on two main functions: (a) estimating the joint distribution of the signal strength values received from access points at each location and (b) grouping the locations into clusters. The joint distributions are used to find the most probable location given the observation sequences of signal strength values. The JC technique also performs location clustering, by grouping locations that have a common feature, to reduce the size of the search space and, hence, reducing the computational requirements of the algorithm. Therefore, the operation of the JC technique can be divided into two phases: (a) offline phase, in which we perform the joint distribution estimation and locations clustering and (b) location determination phase, in which we run the location determination technique to infer the user location.

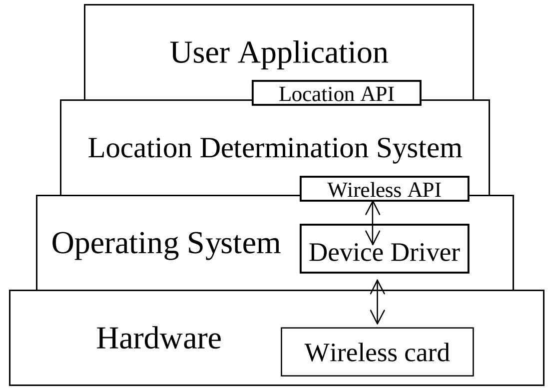

Figure 1. Location determination system architecture

Below, we introduce some notations and then describe the two phases in more details.

We define the following notations:

- | . | denotes the number of elements in a given set or sequence.
- '*' denotes all possible values for a given index.
- For any sequence *x*, *x(i)* denotes the i^{th} element of *x*.
- *SS* is the discrete signal strength space.
- *TrLocs* is a set of locations for which we build the radio map.
- *TsLocs* is a set of locations for which we test the performance of the algorithms.
- $TrSamples_{l,a}$ is a sequence of training signal strength values at location $l \in TrLocs$ from access point *a*.
- $TsSamples_{l,a}$ is a test sequence of signal strength values for a location $l \in TsLocs$ from access point *a*.
- $TrAP_l = \{a: TrSamples_{t,a}(n)>0 \text{ for some } n\}$ is the set of access points heard in the training set at location *l*.
- $TsAP_l = \{a: TsSamples_{t,a}(n)>0 \text{ for some } n\}$ is the set of access points heard in the test sequnce at location *l*.
- $Hist_{l,a}$ is the normalized histogram for signal strength values at location $l \in TrLocs$ from access point $a \in TrAP_l$.

By definition, $Hist_{l,a}(s) = \dfrac{|\{n: TrSamples_{l,a}(n) = s\}|}{|TrSamples_{l,a}|}$ for any $s \in SS$.

- *SortedAP(l, n, AP, Samples)* is the function that sorts the set of access points in *AP* at location *l*, according to the average signal strength value calculated from *Samples*, and returns the first *n* elements of the sorted *AP* set as a sequence. If |*AP*| is less than *n*, the function returns the sorted *AP* set as a sequence.
- *Cluster(key, q)* is a function that returns $\{l \in TrLoc : SortedAP(l, q, TrAP_l, TrSamples_{l,a}) = key\}$. The parameter *key* represents a common set of access points that is shared between all the locations in the cluster.

3.1 Offline Phase

During the offline phase we perform two tasks: joint probability distribution estimation and location clustering.

3.1.1 Estimating the Joint Signal Strength Distribution

At each location in the set of training locations, we store a model for the joint probability distribution of the access points at this location. Therefore, our radio map is stored as a collection of models for joint probability distributions.

The problem of estimating the joint distributions can further be divided into three sub-problems:

1- How to choose a value (*k*) for the dimension of the joint distribution?
2- Which *k* access points, from the set of access points covering a certain location, to choose to be included in the joint distribution?
3- How to estimate the joint distribution between the chosen *k* access points?

In determining the best value for *k* we need to take into account 2 factors: (a) as *k* increases, the process of estimating the joint probability distribution (sub-problem 3) becomes more complex and (b) we need a value for *k* such that all locations are covered by at least *k* access points most of the time.

The second factor is important because at the online phase, we get a number of samples from some of the access points and some of the access points that cover a certain location may be missing from the samples due to the noisy nature of the wireless channel and hence the number of access points covering a location is varying with time. The second factor lessens the affect of variability in the number of access points and hence should lead to better accuracy. Typical values for the parameter *k* can be found in Section 5.

The solution to sub-problem 2 is related to the solution of sub-problem 1. If the number of access points covering a location is varying with time, which access points should we choose? Intuitively, we should choose the access points that appear most of the time in the samples. We did some analysis of the data and found that the access points with the largest signal strength are those that appear in most of the samples. This is expected as the access points with weak signal strength are less probable to be heard than the ones with the strong signal strength.

To summarize, for a given location $l \in TrLocs$, we choose the first *k* access points from the set $TrAP_l$ when sorted according to the average signal strength values, i.e. we use $SortedAP(l, k, TrAP_l, TrSamples_{l,*})$.

The problem of estimating the joint probability distribution can be done in different ways with different accuracy levels. The problem can be stated as: given *k* access points $AP_1 \ldots AP_k$, we want to estimate $P(AP_1 = s_1, AP_2 = s_2, \ldots, AP_k = s_k)$ where s_i is a signal strength value from AP_i. One good way to estimate this joint distribution is to use the Maximum Likelihood Estimation (MLE) method which estimate the joint probabilities as follows:

$$P(AP_1 = s_1, AP_2 = s_2, \ldots, AP_k = s_k) = \frac{Count(s_1, s_2, \ldots, s_k)}{SizeofTrainingData} \quad (1)$$

i.e. the number of times that the signal strength values tuple $(s_1, s_2, \ldots, s_k)$ appeared in the entire training set divided by the size of the training set.

The problem of this approach is that it requires a large training set to obtain good estimate of the joint distribution and the required size increases exponentially with *k*. For example, if we have 3 access points each having a range of 11 signal strength values, then the number of different possible tuples for the joint distribution is $11^3 = 1331$, and hence the training data size cannot be less than this number (actually it must be much bigger).

Since our goal was to use a method that gives a good accuracy and, at the same time, requires reasonable amount of training data and computational power, this approach can only be used with small values of *k*, which may affect the technique accuracy. Instead, we chose to make an approximation that the access points are independent. In this case, the problem of estimating the joint probability distribution becomes the problem of estimating the marginal probability distributions as:

$$P(AP_1 = s_1, AP_2 = s_2, \ldots, AP_k = s_k) = P(AP_1 = s_1).P(AP_2 = s_2)\ldots P(AP_k = s_k) \quad (2)$$

since the random variables $AP_1, \ldots, AP_k$ are independent. For a given location $l \in TrLocs$, $P(AP_i = s_i) = Hist_{l,AP_i}(s_i)$.

Figure 2 gives a typical example of the signal strength normalized histogram from an access point.

This approach reduces the size of the training set required. Using the same example as before, the number of distinct values for each access point is 11, and a small size training set can be used to estimate the marginal distributions. The independence assumption has other advantages as will be described in the discussion section.

3.1.2 Locations Clustering

To reduce the computation overhead, we group the locations into clusters according to the access points that cover the locations. The problem can be stated as follows: Given a location *l*, we want to determine the cluster to which *l* belongs.

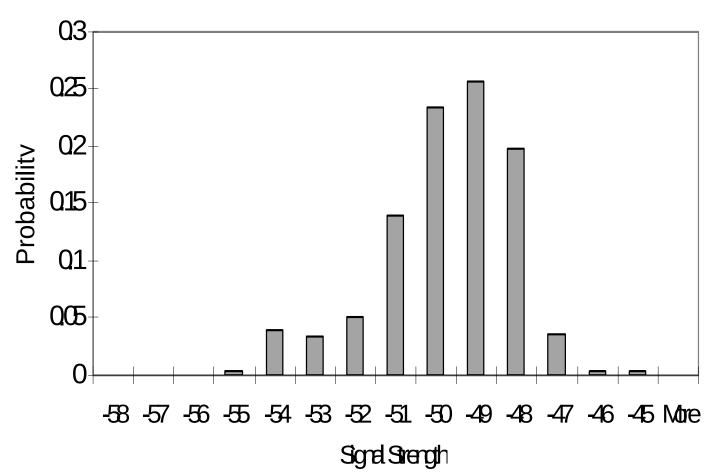

Figure 2. An example of a histogram of the signal strength of an access point

The most obvious way to do clustering is to group locations according to the access points that cover them. i.e. two locations l_1 *and* l_2 are placed in the same cluster *iff* $TrAP_{l_1} = TrAP_{l_2}$. However, this approach for clustering has problems when applied in a real environment. Since the wireless channel is noisy, an access point may be missing from some of the samples and, therefore, using the entire set of access points that cover a location for clustering may fail to find the correct cluster due to the missing access point.

Instead of using the entire set ($TrAP_l$) that covers a location l for clustering, we use a subset of this set containing only q elements and the problem becomes: Given a number q, we want to put all the locations that share q access points in one cluster. Therefore, we have 2 sub-problems:

1- How to determine the value of q ?
2- Which q access points to choose for clustering ?

For the first sub-problem, we need to choose q such that all locations are covered by at least q access points most of the time. This factor is important due to similar reasons as in the discussion of the choice of a value for the parameter k. This suggests that the value of q should be less than or equal to $min(|TrAP_l|)$ for all $l \in TrLocs$. Moreover, we need a value for q that distributes locations evenly between the clusters to reduce the required computations. Determining the value for q is discussed in the Section 5.

For sub-problem 2, we chose to use the q access points with the largest signal strength values at each location, again for similar reasons as in the previous section.

During the data analysis we found that, at some locations, the order of the access points with the largest signal strength values changes when the signal strength values from these access points are near to each other, especially when we take a small number of samples at the online phase. Therefore, we chose to treat the q access points as a set and not as an ordered tuple. For example, if $q=2$ and the two access points with the largest and second largest signal strength value at location l_1 are (AP_1, AP_2) respectively, and (AP_2, AP_1) for another location l_2, then we place location l_1 and location l_2 in the same cluster regardless of the order of the access points.

To summarize, for a given location $l \in TrLocs$, we use the set $\{a$: a is in Sorted$AP(l, q, TrAP_l, TrSamples_{l,*})\}$ to determine the cluster to which l belongs.

We want to emphasize here that the values of the parameters k (dimension of the joint distribution) and q (number of access points to use in clustering) are independent. For example, we can use one access point ($q=1$) for clustering and use a 3-dimenstional ($k=3$) joint distribution.

The next subsection describes the location determination phase.

3.2 Location Determination Phase

The general idea of what happens during the location determination phase is as follows: we get samples from some access points at an unknown location. We use the q access points with the largest signal strength values to determine one cluster to search within for the most probable location. We, then, use Baye's theorem to estimate the probability of each location within the cluster given the observed sample sequences and the radio map built during the offline phase. The most probable location is reported as the estimated user location.

The above algorithm works assuming ideal wireless channel. However, for a practical environment, we need to tackle two problems:

1- The number of access points in a test sample at a location t, $|TsAP_t|$, may be less than q.
2- $|TsAP_t|$ may be less than k, the dimension of the joint distribution.

We first use an example to demonstrate the first problem and our approach to solve it. Assume that number of access points to use in clustering, q, *was set to 3* and assume further that during the location determination phase we got samples from two access points only: AP_1 *and* AP_2. The problem here is that we cannot find a cluster whose key is $\{AP_1, AP_2\}$. To solve this problem, we search for all clusters whose key has $\{AP_1, AP_2\}$ as a subset. We use the union of all the locations in these clusters as our target locations set.

More formally, we define the set of target locations as: $TargetLocs = \bigcup_{SortedAP(t,q,TsAP_t,TsSamples_{t,*}) \subseteq s} Cluster(s)$. The set of target locations reduces to the locations within one cluster if $|TsAP_t|$ is greater than or equal to q.

For the second similar problem, we use the same approach to solve it by reducing the dimension of the joint distribution to $min(k, |TsAP_t|)$.

The only thing that remains to be explained is how to use Baye's theorem to calculate the most probable location out of the target locations set given the observation sequences $TsSamples_{l,*}$. We want to find $l \in TargetLocs$ such that $P(l/TsSamples_{t,a})$ for all $a \in SortedAP(t, k, TsAP_t, TsSamples_{t,*})$, is maximized. i.e. we want

$$\arg\max_l [P(l / TsSamples_{t,a})] \quad (3)$$

Using Baye's theorem, this can be rewritten as:

$$\arg\max_l [P(l / TsSamples_{t,a})] = \arg\max_l [\frac{P(TsSamples_{t,a} / l)P(l)}{P(TsSamples_{t,a})}]$$

(4)

since $P(TsSamples_{t,a})$ is constant for all l, we can rewrite equation (4) as:

$$\arg\max_l [P(l / TsSamples_{t,a})] = \arg\max_l [P(TsSamples_{t,a} / l)P(l)]$$

(5)

$P(l)$ can be determined from the user profile based on the fact that if the user is at a given location, it is more probable that he will be at an adjacent location in the future. If the user profile information is not known, or not used, then we can assume that all the locations are equally likely and the term $P(l)$ can be factored out from the maximization process. Equation (4) becomes:

$$\arg\max_{l}[P(l/TsSamples_{t,a})] = \arg\max_{l}[P(TsSamples_{t,a}/l)]$$
(6)

The remaining term $P(TsSamples_{t,a}/l)$ can be calculated by using:

$$P(TsSamples_{t,a}/l) = \prod_{n=1}^{|TsSamples_{t,a}|} \prod_{a\in SortedAP(t,k,TsAP_t,TsSamples_{t,*})} Hist_{l,a}(TsSamples_{t,a}(n))$$
(7)

assuming independence of access points and samples. The details of the algorithm are given in Figure 3.

The next section presents a discussion of the Joint Clustering technique.

3.3 Discussion

Many operations of the algorithm can be optimized: For example, we do not need to calculate the actual average signal strength of each access point. All we need is just to calculate the sum of the signal strength values because we need to compare the averages and the number of samples is constant. The sorting operations in the algorithm do not take a long time. Sorting the access points according to the average signal strength takes a short time as the typical number of access point at any location is 4 (average number for the specific experiment we performed was 4 access points per location). The independence assumption helps reduce the computations required by converting the multiplications to additions, if we use the logarithms of the probabilities instead of the probabilities themselves. The clustering performed by the algorithm makes the list candidate locations typically small, so sorting the list of candidate locations according to their probabilities should be a fast process.

The memory requirements of the algorithm are limited. If the average number of access points per location is 4 and average range of each access point is 11 distinct values, then for each location we need to store 11* 4 parameters, corresponding to the histograms of each access point, which is a small number. We could instead approximate the histogram by a continuous distribution, e.g. a Normal distribution, and save only the mean and variance of the distribution for each access point. However, this approximation affects the accuracy of the system and the saving of the memory requirement does not justify it.

4 The Incremental Triangulation Technique

The JC technique introduced in the previous section calculates the probability of a location using *k* access points all at the same time, using *k* operations per sample. The Incremental Triangulation technique uses a different approach to calculate the probabilities. It tries to use the access points incrementally, one after the other, until it can estimate the location with certain accuracy, using a predetermined threshold. As we will explain, the IT technique performs implicit clustering at multi-levels leading to a more reduced search space than the JC approach, and hence fewer number of operations, on the average, per sample. However, treating each access point incrementally, instead of using the joint distribution, leads to the loss of some information and thus one should expect that the accuracy of the IT should be lower than the JC technique.

The IT technique works in two phases, in the same way as the JC technique: (a) offline phase, in which we estimate the signal strength distribution from each access point and (b) location determination phase, in which we run the location determination technique to infer the user location.

Note that in the IT technique, we do not need to do clustering in the offline phase as clustering is performed in an implicit way as will be explained in the location determination phase.

- Input:
  - t: Unknown user location.
  - q: Number of access points to use in clustering.
  - k: Number of access points in the joint distribution.
  - *TrLocs*: Set of locations in the radio map.
  - $Hist_{l,a}$: Histogram of each access point at each location (radio map)
  - *Cluster*: Clustering function.
  - $TsSamples_{t.a}$: Test sequence at location *t*.
  - $TsAP_t$: Set of access points heard in the test sequence at location *t.*
- Output:
  - The most probable location in *TrLocs* assigned to t.

1. Set $TargetLocs = \bigcup_{SortedAP(t,q,TsAP_t,TsSamples_{t,*})\subseteq s} Cluster(s)$ ·
2. If *TargetLocs* is empty then set *TargetLocs= TrLocs.*
3. Calculate $X = \{P(TsSamples_{t,a}/l) = \prod_{n=1}^{|TsSamples_{t,a}|} \prod_{a\in SortedAP(t,k,TsAP_t,TsSamples_{t,*})} Hist_{l,a}(TsSamples_{t,a}(n)), \forall l \in TargetLocs\}$ ·
4. Sort the elements of *X* in a descending order. Let *OrderedL* be the sequence of *TargetLocs* corresponding to the sorted *X*.
5. Assign OrderedL*(1)* to *t*.

Figure 3: Detailed inference algorithm for the Joint Probability Distribution with Explicit Clustering technique.

In the rest of this section, we describe the 2 phases followed by introducing the implicit clustering performed by the algorithm in the online phase, and finally a discussion of the algorithm.

4.1 *Offline Phase*

In this phase, we estimate the discrete distribution for each access point at a given location using the histogram and store this information in the radio map. So the radio map for the JC technique and the IT technique are identical. Recall that in the JC case we use the marginal distribution of each access point to approximate the joint distribution.

4.2 *Location Determination Phase*

We start with an example to motivate the algorithm. Given a sequence of observations from each access point, we start by sorting the access points in a descending order according to the average signal strength values received from them. For the first access point, the one with the strongest average signal strength, we calculate the probability of each location in the radio map set (*TrLocs)* given the observation sequence from this access point alone. This will give us a set of candidate locations (locations that have non-zero probability). If the probability of the most probable location is "significantly" higher, according to a measure defined in the algorithm, than the probability of the second most probable location, we return the most probable location as our location estimate, after consulting only one access point. If this is not the case, we go to the next access point in the sorted access point list. For this access point, we repeat the same process again, but only for the set of candidate locations obtained from the first access point. This process of calculating the probabilities and determining the significance of the most probable location is repeated incrementally, for each access point in order, until the location can be estimated or all access points are consulted. In the latter case, the algorithm returns the most probable location in the candidate list that remains after consulting all the access points.

It should be now clear why we call our approach the Incremental Triangulation technique. The reason is that we start by a set of candidate locations using the first access point and reduce this set using other access points iteratively. In contrast, the standard triangulation approach starts by an infinite number of locations on a circle and reduces this number to 2 points using another circle and finally reduces these two points to only one point using a third circle (assuming every thing is perfect). However, typically this is done by solving a set of nonlinear equations and not in an iterative manner.

Figure 4 shows the details of the algorithm.

4.3 *Implicit Clustering*

The algorithm performs implicit clustering using the access points. Starting with the access point that has the strongest average signal strength value, the algorithm restricts itself to calculating the probability for locations inside the range of this access point only, as those are the locations that have histograms for this specific access point. Therefore, depending on the access point that has the strongest average signal strength value, the algorithm examines a different set of locations in its initial step.

Moreover, in the iterative process, the algorithm checks only locations that lie in the coverage area of the first access point and then the locations within those locations that lie in the coverage area of the second access point and so on, leading to a multi-level clustering. This multi-level clustering approach reduces the search space significantly at each iteration, and hence leads to less computation.

4.4 *Discussion*

The parameter *Threshold is used* to determine if the information obtained from consulting an access point is significant enough to make a judgment or not. The value of this parameter ranges from 0 to 1. A value of 0 leads to consulting only one access point, reducing the algorithm accuracy while a value of 1 leads to consulting the entire set of access points at a given location, and hence, increased accuracy.

We use the parameter *Window* in the algorithm to select a subset of all the candidate locations after consulting the first access point, if the set of candidate locations is too large.

The *NAP* parameter is used to set a maximum on the number of access points consulted by the algorithm. The max number of access points parameter is important to see how the technique will perform if the number of access points is limited. Section 4 provides more detailed analysis of the effect of the parameters on performance.

Sorting the access points in a descending order according to the average signal strength has an intuitive sense for the IT technique. We want to sort the access points according to the amount of information we can get from each of them. Using information theory concepts, the access point that has the most variability in its signal strength values should give us the maximum amount of information. From the analysis of the data collected, we found that the access point that has the greatest variability is the one that has the strongest average signal strength. Also, as we mentioned before, the access points that have the largest signal strength appear more often in the samples than the access points with weak signal strength, as will be explained in Section 5, and hence taking the decision based on the access points with the strongest signal strength should give better results.

The implicit clustering performed by the technique reduces the required computations. In addition, using an iterative approach can make the algorithm terminate without examining the entire set of access points, again reducing the required computation.

Comparing the JC technique with the IT technique one expects that the former should lead to better accuracy as it takes into account more information in one step instead of iteratively going through the different access points. However, the computation requirement of the Incremental Triangulation approach may be less as at each iteration we perform the computation for one access point compared to

- Input:
  - t: Unknown user location.
  - *Window*: Window size parameter.
  - *Threshold*: Stopping threshold.
  - *NAP*: Maximum number of access points to be consulted.
  - *TrLocs*: Set of locations in the radio map.
  - $Hist_{l,a}$: Histogram of each access point at each location (radio map)
  - $TsSamples_{t.a}$: Test sequence at location t.
  - $TsAP_t$: Set of access points heard in the test sequence at location t.
- Output:
  - The most probable location.

1. Set $OrderedAP_t$= $SortedAP$(t, ∞, $TsAP_t$, $TsSamples_{t,*}$)
2. Let a= $OrderedAP_t(1)$. Calculate $X = \{P(TsSamples_{t,a} / l) = \prod_{j=1}^{|TsSamples_{t,a}|} Hist_{l,a}(TsSamples_{t,a}(j)), l \text{ is in } TrLocs_l\}$.
3. Sort the elements of X in a descending order. Let *OrderedL* be the sequence of *TrLocs* corresponding to the sorted X.
4. Let $Confidence = \frac{X(1) - X(2)}{X(1)}$.
5. If *Confidence* > *Threshold*, assign Ordered*L(1)* to t and return.
6. Let N be the number of non-zero elements of X. *Set Window= minimum(Window, N)*.
7. Set *CandidateL* to the first *Window* elements of *OrderedL*.
8. For *Count*= 2 to min (| $TsAP_t$ |, *NAP)*
   9. Let a= $OrderedAP_t(Count)$. Calculate $X = \{P(TsSamples_{t,a} / l) = \prod_{j=1}^{|TsSamples_{t,a}|} Hist_{l,a}(TsSamples_{t,a}(j)), l \text{ is in } CandidateL\}$.
   10. Sort the elements of X in a descending order. Let *OrderedL* be the sequence of *CandidateL* corresponding to the sorted X.
   11. Let $Confidence = \frac{X(1) - X(2)}{X(1)}$.
   12. If *Confidence* > *Threshold*, assign *OrderedL(1)* to t and return.
   13. Let N be the number of non-zero elements of X.
   14. Set *CandidateL* to the first 'N' elements of *OrderedL*.
15. End
16. Assign *OrderedL(1)* to t and return.

Figure 4: Detailed inference algorithm for the IT technique.

the JC technique which performs the computation for all the access points but to locations inside the cluster only.

A detailed comparison of the performance of the two algorithms is given in the next section.

5 Experimental Evaluation

In this section, we discuss the experimental testbed, describe the data collection process, discuss the effect of the parameters of the Joint Clustering and Incremental Triangulation techniques on performance, compare the two proposed techniques and, finally, present the performance evaluation of both techniques under an independent test set.

5.1 Experimental Testbed

We performed our experiment in the south wing of the fourth floor of the Computer Science Department building. The layout of the floor is shown in Figure 5. The wing has a dimension of 224 feet by 85.1 feet. Both techniques were tested in the Computer Science Department wireless network. The entire wing is covered by 12 access points installed in the third and fourth floors of the building.

For building the radio map, we took the radio map locations on the corridors on a grid with cells placed 5 feet apart (the corridors width is 5 feet). We have a total of 110 locations along the corridors. On the average, each location is covered by 4 access points.

5.2 Data Collection and Analysis

According to the general location determination system architecture described in Section 2, we modified the Lucent Wavelan driver for Linux to return all the access points in range associated with the current signal strength value from each access point using the active scanning technique [15] (our driver was the first driver to support this feature under Linux). We also developed a wireless API [15] that interfaces with any device driver that supports the wireless extensions [16]. The device driver and the wireless API have been available for public download and have been used in other wireless research.

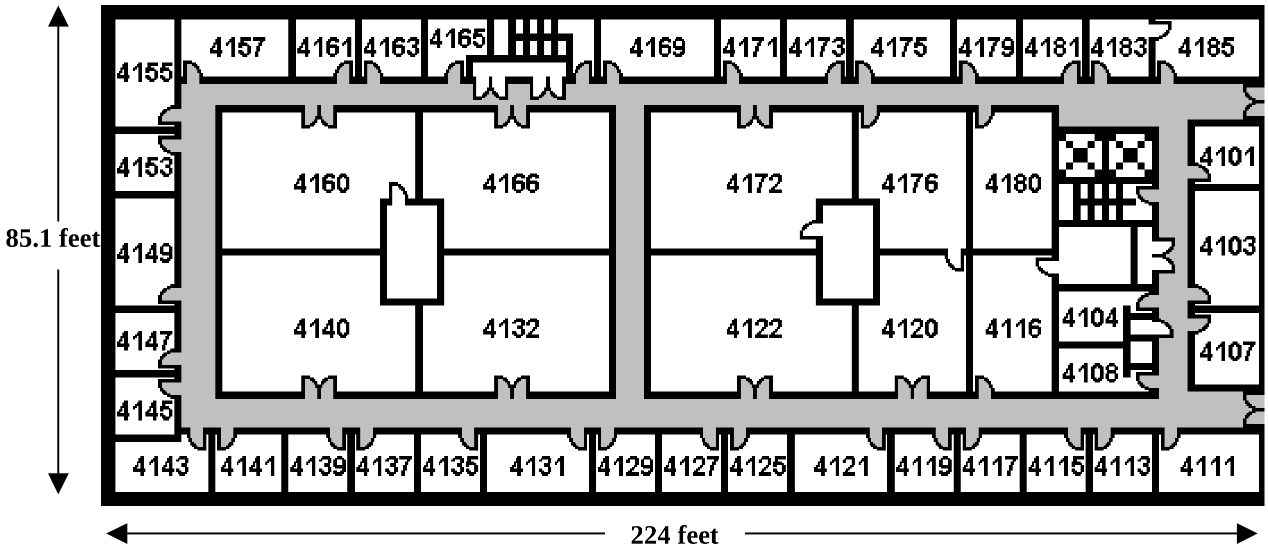

Figure 5: Plan of the south wing of the 4th floor of the Computer Science Department building where the experiment was conducted. Readings were collected in the corridors (shown in gray).

Using the device driver and the API, we collected 300 samples at each location, one sample per second. We divided this data at random into two sets: training set and development test set. The training set constituted 80% of the 300 samples and was used to estimate the distribution of each access point at each location using the method previously described. The development test set constituted the remaining 20% and we used it to estimate the initial performance of the algorithms and tune the models parameters. We also used an **independent** test set, different from the entire training set, to test the performance of the algorithms. Unless otherwise specified, we take the length of the testing sequences to be 3 samples in the rest of the paper.

Both the JC and the IT techniques depend on the property that the access points with the strongest signal strength values are the ones from which we receive samples most of the time. Figure 6 shows the relation between the average signal strength received from an access point and the number of samples we receive from it during a period of 5 minutes (300 samples). The figure shows that the number of samples collected from an access point is a monotonically increasing function of the average signal strength of this access point, which justifies the use of the strongest access points in our techniques.

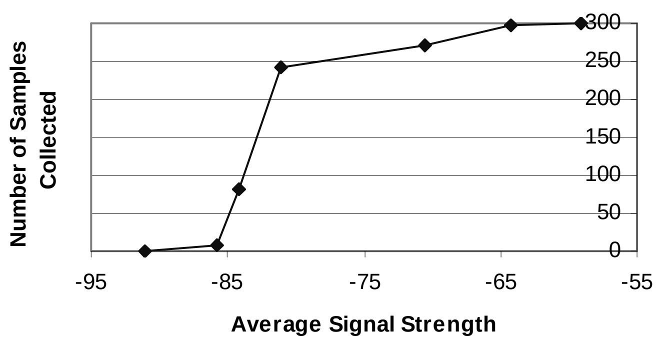

Figure 6: Relation between the average signal strength value from an access point and the number of samples received from it during a 5 minutes interval.

5.3 *Effect of the Parameters on Performance*

Each of the proposed techniques has a number of tunable parameters. In this section, we study the effect of these parameters on the performance of the techniques. In Section 5.3.1, we define the performance measures that will be used to compare the techniques. Section 5.3.2 discusses the effect of the Joint Clustering parameters on the performance measures. The effect of the parameters of the Incremental Triangulation technique on performance is discussed in Section 5.3.3. Finally, the effect of the length of the observation sequence on performance is discussed in Section 5.3.4.

5.3.1 *Performance Measures*

- *Accuracy:* This measure is defined as the percentage of time in which the technique gives the correct location estimate. However, we give the complete CDF of the error in distance in Sections 5.4 and 5.5.
- *Average number of access points consulted for each location estimate*: This measure is important because it shows a practical aspect of the technique. For example, there may be two techniques that give the same accuracy, but one uses information from 3 access points while the other requires information from 5 access points. In such as situation, the first technique should be preferred, as it requires less information and hence less computation.
- *Number of operations per location estimate*: This measure is defined as the total number of operations (additions when using the logarithm of the probabilities) performed for a single location estimate. Combined with the previous measure, this measure indicates the required computation needed for each access point consulted.

This is important in minimizing the computation time, but more so in minimizing the power consumption.

5.3.2 Joint Clustering Technique

The Joint Clustering technique has two control parameters. In this section, we study the effect of these parameters, specifically *k* (dimension of the joint distribution) and *q* (number of access points to use in clustering) on its performance.

We start by showing the effect of changing *q* on the clustering process. For this experiment, we changed the value of *q* from 1 to 4 and calculated the number of clusters, the average size of each cluster, and the standard deviation of the cluster size. This is shown in Figure 7. From the figure we can see that as *q* increases, the number of clusters increases and the average size of each clusters decreases until we reach a saturation point at *q=2.* For the standard deviation, the variation of the size of the clusters decreases until we reach a minimum value, at *q=3,* and it increases again. A small value for the standard deviation means that the sizes of the clusters are more uniform, which is a desirable property. The minimum value at *q=3* can be explained by noting that as *q* increases from 1 to 3, more locations are differentiated into different clusters due to the addition of new access points. When *q* is increased past 3, i.e. *q=4,* different locations start to share the same *4* access points, especially for locations close to each other (recall that the average number of access points per location was 4 in our experiment), and thus the number of locations per cluster starts to deviate from being uniform across clusters leading to increased standard deviation.

Figures 8 and 9 show the effect of parameters *q* and *k* together on performance. From the figures we see that as dimension *k* increases, the accuracy increases as we have more information due to the addition of access points and, due to the same reason, the number of operations required per location estimate increases. As the number of access points used in clustering (*q*) increases, the number of elements per cluster decreases leading to increased accuracy and less number of operations per location estimate.

For the rest of the paper, we chose to take the values of the parameters as *q=3* and *k=4* as these values lead to the best performance for the JC algorithm for our experiment.

5.3.3 Incremental Triangulation Technique

The Incremental Triangulation technique has three parameters: Threshold, Window, and the maximum number of access points (NAP). We consider the effect of each of these parameters on the performance of the technique.

Figure 10 through 12 show the effect of the *Threshold* parameter on performance. For this experiment, we fixed the value of the *Window* parameter at 12 and the value of the *NAP* parameter at 10 (equivalent to examining all the access points that the technique can use, if needed). For small values of the *Threshold* parameter, the decision is taken quickly after examining a small number of access points. As the threshold value increase, more access points are consulted to reach a decision. As the number of access points consulted increases, the number of operations per location estimate increase and the so does the accuracy. It is important to note here that the average number of access points consulted and the average number of operations per location estimate is small which support our previous discussion that the computation requirements of the Incremental Triangulation technique is modest.

The effect of the *Window* parameter on performance is shown in Figures 13 through 15 (*Threshold*= 0.4, *NAP*= 10). A large value of the window parameter leads to a wider set of candidate locations to work on, if the decision cannot be taken based on consulting the first access point alone. Therefore, as the value of the window parameter increases, the set of candidate location from the first access point increases leading to consulting more access points, more operations per location estimate, and better accuracy. However, the average number of operations per location estimate does not increase significantly. This suggests that, in most of the time, the number of candidate locations (i.e. locations with non-zero probability) is small that we do not reach the upper bound provided by the *Window* parameter.

The maximum number of access points parameter is important to see how the technique will perform if the number of access points is limited. The effect of changing the *NAP* parameter is shown in Figures 16 through 18 (*Threshold*= 0.4, *Window*= 12). It is shown from the figures that 3 access points per location are sufficient to obtain good performance (94% accuracy). This makes intuitive sense as the triangulation technique requires 3 access points.

Unless otherwise specified, we chose to take the values of the parameters as *Threshold= 0.4 Window=12* and *NAP=4* as these values lead to the best performance for the IT algorithm in our experiment.

5.3.4 Effect of the Length of the Observation Sequence on Performance

This section studies the effect of increasing the length of the observation sequence, used in location determination phase, on the performance of the algorithms. Figures 19 through 21 show the results.

As the length of the observation sequence increases, the accuracy of the both technique increases till it reaches a saturation point at 3 samples. This is expected as the more samples we have the more information we have about the signal strength distribution and hence better the accuracy.

The number of operations per location estimate increases linearly with the increase of the length of the observation sequence for the JC technique. The curve for the number of operations per location estimate of the IT technique is interesting. The minimum point at 2 samples can be explained by noting that as the length of the observation sequence increases we have 2 conflicting factors: (a) the number of operations performed per access point for each location increases (linearly with the sequence length) and (b) the technique has more information from the samples and

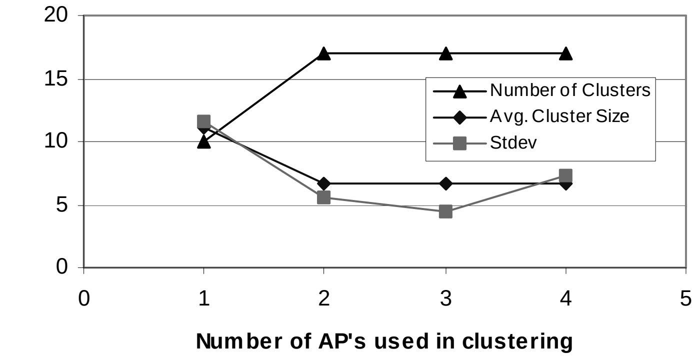

Figure 7: Effect of *q* on the clustering process.

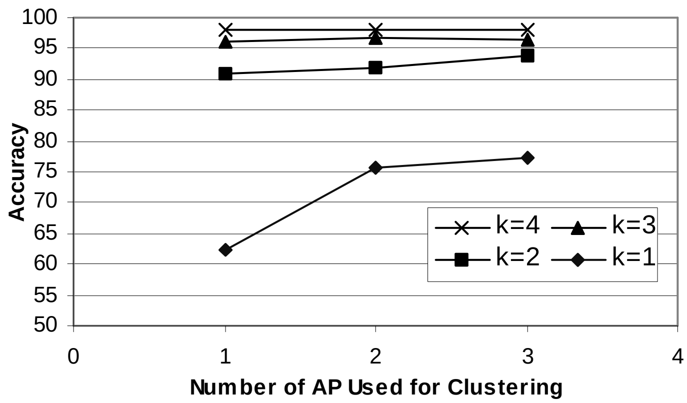

Figure 8: Effect of parameters q and k on accuracy.

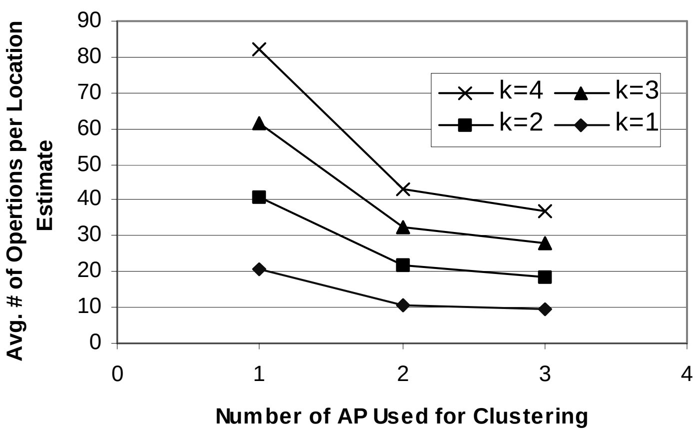

Figure 9: Effect of parameters q and k on average number of operation per estimate.

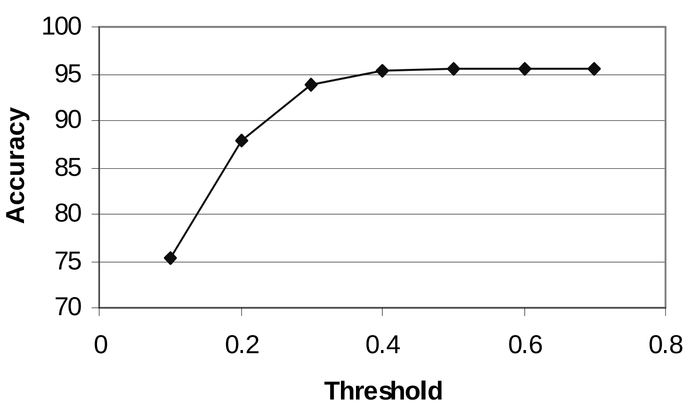

Figure 10: Effect of the Threshold parameter on accuracy.

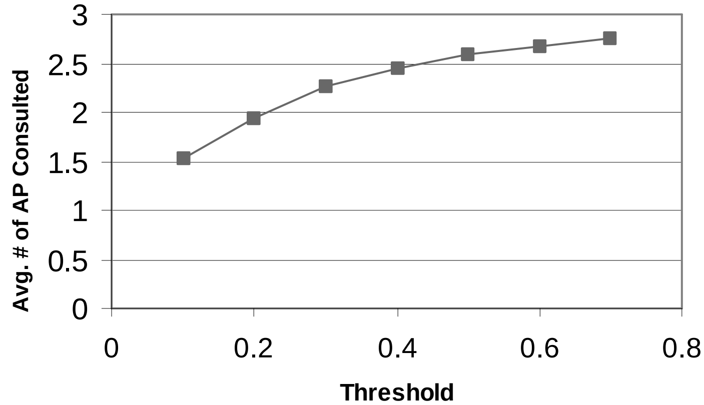

Figure 11: Effect of the Threshold parameter on average number of AP consulted.

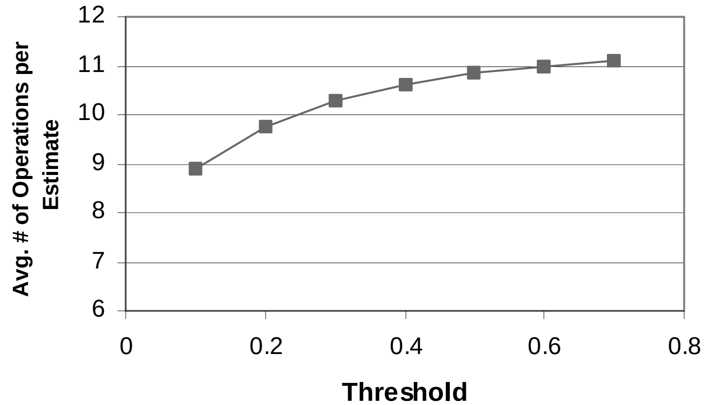

Figure 12: Effect of the Threshold parameter on average number of operation per estimate.

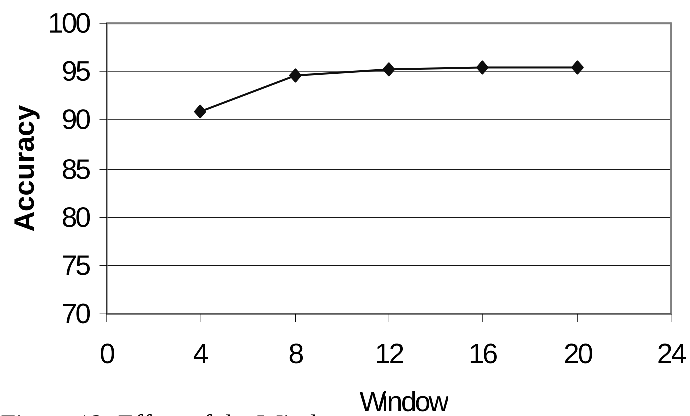

Figure 13: Effect of the Window parameter on accuracy.

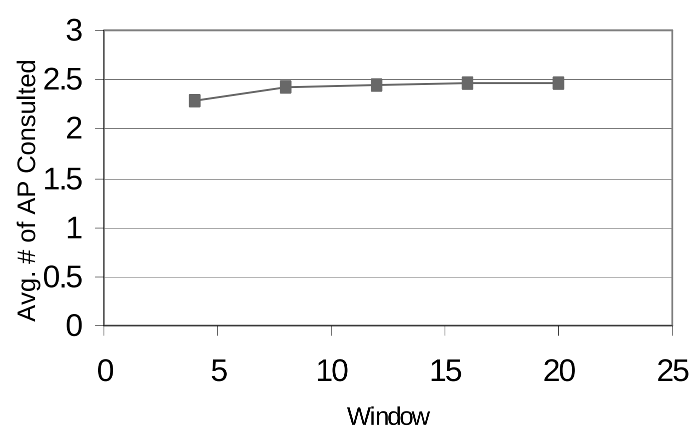

Figure 14: Effect of the Window parameter on average number of AP consulted.

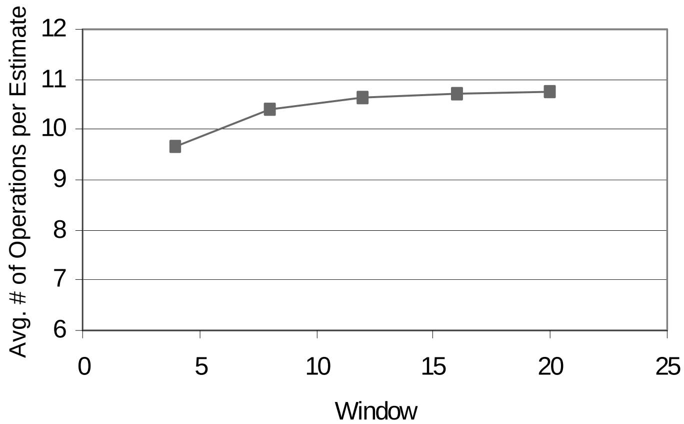

Figure 15: Effect of the Window parameter on average number of operation per estimate.

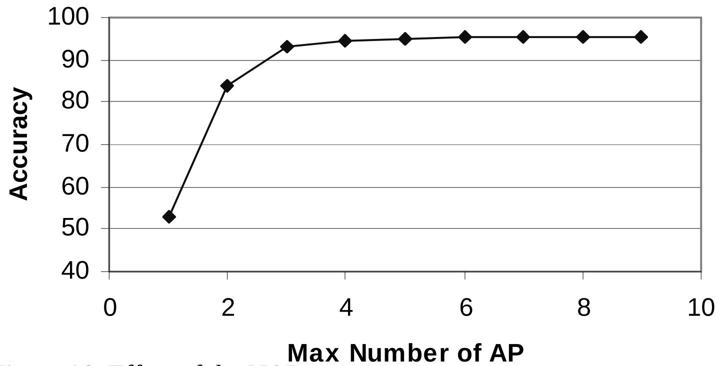

Figure 16: Effect of the NAP parameter on accuracy.

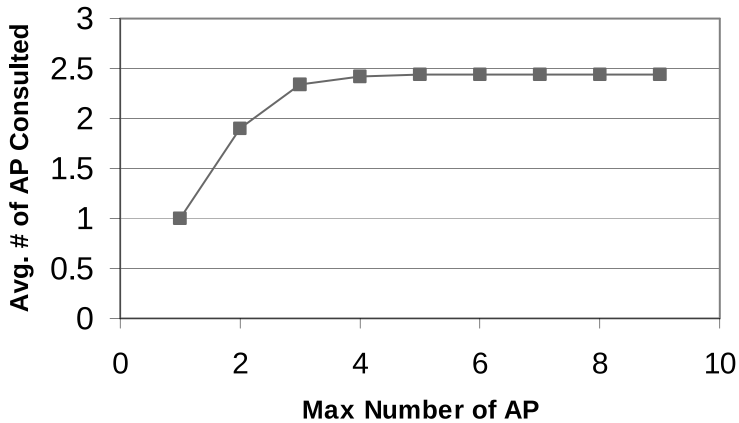

Figure 17: Effect of the NAP parameter on average number of AP consulted.

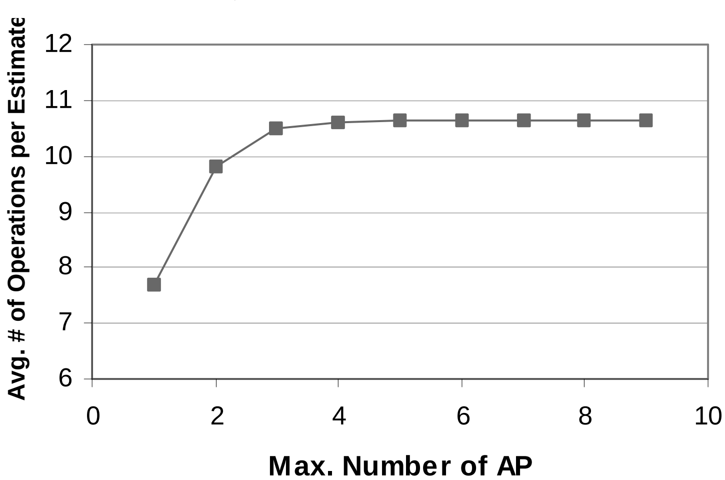

Figure 18: Effect of the NAP parameter on average number of operation per estimate.

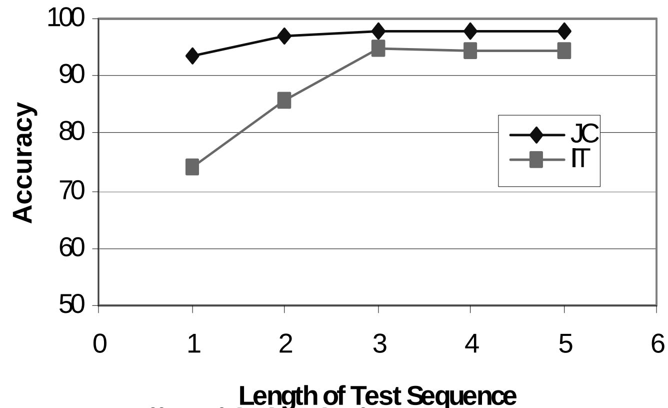

Figure 19: Effect of the length of testing sequence on accuracy.

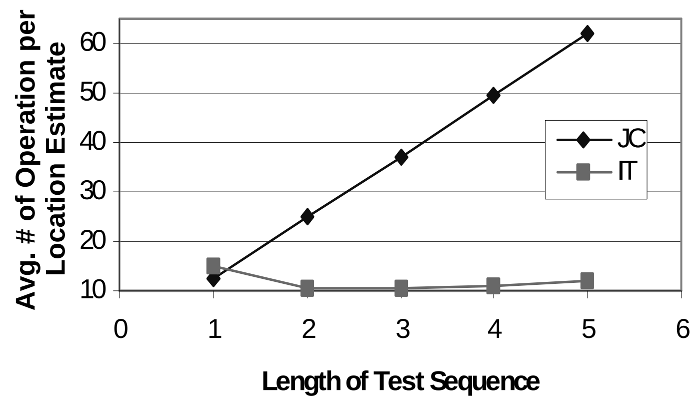

Figure 20: Effect of the length of testing sequence on the average number of operations per location estimate.

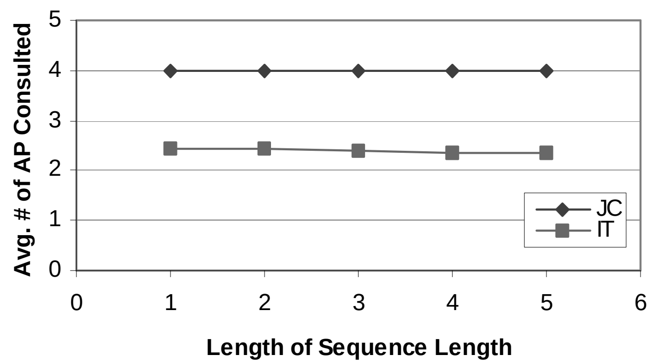

Figure 21: Effect of the length of testing sequence on the average number of AP's consulted.

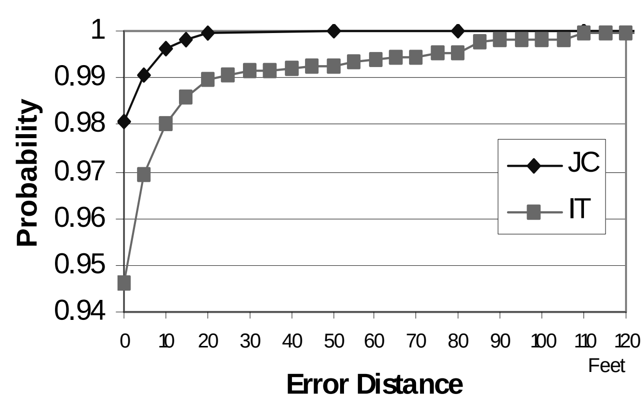

Figure 22: CDF for the JC and IT techniques for the development test set

hence should consult fewer number of access points. This explains the minimum point when the sequence length is 2.

For Figure 21, the number of access points consulted for the JC technique is constant, equals the dimension of the joint distribution *k*, and independent of the length of the observation sequence. For the IT technique, the average number of access points consulted is slightly decreasing due to the availability of more information with the increase in the observation sequence length.

5.4 Comparison Between the Proposed Techniques

Figure 22 shows the CDF of the error distance for the two techniques (note that the y-axis starts from 94%). It is interesting to note that both techniques give more than 94% accuracy for the exact position. This can be explained by noting that our development test set was taken at the same grid positions as the training test set, and hence the exact match (0 error distance). This is different from the independent test set where all locations were off the grid as will be discussed in Section 5.5.

A comparison of the two techniques in terms of the performance measures is shown in Figure 23. From the figures we note that the Joint Clustering technique gives better accuracy than the Incremental Triangulation technique and its tail is much lower than the Incremental Triangulation technique. However, the average number of operations performed per location estimate for the Incremental Triangulation technique is much lower than the corresponding number of Joint Clustering technique. Therefore we have a tradeoff here, if one is interested in accuracy more than power consumption then the Joint Clustering technique is the one to use. If on the other hand the power consumption is the key factor then one should choose the Incremental Triangulation technique as it leads to less computation and hence better power consumption.

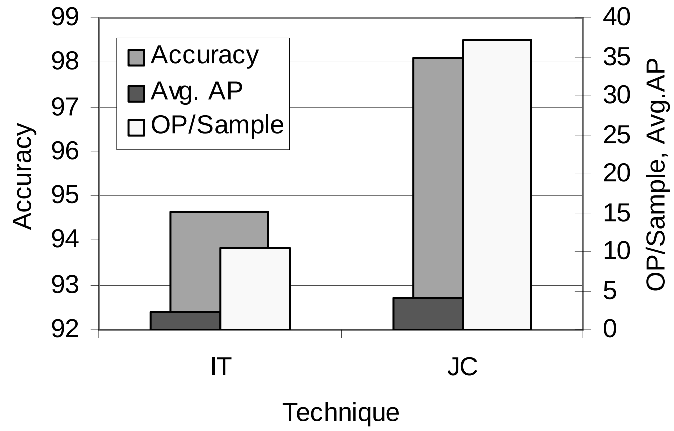

Figure 23: A comparison between the IT and JC techniques.

5.5 Using an Independent Test Set

To better test the proposed techniques, we ran the techniques against an independent test set. This test set was collected at different days and times of the day than the original training sample set and, hence, should give good indication of the performance of the algorithm in different environments. To collect the testing set, we moved along the corridors and selected locations randomly for test. The coordinates of each location along with the test samples were collected. We compare the results of running the techniques with the coordinates stored in the files to determine the error distance.

Figure 24 shows the CDF of the error distance for the Joint Clustering and Incremental Triangulation techniques. The figure shows that both techniques give an accuracy of 7 feet for more than 90% of the time, lower than the results

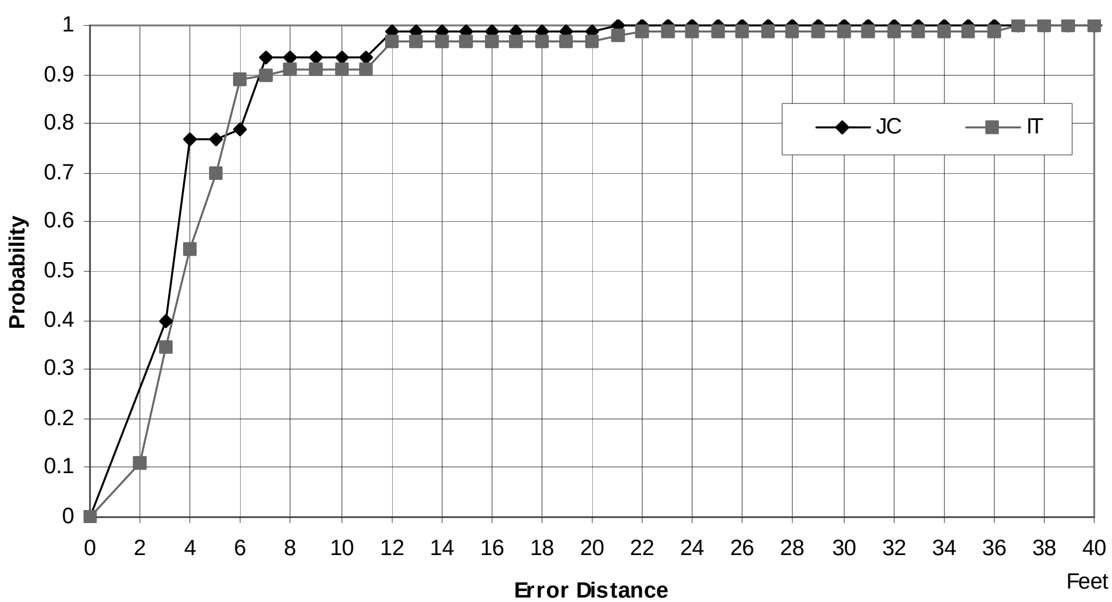

Figure 24: Error Distance CDF for the independent test set.

obtained by the development test set by approximately 10%. It is worth mentioning here that, to the best of our knowledge, the best results reported using other RF-based indoor location determination systems was 75% to within 3 meters (approximately 9.6 feet) [13], which is an indication that using a probabilistic approach for RF-based indoor location determination leads to better results.

6 Applications

We have implemented the location determination system architecture described in Section 2 in two applications. Both applications are implemented on Compaq iPAQ Pocket PC's (model H3650) running the Familiar distribution (release version 0.4 and 0.5) of Linux for PDA's. The iPAQ was running a modified version of our device driver, designed specially for iPAQ's, and our wireless API.

One application, called **User Positioning System,** provides a user moving in a building with the current position as well as directions to specific places of interest. The iPAQ collects the signal strength measurements and sends them to a central monitoring system, which determine the user location and displays it. The user can request directions from his current location to places of interests.

The other application, called **Rover** [20], provides location-based services, as well as the traditional time-aware, user-aware and device-aware services. Examples of such services are displaying the user location on a map, giving the user directions from one place to another and identifying places of interests near the user. Rover also allows the user to see the positions of other people of interest, e.g. members of his group.

Rover has been tested in indoor and outdoor environments. At startup, the clients determine their position and request position-appropriate services, e.g., map, from the Rover server. Unlike the User Positioning System, the client estimates the user position and forwards it to the server. This assures that the proposed location determination techniques are lightweight enough to be implemented on iPAQ's.

7 Related Work

There have been many systems over the years tackling the problem of user positioning and tracking. Examples include GPS, wide-area cellular-based systems, infrared-based systems, magnetic tracking systems, various computer vision systems, physical proximity systems, and radio frequency (RF) based systems.

The GPS system [2] is very useful in outdoors environments. However, the line-of-sight to GPS satellites is not available inside buildings and hence the GPS system cannot be used indoors.

Locating users in the wide-area cellular-based systems has been motivated in recent years by the FCC 94-102 order [17], mandating wireless E911 (automatically locating 911 callers). The two most widely known location technologies used in the wide-area celluar-based systems are Time Difference of Arrival (TDOA) and Angle of Arrival (AOA). TDOA systems use the principle that the emitter location can be estimated by intersection of the hyperbolae of constant differential Time of Arrival (TOA) of the signal at two or more pairs of base stations. AOA systems use simple triangulation based on estimated AOA of a signal at two or more base stations to estimate the location of the desired transmitter [3]. While these systems are promising in outdoor environments, their effectiveness in indoor environments is limited by the multiple reflections suffered by the RF signal, which leads to inaccurate estimate of the TOA or AOA, and the lack of off-the-shelf and inexpensive hardware to provide fine-grain time synchronization.

Many infrared-based (IR) based systems have been proposed and reported. In the Active Badge system [4], a badge worn by a person emits a unique IR signal. Fixed IR receivers pick up this signal and relay it to the location manager software. The walls of a room blocks the IR signal, thus the user can be identified accurately within a room. In [5] IR transmitters attached to known positions in the ceiling emit beacons. A head mounted optical sensor senses these beacons. This enables the system software to determine the user location.

IR based techniques suffer from several drawbacks: (a) they scale poorly due to the limited range of IR, (b) incur significant installation and maintenance costs and (c) perform poorly in the presence of direct sunlight.

Magnetic tracking has been used to support virtual reality and motion capture for computer animation. For example, Ascension [6] offers a variety of motion capture solutions such as the MotionStar DC magnetic tracker. These tracking systems generate axial DC magnetic field pulses from a transmitting antenna in a fixed location. The system computes the position and orientation of the receiving antennas by measuring the response in three orthogonal axes to the transmitter field pulse, combined with the fixed effect of the earth's magnetic field. Such systems suffer from the steep implementation costs and the need to tether the tracked object to a control unit. Furthermore, the sensors must remain within 1 to 3 meters of the transmitter, and accuracy degrades with the presence of metallic objects in the environment.

Several groups have explored using computer vision technology for locating objects. Microsoft research's Easy Living [7] provides one example of this approach where real-time 3D cameras provide a stereovision positioning capabilities in a home environment. Computer-vision based techniques have two drawbacks: (a) they use substantial processing power to analyze frame captured with comparatively low-complexity hardware; (b) the analysis becomes more complex when the scene complexity increases or more occlusive motion occurs.

In Georgia Tech's Smart Floor proximity location system [8], embedded pressure sensors capture footfalls, and the system uses Hidden Markov Models to recognize the users according to their profiles. The system has the disadvantages of poor scalability and high incremental cost as the floor of each building in which Smart Floor is deployed must be physically altered to install the pressure sensor grids.

Recently, there has been ongoing research on RF based techniques. These techniques can be categorized into two broad categories. One that uses specialized hardware and another that uses the underlying data network.

Many systems fall into the first category: The Active Bat [10], [11] system is based on combining the RF and the ultrasonic technologies. A short pulse of ultrasound is emitted from a transmitter (a *Bat*) attached to the object to be located in response to an RF request from a local controller. The local controller sends, at the same time as the request packet, a synchronized reset signal to the ceiling sensors using a wired serial network. The system measures the times-of-flight of the pulse to the mounted receivers on the ceiling. The system uses the speed of sound in air to calculate the distances from the Bat to each receiver. The local controller forwards these distances to a central controller that performs the location determination computation. The scalability and ease of deployment are disadvantages to this approach.

The Cricket location support system [9] uses a combination of RF and ultrasound technologies to provide a location-support service to users and applications. Wall- and ceiling-mounted beacons are spread through the building, publishing information on an RF signal operating in the 418 MHz AM band. With each RF advertisement, the beacon transmits a concurrent ultrasonic pulse. Listeners attached to devices and mobiles listen for RF signals, and upon receipt of the first few bits, listen for the corresponding ultrasonic pulse. When this pulse arrives, they obtain a distance estimate for the corresponding beacon. The listeners run maximum-likelihood estimators to correlate RF and ultrasound samples and pick the best pair. The disadvantages lie in the lake of centralized management or monitoring and the computational, and hence the power consumption burden, at the receiver due to timing and processing the RF data and ultrasound pulses.

Another indoor RF system is the 3D-iD RF tag built by PinPoint Corporation [18]. Antennas planted around a facility emit RF signals. Tags, acting like RF mirrors, transmit a response signal along with an identification code. Various antennas receive the response signal and send the results to a central controller that triangulates the user location. The cost of the entire system is quite high.

All the techniques that fall in the specialized hardware category have common disadvantages: (a) requirement of specialized hardware leading to more deployment and maintenance cost, and (b) poor scalability.

In the last few years, many techniques have been proposed that fall into the second category, RF-based systems that do not require additional hardware. The Daedalus project [19] developed a system for coarse-grained user location. A mobile host estimates its location to be the same as the base station to which it is attached. Therefore, the accuracy of the system is limited to the coverage area of the access point.

The RADAR system [12] uses the RF signal strength as an indication of the distance between the transmitter and receiver. During an offline phase, the system builds a radio map for the RF signal strength from a fixed number of receivers. During normal operation, the RF signal strength of the transmitter is measured by a set of fixed receivers and is sent to a central controller. The central controller uses a K-nearest approach to determine the location from the radio map that best fits the collected signal strength information.

The CMU system proposed in [13] uses two techniques: pattern matching (PM) and triangulation, mapping and interpolation (TMI). The PM approach is very similar to the RADAR approach. In the TMI technique, the physical position of all the access points in the area needs to be known and a function is required to map signal strength onto distances. They generate a set of training points at each trained position. The interpolation of the training data allows the algorithm to use less training data than the PM approach. During user location determination phase, they use the approximate function they got from the training data to generate contour and they calculate the intersection between different contours yielding the signal space position of the user. The nearest set of mappings from the signal space to the physical space is found by applying a weighted average, based on proximity, to the signal space position.

Our proposed techniques differ from the RADAR and the CMU approaches in many ways: (a) we use a probabilistic approach to rank the candidate user locations reducing the search space; (b) both of our techniques perform clustering, either explicit or implicit, which further reduces the search space; (c) user profile information can be easily added to the probabilistic models.

The Nibble location system from UCLA uses a Bayesian network to infer a user location [14]. Their Bayesian network model include nodes for location, noise, and access points (sensors). The signal to noise ratio observed from an access point at a given location is taken as an indication of that location. The system also quantizes the SNR into four levels: high, medium, low, and none. While our approaches and the Nibble approach are similar in some ways, they also differ in significant ways: (a) our approaches do not store the joint distribution between all the random variables of the system. Instead, we store only the marginal distribution of each access point, which reduces the computation significantly and enhances system scalability; (b) we use the received signal strength instead of the SNR as the signal strength is a stronger function of location than the SNR [12]. (c) we have a much finer quantization of the received signal strength, which gives us more information, and thus should lead to better accuracy without affecting performance or scalability; (d) our techniques perform clustering.

Table 1 gives a comparison between the previous systems in the area of RF indoor location determination with our proposed techniques.

8 Conclusions and Future Work

In this paper, we presented the design, implementation, and evaluation of two novel probabilistic indoor location determination techniques: the Joint Clustering technique and the Incremental Triangulation technique. Both techniques depend on (a) probability distributions to handle the noisy

System	RADAR	Nibble	CMU	JC	IT
Technique	Pattern matching	Bayesian Network	Triangulation-Mapping and Interpolation, Pattern matching	Probability with joint distribution.	Probability with incremental triangulation
Clustering	None	None	None	Explicit	Implicit
Feature used	Signal strength	SNR	Signal strength	Signal strength	Signal strength
Quantization of feature	No	Yes	No	No	No
Ease of adding user profile	Not part of the model.	Part of the model.	Not part of the model	Part of the model	Part of the model
Position of AP's	Needed for Radio propagation model.	Not needed	Needed for Triangulation.	Not needed	Not needed

Table 1: Qualitative comparison between other RF-location based techniques with IPS.

characteristics of the wireless channel, and (b) clustering to manage the computational cost.

The Joint Clustering technique gives better accuracy than the Incremental Triangulation technique. However, the average number of operations performed per location estimate for the Incremental Triangulation technique is much lower than the corresponding number of the Joint Clustering technique. Therefore a tradeoff exists between accuracy and computation power. Both techniques lead to accuracy of more than 90% to within 7 feet in our experiments.

During the course of our implementation, we developed a new device driver for the Lucent Wavelan card and a new wireless API. Both software pieces are available for public download and are being used by many researchers throughout the world.

Currently we are working to enhance accuracy and reduce computational cost. By using the user history profile and better clustering techniques, the accuracy of the location determination techniques can be enhanced. Interpolating between a number of the most probable locations is another direction that we are looking into to improve the accuracy. We believe that understanding the nature of the radio channel and building accurate models for it are important for building more accurate location determination systems for the indoor environments and for reducing the overhead of building the radio map.

Our results gave us confidence that, despite the hostile nature of the wireless channel, we can infer the user location with a high degree of accuracy enabling a set of context-aware applications for the indoor environments.

References

[1] Guanling Chen and David Kotz, "A survey of context-aware mobile computing research," Dartmouth Computer Science Technical Report TR2000-381.

[2] P. Enge, and P. Misra, "Special issue on GPS: The Global Positioning System," Proc. of the IEEE, pp. 3-172, Jan. 1999.

[3] Sirin Tekinay, "Special issue on Wireless Geolocation System and Services," IEEE Communications Magazine, April 1998.

[4] Roy Want, Andy Hopper, Veronica Falcão, and Jonathan Gibbons, "The Active Badge Location System," ACM Transactions on Information Systems, Vol. 10, No. 1, January 1992, pp 91-102.

[5] Ronald Azuma, "Tracking requirements for augmented reality," Communications of the ACM, Jul. 1995, Vol. 36, No. 7.

[6] http://www.ascension-tech.com

[7] J. Krumm et al., "Multi-camera multi-person tracking for Easy Living," Proc. 3rd IEEE Int'l Workshop Visual Surveillance, IEEE Press, Piscataway, N.J., 2000, pp. 3-10.

[8] Robert J. Orr and Gregory D. Abowd, "The Smart Floor: A Mechanism for Natural User Identification and Tracking," Proceedings of the 2000 Conference on Human Factors in Computing Systems (CHI 2000), The Hague, Netherlands, April 1-6, 2000.

[9] Nissanka B. Priyantha, Anit Chakraborty, Hari Balakrishnan, "The Cricket Location-Support system", Proc. 6th ACM MOBICOM, Boston, MA, August 2000.

[10] Andy Ward, Alan Jones, Andy Hopper. "A New Location Technique for the Active Office." IEEE Personal Communications, Vol. 4, No. 5, October 1997, pp. 42-47.

[11] A. Harter, A. Hopper, P. Steggles, A. Ward, and P. Webster, "The anatomy of a context-aware application," In Proc. 5th ACM MobiCom Conf., Aug. 1999.

[12] Paramvir Bahl and Venkata N. Padmanabhan, "RADAR: An In-Building RF-based User Location and Tracking System," IEEE Infocom 2000.

[13] Asim Smailagic, Daniel P. Siewiorek, Joshua Anhalt, David Kogan, and Yang Wang, "Location Sensing and Privacy in a Context Aware Computing Environment," Pervasive Computing, 2001.

[14] Paul Castro, Patrick Chiu, Ted Kremenek, and Richard Muntz, "A Probabilistic Location Service for Wireless Network Environments", Ubiquitous Computing 2001, Atlanta, GA, September 2001.

[15] http://www.cs.umd.edu/users/moustafa/Downloads.html

[16] http://www.hpl.hp.com/personal/Jean_Tourrilhes/

[17] FCC Docket No. 94-102, "Revision of the commission's rules to ensure compatibility with Enhanced 9111 emergency call systems," RM-8143, July 26 1996.

[18] J. Werb and C. Lanzl, "Designing a positioning system for finding things and people indoors, " IEEE Spectrum, 35(9), pp71-78, September 1998.

[19] T. D. Hodes, R. H. Katz, E. S. Schreiber, and L. Rowe, "Composable ad hoc mobile services for universal interaction," MobiCom '97 Proceedings, pp:1-12, Sep. 1997.

[20] S. Banerjee, S. Agarwal, K. Kamel, A. Kochut, C. Kommareddy, T. Nadeem, P. Thakkar, B. Trinh, A. Youssef, M. Youssef, R. Larsen, A. U. Shankar, A. Agrawala, "Rover Technology: Enabling Scalable Location-Aware Computing," UMIACS-TR 2001-89 and CS-TR 4312, December 2001.