The Horus Location Determination System

Moustafa Youssef

1 Overview

Figure 1 shows the overall system. The *Horus* system works in two phases:

1. *Offline phase*: to build the radio map, cluster radio map locations, and do other preprocessing of the signal strength models.

2. *Online Phase*: to estimate the user location based on the received signal strength from each access point and the radio map prepared in the offline phase.

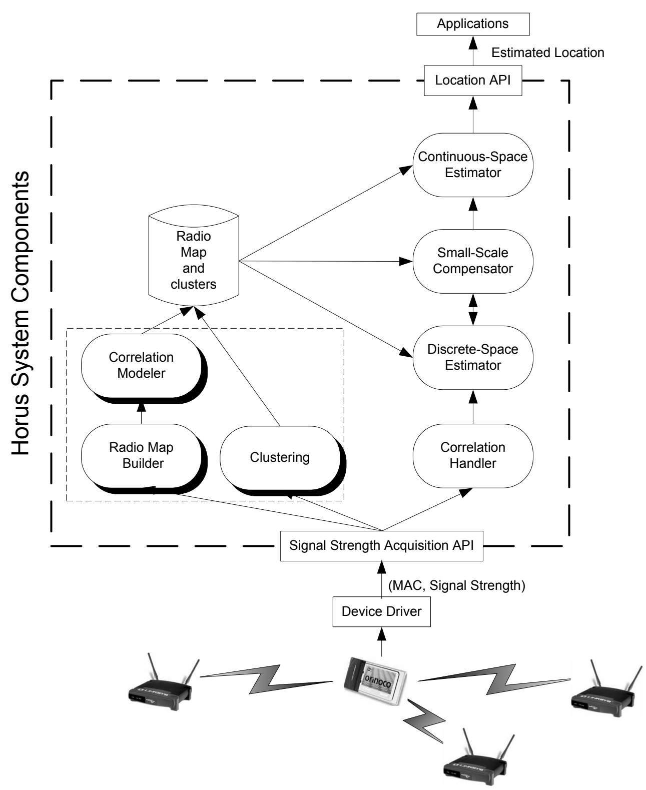

Figure 1: *Horus* Components: the arrows show information flow in the system. Shadowed blocks represent modules used during the offline phase.

The radio map stores the distribution of signal strength received from each access point at each sampled location. There are two modes for operation of the *Horus* system: one uses non-parametric distributions and the other uses parametric distributions.

The *Clustering* module is used to group radio map locations based on the access points covering them. Clustering is used to reduce the computational requirements of the system and, hence, conserve power.

The *Discrete Space Estimator* module returns the radio map location that has the maximum probability given the received signal strength vector from different access points.

The *Correlation Modelling and Handling* modules use an autoregressive model to capture the correlation between consecutive samples from the same access point. This model is used to obtain a better discrete location estimate using the average of n correlated samples.

The *Continuous Space Estimator* takes as an input the discrete estimated user location, one of the radio map locations, and returns a more accurate estimate of the user location in the continuous space.

The *Small-Scale Compensator* module handles the small-scale variation characteristics of the wireless channel.

We start by laying out the mathematical framework for the approach then give details about different components of the system.

2 Mathematical Model

Without loss of generality, let $\mathbb{X}$ be a 2 dimensional physical space. At each location $x \in \mathbb{X}$, we can get the signal strength from k access points. We denote the k-dimensional signal strength space as $\mathbb{S}$. Each element in this space is a k-dimensional vector whose entries represent the signal strength readings from different access points. We denote samples from the signal strength space $\mathbb{S}$ as s. We also assume that the samples from *different* access points are independent.

The problem becomes, given a signal strength vector $s = (s_1, ..., s_k)$, we want to find the location $x \in \mathbb{X}$ that maximizes the probability $P(x/s)$.

In the next section, we assume a discrete $\mathbb{X}$ space. We discuss the continuous space case in Section 5.

3 Discrete Space Estimator

During the offline phase, the *Horus* system estimates the signal strength histogram for each access point at each location. These histograms represent the *Horus* system's radio map. Now consider the online phase. Given a signal strength vector $s = (s_1, ..., s_k)$, we want to find the location $x \in \mathbb{X}$ that maximizes the probability $P(x/s)$, i.e., we want

$$\text{argmax}_x[P(x/s)] \tag{1}$$

Using Bayes' theorem, this can be shown to be equivalent to:

$$\text{argmax}_x[P(x/s)] = \text{argmax}_x[P(s/x)] \tag{2}$$

$P(s/x)$ can be calculated using the radio map as:

$$P(s/x) = \prod_{i=1}^{k} P(s_i/x) \tag{3}$$

An outline of the algorithm used is given in Algorithm 1.

The signal-strength histogram can be approximated by a parametric distribution such as the Gaussian distribution.

Alg. 1 x= Horus_GetLocation (s, $\mathbb{X}$, P_RM)

Input:

s : Measured signal strength vector from k access points ($s = (s_1, ..., s_k)$).

$\mathbb{X}$: Radio map locations.

RM : A radio-map function, where $P\_RM(s_a, a, x)$ returns the probability of receiving signal strength s_a from access point a at location $x \in \mathbb{X}$.

Output:

The location $x \in \mathbb{X}$ that maximizes $P(x/s)$.

Max $\leftarrow 0$
for $l \in \mathbb{X}$ **do**
  $P \leftarrow \prod_{i=1}^{k} P\_RM(s_i, i, l)$
  if $P >$ Max **then**
    $x \leftarrow l$
    $Max \leftarrow P$
  end if
end for

4 Correlation Handling

To account for the temporal signal-strength variations, it is important to average multiple signal strength samples from the same access point. the autocorrelation of successive samples collected from one access point is as high as 0.9. Assuming independence of samples from the same access point leads to the undesirable result of degraded system performance as the number of averaged samples is increased.

4.1 Mathematical model

We use an autoregressive model to capture the correlation between different samples from the same AP.

Let s_t be the *stationary* time series representing the samples from an access point, where t is the discrete time index. s_t can be represented as a *first order* autoregressive model as:

$$s_t = \alpha s_{t-1} + (1-\alpha)v_t \quad ; 0 \leq \alpha \leq 1 \tag{4}$$

where v_t is a noise process, independent of s_t, and α is a parameter that determines the degree of autocorrelation of the original samples. Moreover, different samples from v_t are independent and identically distributed (i.i.d.).

The model in Equation 4 states that the current signal strength value (s_t) is a linear aggregate of the previous signal strength value (s_{t-1}) and an independent noise value (v_t). The parameter α gives flexibility to the model as it can be used to determine the degree of autocorrelation of the original process. For example, if α is zero, the samples of the process s_t are i.i.d.'s, whereas if α is 1 the original samples are identical (autocorrelation=1).

Assuming that the signal strength distribution of samples from an access point is Gaussian with mean μ and variance σ^2, the distribution of the average of n correlated samples is a Gaussian distribution with mean μ and variance given by:

$$\frac{1+\alpha}{1-\alpha}\sigma^2 \tag{5}$$

4.2 Correlation modeler

The purpose of the correlation modeler component is to estimate the value of α in the autoregressive model and to pre-calculate the parameters of the distribution of the average of n correlated samples during the offline phase. α can be approximated using the autocorrelation coefficient with lag 1. The variance of the distribution can be calculated using Equation 5. These distribution parameters (μ, σ, and α) are then stored in the radio map.

4.3 Correlation handler

During the online phase, the correlation handler module averages the value of n consecutive samples from an access point and passes this information to the discrete-space estimator, which uses the distributions stored in the radio map.

5 Continuous Space Estimator

The discrete-space estimator returns a single location from the set of locations in the radio map. To increase the system accuracy, the *Horus* system uses two techniques to obtain a location estimate in the continuous space.

5.1 Technique 1: Center of Mass of the Top Candidate Locations

This technique is based on treating each location in the radio map as an object in the physical space whose weight is equal to the normalized probability[1] assigned by the discrete-space estimator. We then obtain the center of mass of the N objects with the largest mass, where N is a parameter to the system, $1 \leq N \leq ||\mathbf{X}||$.

More formally, let $p(x)$ be the probability of a location $x \in \mathbb{X}$, i.e., the radio map, and let $\bar{\mathbb{X}}$ be the list of locations in the radio map *ordered* in a descending order according to the normalized probability (the location with lower ID comes first for locations with equal probability). The center

[1]The normalization is used to ensure that the sum of the probabilities of all locations equals one.

of mass technique estimates the current location x as:

$$x = \frac{\sum_{i=1}^{N} p(i)\bar{\mathbb{X}}(i)}{\sum_{i=1}^{N} p(i)} \tag{6}$$

where $\bar{\mathbb{X}}(i)$ is the i^{th} element of $\bar{\mathbb{X}}$

Note that the estimated location x need not be one of the radio map locations.

5.2 Technique 2: Time-Averaging in the Physical Space

The second technique uses a time-average window to smooth the resulting location estimate. The technique obtains the location estimate by averaging the last W locations estimates obtained by either the discrete-space estimator or the continuous-space estimator discussed in the previous section.

More formally, given a stream of location estimates $x_1, x_2, ..., x_t$, the technique estimates the current location $\bar{x_t}$ at time t as:

$$\bar{x_t} = \frac{1}{\min(W,t)} \sum_{t-\min(W,t)+1}^{t} x_i \tag{7}$$

6 Small-Scale Compensator

Dealing with small-scale variations is challenging. Since the selected radio map locations are typically placed more than a meter apart, to limit the radio map size, the radio map does not capture small-scale variations. This contributes significantly to the estimation errors in the current systems. The *Horus* system uses the *Perturbation* technique to handle the small-scale variations. The technique is based on two sub-functions: detecting small-scale variations and compensating for small-scale variations.

6.1 Detecting small-scale variations

In order to detect small-scale variations, the *Horus* system uses the heuristic that users' location cannot change faster than their movement rate. The system calculates the estimated location using the standard radio map and the inference algorithm, then calculates the distance between the estimated location and the previous user location. If this distance is above a threshold, based on the user movement rate and estimation frequency, the system decides that there are small-scale variations affecting the signal strength.

6.2 Compensating for small-scale variations

To compensate for these small-scale variations, the system perturbs the received vector entries, re-estimates the location, and chooses the nearest location to the previous user location as the final location estimate. For example, if one sample includes a signal-strength observation from each of k access points $(s_1, s_2, \ldots, s_k)$, the system tries all 3^k combinations in which each of the k

observations i is replaced by one of three values, s_i, $s_i(1+d)$, or $s_i(1-d)$; An enhancement of this approach is to perturb a subset of the access points.

7 Clustering Module

This section describes the *Incremental Triangulation* (IT) clustering technique used by the *Horus* system to reduce the computational requirements of the location determination algorithm. We define a *cluster* as a set of locations sharing a common set of access points. We call this common set of access points the *cluster key*. The problem can be stated as: Given a location x, we want to determine the cluster to which x belongs. The noisy characteristics of the wireless channel make clustering a challenging problem because the number of access points covering a location varies with time.

The *IT* approach is based on the idea that each access point defines a subset of the radio map locations that are covered by this access point. These locations can be viewed as a cluster of locations whose key is the access point covering the locations in this cluster. If during the location determination phase we use the access points incrementally, one after the other, then starting with the first access point, we restrict our search space to the locations covered by this access point. The second access point chooses only the locations in the range of the first access point and covered by the second access point and so on, leading to a multi-level clustering process.

Notice that no preprocessing is required in the offline training phase. During the online phase, a location x belongs to a cluster whose key is access point a if there is information about access point a at location x in the radio map.

The algorithm works as follows. Given a sequence of observations from each access point, we start by sorting the access points in descending order according to the average received signal strength. 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 given the observation sequence from this access point alone. This gives 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 threshold) than the probability of the second most probable location, we return that 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.