MIT 6.808 Mobile and Sensor Computing (The IoT Class)
Spring 2022

Hari Balakrishnan

Notes for Lecture 3: Indoor Location (RADAR, Cricket) Feb 7 2022

Plan for lectures 2 through 4: Principles, GPS, Wi-Fi loc, ultrasonic loc, Wi-Fi sensing

Goal: Fundamentals of wireless localization and positioning and systems to achieve localization

In the previous lecture we studied these issues:

1. Motivating applications: why do we care about location, localization, positioning?
2. Key approaches: a) sensing modalities for location, b) principles behind the methods
   a. Modalities: radio, ultrasound/acoustics, inertial (accel, gyro), light (camera, LIDAR), and assisting info from barometers and magnetometers
   b. Five broad principles: proximity sensing (aka identity-based localization), lateration (solve using distances), angulation (solve using angles), fingerprinting (pattern matching), dead reckoning
3. Key ideas in GPS: time-sync'd satellites, need four satellites (three position and time offset between GPS receiver and satellite is fourth unknown), orthogonal codes.

Today's plan

1. **Wireless LAN-based localization with RF[1] fingerprinting (RADAR, Infocom 2000)**
2. **Ultrasonic (+RF) localization (Cricket, MobiCom 2000)**

Both papers have won "test of time" awards from SIGMOBILE (and thousands of citations).

RADAR Wireless LAN localization

Q: What are the key ideas in this system?

1. There is a correlation between received RF signal strength (today called RSSI, the received signal strength indicator) and distance from the transmitter. The reason for that is the physics of signal propagation where the strength of a signal goes down with distance as 1 / dn, where usually 2 <= n < 4. In free space, n = 2; indoors and with obstacles, n increases.
2. But the physics of RF signal propagation is quite complicated and hard to model, so we should view these formulas and models as guidelines.

---

[1] RF = Radio Frequency

3. Side note: how is RSSI measured? Typically, in Wi-Fi, most chipsets measure it on the packet's preamble. It provides an estimate of the power level of the received signal and is typically expressed in dBm.
    a. What is dBm? To answer that, first let's understand what the dB scale is. It's a way to compare relative power levels. If you have two signals of strengths S1 and S2, then the power difference between them, in dB, is $10 \log_{10} (S1/S2)$. It's used, for example, to express the signal-to-noise ratio (SNR).
    b. dBm simply means that the reference signal, or denominator, is 1 milliwatt. Thus, if you receive a signal with power *S* Watts, it's dBm is $10 \log_{10} (S/0.001)$, or equivalently, if the received power is *m* milliwatts, then its dBm is $10 \log_{10}(m)$.
4. The challenge in practice is that there is no specific standard on how these numbers are reported; they tend to be hardware-specific. Most cards will rescale it to a certain number of bits (e.g., 8 bits).

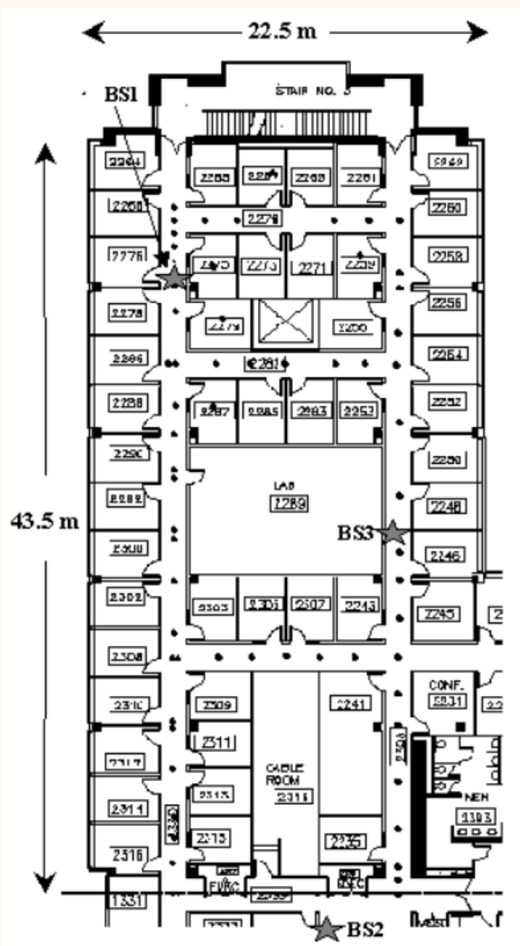

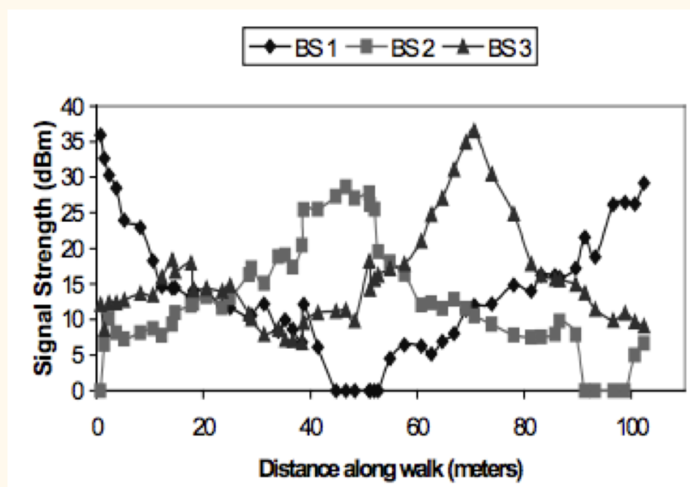

The pic above (left) shows the Microsoft experimental setup deployment from the RADAR paper. The pic on the right shows how the signal strength varies as a user moves starting from the BS1 location along the corridor loop. You can see that it's correlated with the distance from the base station, roughly speaking.

Method

Suppose we have a building with some wireless LAN (Wi-Fi / Bluetooth / something else) RF access points (APs) / base stations. In the picture below, the APs are shown here as small black squares. Then, later, users may want to know their locations at unknown uncalibrated locations, e.g., those shown in the small black circles.

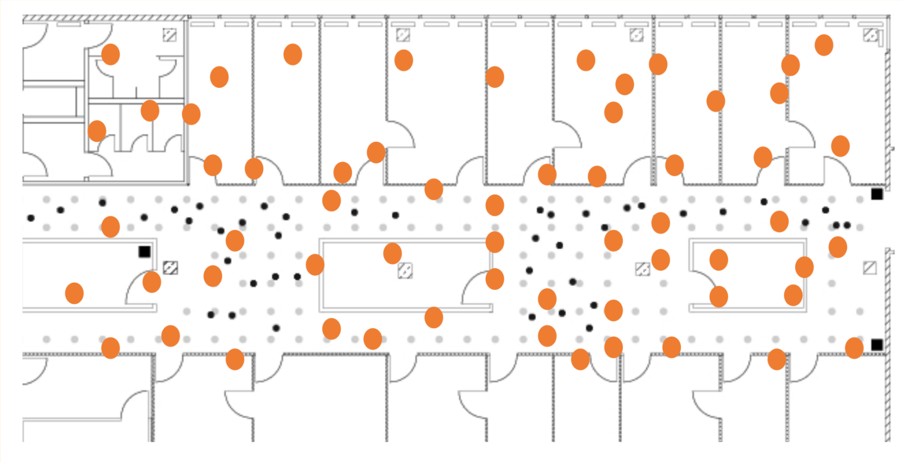

Our job is to obtain the best answer. To do that, we will:

1. Obtain training data for modeling. For example, in this pic below the orange circles are calibrated locations where we will gather RSSI information. We will produce a training database of RSSI samples at several locations, e.g., (AP1 RSSI 7, AP2 RSSI 14, AP3 RSSI 24, AP6 RSSI 8) for location L with coordinates (x, y, z) – we can think of the problem as in two dimensions on a given floor or in three dimensions. The calibration data needs to include the *direction* or *heading* of the user, especially at 2.4 GHz (but even otherwise) as antenna orientations affect signal reception.
2. For each location L, we can use this data to determine what the signal strength would be. The paper discusses two ways to do this:
    a. Empirical, i.e., use the measured data from the calibration phase as-is
    b. Physical modeling, i.e., use a signal propagation model
3. At runtime, the system infers the best location for the user or device given an observed signal strength vector: (AP1:45, AP2:3) → what is the best location?
    a. Need a nearest-neighbor metric and search
        i. Minimize Euclidean distance in signal strength space \sum_i (received_ss_i^2 - calibrated_ss_i^2)

ii. Should we weight the locations of k nearest APs?
Q: Why do we want any weighting?
A: Because we aren't going to be at the exact location where we were calibrated.
Some sort of weighted averaging would help (not in paper).

b. What do we do about partial AP coverage (not in paper)
   i. Could **maximize** a metric such as A*M + (dmax - euclidean_dist), where A is some constant scaling factor, M the number of matching APs between the calibration set and the observed set, and we compute the Euclidean distance over the matches and subtract that from the max possible Euclidean distance.

RADAR evaluation

What does this picture show? What does it mean? These are CDFs, or cumulative distribution functions. THis was an eval where 1 of the 70 calibrated points was queried in each trial against the other 69. Each CDF has 70 contributing data points.

Is there a better evaluation approach?

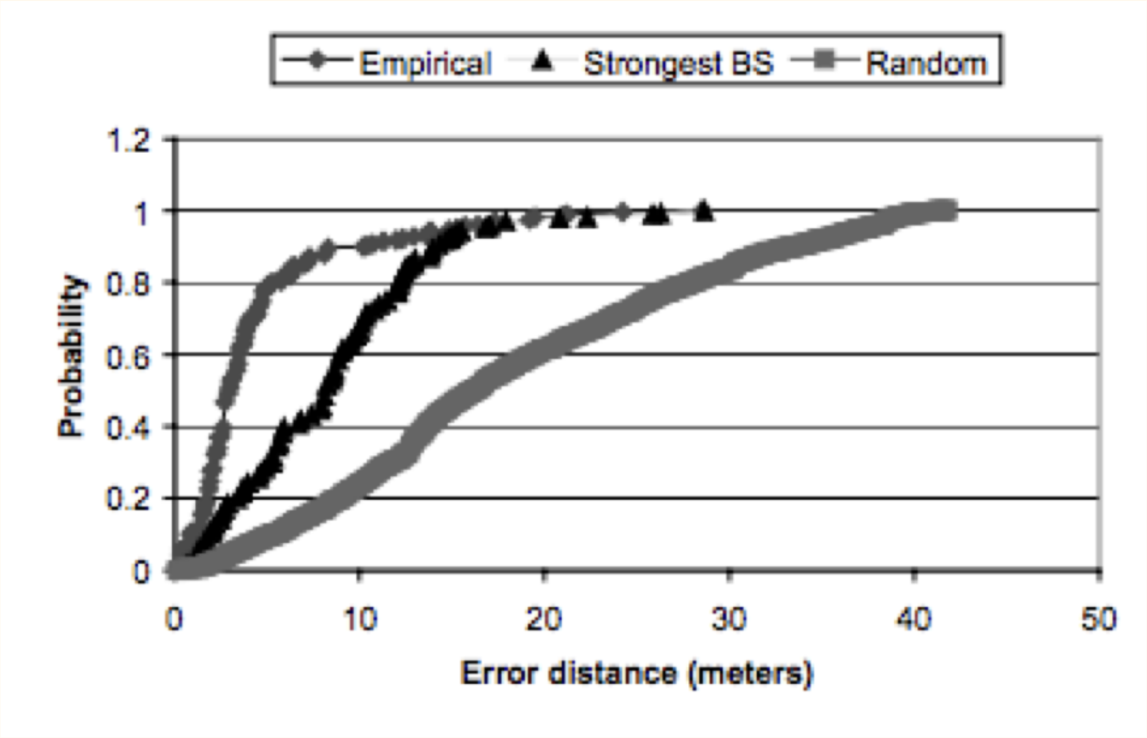

Cricket

Goals: Precise indoor localization of spaces and coordinates, scalability, privacy, easy deployment, low energy consumption.

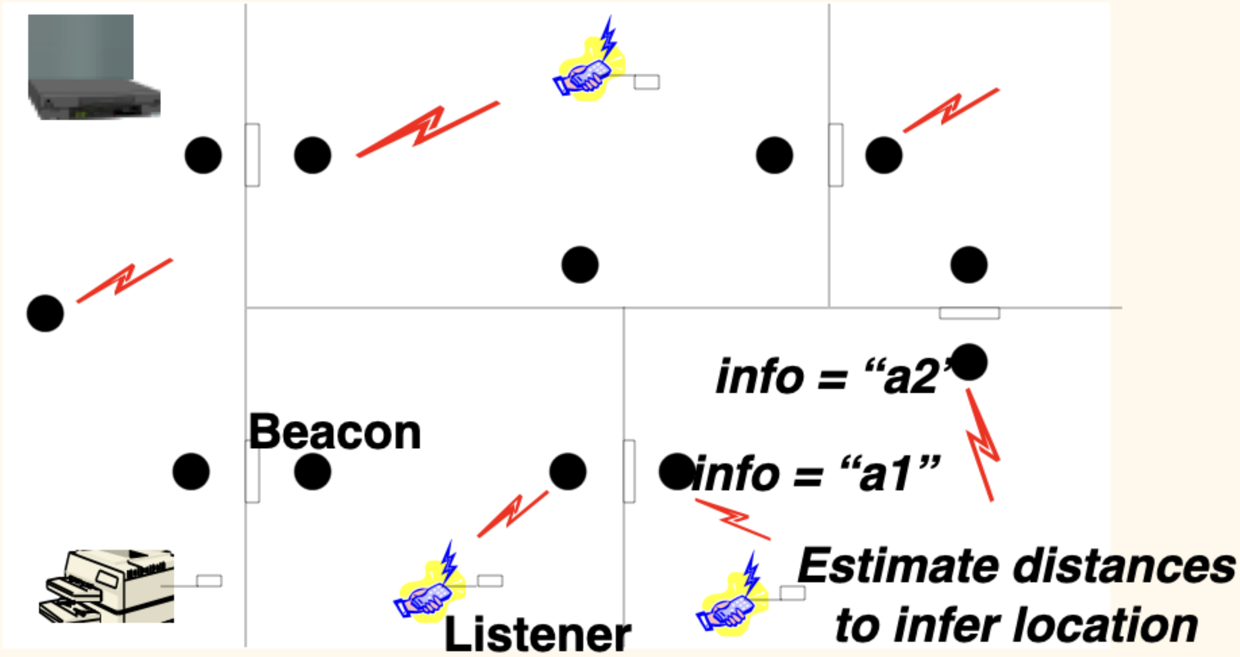

Passive listeners, active beacons. Decentralized and self-configuring coordinate system in beacons (not in this paper).

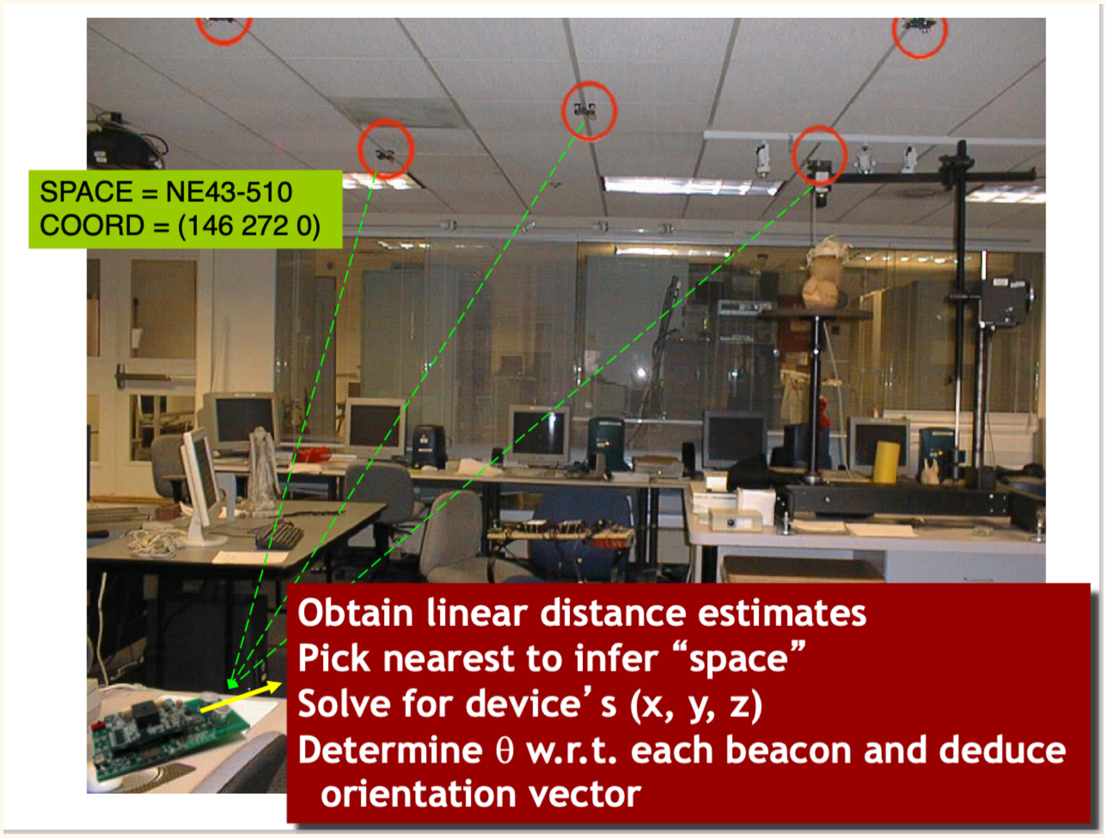

Note: determination of the angle is not in this paper. That was in a MobiCom 2001 paper on the Cricket Compass.

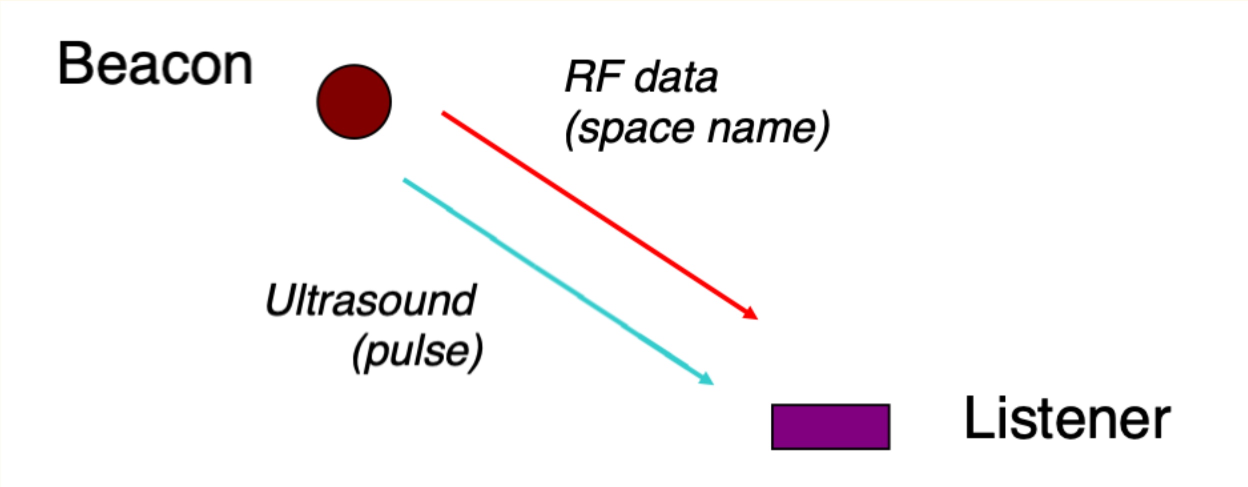

How does the system handle interference?

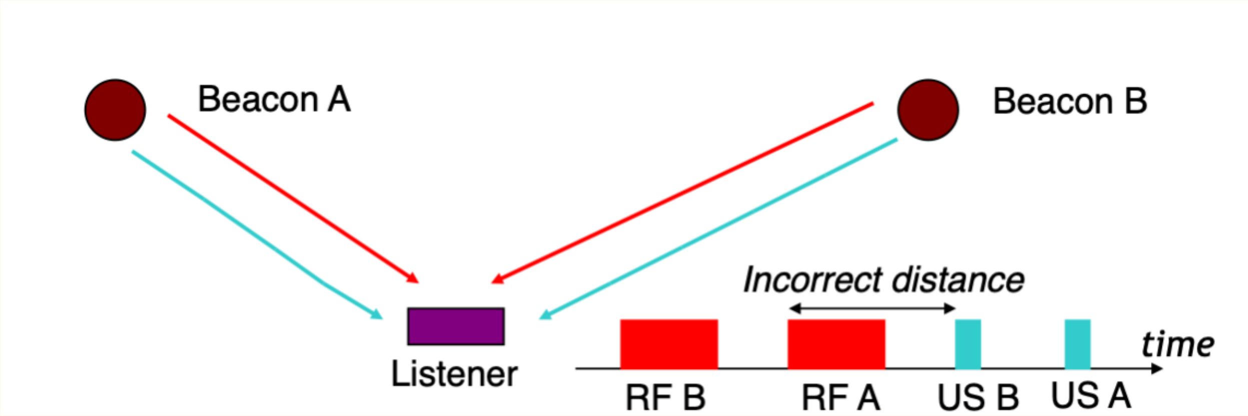

Ultrasound has no messages encoded on it, just pulses. We will revisit this point in a bit. But for now assume it's a pulse and has significant reflections.

1. RF always envelops US
2. RF-A, US-A, US-RA (ultrasonic reflection of beacon A) and we also have RF-I, US-I, US-RI of interferer (shown as B above).
3. What are the different possible interference patterns?

Beacon deployment: equidistant from boundary between spaces.

Distance estimation: use time estimates, multiply by speed of sound. Each sample produces a distance estimate. Produce a distribution of these. Estimate the distance as the mode of the distribution. This is robust to reflections causing higher estimates and to spurious ultrasonic pulses causing occasional under-estimates.

Space estimation: **MinMode,** i.e., the distance whose statistical mode is smallest.

Impact: Many hundreds of thousands of devices were made and deployed. Influenced subsequent “echolocation” systems that used acoustic/ultrasound. Also had an influence on recent approaches to precise contact tracing with bi-directional ultrasonic ranging.