Location and Time
in
Wireless Environments

Ashok K. Agrawala
Director, MIND Lab
Professor, Computer Science
University of Maryland

UNIVERSITY OF MARYLAND 18 56

MIND The MIND Lab
The Maryland Information and Network Dynamics Lab

Environment

- N nodes
  - local clock
    - Stable
  - Wireless Communications
  - Computation
  - Storage
  - Sensors
- Deployed in 2D/3D region
  - Regularly spaced
  - Randomly placed
  - Static or mobile
- Deployment
  - Infrastructure mode
  - Ad hoc mode
- Indoors/Outdoors

What can such a group of nodes do ?

Applications

- Location-Aware Applications
  - Shopping Center
  - Amusement Park
  - Museum
  - Hospital
  - First Responders
  - …
- Location-Aware Security
- Location-Aware Routing
- …

- Synchronized Actions
  - By group of people
  - By devices
- Information Fusion
- Ad-hoc Phased Array
  - Transmitter
  - Receiver
- Time-based Management of Resources
- …

System Synchronization

- Coordinated action by N-nodes
- Are synchronized clocks essential ?
  - Sufficient, not necessary and sufficient
  - If clocks are not synchronized and no information about clocks of each node is used, lower bound on synchrony is the signal transit delay.
- Stable Clocks
  - Clock characteristics do not change rapidly
    - Drift rate remains constant
  - Can lead to system synchronization without clock synchronization !!

Outline

- Localization – Active Techniques
- RSSI Based
  - Characteristics of 802.11b signals
  - Horus
- Transit Time Based/ Synchronization
  - PinPoint
  - System Synchronization
- Localization – Passive Techniques
  - Nuzzer
- Concluding Remarks

Location Determination or Localization

- Indoors/Outdoors
- Active
  - Node actively participates in determining the location – participates in sending/receiving/processing messages
- Passive
  - Node, held by a human, does not participate in location determination
    - Essentially locating a human being.

Active Localization

- Measure
  - Distance
  - Some function of distance $\tau(d)$
  - Some function of Location

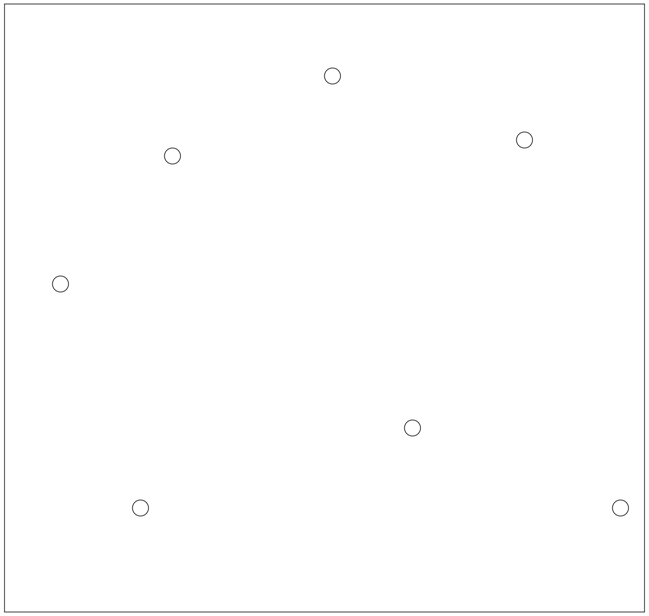

$$R(x,y,z) = \begin{bmatrix} r_1 \\ r_2 \\ \cdots \\ r_n \end{bmatrix} (x,y,z)$$

Signal Strength Function

- If we know the function

$$R(x, y, z)$$

- Measure R at a location and invert the function
  - Easy??
  - Practical Realities are complex

Outline

- Localization – Active Techniques
- RSSI Based
  - Characteristics of 802.11b signals
  - Horus
- Transit Time Based/ Synchronization
  - PinPoint
  - System Synchronization
- Localization – Passive Techniques
  - Nuzzer
- Concluding Remarks

Spatial Variations: Small-Scale

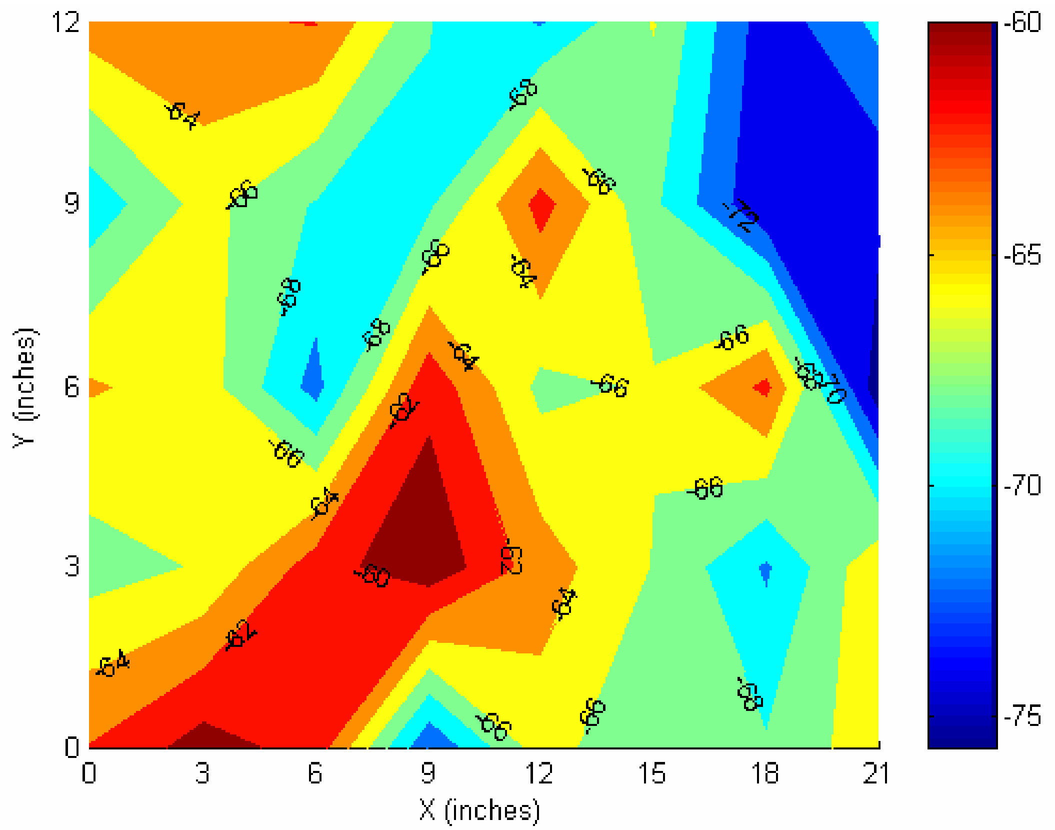

- Multipath effect

Spatial Variations: Large-Scale

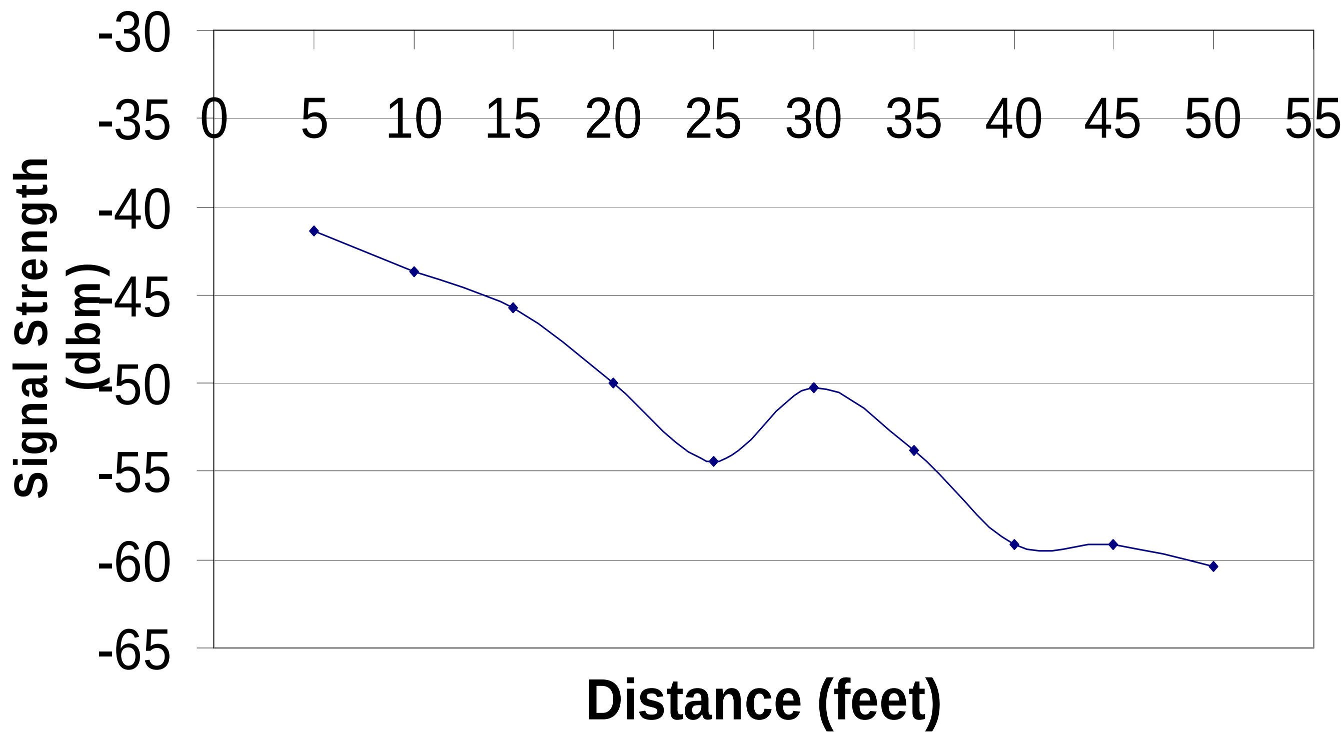

- Desirable !

Temporal Variations: One Access Point

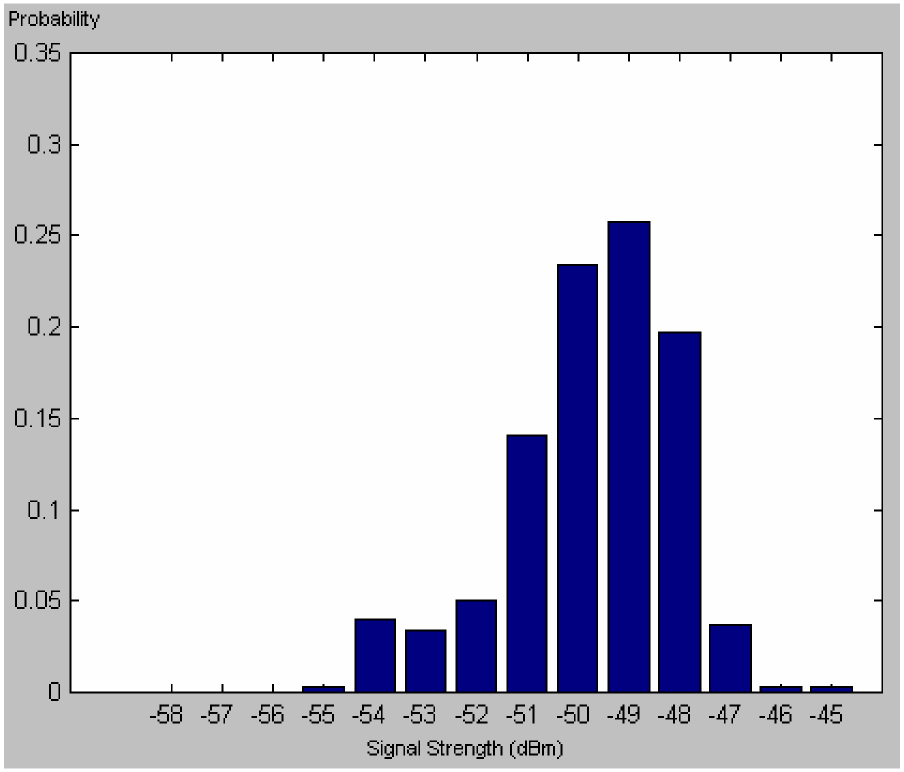

- Environment changes
- Using average only leads to loss of information

The Maryland Information and Network Dynamics Lab

Temporal Variations: Multiple Access Points

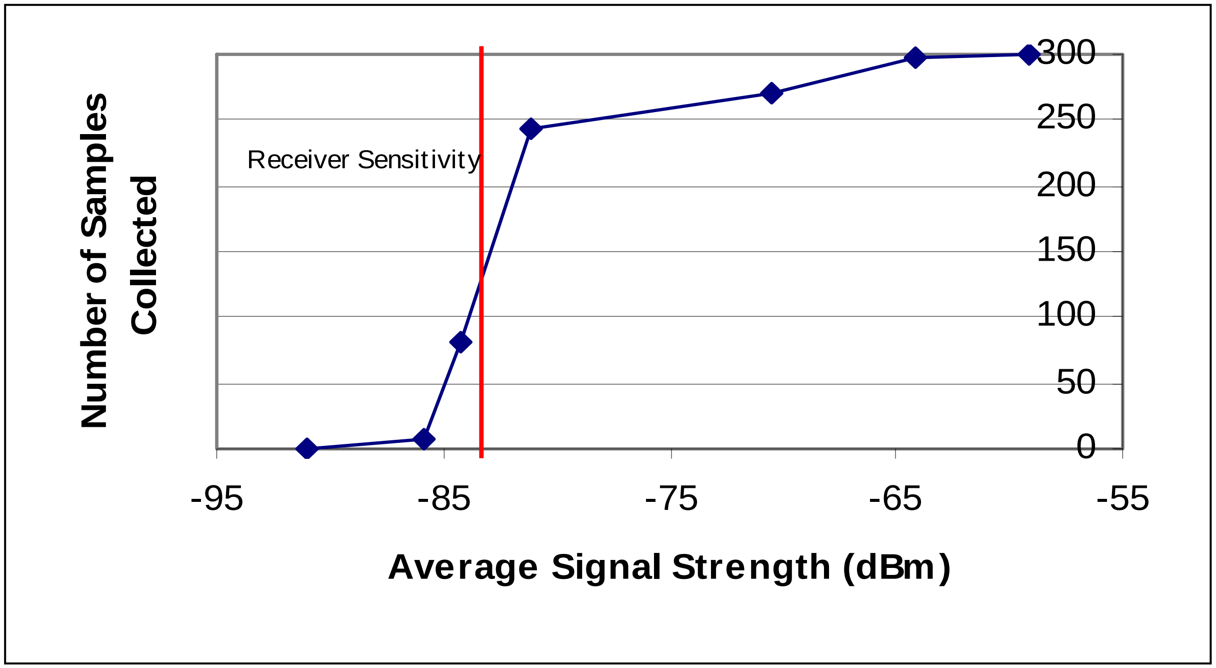

- Number of access points changes over time
- Choose the strongest access points

Temporal Variations: Correlation

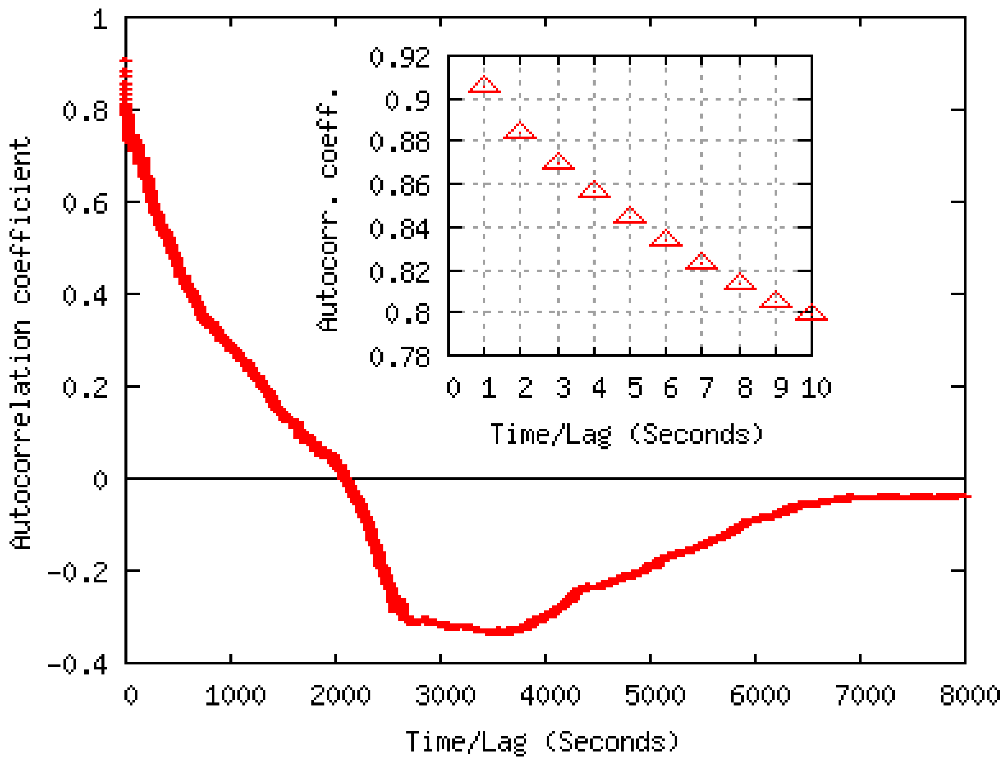

- Independence assumption is wrong

The Maryland Information and Network Dynamics Lab

Environmental Factors

- Distance
  - Used in determining location – Horus Technology
- Multipath
  - Always there indoors
- Objects
  - May effect
    - Door open vs closed
- People
  - Presence and movement always affects the signal

- Can we use the infrastructure to determine the presence of people ?

Vault Measurements

- Does the RSSI vary in controlled environments?
- Bank Vault
  - CISCO AP
  - Measure RSSI in controlled environment

Measurement Example

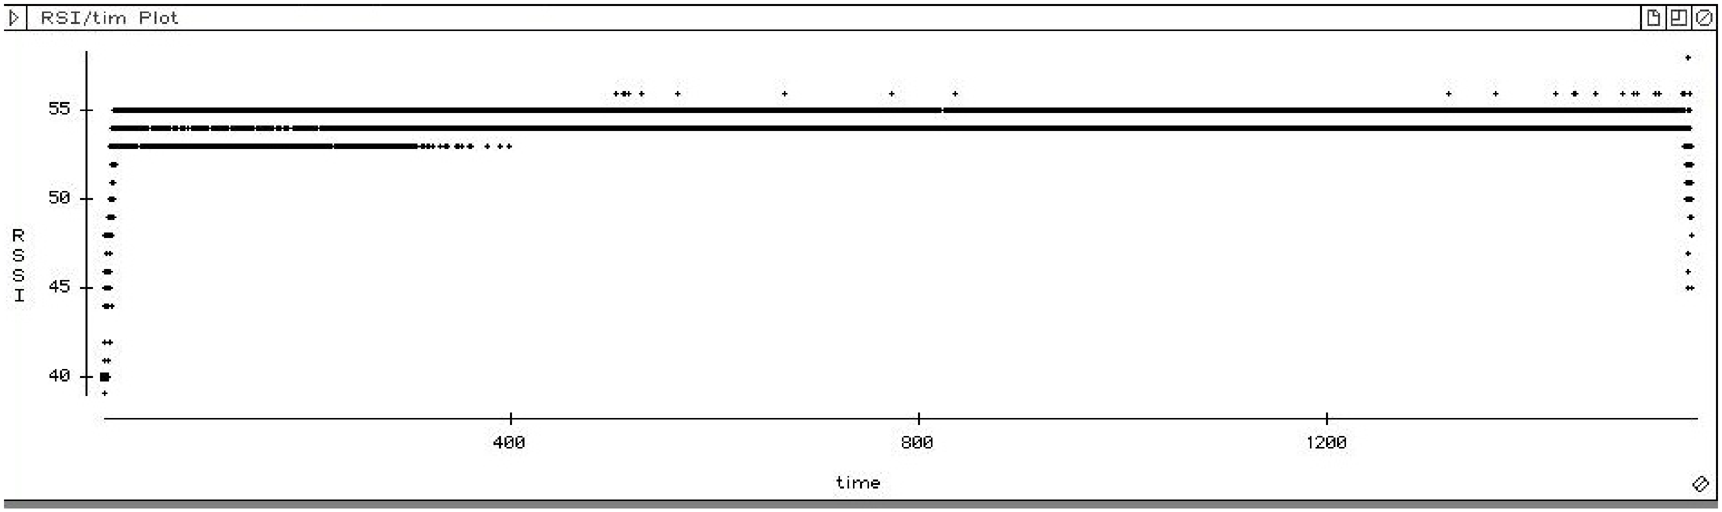

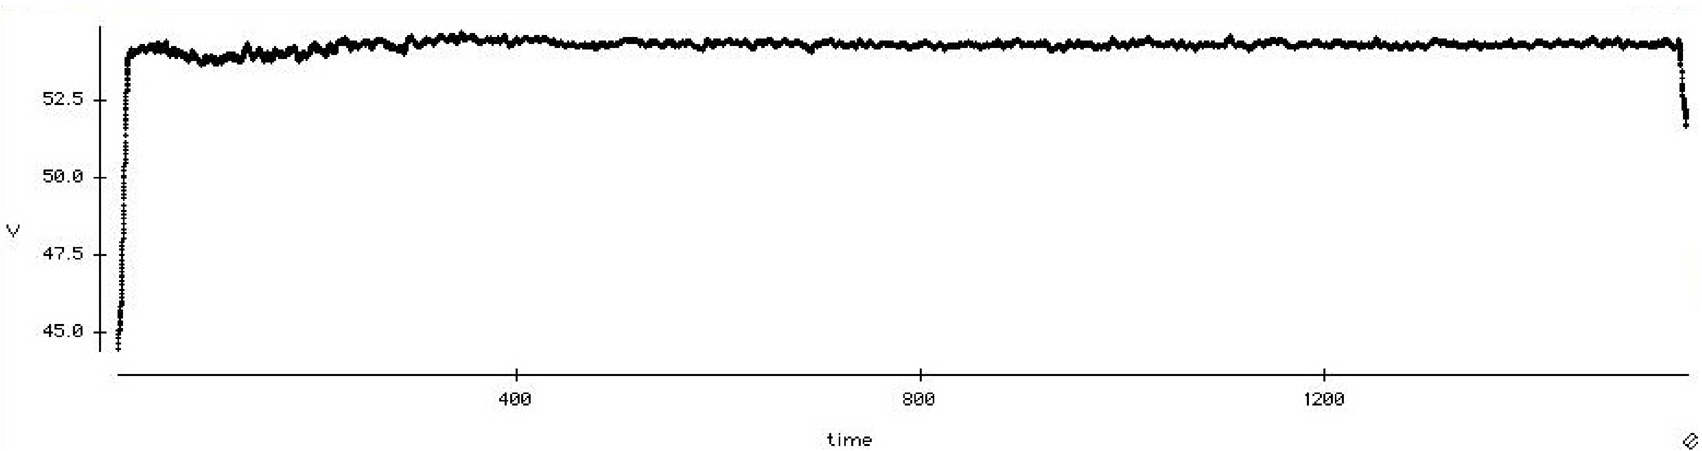

The Maryland Information and Network Dynamics Lab

Noise in NICs

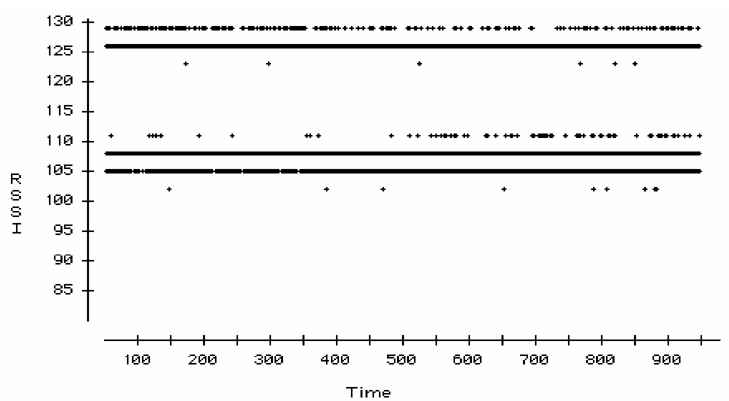

Figure 1. Observed RSSI values for Compaq Wireless LAN Adapter during 15 minutes period.

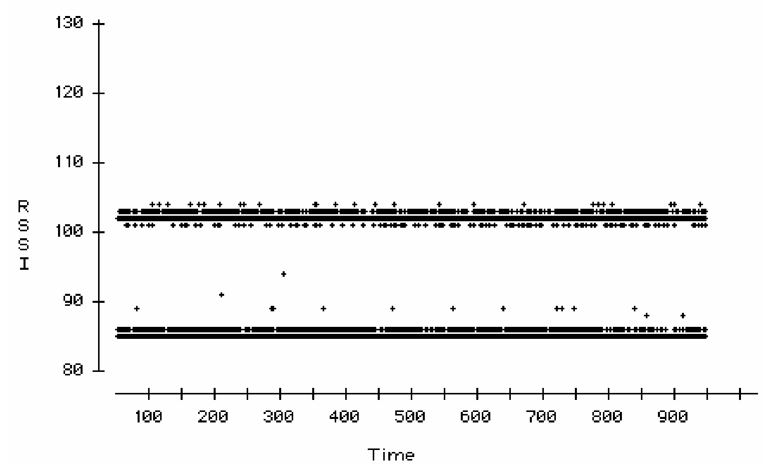

Figure 2. Observed RSSI values for Demarc Wireless LAN Adapter during 15 minutes period.

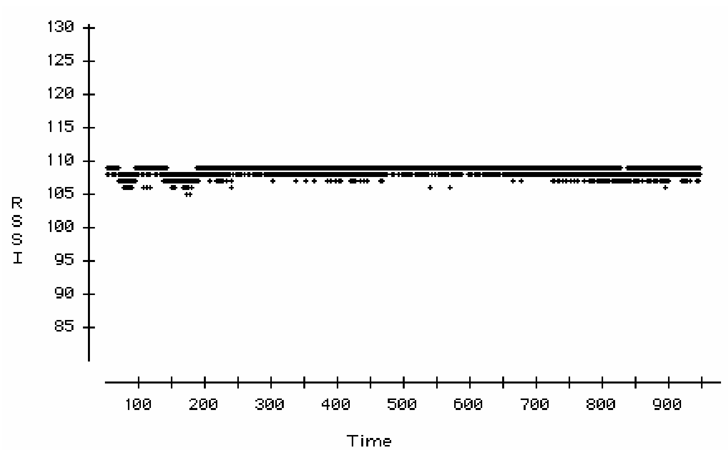

Figure 3. Observed RSSI values for Orinoco Wireless LAN Adapter during 15 minutes period.

NIC	RSSI	SD	NormSD
Orinoco8	108.57	0.618	0.569
Compaq24	116.75	9.598	8.221
Generic	94.24	8.438	8.954
ZoomAir-Ant	42.299	5.816	13.750
ZoomAir-NoAnt	31.74	5.113	16.108
LinkSYS	45.506	0.566	1.243
Orinoco12	105.03	0.839	0.799

Table 1. RSSI Measurements Comparisons based on Ethereal network analyzer.

UNIVERSITY OF MARYLAND 1856

MIND The MIND Lab
The Maryland Information and Network Dynamics Lab

NIC Performance

- NICs available in the market vary in performance

NIC	RSSI	SD	NormSD
Orinoco8	108.57	0.618	0.569
Compaq24	116.75	9.598	8.221
Generic	94.24	8.438	8.954
ZoomAir-Ant	42.299	5.816	13.750
ZoomAir-NoAnt	31.74	5.113	16.108
LinkSYS	45.506	0.566	1.243
Orinoco12	105.03	0.839	0.799

NIC	RSSI	SD	NormSD
Orinoco3	53.35	0.673	1.262
Orinoco8	51.38	0.488	0.950
Orinoco7	51.21	1.683	3.286
Orinoco17	52.57	0.703	1.338
Orinoco12	53.48	0.598	1.118
Cisco249	100.00	0.000	0.000
Cisco175	100.00	0.000	0.000
CISCO872	100.00	0.000	0.000
CiscoMind11	97.99	0.233	0.237
Cisco138	99.89	0.453	0.453

Vault Measurement Results

- AP power does not vary
  - Measured using two sniffers
  - No correlation between the two measurements
    - Implies AP power variability is not there
- Noise introduced by NIC can be significant
  - ZoomAir
- Some NICS introduce very little or No Noise.
  - CISCO

Another Measurement

- In AVW
- Over 12 hour period
- From 6:30 PM on
- 50,000 Seconds

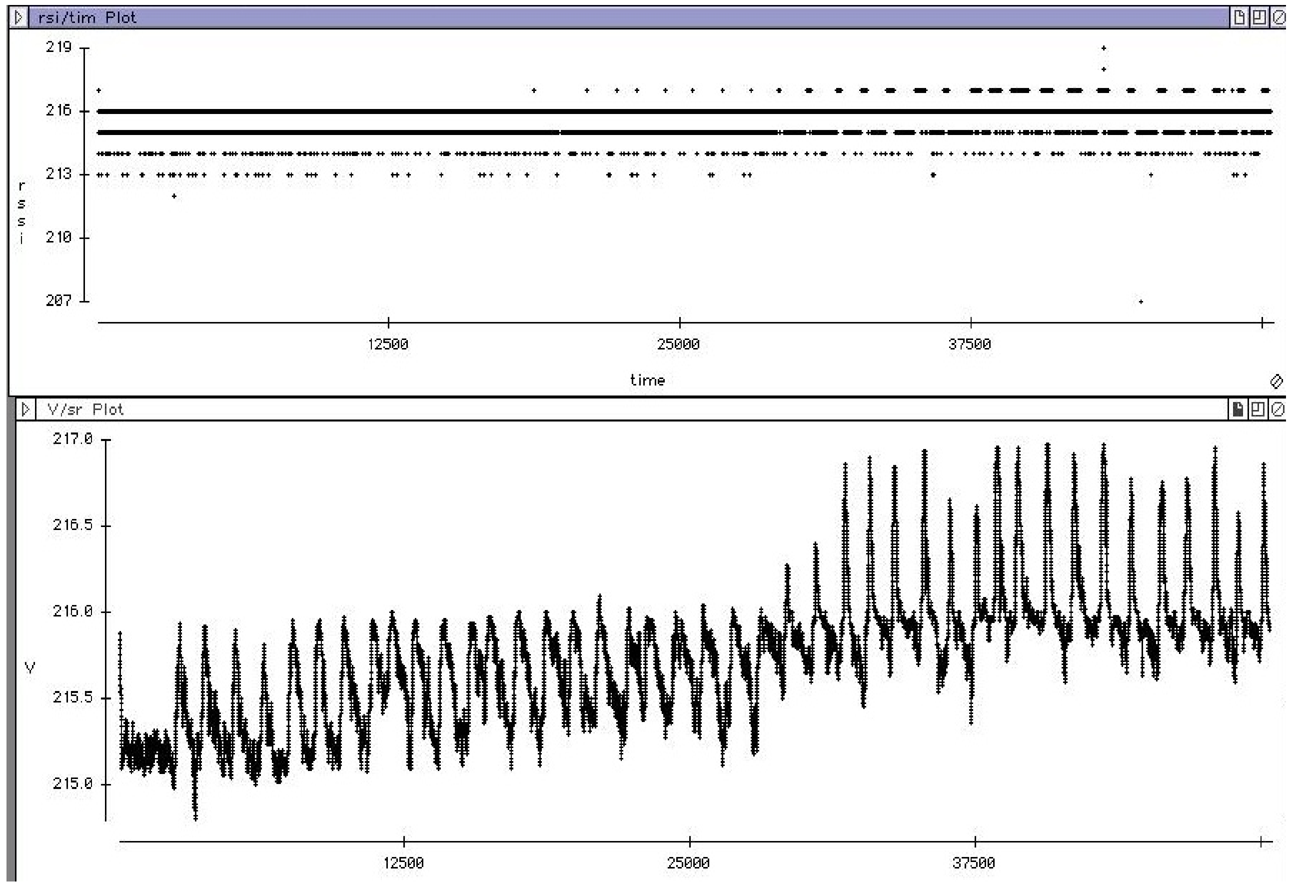

The Maryland Information and Network Dynamics Lab

Outline

- Localization – Active Techniques
- RSSI Based
  - Characteristics of 802.11b signals
  - Horus (PhD. Work of Moustafa Youssef)
- Transit Time Based/ Synchronization
  - PinPoint
  - System Synchronization
- Localization – Passive Techniques
  - Nuzzer
- Concluding Remarks

HORUS Technology
Basic Algorithm: Mathematical Formulation

- x: Position vector
- s: Signal strength vector
  - One entry for each access point
- s(x) is a stochastic process
- P[s(x), t]: probability of receiving s at x at time t
- s(x) is a stationary process
  - P[s(x)] is the histogram of signal strength at x
- $\text{Argmax}_x[P(x/s)]$
- Using Bayesian inversion
  - $\text{Argmax}_x[P(s/x).P(x)/P(s)]$
  - $\text{Argmax}_x[P(s/x).P(x)]$
- P(x): User history

Horus Components

- Basic algorithm
- Correlation handler
- Continuous space estimator
- Small-scale compensator
- Locations clustering

Basic Algorithm: Radio map

[Percom03] [CNDS04]

- Offline phase
  - Radio map: signal strength histograms
- Online phase
  - Bayesian based inference

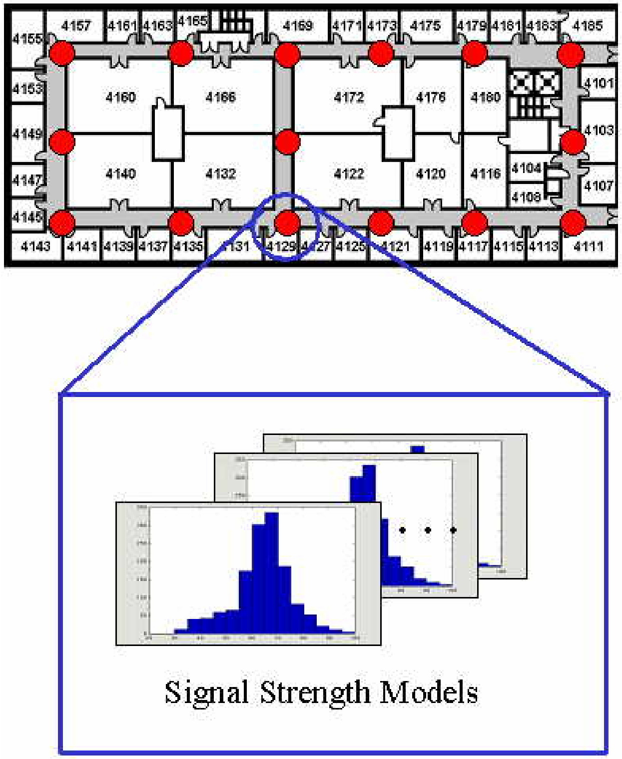

The Maryland Information and Network Dynamics Lab

Basic Algorithm: Example

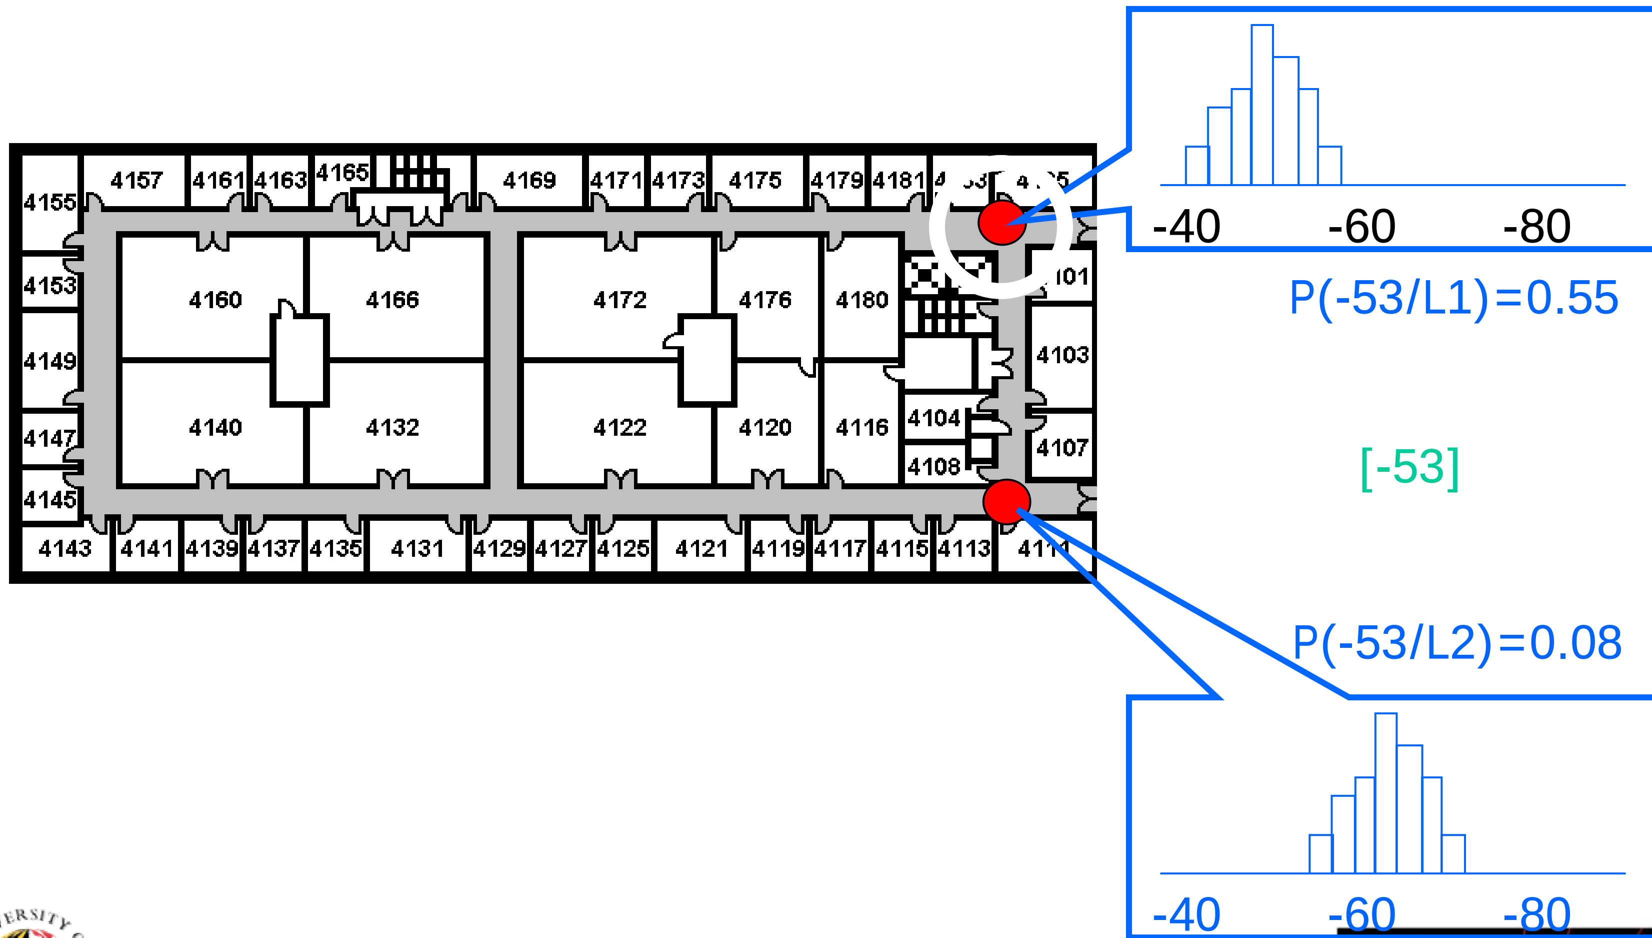

Basic Algorithm:
Parametric Distributions

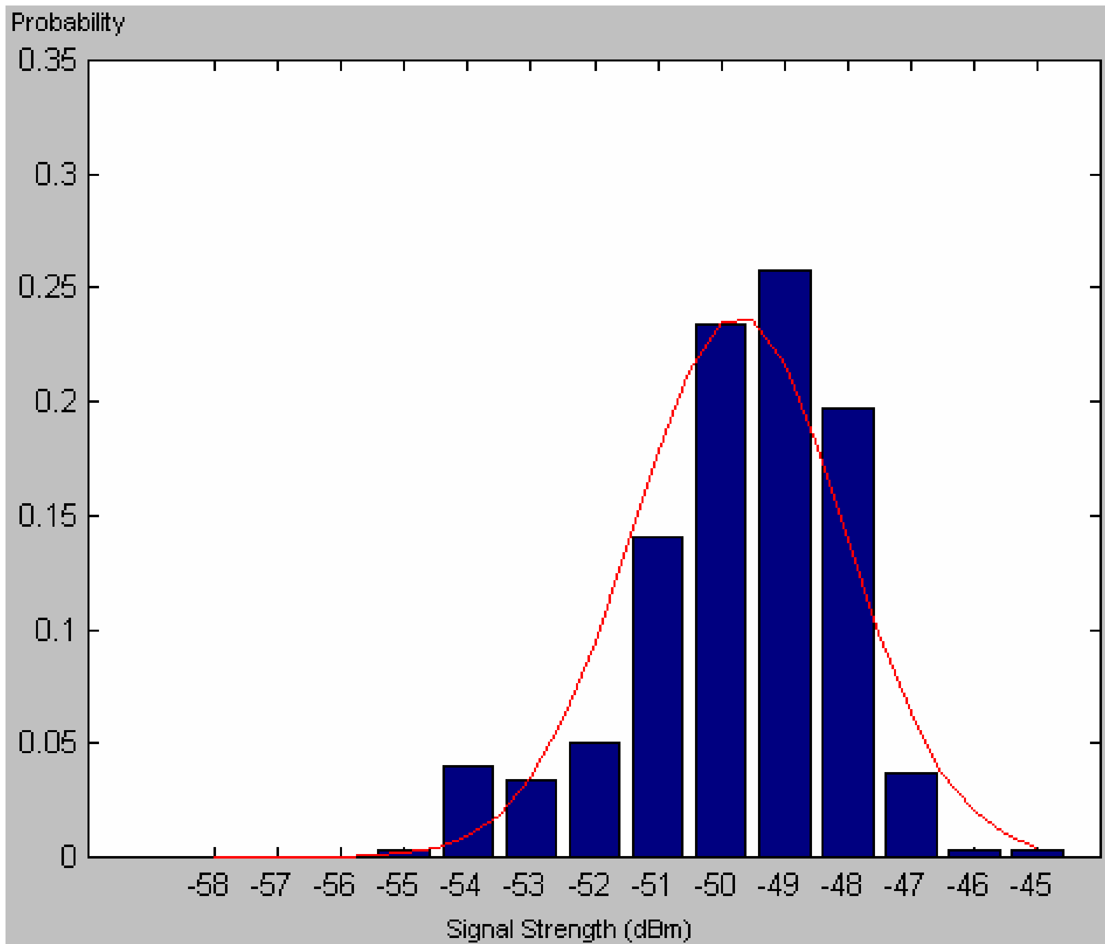

The Maryland Information and Network Dynamics Lab

Basic Algorithm:
Results

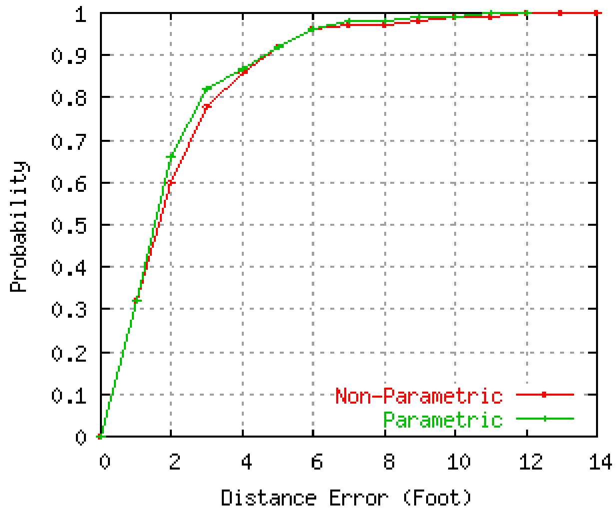

- Accuracy of 5 feet 90% of the time
- Slight advantage of parametric over non-parametric method
  - Smoothing of distribution shape

Correlation Handler
[InfoCom04]

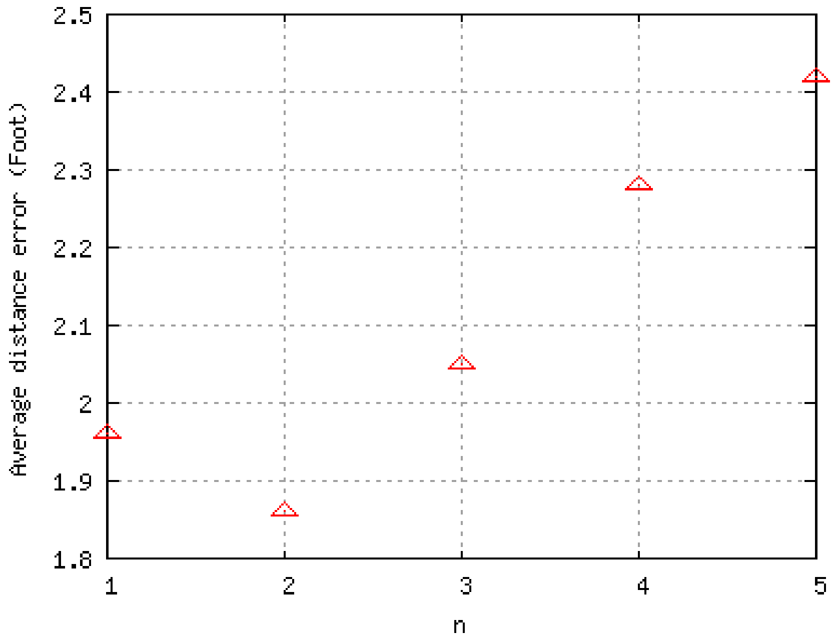

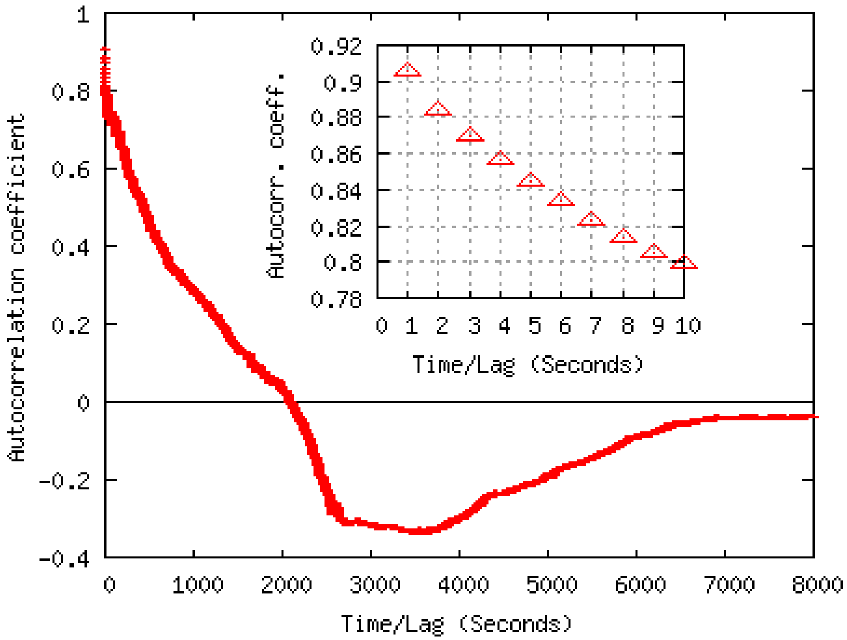

- Need to average multiple samples to increase accuracy
- Independence assumption is wrong

Correlation Handler: Autoregressive Model

- $s(t+1)=\alpha.s(t)+(1-\alpha).v(t)$
- α: correlation degree
- $E[v(t)]=E[s(t)]$
- $Var[v(t)]= (1+\alpha)/(1-\alpha)\ Var[s(t)]$
- $s(t+1)=\alpha.s(t)+(1-\alpha).v(t)$
- $s \sim N(0, m)$
- $v \sim N(0, r)$
- $A=1/n\ (s_1+s_2+...+s_n)$
- $E[A(t)]=E[s(t)]=0$
- $Var[A(t)]= m^2/n^2\ \{\ [(1-\alpha^{n})/(1-\alpha)]^2 + n+ 1-\alpha^{2} * (1-\alpha^{2(n-1)})/(1-\alpha^{2})\ \}$

The Maryland Information and Network Dynamics Lab

Correlation Handler: Var(A)/Var(s)

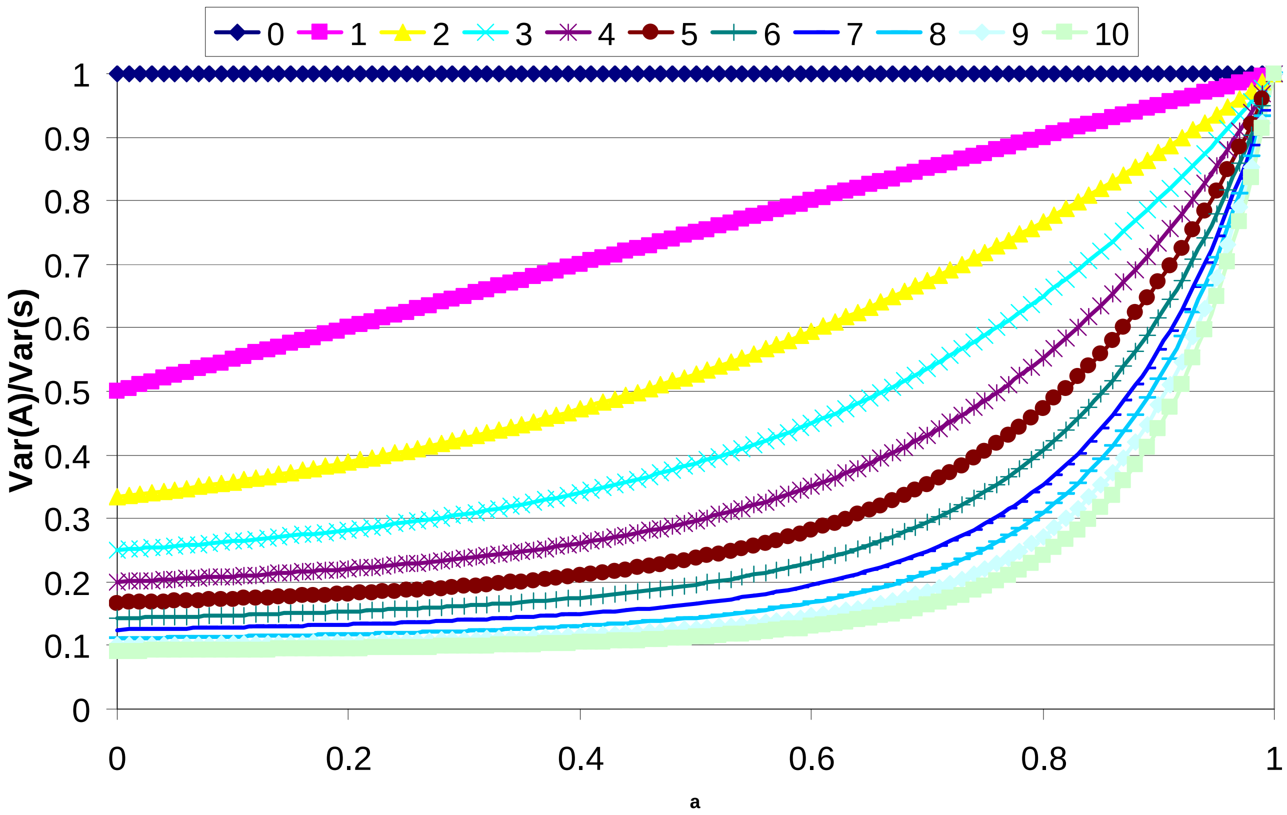

- Independence assumption underestimates true variance

Correlation Handler:
Results

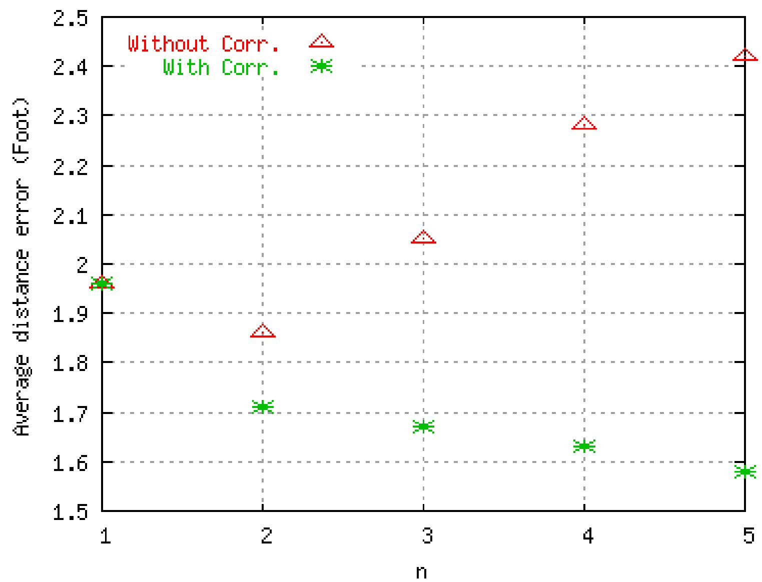

- Independence assumption: performance degrades as n increases
- Two factors affecting accuracy
  - Increasing n
  - Deviation from the actual distribution

Continuous Space Estimator

- Enhance the discrete radio map space estimator
- Two techniques
  - Center of mass of the top ranked locations

$$x \leftarrow \frac{\sum_{i=1}^{\min(N,\|\bar{\mathbb{X}}\|)} p(i) * \bar{\mathbb{X}}(i)}{\sum_{i=1}^{\min(N,\|\bar{\mathbb{X}}\|)} p(i)}$$

  - Time averaging

$$\bar{x}_t = \frac{1}{\min(W,t)} \cdot \sum_{t-\min(W,t)+1}^{t} x_i$$

Center of Mass:
Results

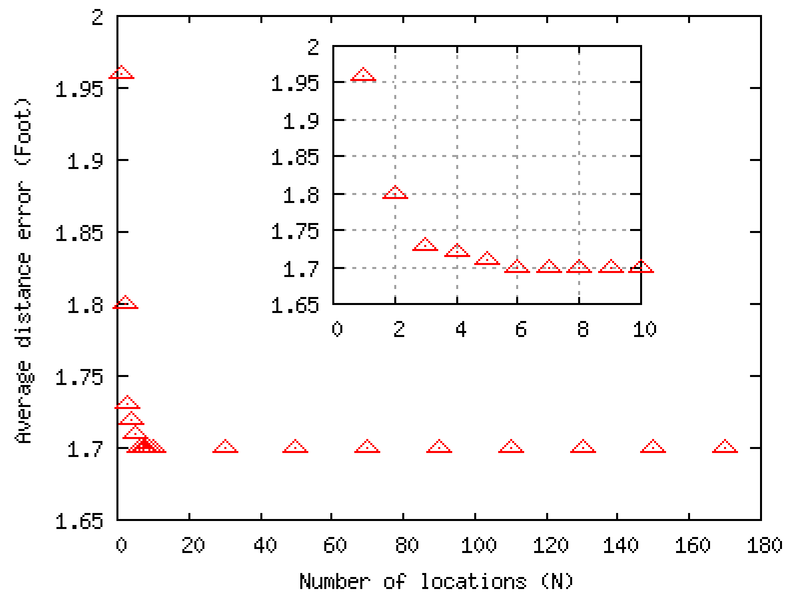

- N = 1 is the discrete-space estimator
- Accuracy enhanced by more than 13%

The Maryland Information and Network Dynamics Lab

Time Averaging Window: Results

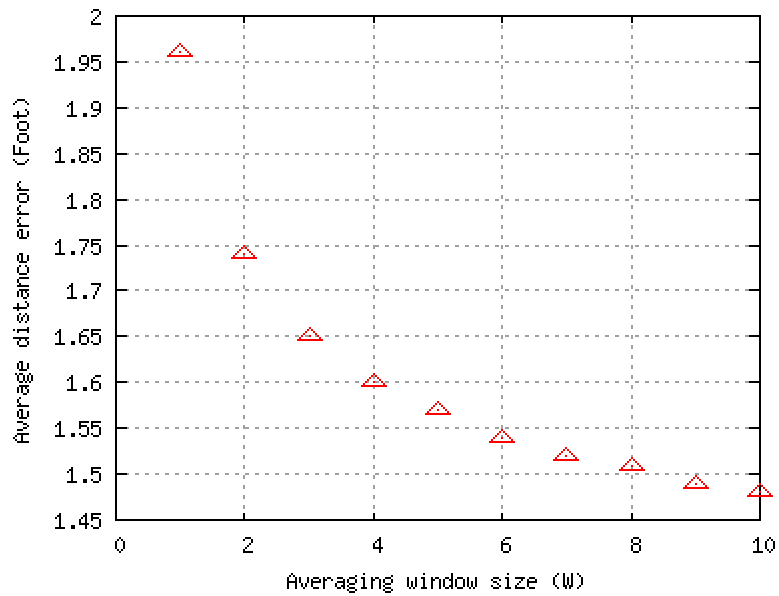

- N = 1 is the discrete-space estimator
- Accuracy enhanced by more than 24%

Small-scale Compensator

- [WCNC03]

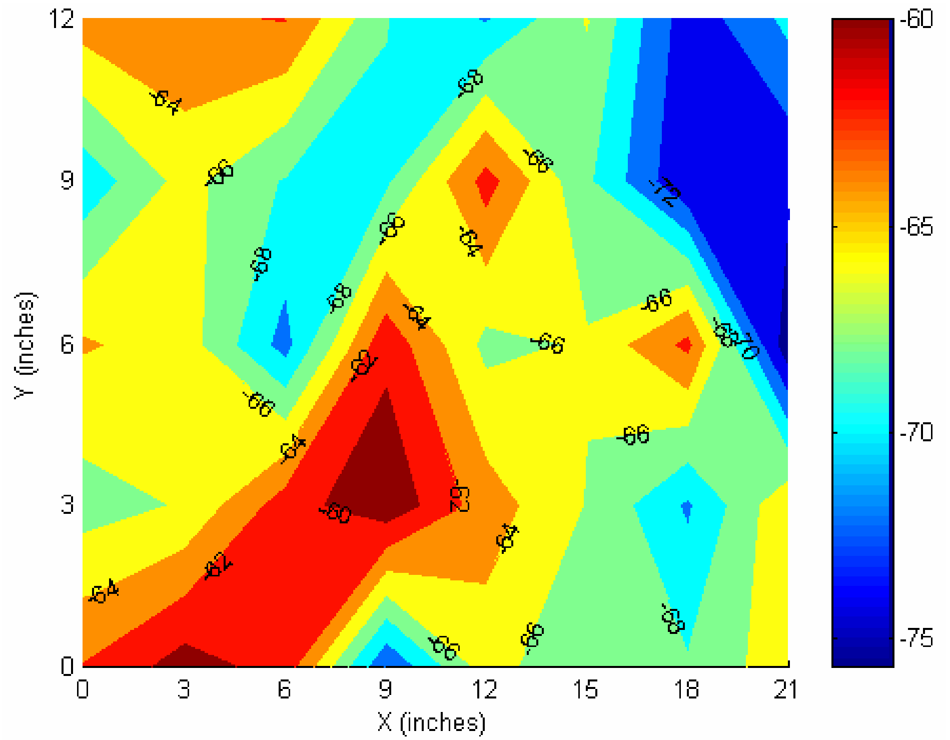

- Multi-path effect
- Hard to capture by radio map (size/time)

The Maryland Information and Network Dynamics Lab

Small-scale Compensator: Small-scale Variations

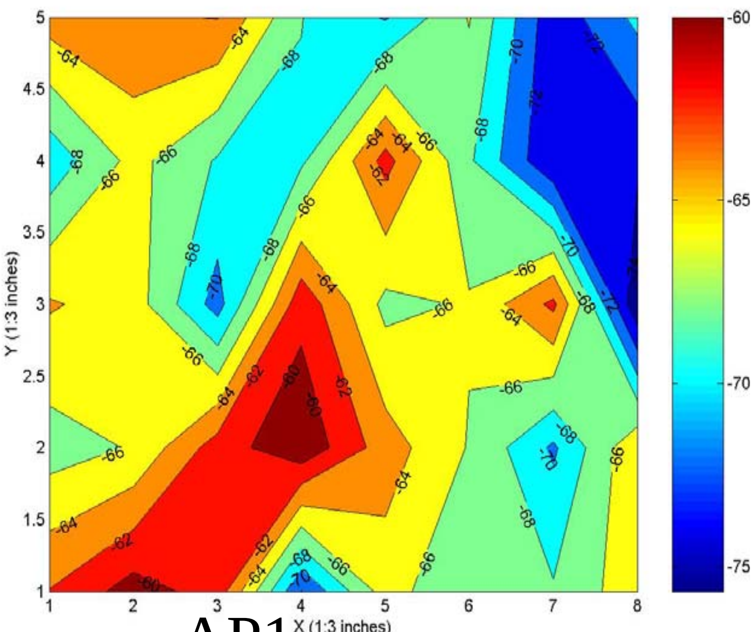

AP1

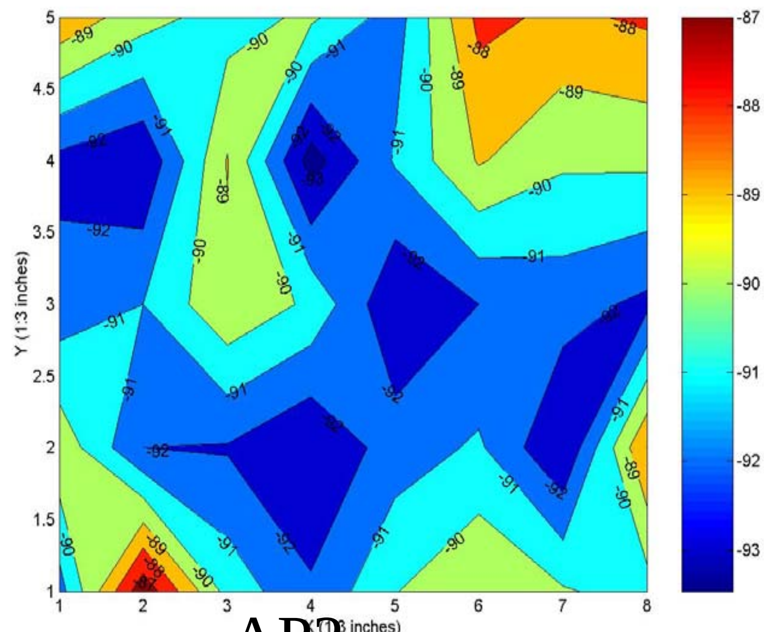

AP2

- Variations up to 10 dBm in 3 inches
- Variations proportional to average signal strength

Small-scale Compensator: Perturbation Technique

- Detect small-scale variations
  - Using previous user location
- Perturb signal strength vector
  - $(s_1, s_2, \ldots, s_n) \rightarrow (s_1 \pm d_1, s_2 \pm d_2, \ldots, s_n \pm d_n)$
  - Typically, n=3-4
- d_i is chosen relative to the received signal strength

Small-scale Compensator: Results

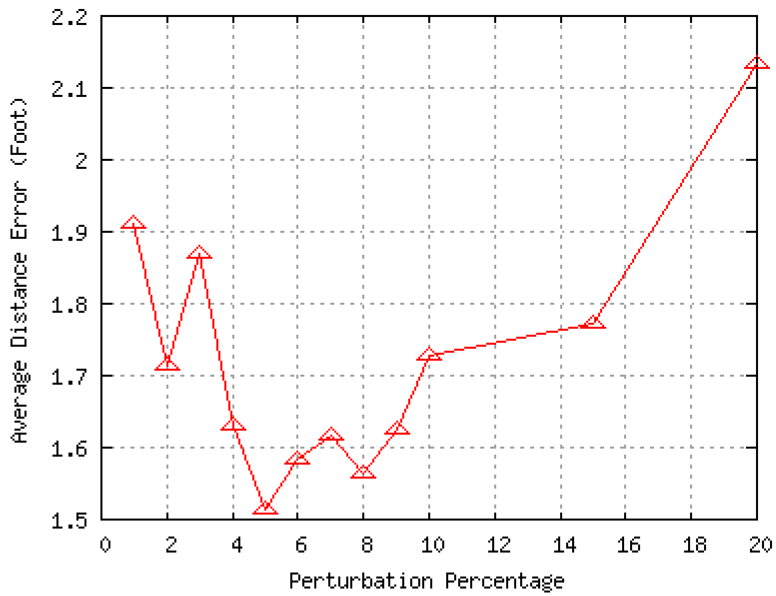

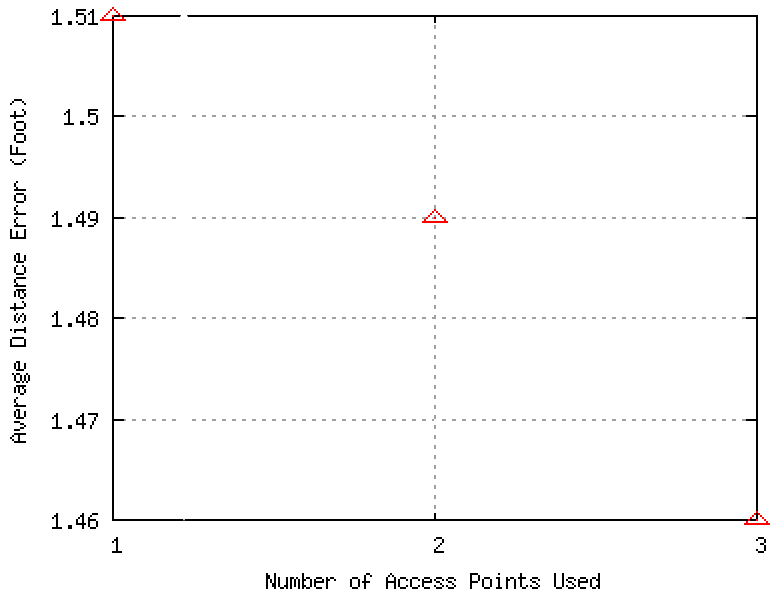

- Perturbation technique is not sensitive to the number of APs perturbed
- Better by more than 25%

Locations Clustering
[Percom03]

- Reduce computational requirements
- Two techniques
  - Explicit
  - Implicit

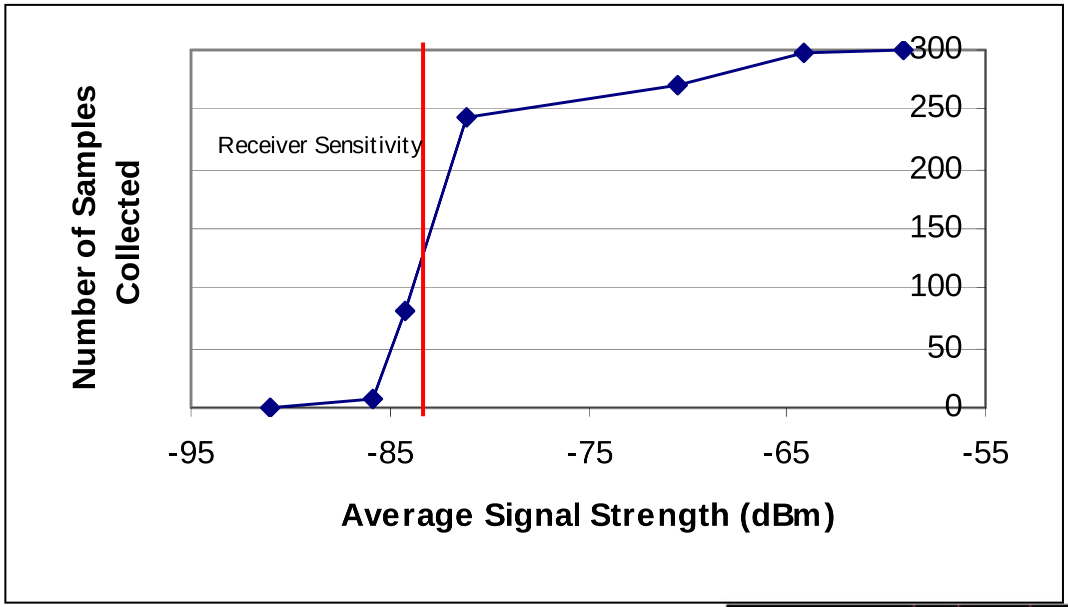

The Maryland Information and Network Dynamics Lab

Locations Clustering: Explicit Clustering

- Use access points that cover each location
- Use the *q* strongest access points

S=[-60, -45, -80, -86, -70]

S=[-45, -60, -70, -80, -86]

q=3

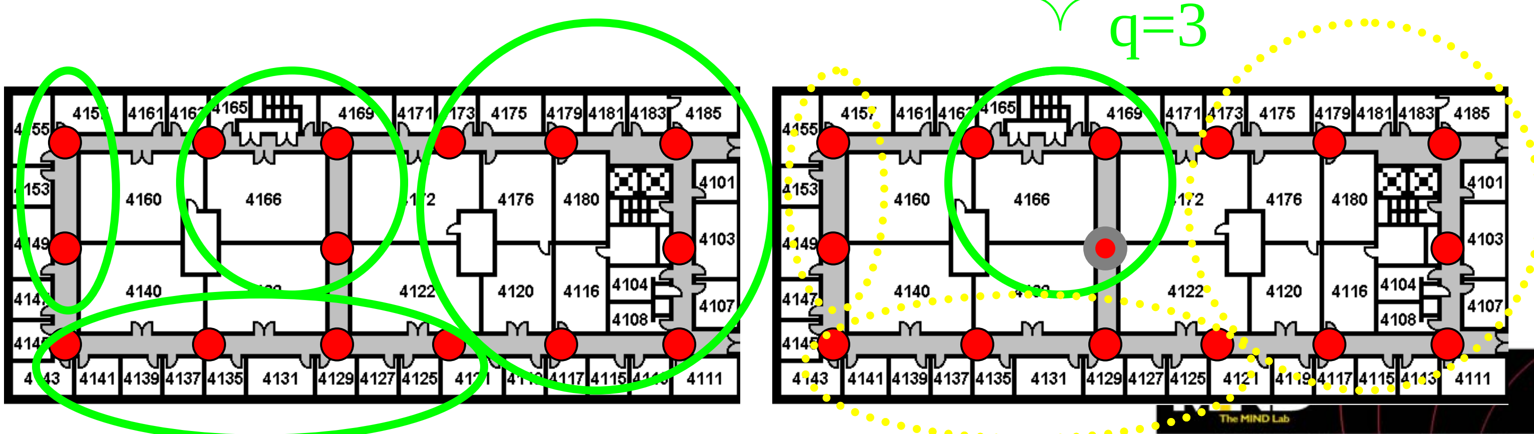

Locations Clustering:

Results- Explicit Clustering

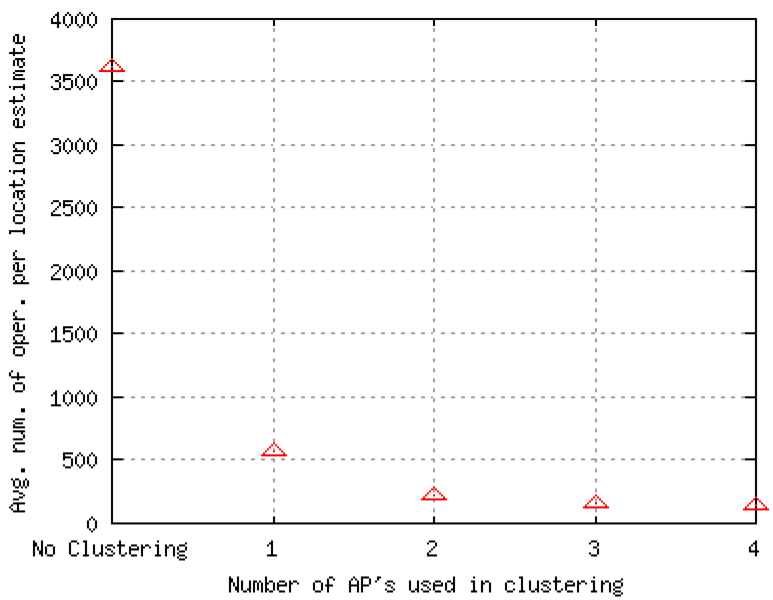

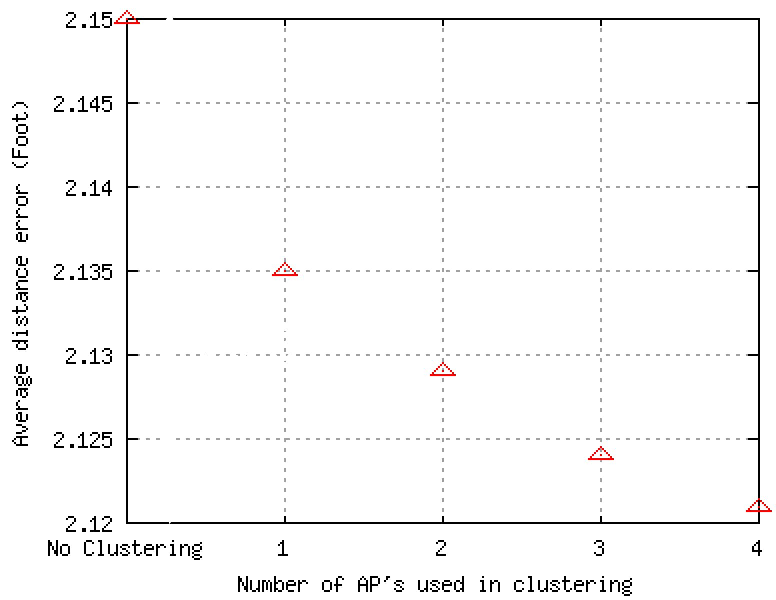

- An order of magnitude enhancement in avg. num. of oper. /location estimate
- As q increases, accuracy slightly increases

Locations Clustering: Implicit Clustering

- Use the access points incrementally
- Implicit multi-level clustering

S=[-60, -45, -80, -86, -70]

S=[-45, -60, -70, -80, -86]

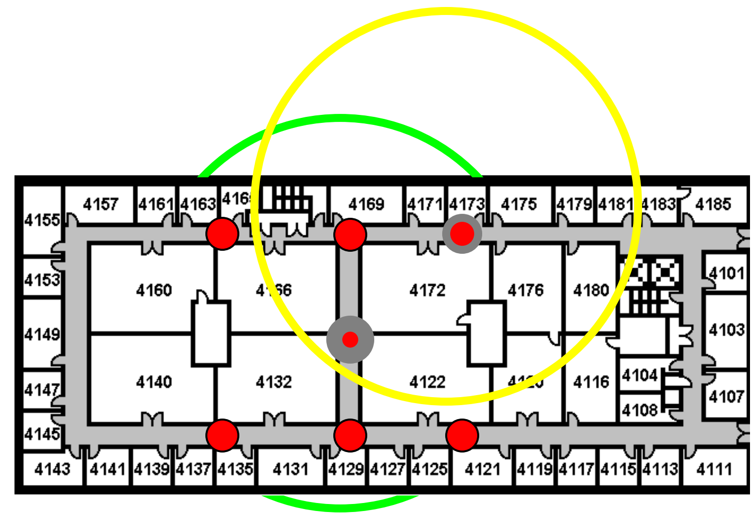

Locations Clustering:

Results- Implicit Clustering

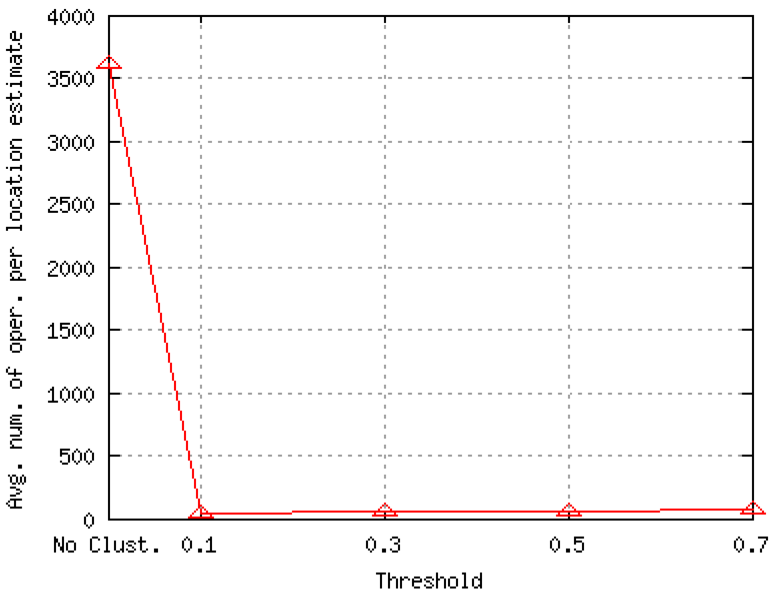

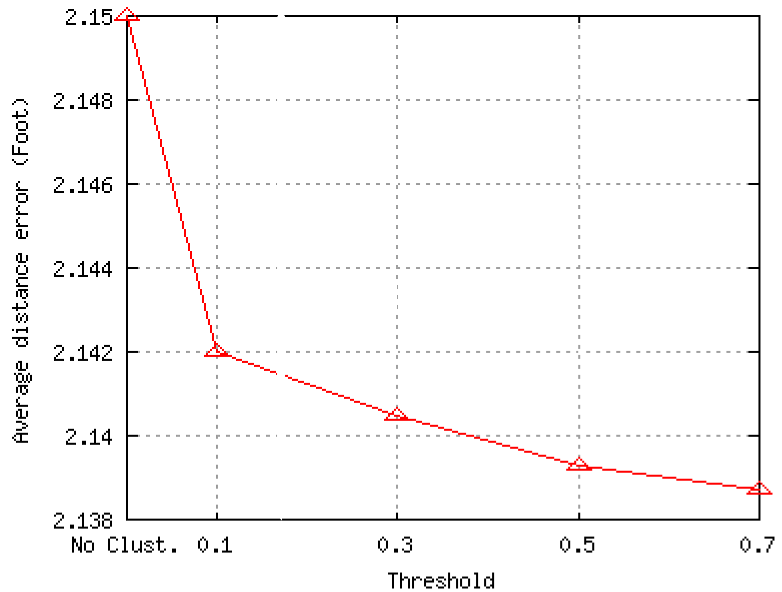

- Avg. num. of oper. /location estimate better than explicit clustering
- Accuracy increases with Threshold

Testbeds

- A.V. William's
  - 4th floor, AVW
  - 224 feet by 85.1 feet
  - UMD net (*Cisco* APs)
  - 21 APs (6 on avg.)
  - 172 locations
  - 5 feet apart
  - *Windows XP Prof.*

- FLA
  - 3rd floor, 8400 Baltimore Ave
  - 39 feet by 118 feet
  - LinkSys/Cisco APs
  - 6 APs (4 on avg.)
  - 110 locations
  - 7 feet apart
  - Linux (kernel 2.5.7)

Orinoco/Compaq cards

Horus-Radar Comparison

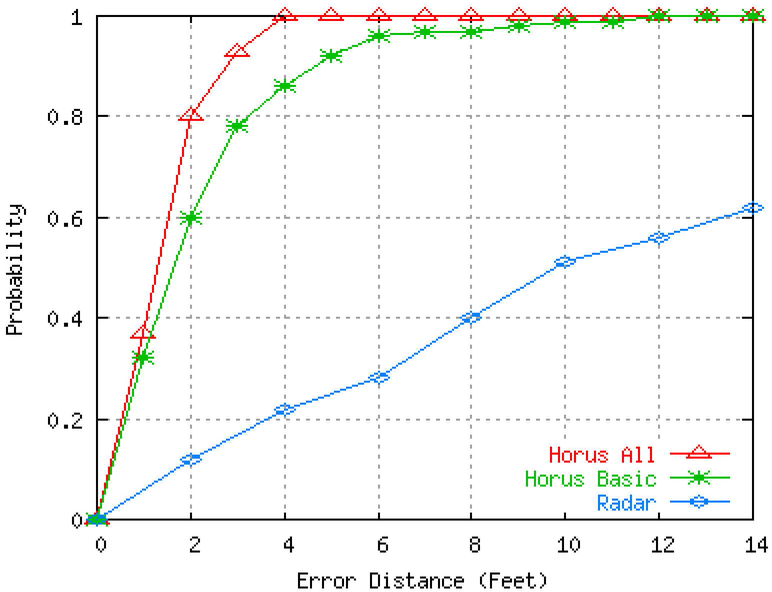

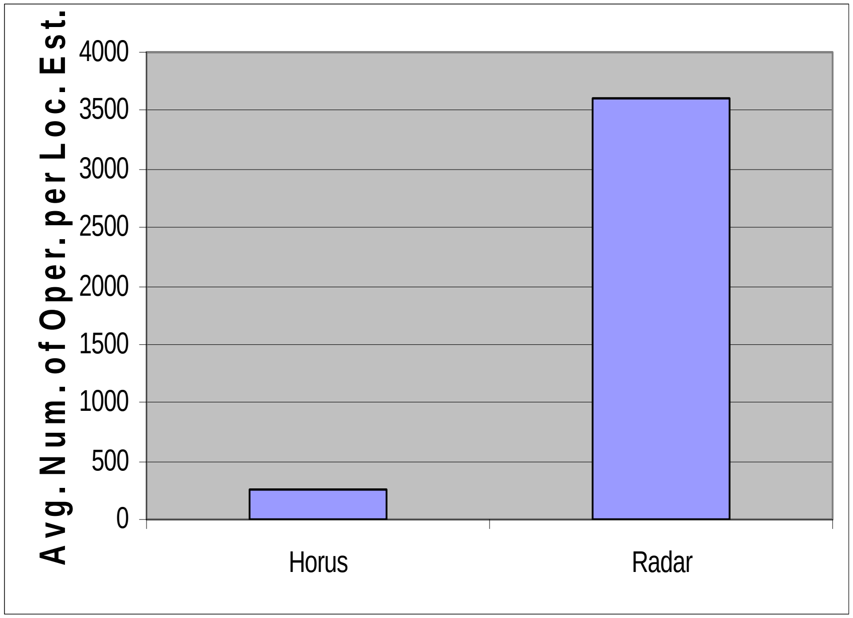

	Median	Avg	Stdev	Max
Horus (all components)	1.28	1.38	0.95	4
Horus (basic)	1.6	2.16	2.09	18.08
Radar	9.74	13.15	10.71	57.67

The Maryland Information and Network Dynamics Lab

Comparison With Other Systems: Ekahau

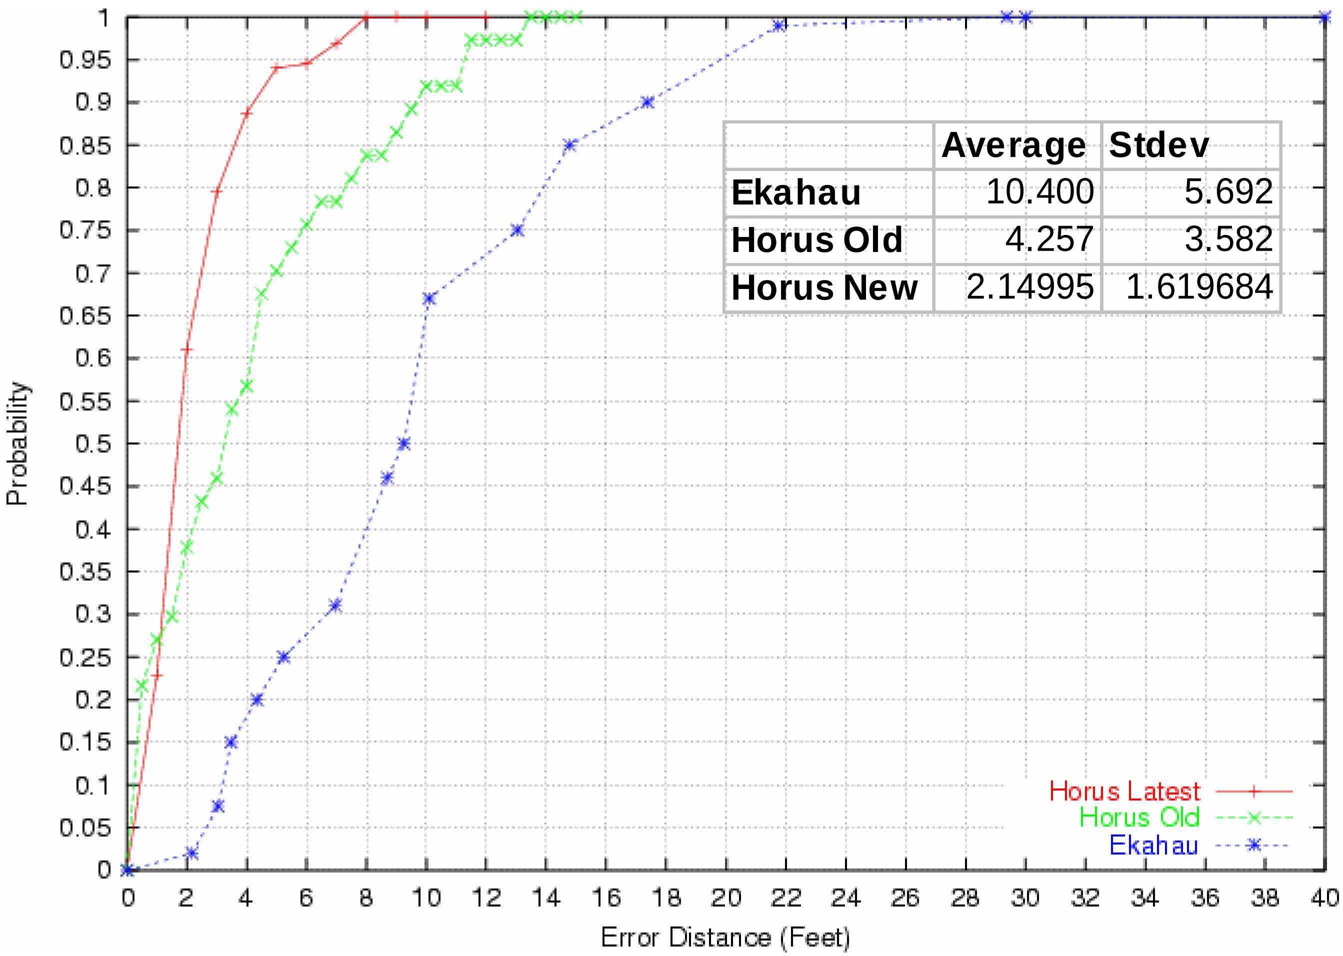

	Average	Stdev
Ekahau	10.400	5.692
Horus Old	4.257	3.582
Horus New	2.14995	1.619684

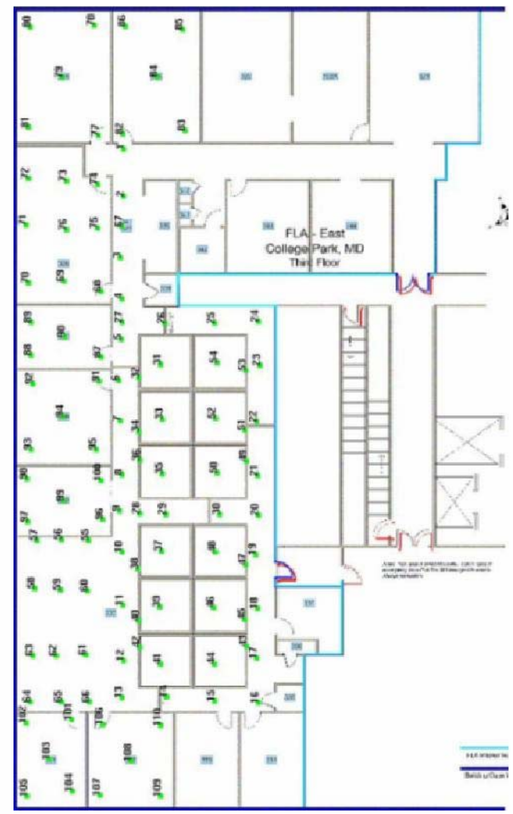

Radar with Horus Techniques

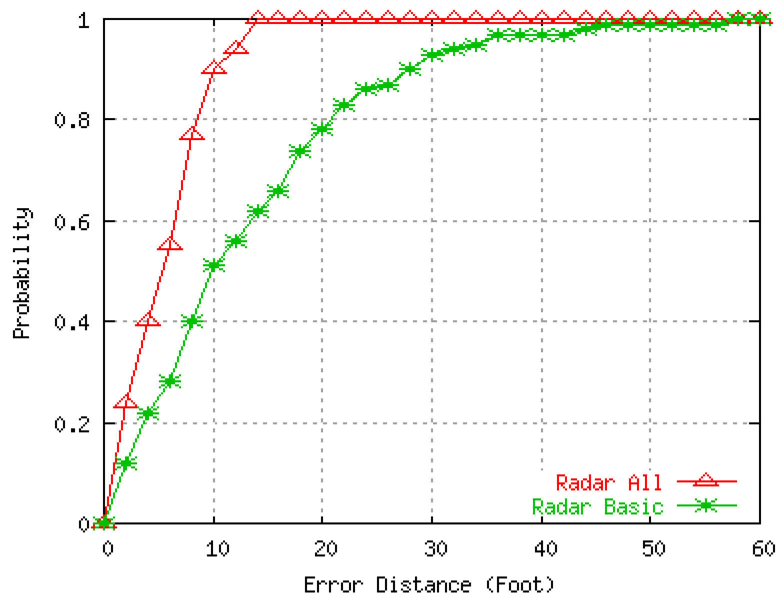

- Average distance error enhanced by more than 58%
- Worst case error decreased by more than 76%

The Maryland Information and Network Dynamics Lab

Horus Status

- The *Horus* system achieves its goals
- High accuracy
  - Through a probabilistic location determination technique
  - Smoothing signal strength distributions by Gaussian approximation
  - Using a continuous-space estimator
  - Handling the high correlation between samples from the same access point
  - The perturbation technique to handle small-scale variations
- Low computational requirements
  - Through the use of clustering techniques
- Scalability in terms of the coverage area
  - Through the use of clustering techniques
- Scalability in terms of the number of users
  - Through the distributed implementation
- Training time of 15 seconds per location is enough to construct the radio-map
- Radio map spacing of 14 feet
- Horus vs. Radar
  - More accurate by more than 11 feet, on the average
  - More than an order of magnitude savings in number of operations required per location estimate
- Horus vs. Ekahau

Outline

- Localization – Active Techniques
- RSSI Based
  - Characteristics of 802.11b signals
  - Horus
- Transit Time Based/ Synchronization
  - PinPoint
  - System Synchronization
- Localization – Passive Techniques
  - Nuzzer
- Concluding Remarks

Time-Based Approach

- Determine the distance by measuring the flight time of signal
- Accuracy of distance measurement depends on the clock resolution
  - 1 ns = 30 cm
- Roundtrip measurement vs. synchronized clocks
- Can we use stable clocks and determine location/time ?
- PinPoint technology
  - Joint work with A.U. Shankar, R.L. Larsen and D. Szajda

Problem

- Consider a collection of nodes
- Each node has
  - Unique ID (10 bits)
  - A clock with one nanosecond resolution
  - Processor and storage capability
- Each capable of
  - Sending and receiving digital information using UHF
  - Time Stamping using 64 bit time stamp with ns resolution

Can each node know the topology of all nodes it can talk to?

Can each node know enough to carry out a synchronous action with other nodes?

Node Structure

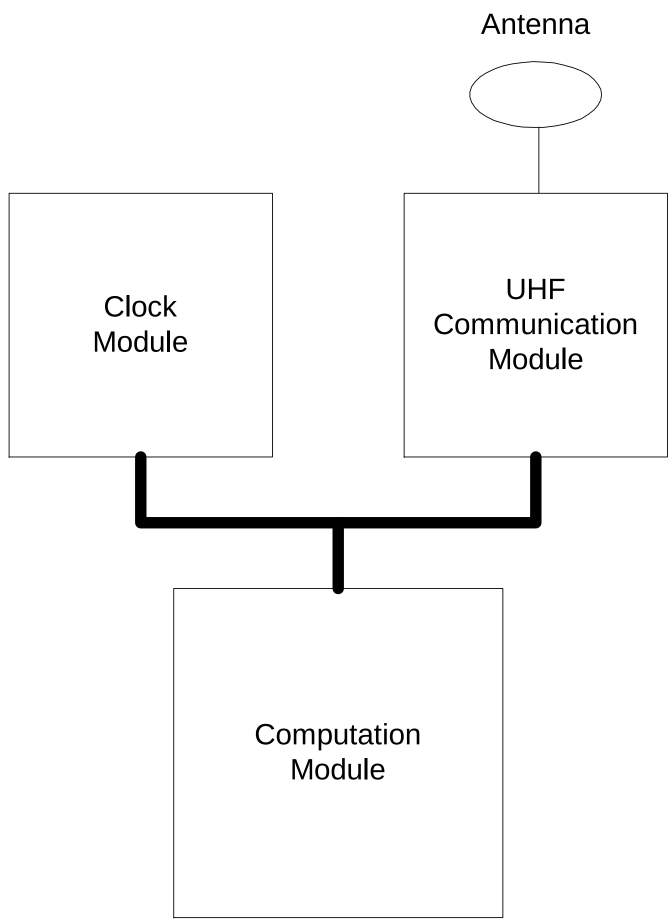

Clock Module

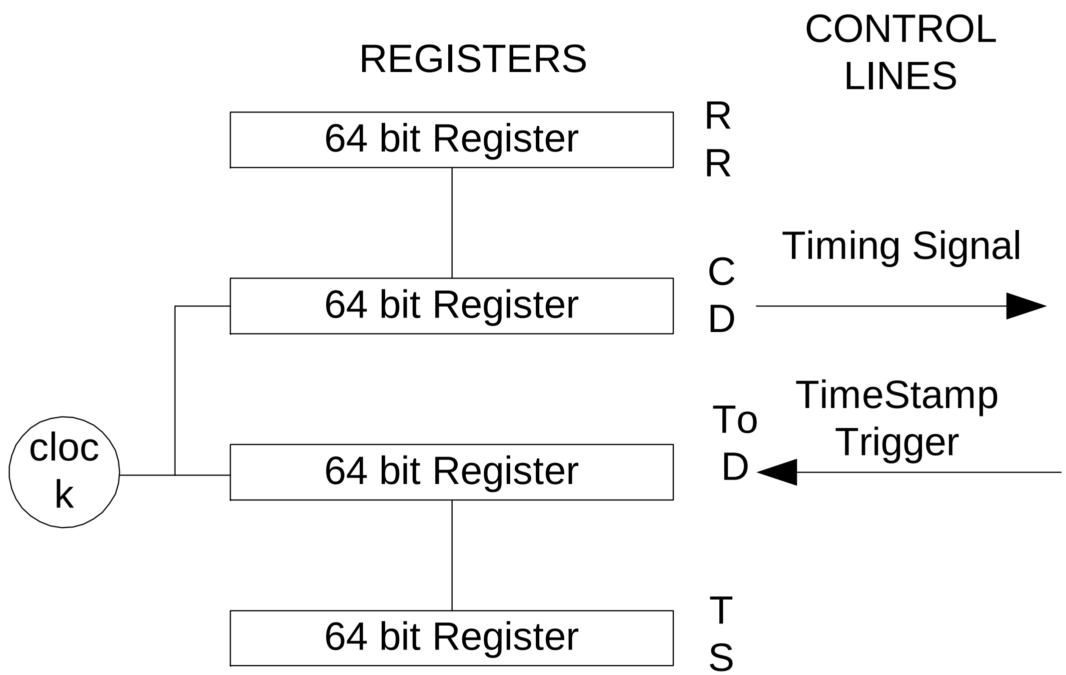

Communications Module

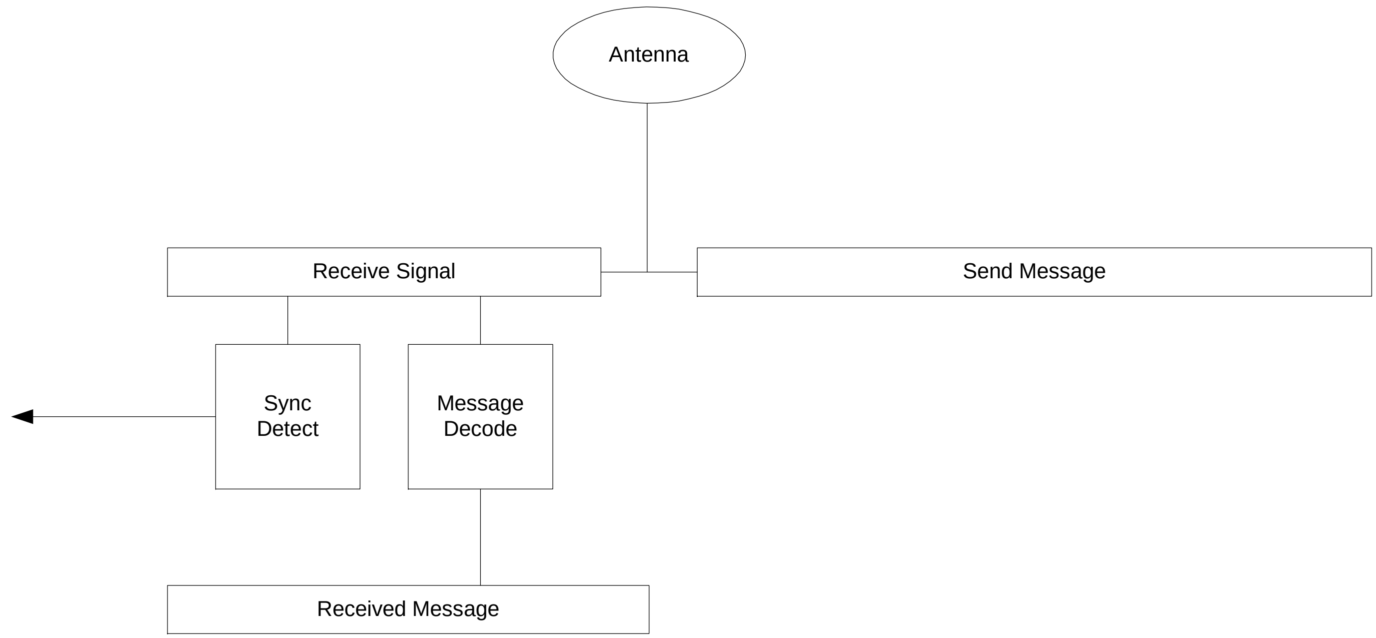

Approach

- Three Phases
  - Measurement Phase
  - Information Exchange Phase
  - Computation Phase

Measurement Phase

- Each node sends (a, t1) where
  - a is its 10 bit ID, and
  - t1 is 64 bit time stamp of when it started sending this message
- All nodes listen to all the messages and keep them after adding a time stamp according to their clocks for the receive time for the first bit.
- After some time a second round of the transmission is started
- The measurement phase ends when all nodes have sent the (a, t) message twice
  - Note that (a,t) message is 74 bits long

Information Exchange phase

- In this phase nodes take turn in broadcasting their receive time stamps for all the messages they have received.

{ (a, ta),(b, tb1,tb2),(c,tc1,tc2)…}

- In this message all receive timestamps, tb1,tb2,tc1, etc. are offset from ta which is 64 bit long while all others are 32 bit long.

Computation Phase

- Each node has a set of nodes {na} in its receive zone
- In this phase using the information it has which includes,
  - send times and receive times for its messages as well as messages among the nodes in {na}.
- A node calculates
  - Distance to all nodes in {na}
  - Clock characteristics of clocks of all nodes in {na}
  - Location of all nodes in {na} in 3-d space

Clock model

- The calculations are based on the clock which is assume to remain stable for short periods of time in that the clock time τ is related to the real time t as follows:

$$\tau_a(t) = \beta_a(\alpha_a + t)$$

- We assume that
  - α and β remain constant for the measurement phase.
  - β , the drift rate of the clock is no worse than 100 parts per million
  - τ is measured with a nanosecond resolution

The Maryland Information and Network Dynamics Lab

Time at Two Node

- At time t the clock reading at node a and node b are:

$$\tau_a(t) = \beta_a(\alpha_a + t)$$
$$\tau_b(t) = \beta_b(\alpha_b + t)$$

- Each node has its own offset and drift rate

Measurement Cycle

In the first measurement cycle, node *A* broadcasts, at global time t_1, a message (A, τ_{a1})

$$\tau_a(t_1) = \beta_a(\alpha_a + t_1)$$

Node *B* receives it at global time t_1+d and records the receive timestamp as equaling $\tau_b(t_1+d) = \beta_b(\alpha_b+t_1+d)$.

University of Maryland 1856

The MIND Lab
The Maryland Information and Network Dynamics Lab

Measurement Cycle

This is repeated in the second measurement cycle

$$\tau_{a3} = \beta_a \left(\alpha_a + t_3 \right) \qquad \tau_{b3} = \beta_b \left(\alpha_b + t_3 + d \right)$$

$$\tau_{a4} = \beta_a \left(\alpha_a + t_4 + d \right) \qquad \tau_{b4} = \beta_b \left(\alpha_b + t_4 \right)$$

Measurement Equations

$$\tau_{a1} = \beta_a (\alpha_a + t_1)$$

$$\tau_{b1} = \beta_b (\alpha_b + t_1 + d)$$

$$\tau_{a2} = \beta_a (\alpha_a + t_2 + d)$$

$$\tau_{b2} = \beta_b (\alpha_b + t_2)$$

$$\tau_{a3} = \beta_a (\alpha_a + t_3)$$

$$\tau_{b3} = \beta_b (\alpha_b + t_3 + d)$$

$$\tau_{a4} = \beta_a (\alpha_a + t_4 + d)$$

$$\tau_{b4} = \beta_b (\alpha_b + t_4)$$

Drift Ratio

$$\frac{\tau_{a3}-\tau_{a1}}{\tau_{b3}-\tau_{b1}}=\frac{\beta_a\left(\alpha_a+t_3\right)-\beta_a\left(\alpha_a+t_1\right)}{\beta_b\left(\alpha_b+t_3+d\right)-\beta_b\left(\alpha_b+t_1+d\right)}=\frac{\beta_a}{\beta_b}$$

Propagation Delay

$$\beta_b d = \frac{(\tau_{b1} - \tau_{a1}) + (\tau_{a2} - \tau_{b2})}{2} + \frac{1}{2}\left(\frac{\beta_a}{\beta_b} - 1\right)(\tau_{a2} - \tau_{a1})$$

Remote Clock Reading

$$\tau_b(t) = \tau_{b1} - \beta_b d - \frac{\beta_b}{\beta_a}\tau_{a1} + \frac{\beta_a}{\beta_b}\tau_a(t)$$

$$t = \frac{\tau_a(t)}{\beta_a} - \alpha_a$$

The Maryland Information and Network Dynamics Lab

Point Set Determination

- Each node can determine the distance to all other nodes within its listening range
- Based on this information each node can determine the relative location of all these nodes

Point Set Determination

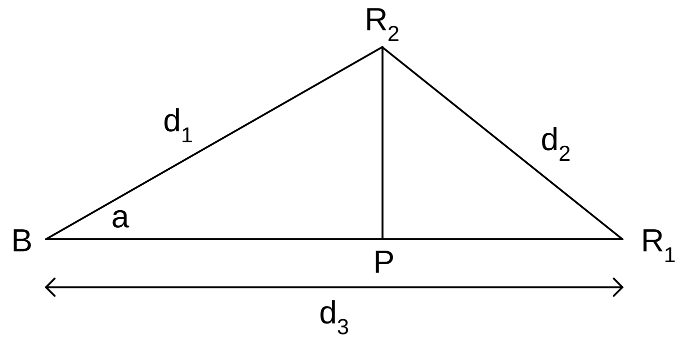

$$\cos(a) = \frac{d_1^{\,2} + d_3^{\,2} - d_2^{\,2}}{2d_1d_3}$$

- Can determine BP and R_2P

Combining Point Sets

- Each node may have different set of nodes in its listening range.
- All calculations are based on common information
- Sets can thus be combined to create a common picture of the whole space

Error Analysis

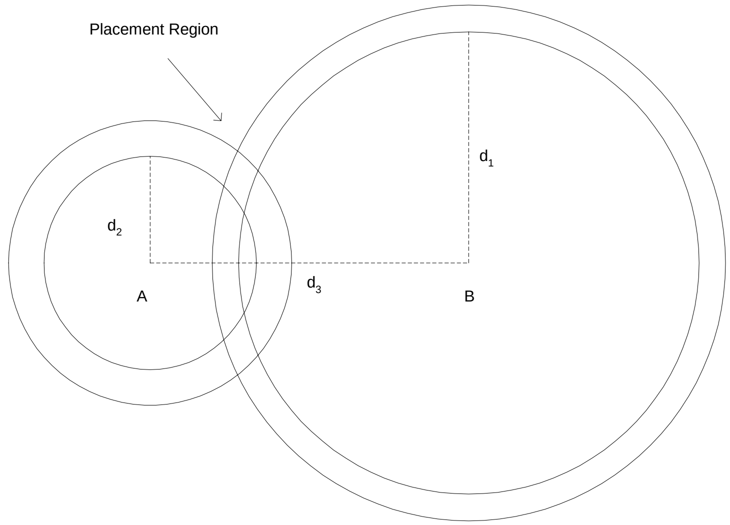

- First order error analysis is based on this geometry

Error Analysis

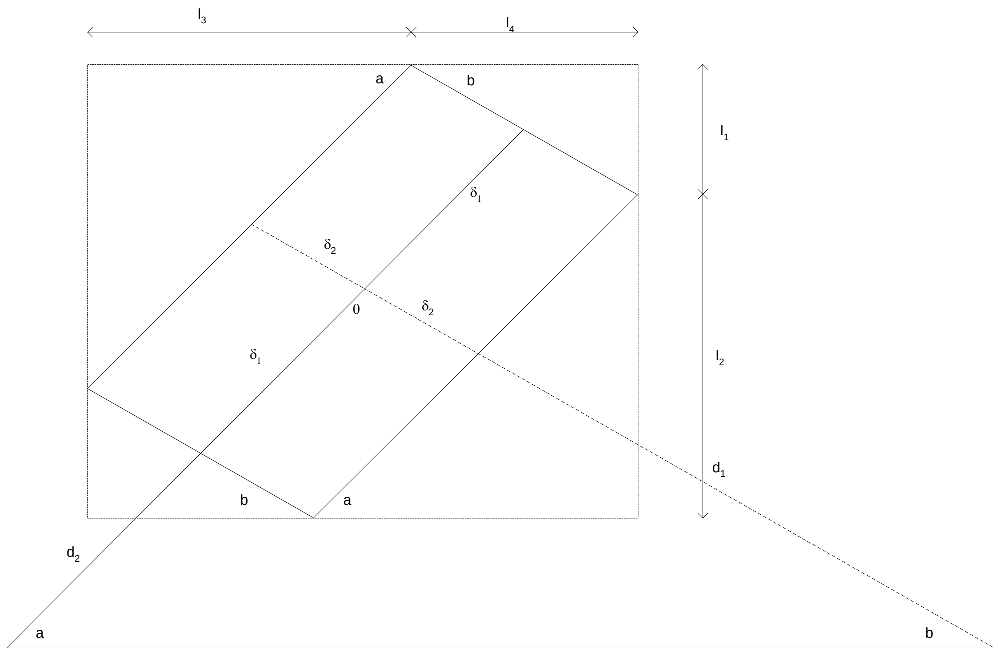

- Can write expressions for the errors
- X variation is given as

$$2\delta_2\sqrt{1-\left(\frac{d_1}{d_3}\sin\theta\right)^2}+2\delta_1\sqrt{1-\left(\frac{d_2}{d_3}\sin\theta\right)^2}$$

Operations

- Timing Diagram

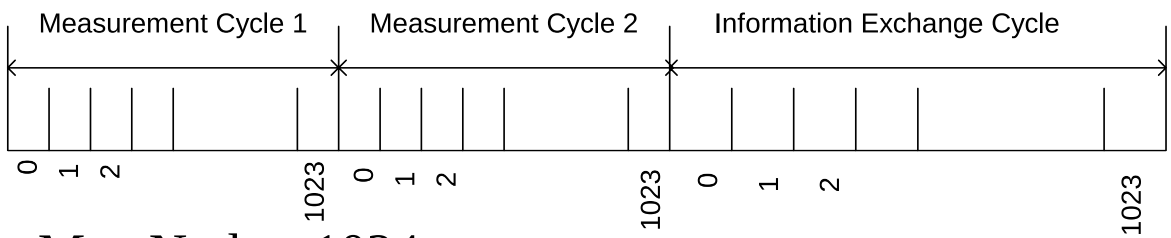

- Max Nodes: 1024
- Mslot : 10 μs
- Islot :   10 ms

Open Issues

- Hardware Implementation
  - Can we have hardware that can give timestamps with the required accuracy?
  - Can that hardware be reduced to a chip?
  - Can that chip be integrated with other systems, e.g. 802.11b
- Accuracy analysis and Improvements
- Algorithmic improvements
- Point Set Integration
  - Multi hop environment
- Operation with a few fixed locations, e.g. Access Points
- …

Outline

- Localization – Active Techniques
- RSSI Based
  - Characteristics of 802.11b signals
  - Horus
- Transit Time Based/ Synchronization
  - PinPoint
  - System Synchronization
- Localization – Passive Techniques
  - Nuzzer
- Concluding Remarks

Passive Localization

- Exploit the variability in the signal seen due to the presence of people
- Can we determine the location of a person or persons?
- Nuzzer Technology
  - Work in Progress – Leila Shahamatdar, Moustafa Youssef

Nuzzer Technology

- Measure RSSI at fixed locations

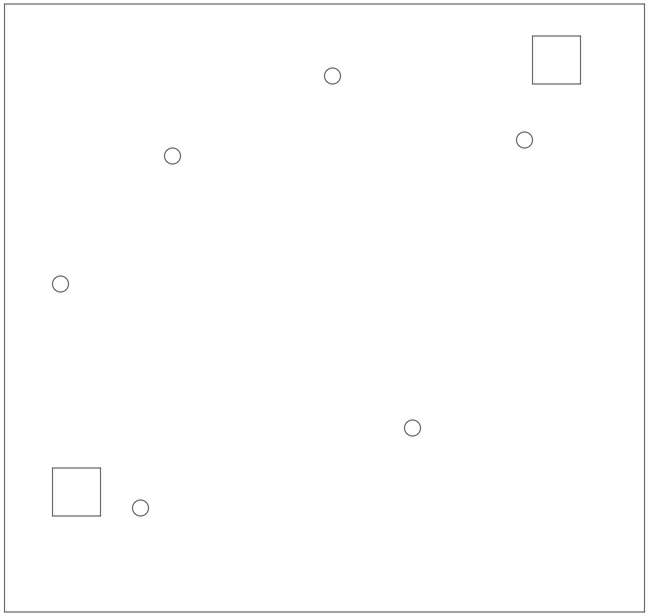

$$\begin{bmatrix} r_1 \\ r_2 \\ \dots \\ r_n \end{bmatrix} (x, y, z)$$

Nuzzer System

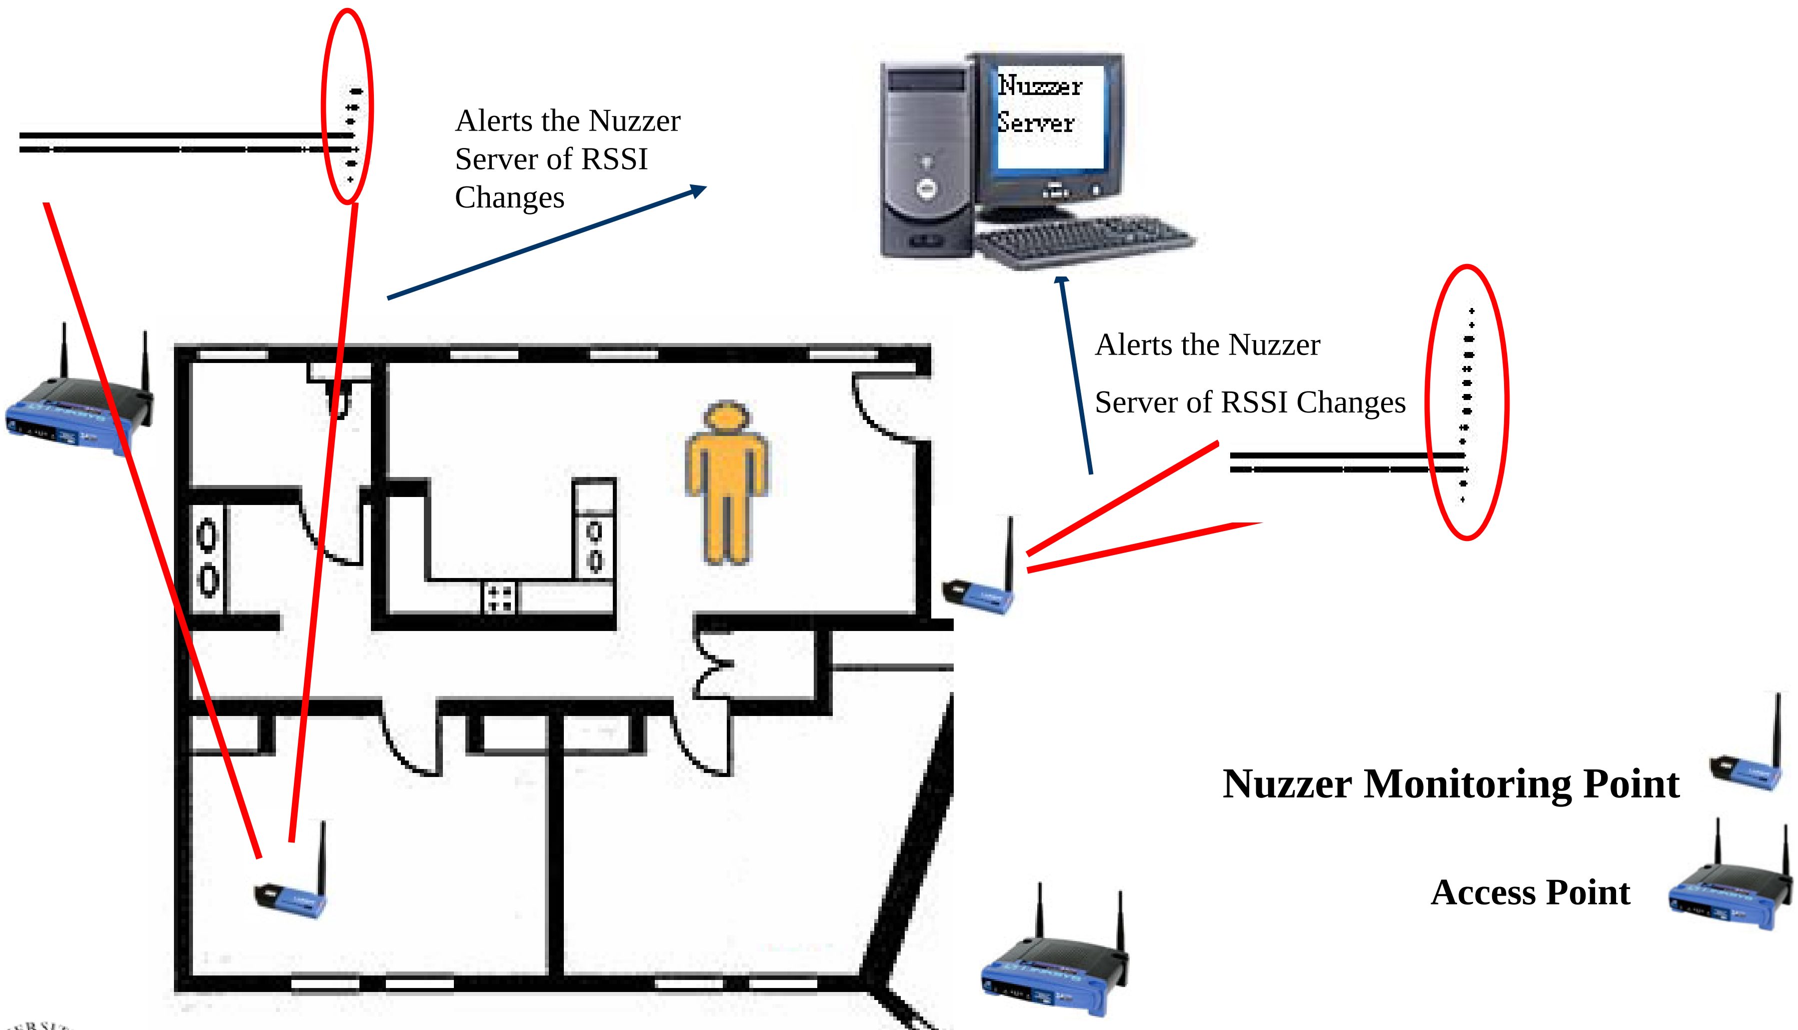

The Maryland Information and Network Dynamics Lab

Nuzzer Steps

- Presence/Absence of a person
- Location of a person
- Location and tracking of multiple people
- …

Experimental Evidence

- RSSI varies as people move around.

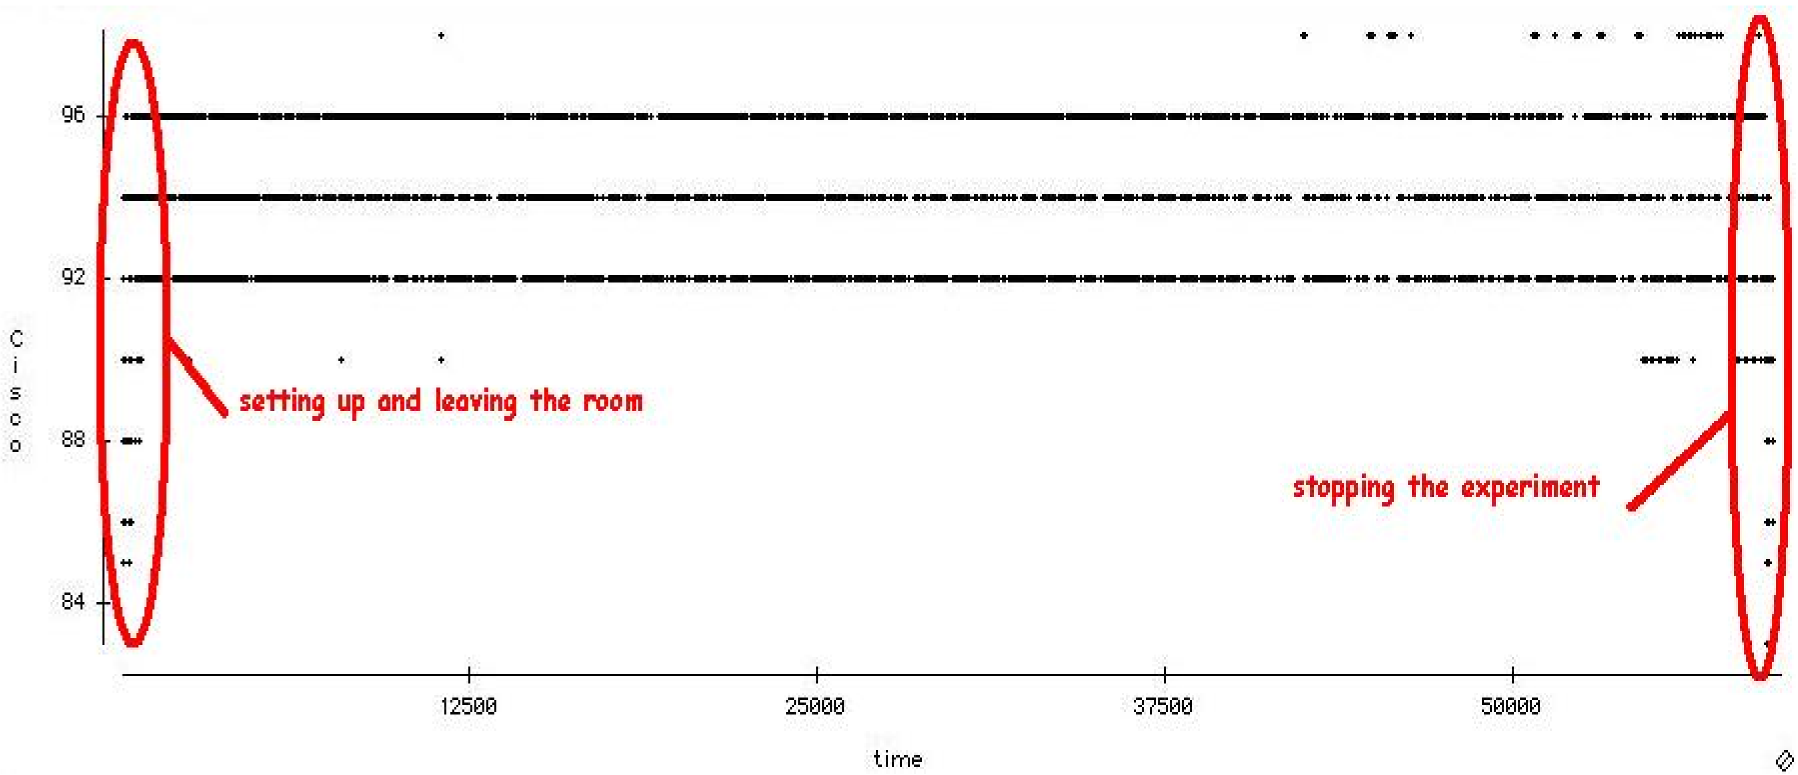

The Maryland Information and Network Dynamics Lab

Concluding Remarks

- Can we realize the applications we talked about in the beginning of the discussion today?
- Location and time in distributed systems of tomorrow are going to play a major role.
- Techniques for location
- System Synchronization with stable clocks