

Understanding Wireless Interference

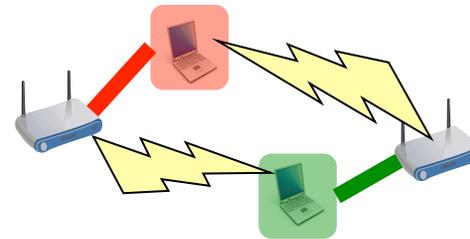
Suman Banerjee

UW Madison

How good is my wireless network?

- Even after 14 years of standardization of IEEE 802.11
 - Not a good understanding of whether my wireless network operates well enough

- WiFi to WiFi interference



- Non-WiFi to WiFi interference

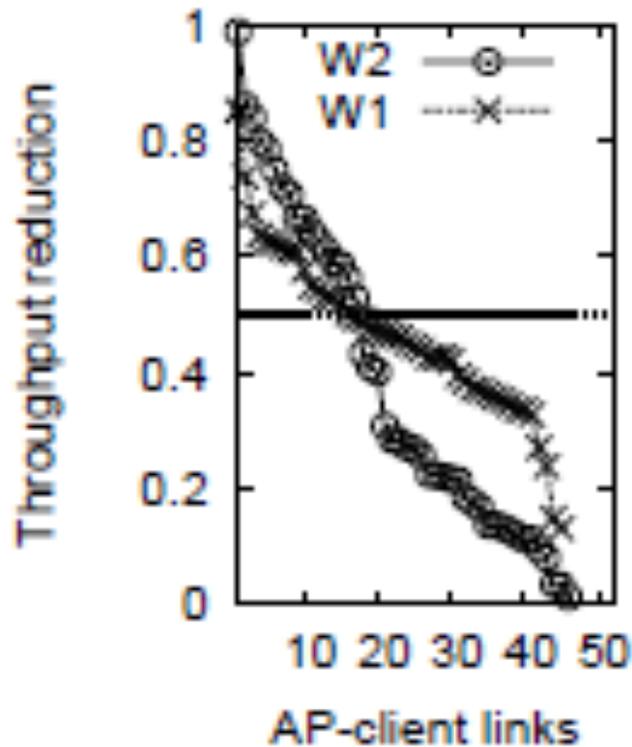


A Snapshot of (WiFi-WiFi) Interference

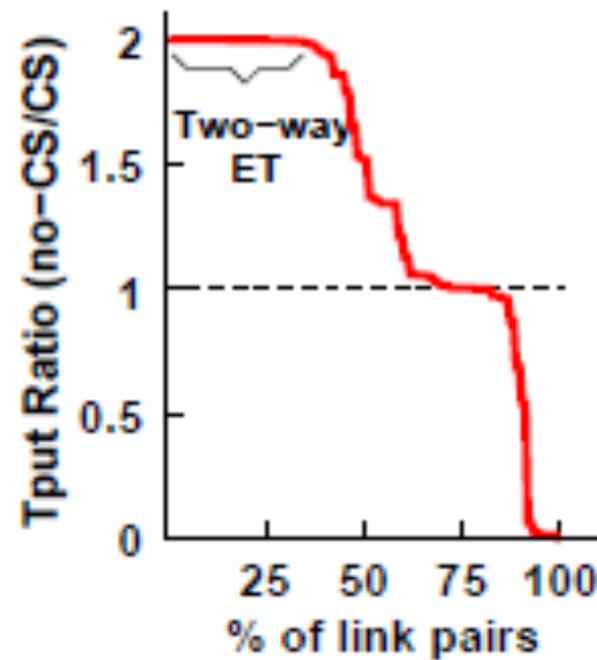
W1: 9 APs, 45 clients

W2: 20 APs, 51 clients

30 node testbed



Hidden terminals



Exposed terminals

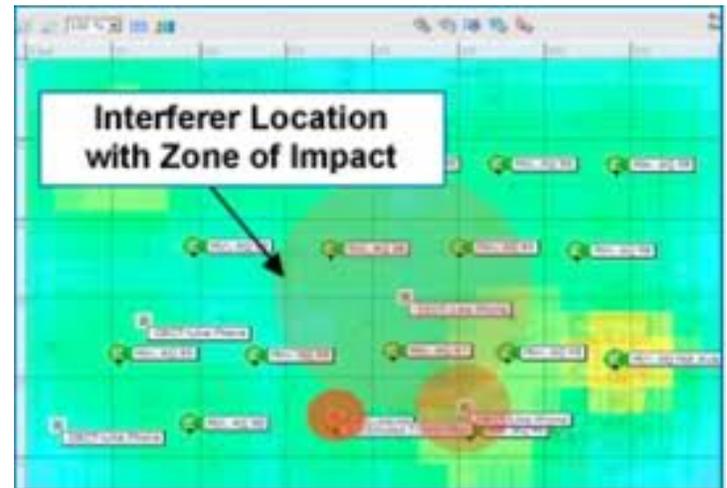
Wireless frustrations

- Users hardly call administrators
- Standard practices
 - Re-boot, Pray, Re-try (Repeat)



Current solutions

- Model based view of interference “potential”
- Does not capture what the likely impact is
- No real-time understanding that evolves with interference



Basic question

- Can I give an overall score to my wireless network performance?
 - In real-time?
 - Without making active measurements?



Real-time Interference Estimators

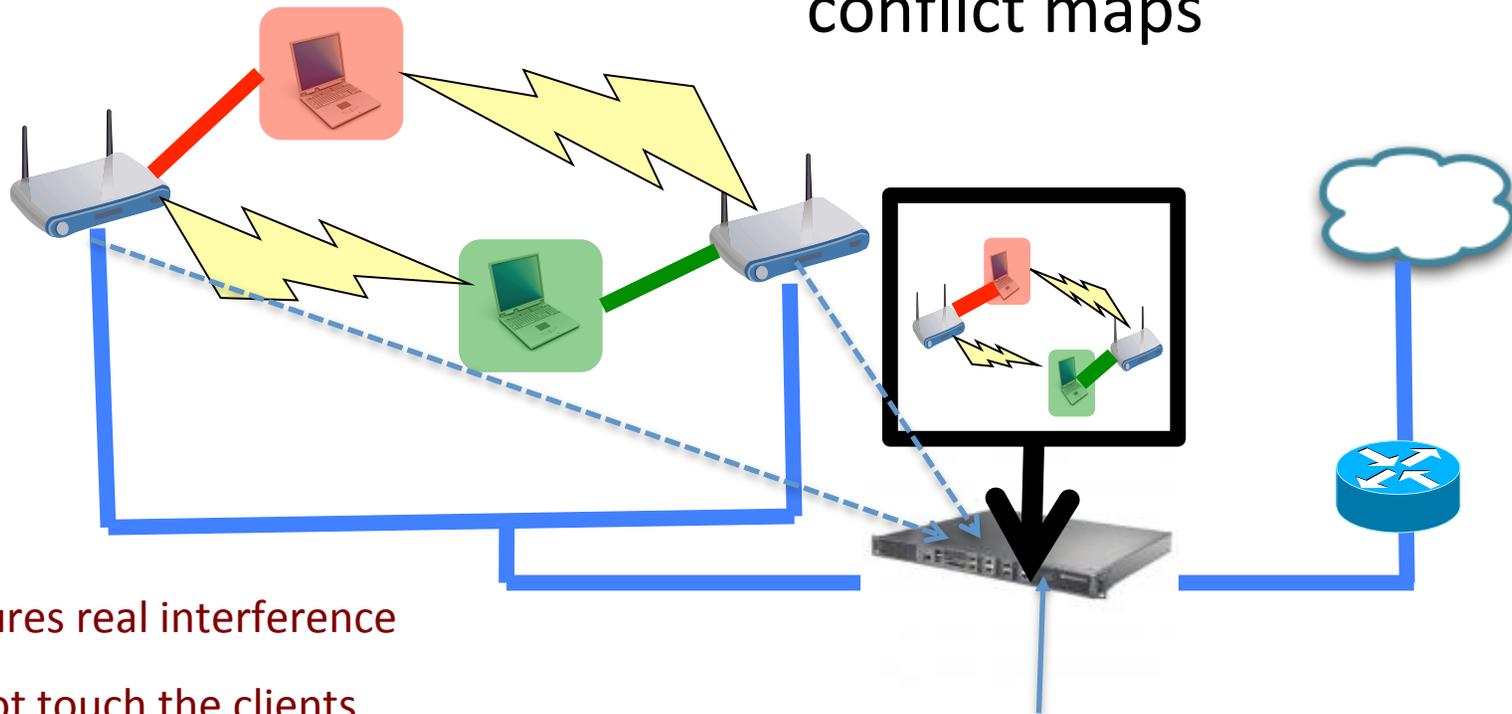
- WiFi to WiFi interference
- Non-WiFi to WiFi interference
 - Special restriction: Use WiFi-only hardware
- In both cases, will try to quantify the impact of interferer

PIE: WiFi-WiFi Interference Estimator

[NSDI 2011]

Interfering links

Controller builds real-time conflict maps



Captures real interference

Do not touch the clients

Passive observations only at controller with explicit interaction with APs (minimal overheads)

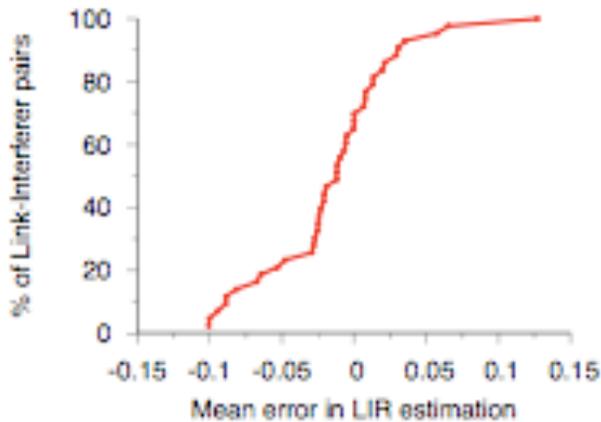
PIE Approach Summary

- Observe events in air at each AP (WiFi level)
 - What packets are transmitted? When?
Success? Failure?
- Send periodic reports to controller
 - < 1 Kbps per AP
- Use statistics to identify different types of interference

Advantages of PIE

- No client support required
- No active wireless measurements involved
- Works in a real-time fashion
- Detects all types of conflicts and their relative severity as and when they occur

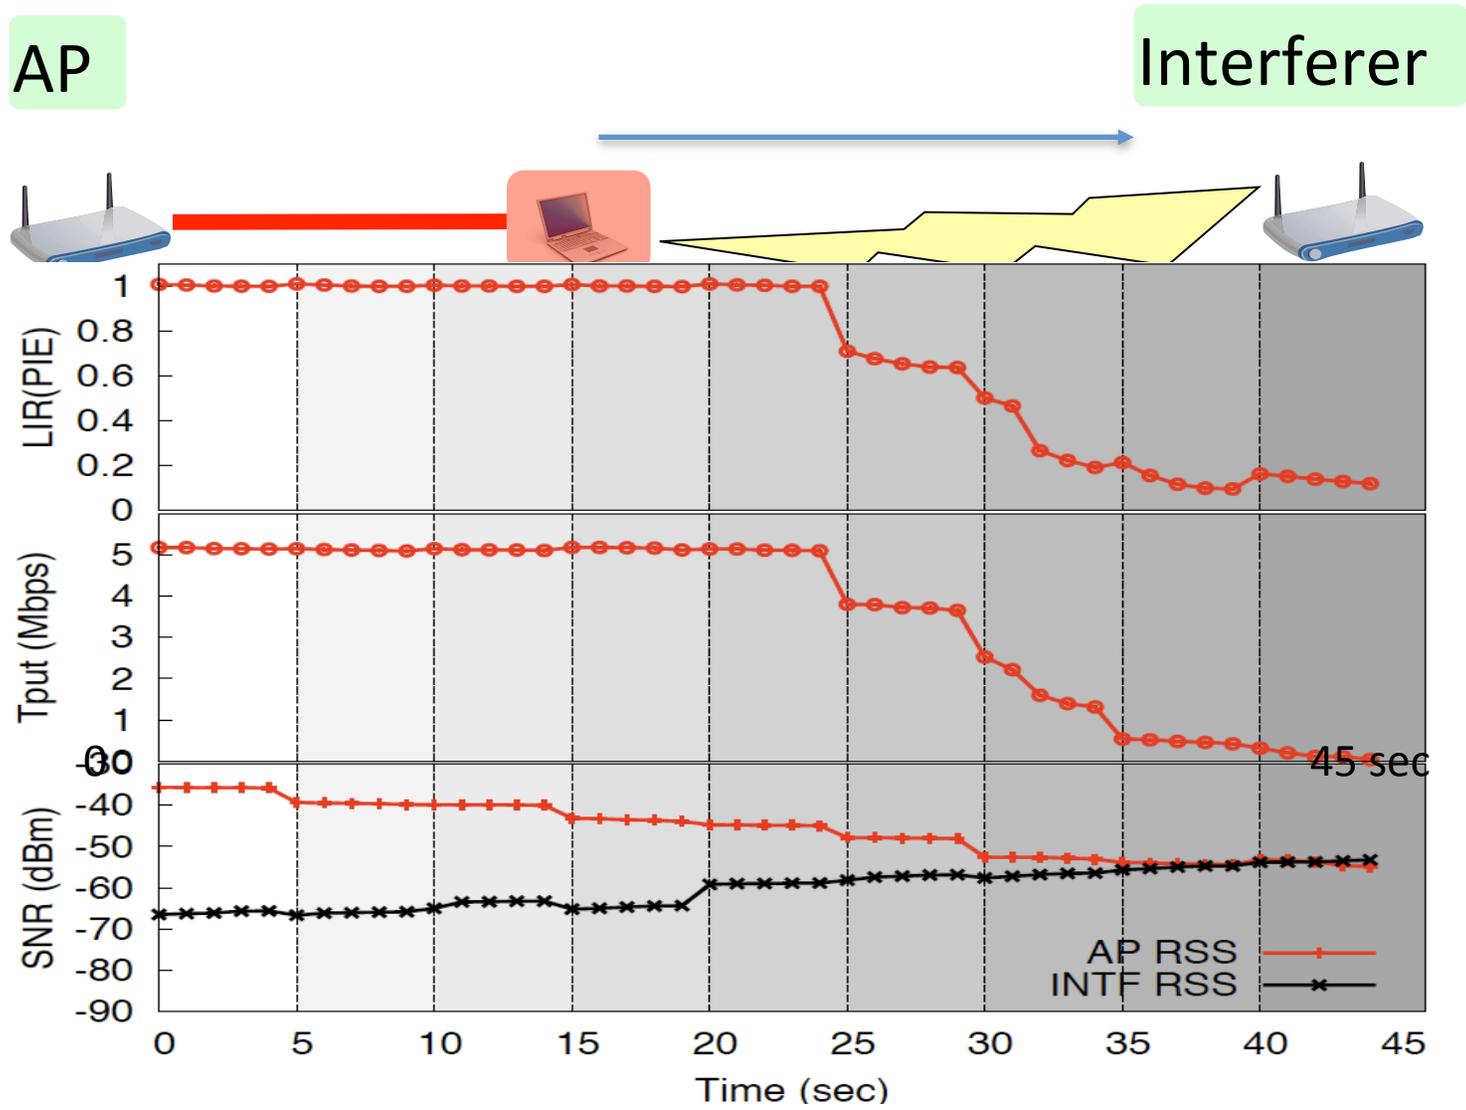
Properties of PIE Performance



Highly accurate

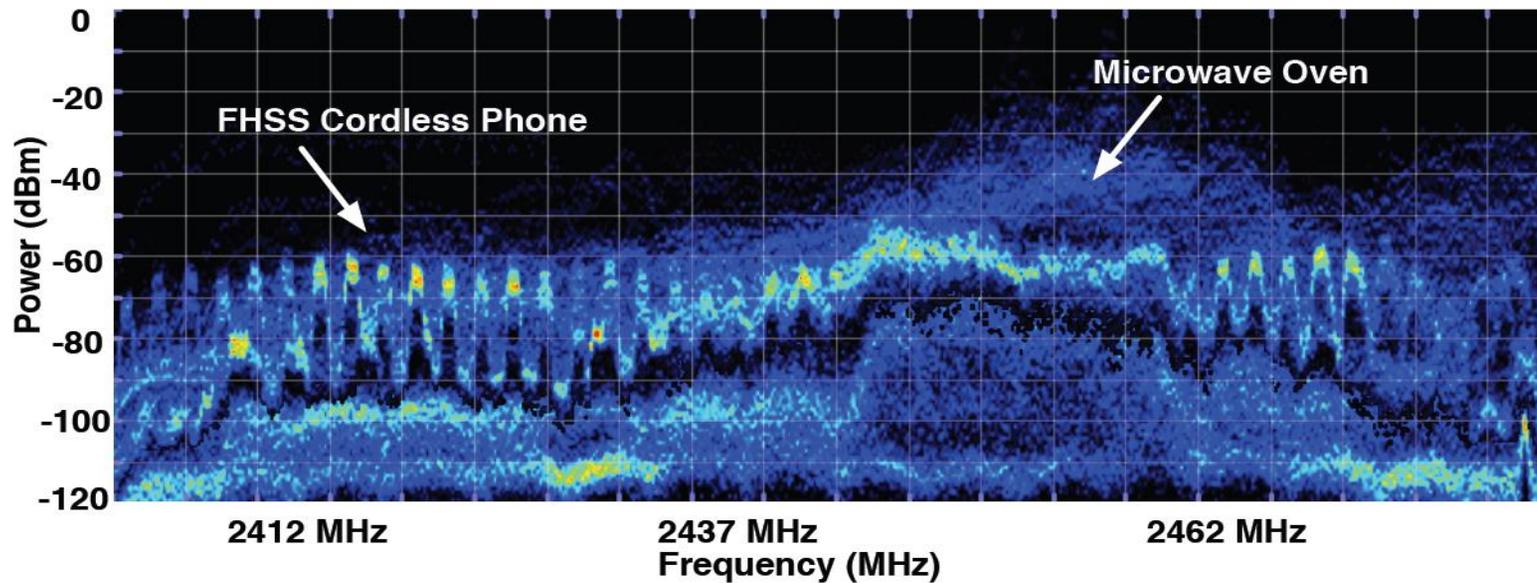
- High accuracy
- Can handle multiple interferers
- Can handle client mobility
- Scales efficiently
- Very agile: ~ few seconds

PIE Agility



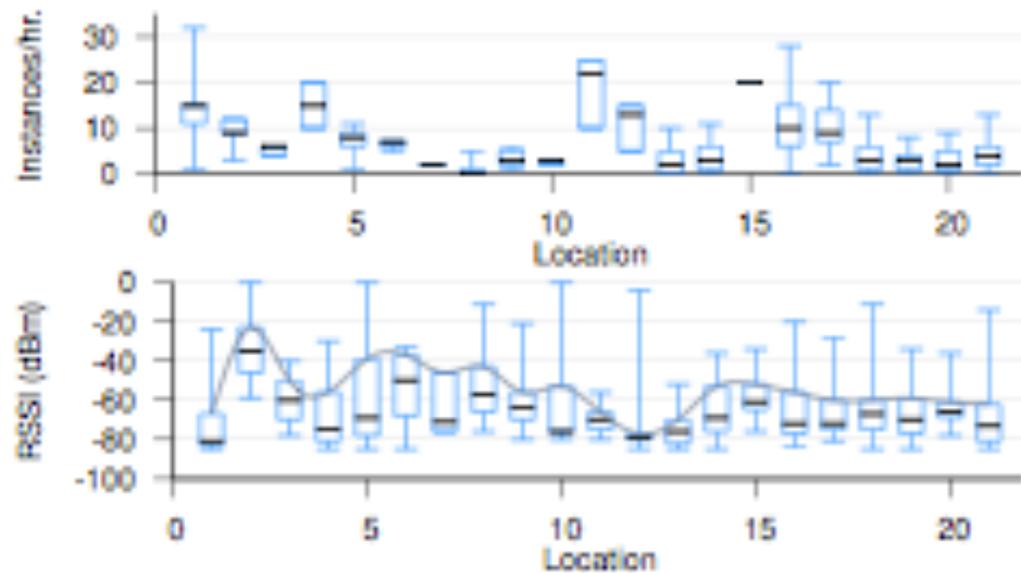
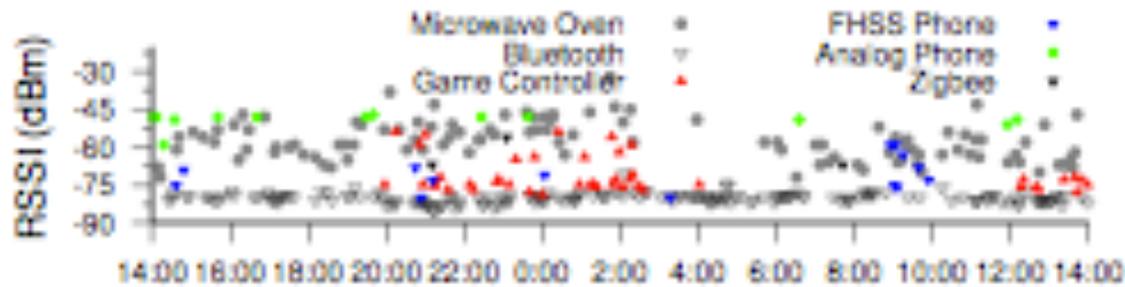
What about non-WiFi interference?

Spectrum at a university cafe



High powered non WiFi devices share the spectrum with WiFi devices

Snapshot of interference



Plethora of devices in the spectrum ...



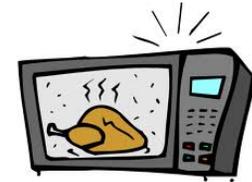
Analog Cordless Phone



Video Camera



ZigBee Device



Microwave Ovens

Narrowband/High-duty devices

Broadband devices



FHSS Cordless Phone



Bluetooth ACL/SCO



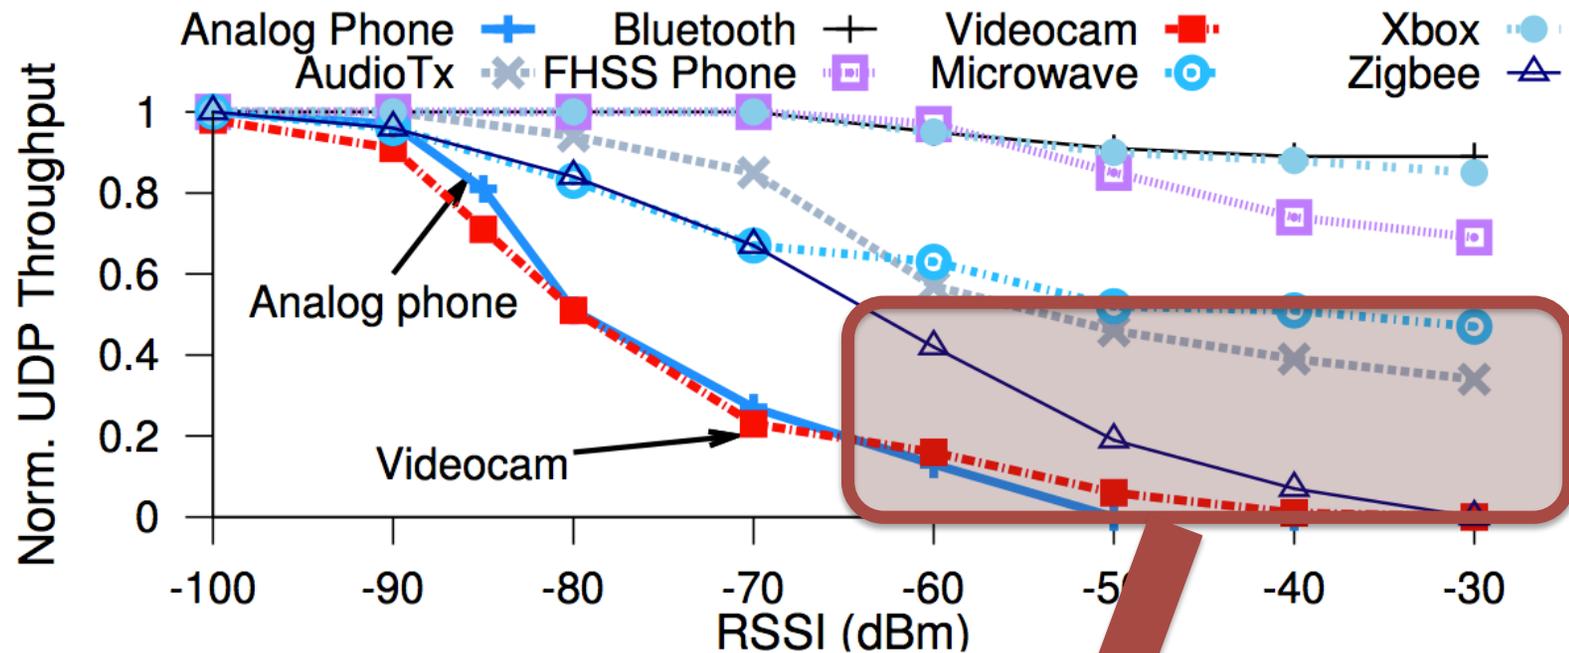
Game controllers
(Wii, PS3)



Wireless headphones

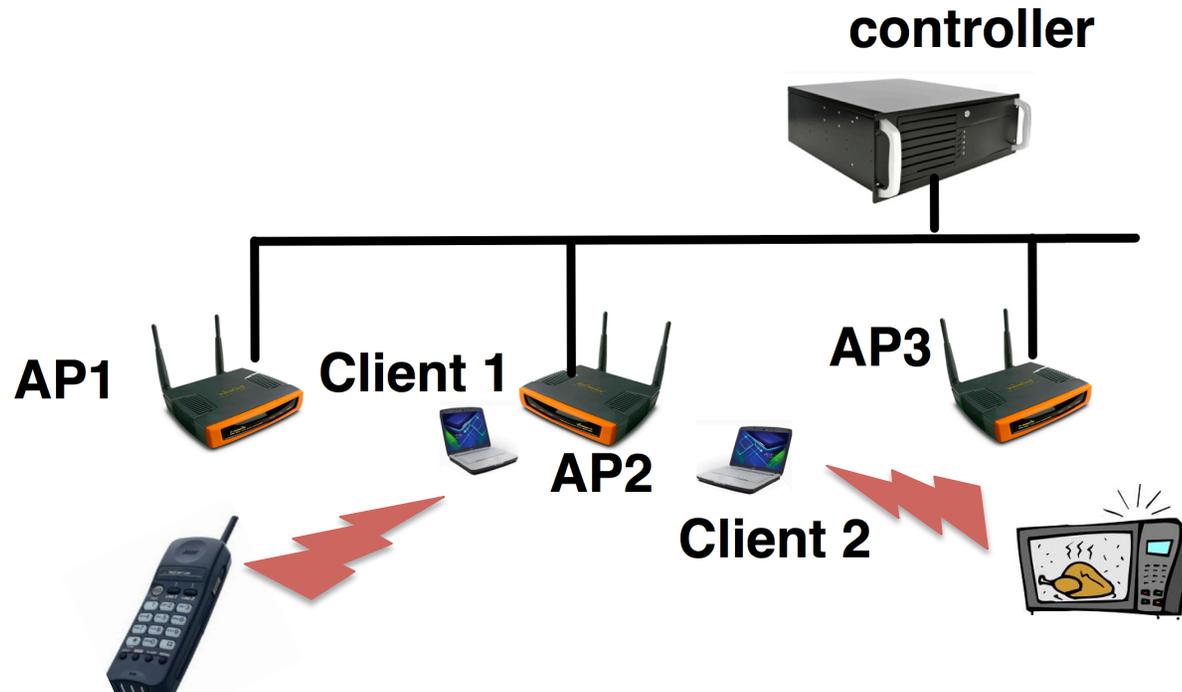
Frequency-hopping devices

Impact on WiFi links



More than 50% throughput reduction, and in some cases, throughput drops to zero!

Enterprise WLAN scenario



QUESTIONS

- Are there any non-WiFi devices in the medium?
- Which non-WiFi devices caused interference?
- How do we locate these non-WiFi devices?

Typical solutions

- A heavy-weight spectrum sensor
- Can we do this with a WiFi card alone?
 - Advantages: integrate into existing WiFi APs and clients, and new class of adaptation becomes possible



Our solution: **WiFiNet**

- Collaborative neighborhood of WiFi-only nodes:
 1. Individually **detect** non-WiFi devices using WiFi hardware (Airshark)
 2. Identify the **interference impact** of each non-WiFi device
 3. Pin point the **physical location** of each non-WiFi device

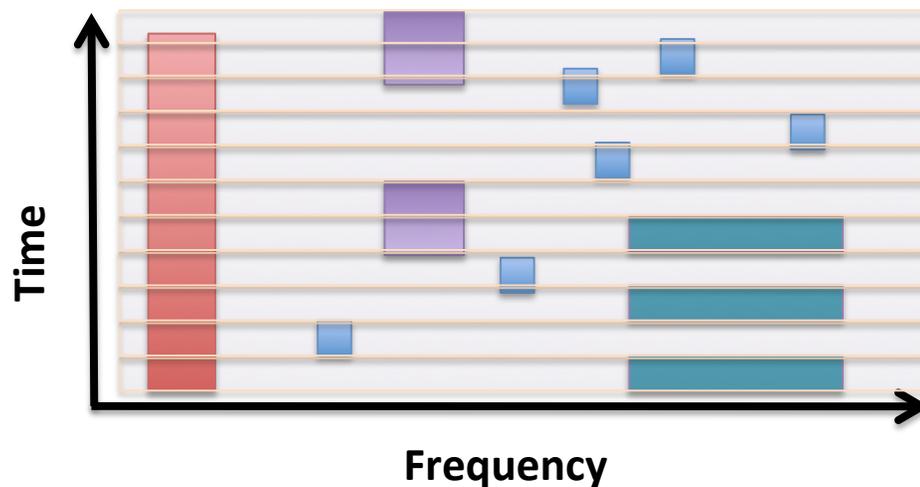
How to detect device types?

- Airshark [IMC 2011]
 - Single WiFi node system
 - Software-only solution
 - Works on top of emerging wireless cards providing fine-grained energy samples (FFTs)
 - Atheros AR9280
 - Benefits
 - Low cost, software solution
 - Does not require sophisticated spectrum analyzers
 - Easily deployable in Access Points and clients

How to detect device types?

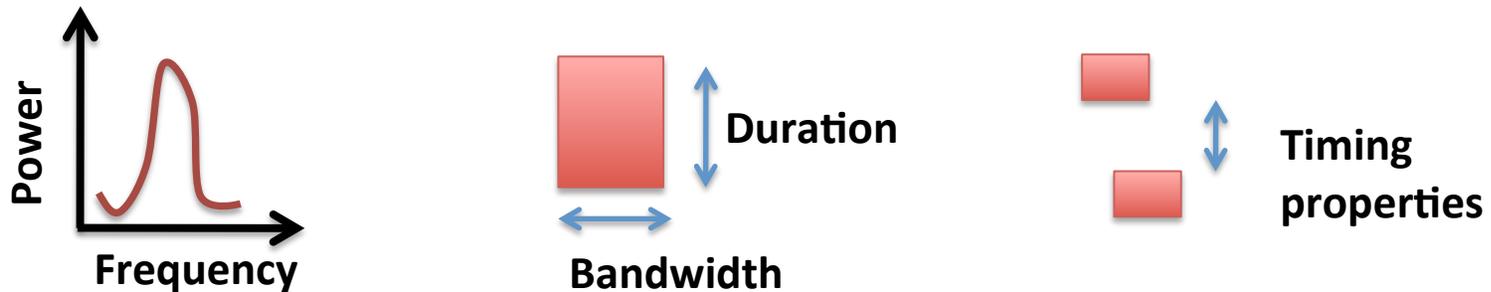
- Capture FFT samples from the WiFi card
- Identify “Pulses”
 - Time-Frequency blocks capturing device transmissions

Analog Phone	
ZigBee device	
Microwave Oven	
Bluetooth	

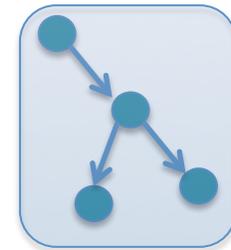


How to detect device types?

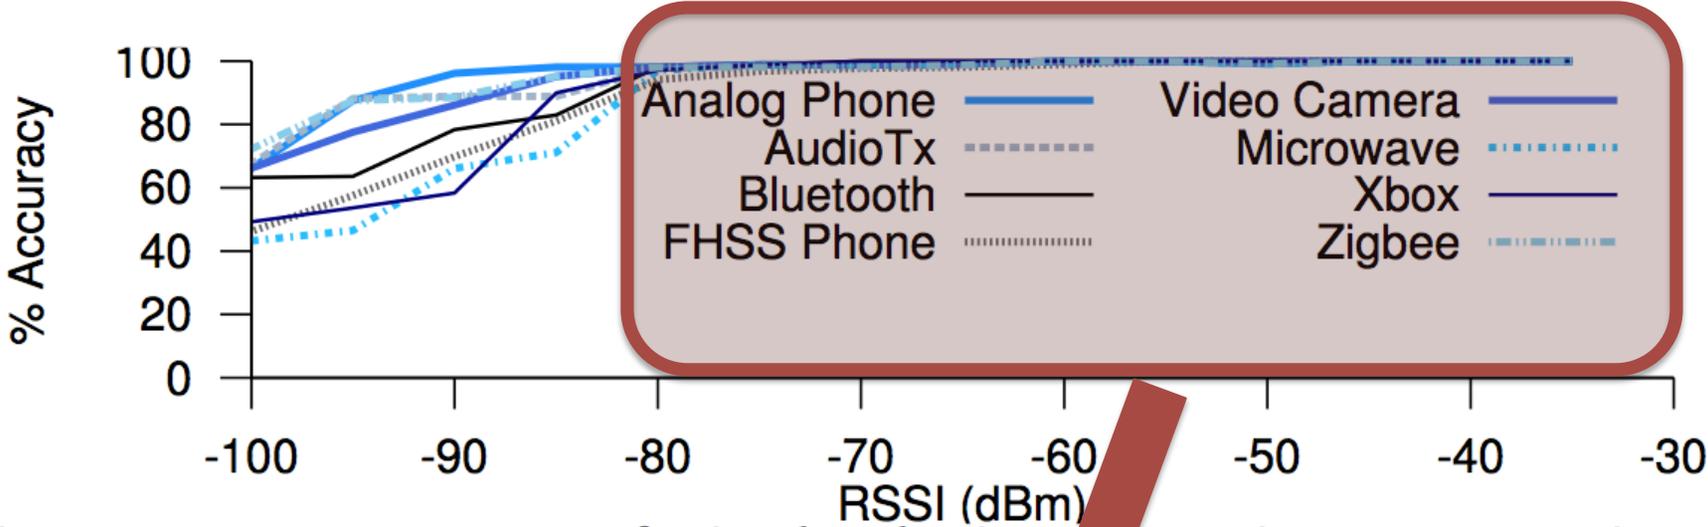
- Extract “Features”
 - Spectral and temporal properties



- Decision Tree based classification
 - Per-device classifier outputs “1” or “0”

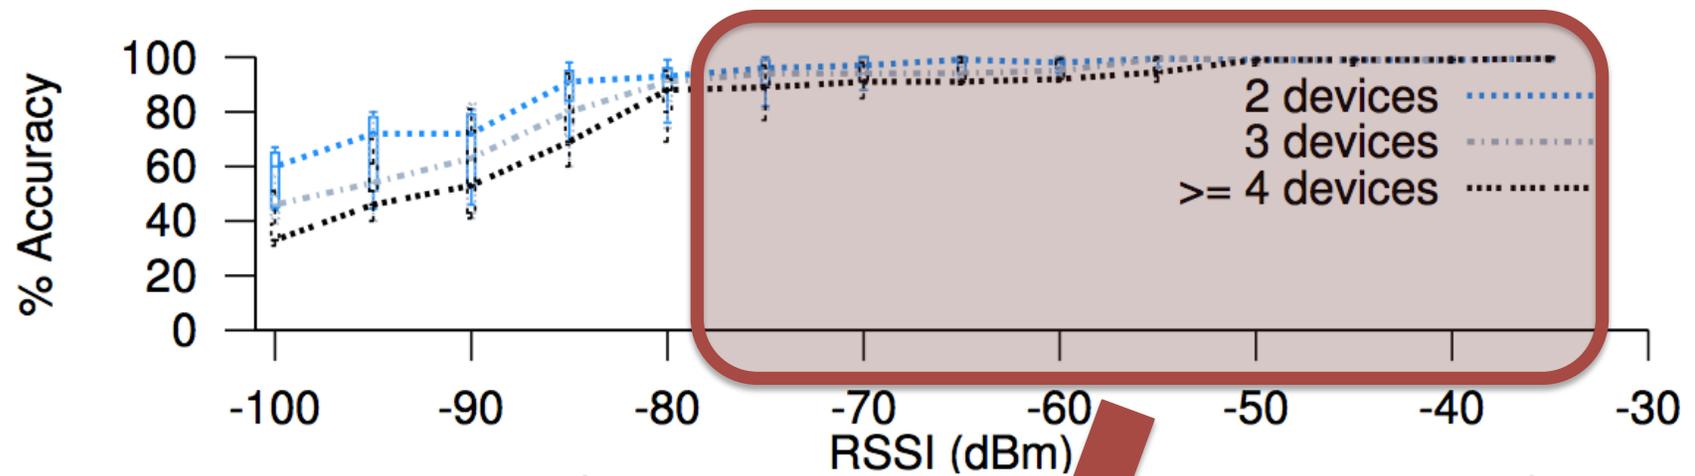


Detection Accuracy: Single device



> 98% accuracy at signal strengths \geq -80 dBm

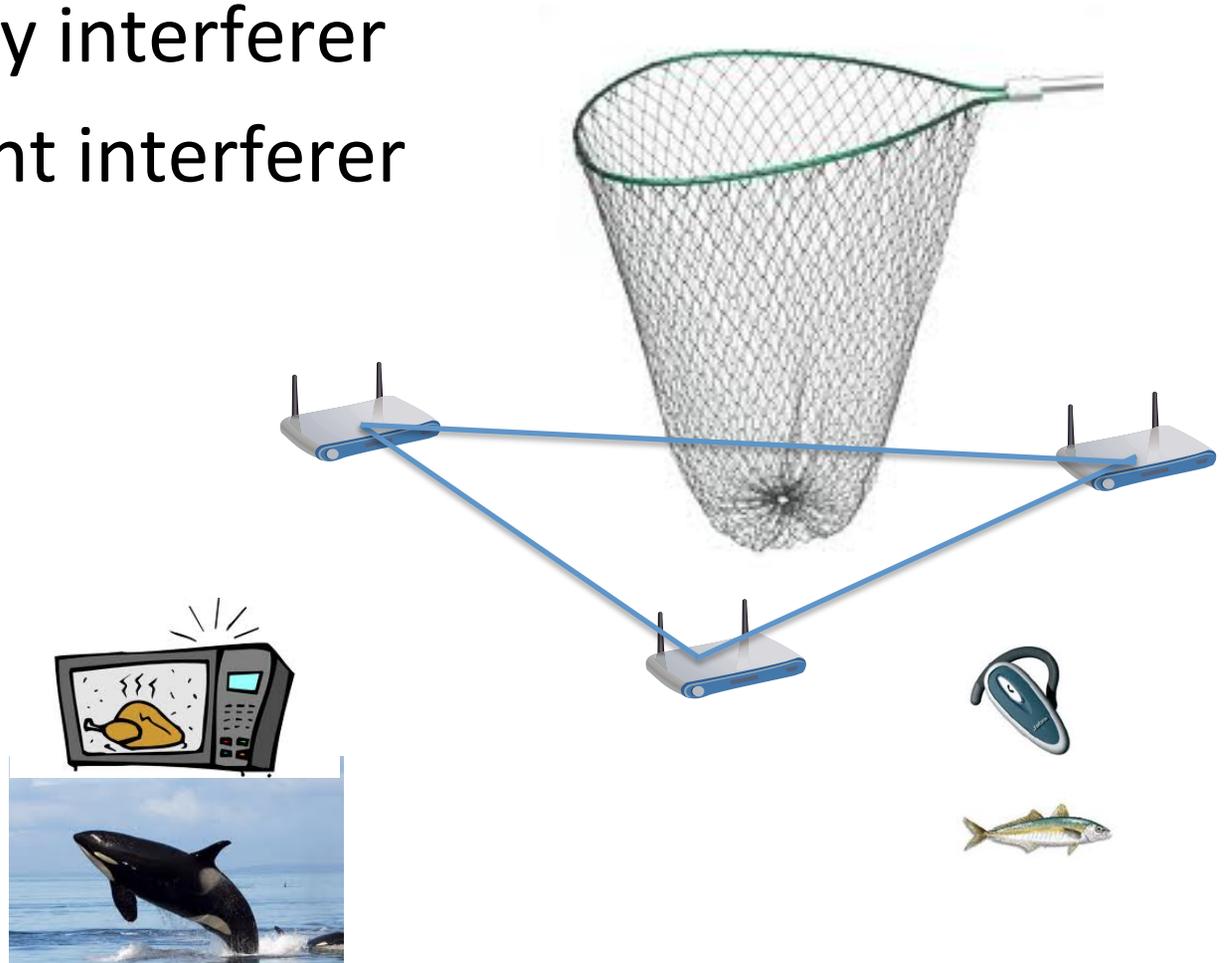
Detection Accuracy: Multiple devices



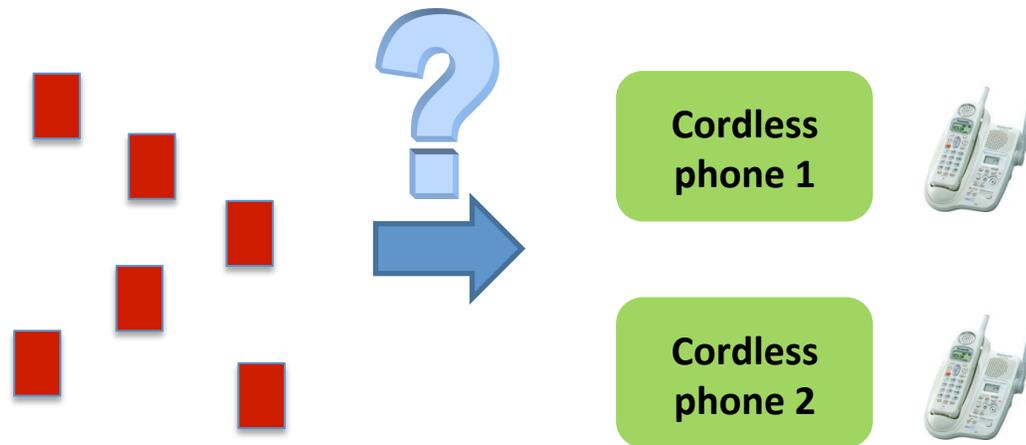
> 91% accuracy at signal strengths ≥ -80 dBm

WiFiNet: Catching whales and minnows

- Whale: Heavy interferer
- Minnow: Light interferer



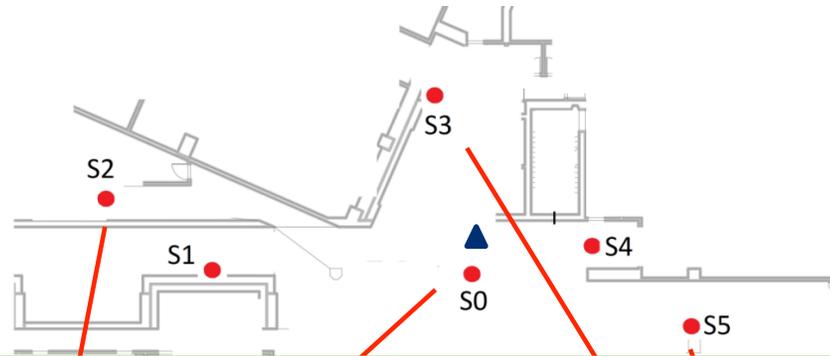
WiFiNet: Key challenges solved



**What about multiple devices of same type?
Which pulses belong to which device?**

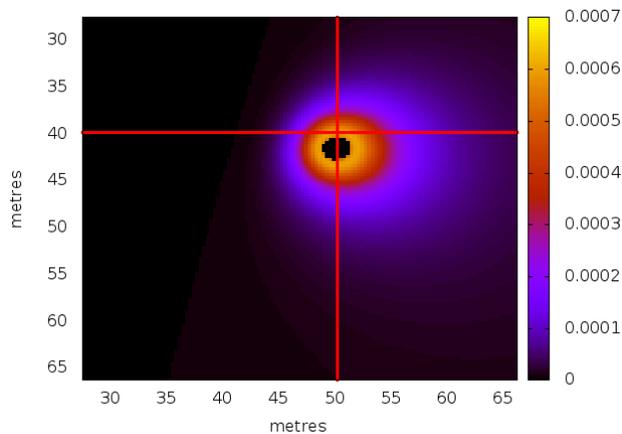
Multiple WiFi observers have to correctly agree to this mapping and then do some triangulation (lots of details in triangulation as well)

Model based localization

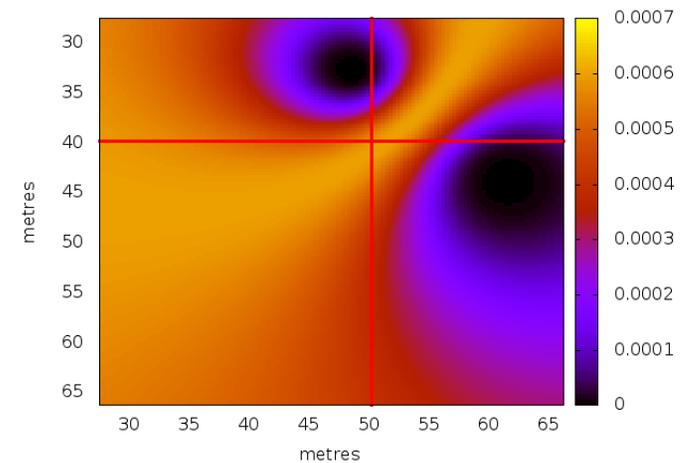


Feedback from multiple pairs of nodes helps narrow down the device location

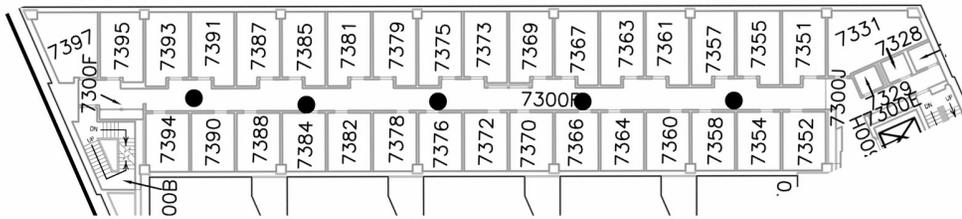
Feedback from: S0,S2



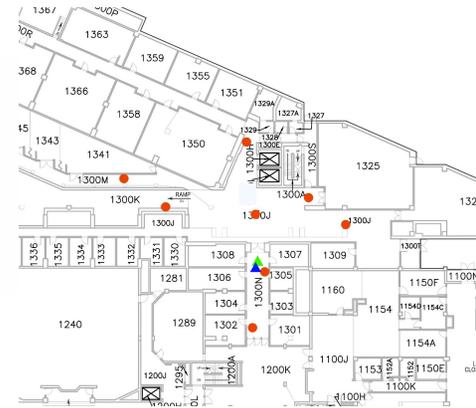
Feedback from: S3,S5



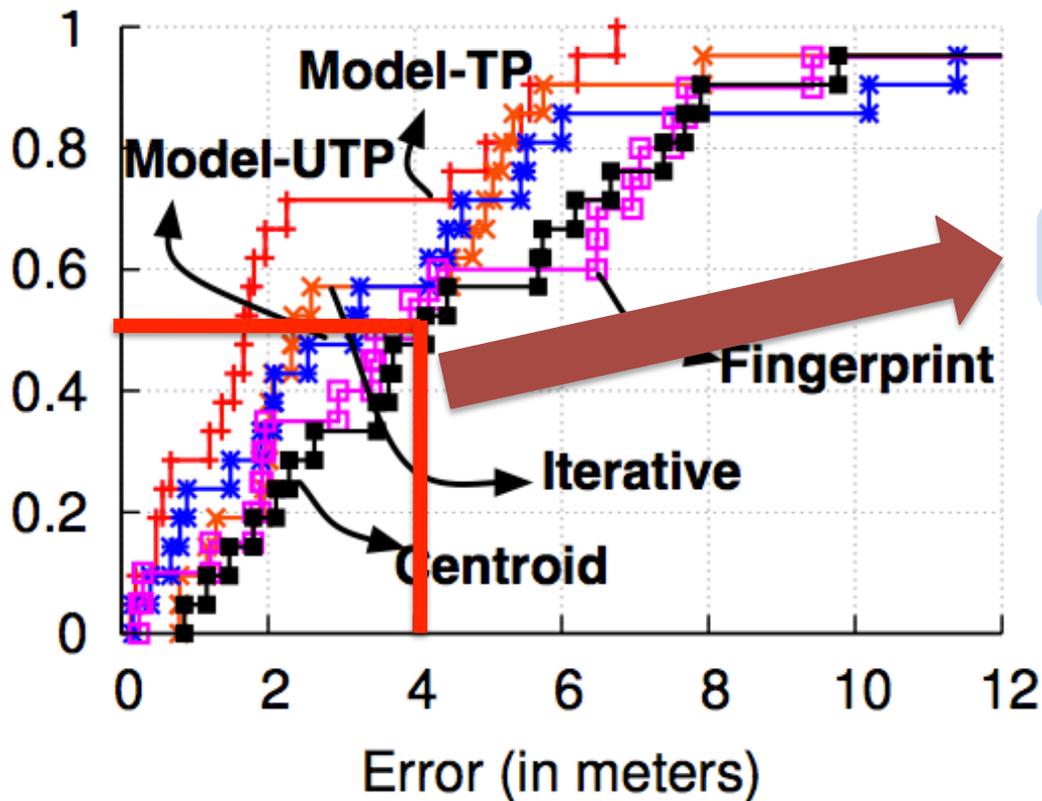
Device Localization Results



Deployment 1

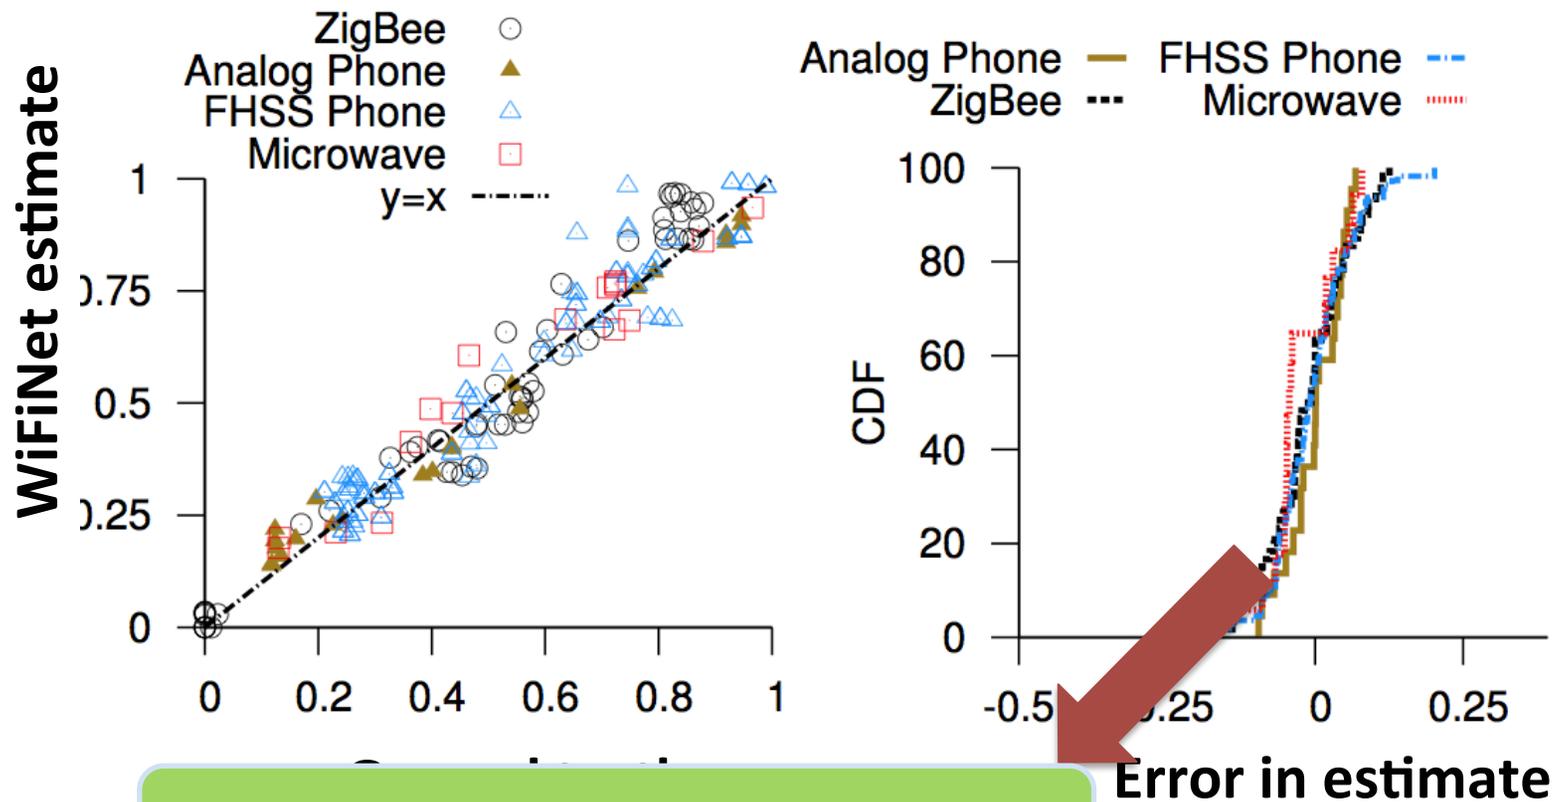


Deployment 2



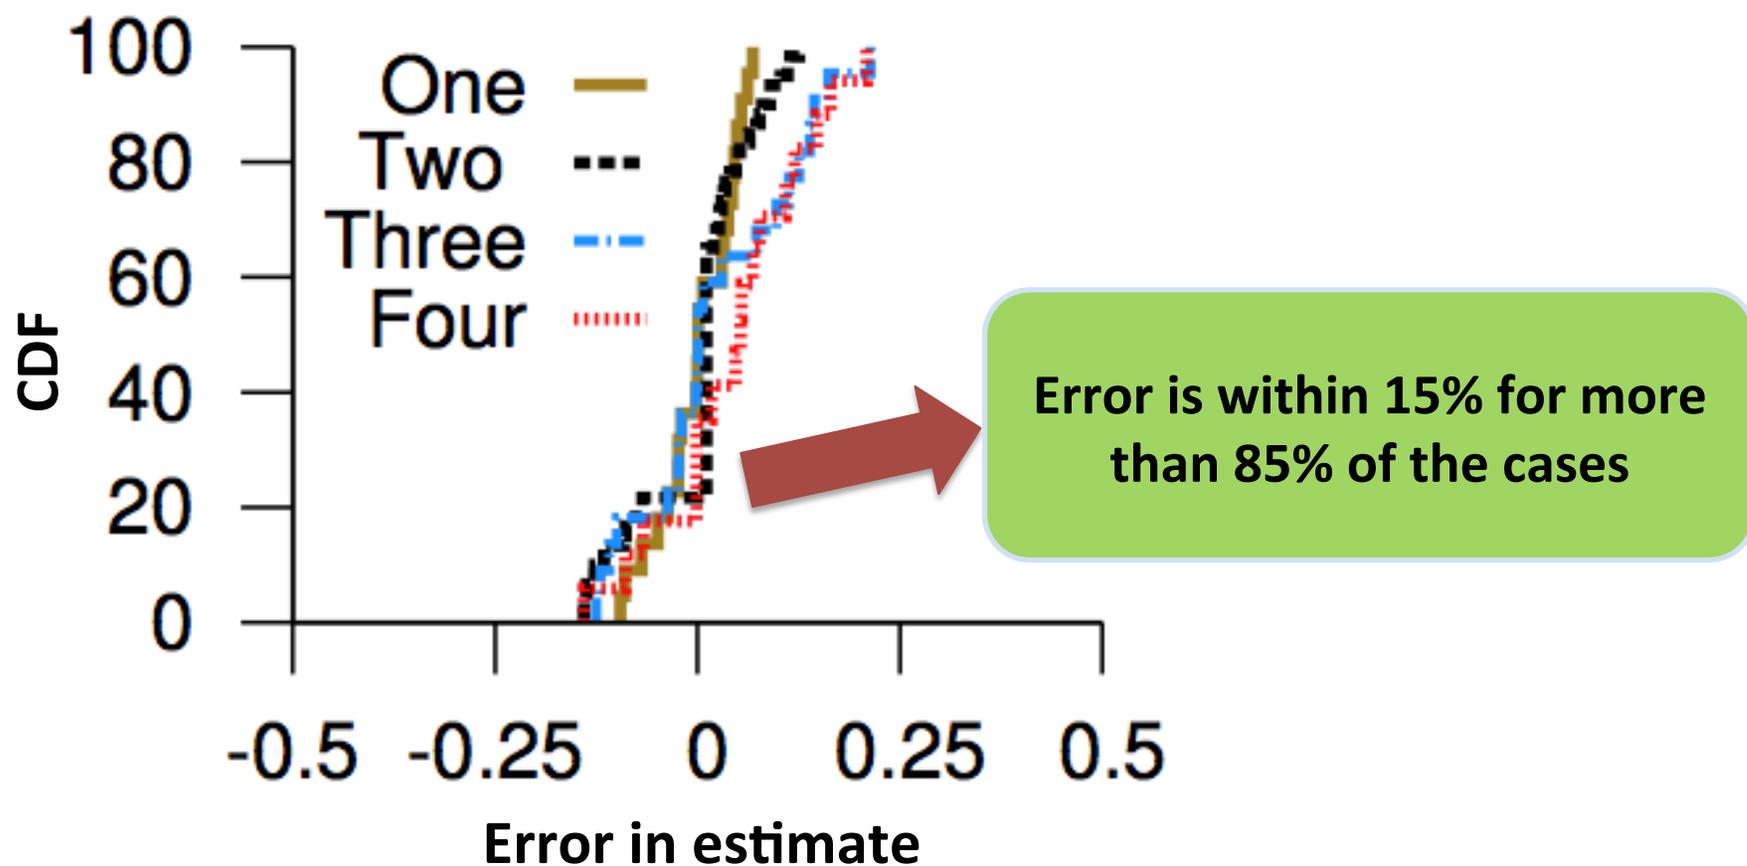
Median error ≤ 4 mts

Interference Estimation Results: single non-WiFi interferer



Error is within 10% for 95% cases

Interference Estimation Results: Multiple non-WiFi interferers



Future challenges

- Scratching the surface of understanding interference
 - Better techniques
 - Better mitigation
 - Better UI for admins and users

