Microsoft Research

Are we there Yet?
Self-Managing Wireless Networks

Victor Bahl
Microsoft Corporation

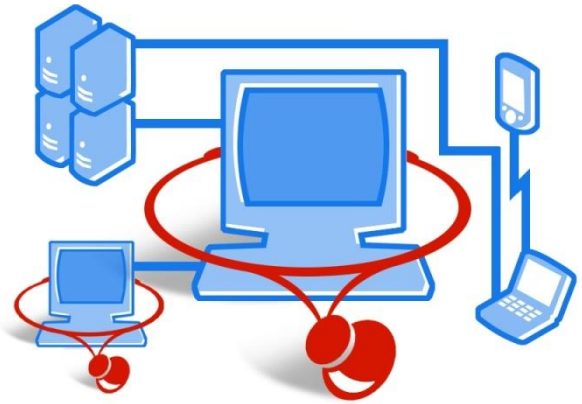

February 2007

Source: Victoria Poncini, MS IT

MS IT Wireless Satisfaction Survey

Wireless networks perceived to be “flaky”, less secure

~7,000 Access Points
~65,000 XP & Vista Clients
~40,000 connections/day
~35,000 handheld devices

December 2006

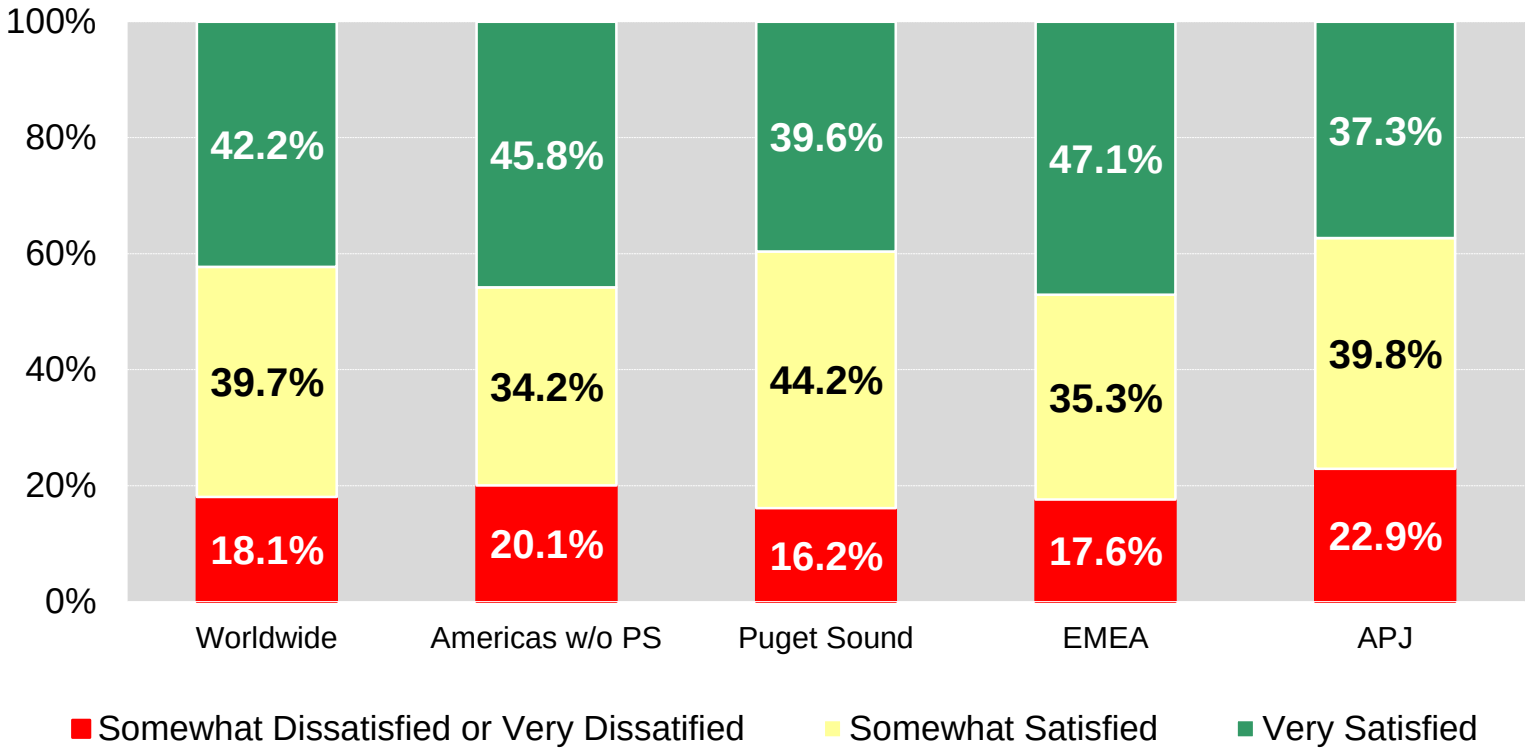

Victor Bahl

Microsoft Research

User Complaints & IT Headaches

Microsoft's IT Dept. logs several hundred complaints / month

- 70% calls are about client connectivity issues (e.g. ping-ponging between APs)
- 30% (and growing) are about performance problems due to interference

End-users complain about

- Lack of RF coverage, performance & reliability
- Connectivity & authentication problems

Network administrators worry about

- Providing adequate coverage, performance
- Security and unauthorized access

Corporations spend lots of $$ on WLAN infrastructure

- WLAN hardware business to reach $2.6 billion in 2007. (Forester 2006)
- Heavy VC funding in this area (e.g. AirTight $36M in the last 16 months)

The Business World

Systems & Management

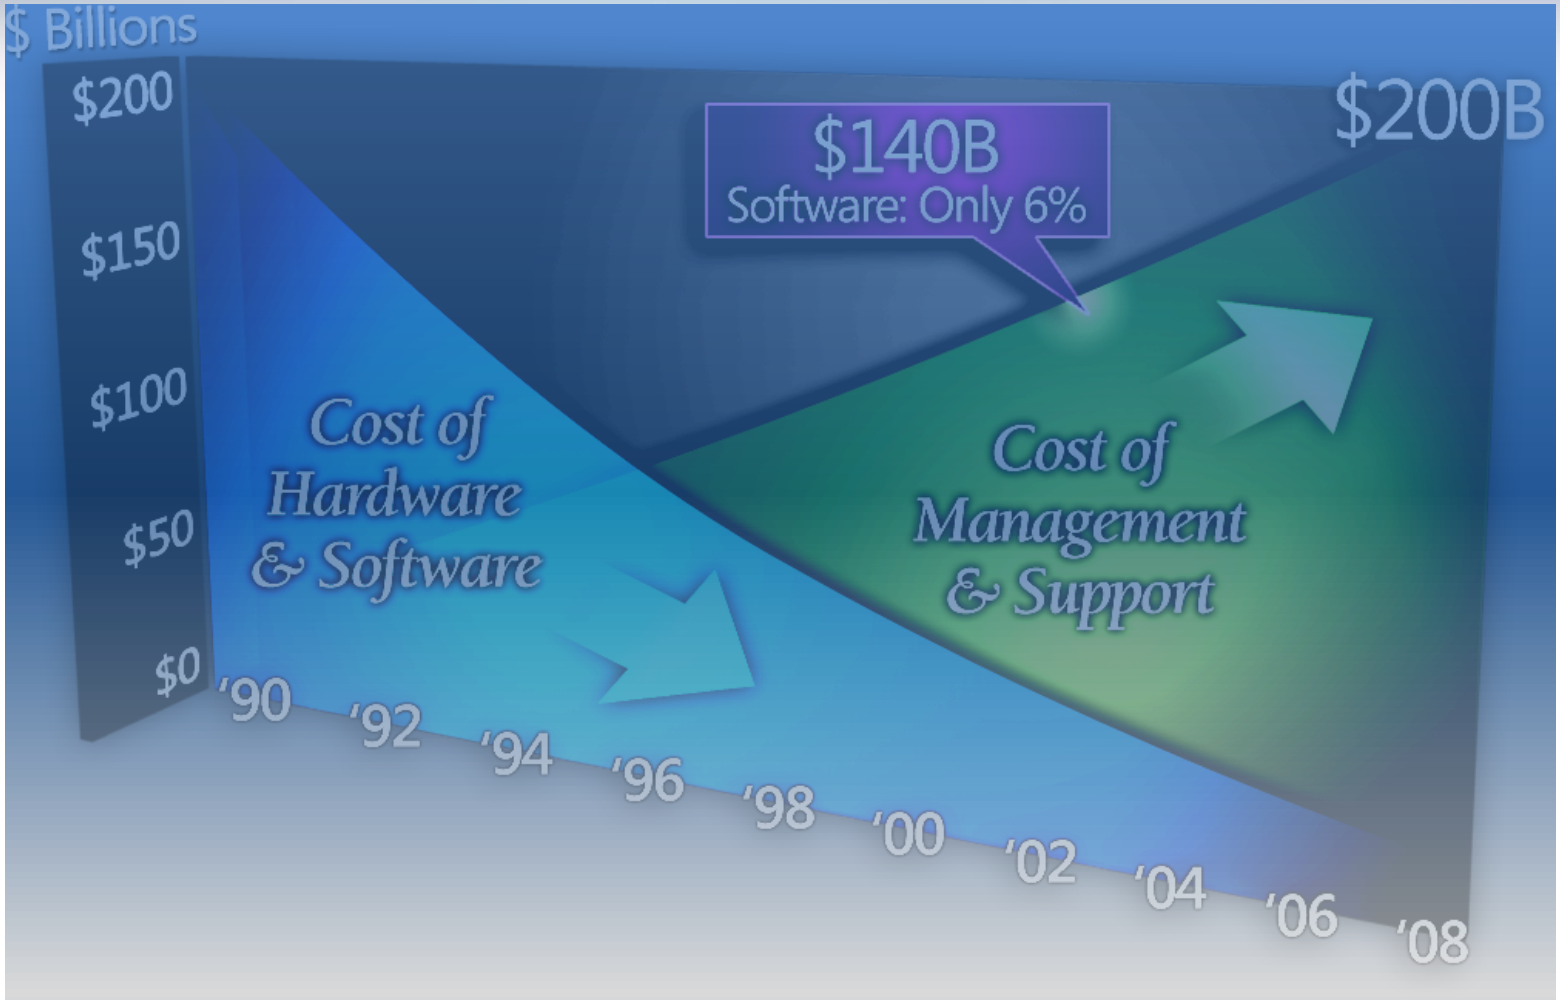

Example: Microsoft IT FY05 $ Expenses

Functional View	FY05 Breakdown
Applications	**60%**
App Development (29%)	
App Support (31%)	
Infrastructure	**40%**
Network (14%)	
Data Center (7%)	
Employee Services (5%)	
Voice (5%)	
Helpdesk (5%)	
Security (3%)	

Cost Element View	
People	72%
Data & Voice	16%
Hardware	5%
Facilities	5%
Software*	2%

* 5% If MS software were included

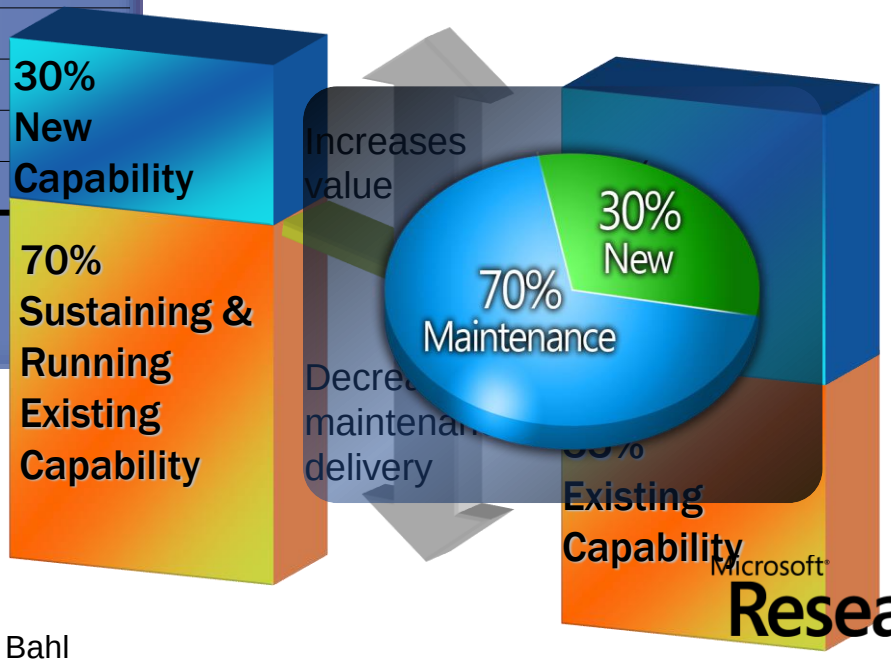

Our March Towards Self Managing Networks

Timeline

HotNets'05, MobiSys'06, NSDI '07

ACM CCR ' 06

Mesh Management

Enterprise Wireless LAN Management (*DAIR*)

2002 2003 2004 2005 2006 2007

Wireless Office

Cooperative WLAN Management (*Client Conduit*)

MobiCom' 04

Cooperative WiFi Diagnostics (*WiFi Profiler*)

MobiSys' 06

EdgeNet2006

sigcomm 2006

ACM SIGCOMM Workshop on Internet Network Management (INM)

Pisa, Italy - September 11, 2006

Access Point Replacement (*Dense AP*)

Microsoft Research

Network Management is Hard!

Heterogeneous world

- Multiple technologies: 802.11 /.15 /.16 /.20 / .22, GPRS, 3G, 1xRTT, EvDO, 4G,…
- Multiple layers: Transport, IP, Ethernet…
- Multiple equipment vendors: Cisco, Juniper, Extreme, Symbol, Aruba,…

Problems can occur anywhere

- Applications, services, first/last hop link, AP, proxy, server, application, switch...

No standard monitoring technique

- What to monitor? Flood of low quality information; Scalability? Cryptic Analyses

Users have very limited understanding & control

- Increased support calls are NOT the answer
- Don’t want to have to call anyone, just want the problem fixed and/or told when it will be fixed

Complexity = expense & slow progress

WLAN Management is Harder

Unpredictable RF Propagation

Many tunable Parameters & Parameter Sensitivity is High

- Frequency band, channel-width, power, rate, multiple radios, ….

Cross-Industry Cooperation is Difficult to Achieve

- Some of them (e.g. cordless phones, baby monitors) may not follow channel discipline
- Some devices such as microwave ovens are incapable of following
- No built in incentive

Topology Discovery is Hard

- Who is affecting my transmission - hidden terminals, mobility, interference,…

Self-interference is rampant

- Multiple host interfaces, multi-hop networks

Root Cause Analysis Techniques are in Their Infancy

- Signature-based techniques do not work - what is normal behavior?

No Standard Metrics for Noise, Power Level etc

Shortcomings of AP based Solutions

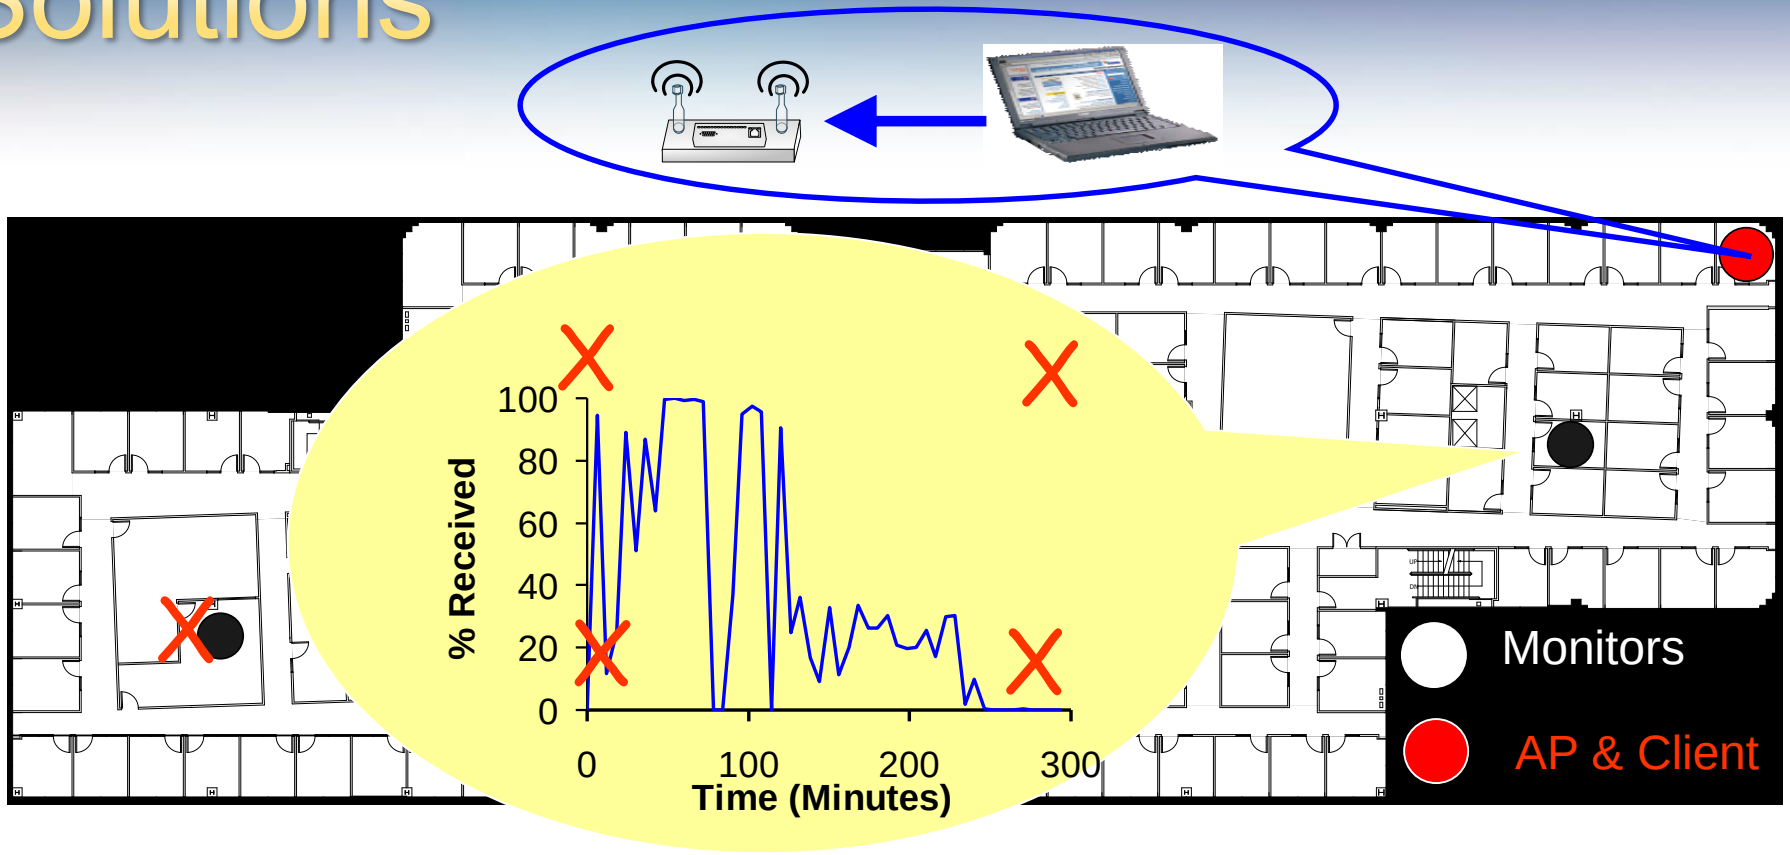

Giving Users Greater Control

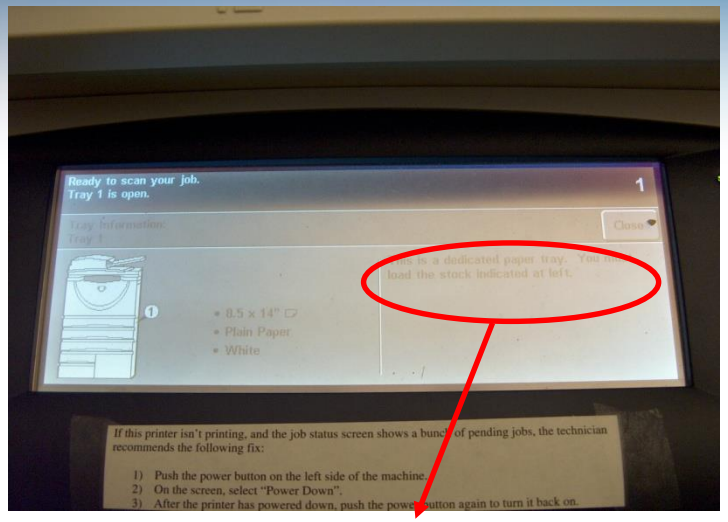

This is a dedicated paper tray. You must load the stock indicated at left

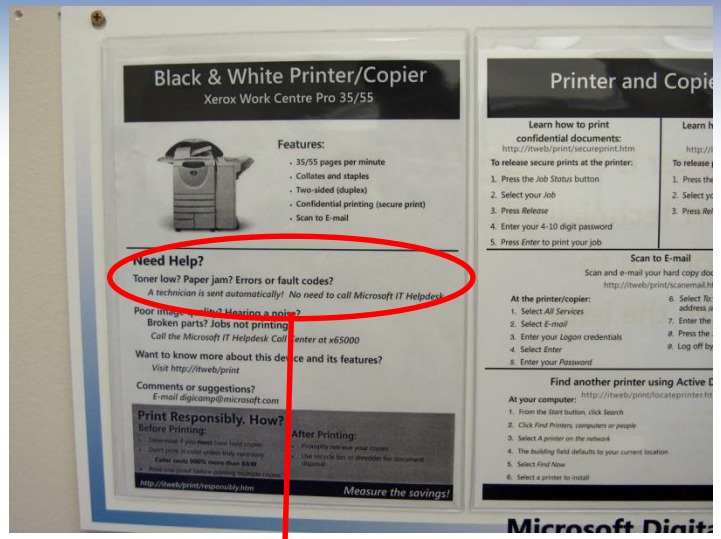

Need Help?

Toner Low? Paper jam? Errors or fault codes?

A technician is sent automatically! No need to call Microsoft IT Helpdesk

Reduce number of support calls
- Help the user/app/network help itself
- Locate the correct party to contact if not

Reduce the time spent on support calls that do occur

Tension between control & automation

Control ⟷ Automation

NetHealth

NetHealth is an end-node based framework for the management of enterprise networks.

Framework

- Integrate end-node view of the network with network services & applications
  - share network experience across end points
  - draw inferences based on automatic correlation
  - automate what expert users do manually
- Integrate peer cooperation
- Compliment existing technologies

Goals

Proactively and reactively:

- Detect, alert, diagnose & repair problems
- Detect, alert & contain security compromises
- Perform root cause analysis of performance problems
- Allow what-if analysis for better resource management

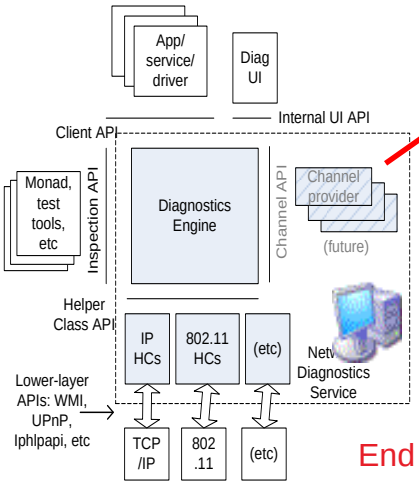

End Host

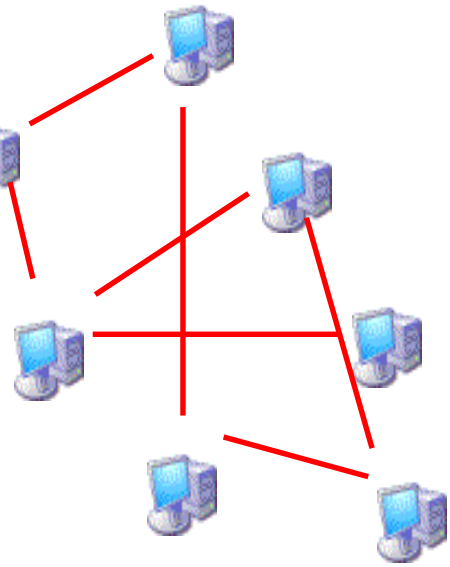

End Host Cooperation

Microsoft Research

NetHealth (Wireless) Projects

Tools to Help Users Help Themselves

- Cooperation between end-nodes for Network Diagnosis & Recovery
  - VirtualWiFi, Client Conduit, WiFiProfiler, SoftRepeater Projects

System & Tools for Managing Enterprise Wireless LAN

- Cooperation between end-nodes and infrastructure servers
  - The DAIR WiFi Network Management Project

Systems & Tools for Managing Wireless Meshes

- Cooperation between end-nodes and infrastructure servers
  - Online simulation based root cause fault analysis
  - What-if Analysis (Time permitting)

Software Infrastructure

Instrumentation
Hooks to look

Naming
Problem identification

Alerting
Getting problem instance (message) to capable agent

Dependency
Learning relationships between distributed application, services & network components

Verifying
Quantifying the user's complaint

Learning & Improving
What is normal/abnormal within a class

Diagnosing & Repairing
Handling faults until they are fixed

Network Visualization

Important:
Must be Complimentary to Existing Technologies

- Network Diagnostic Infrastructure
- SNMP
- Native WiFi
- MOM
- SMS / Event logger
- Operations Manager
- Systems Center Capacity Planner
- Active Directory & Group Policy

Tools to Help Users Help Themselves

Cooperative Peer–to-Peer Network Diagnosis & Recovery

Automate network fault diagnosis and recovery
Reduce user frustration and admin load

Use peer cooperation to improve network health

Details: INFOCOM ‘04

VirtualWiFi

A single wireless NIC appears as multiple cards

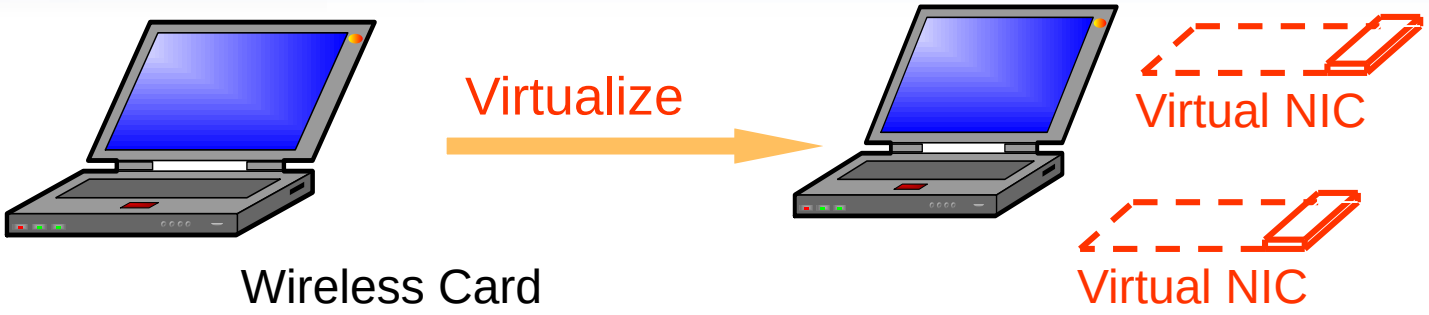

Virtual cards

- Appear as real network interfaces to upper layers
- Each virtual card can connect to any network

Microsoft Research

Details: MobiCom '04

Helping Disconnected Clients

Client Conduit

Possible causes of disconnection:

- Lack of coverage, e.g. In an RF Hole, just outside AP range, …
- Authentication problem, e.g., stale certificates, …
- Protocol problem, e.g., no DHCP address

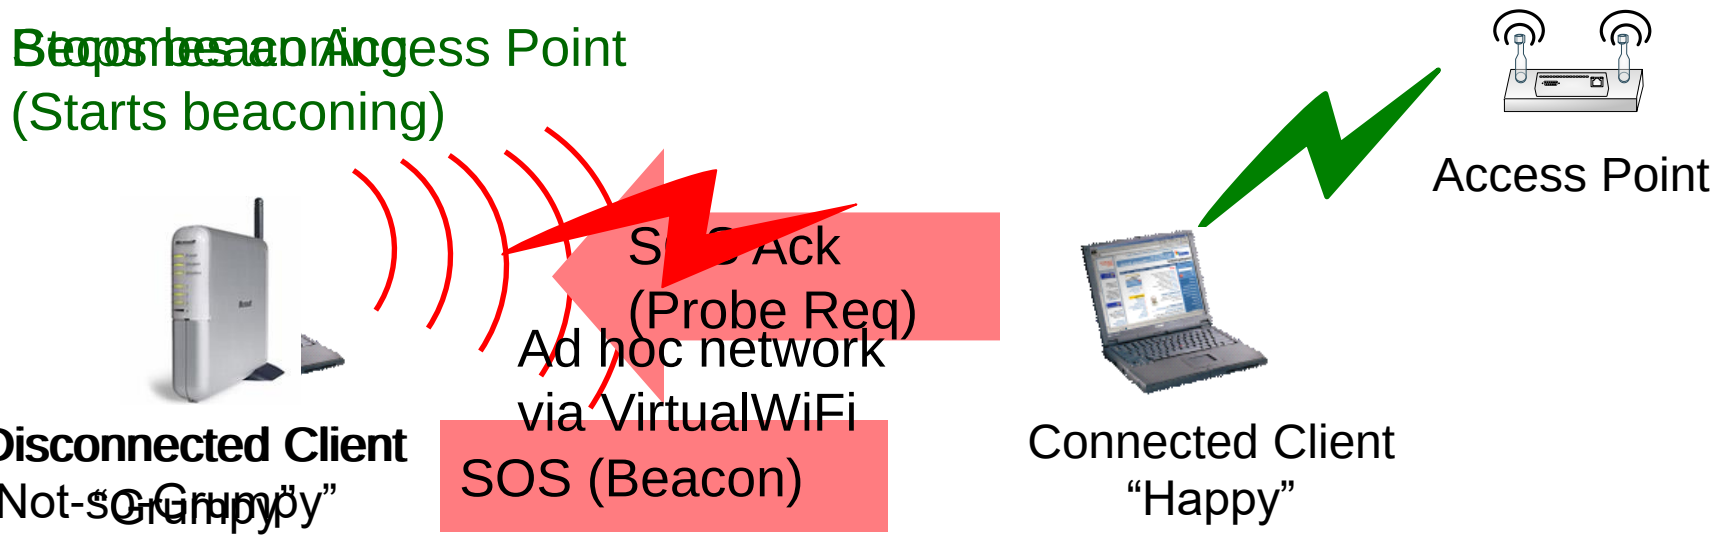

When "Happy" donates only 20% of time; Bandwidth available for diagnosis > 400 Kbps

Microsoft Research

Details: MobiSys '06

WiFiProfiler
Cooperative Diagnosis in WLANs

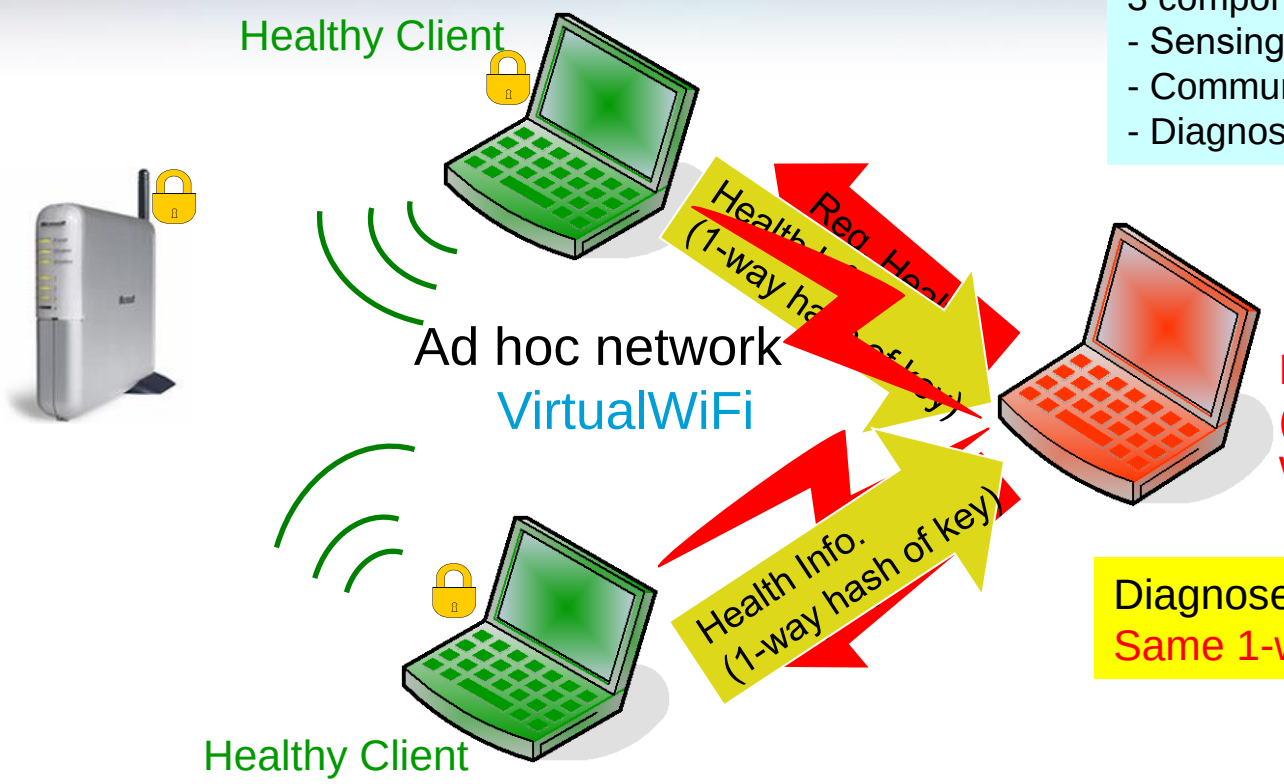

3 components:
- Sensing: Collect local "health" info
- Communication: Send info to peers
- Diagnosis: Use info to diagnose faults

Diagnose Problem:
Same 1-way hash of WEP key?

Diagnose range of problems across layers

- No association due to MAC filtering or driver incompatibility
- No DHCP address due to bad WEP key or bad server
- Poor WAN Performance due to wireless or wired problems
- No Internet connectivity due to incorrect proxy

Microsoft Research

SoftRepeater

Solving Performance Problems

802.11 data rate depends on RF distance

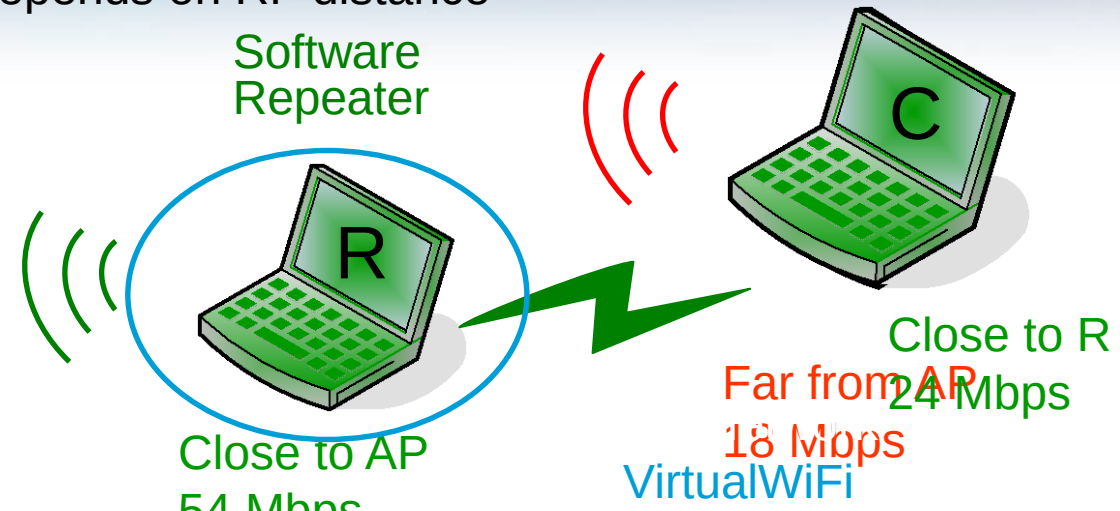

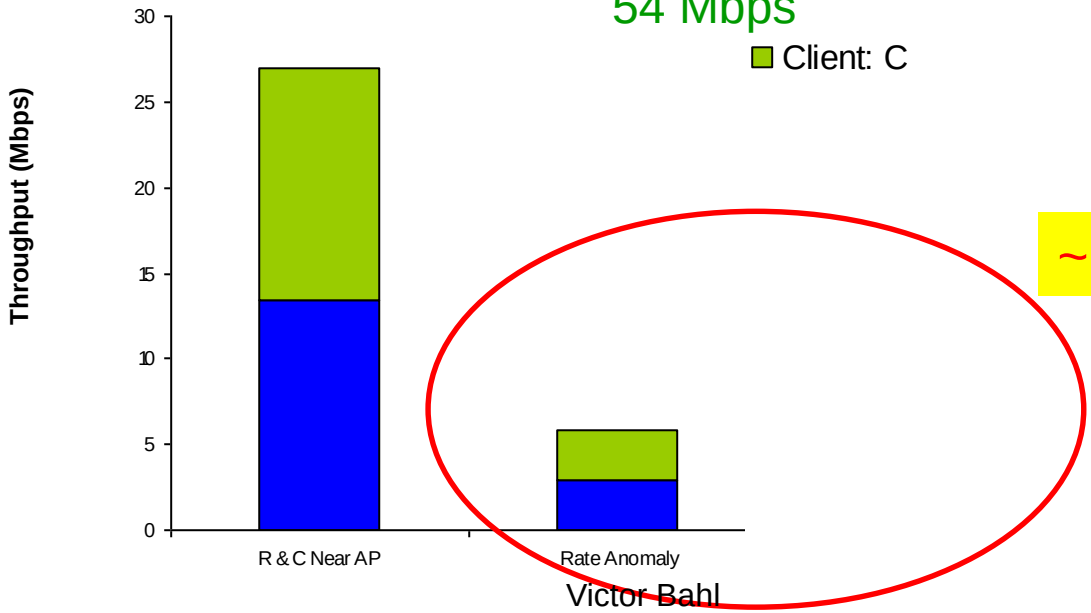

~ double throughput

Victor Bahl

Microsoft Research

SoftRepeater

Solving Performance Problems

Using Network Coding to improve capacity

= ~~3~~ 2 transmissions in the air

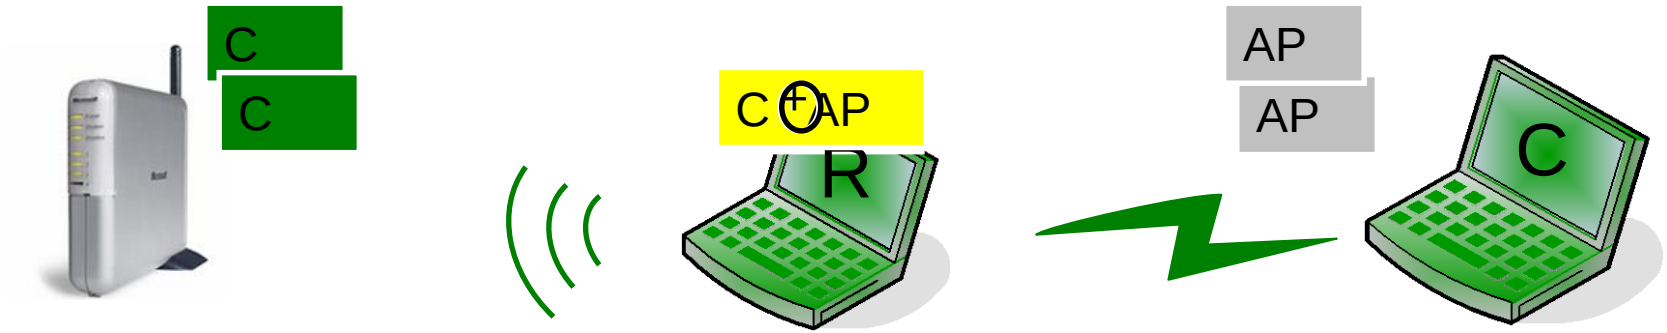

Zero network overhead implementation on Windows XP

- no extra bytes in packet headers

Throughput (in Mbps)	w/o Network Coding	Network Coding
UDP (AP→C, C→AP)	11.02	18.13 **(+64%)**
TCP (AP→C, C→AP)	10.91	13.97 **(+28%)**
TCP (C →AP)	10.55	12.11 **(+15%)**

Summarizing
Using Mobile Hosts for Management

The Good

- No infrastructure required
- Exploits host-view of network
- Provides quick and effective diagnosis
- Incurs low overhead for connected (healthy) clients
  - Use existing 802.11 messages: beacons & probes
- Lets users help themselves

The Bad

- Difficult to provide predictable coverage
- Dependent on battery & energy constraints

….what if we have infrastructure support

Tools for Managing Enterprise Wireless Networks

Cooperative Client-Server Network Diagnosis & Recovery

Automate network fault diagnosis and recovery
Reduce user frustration and admin load

Wireless LAN Management System
Requirements

- Must manage the effects of RF propagation
  - Provide comprehensive spatial coverage
- Must Integrate location into the management system

High Sensor Density

- Should determine performance problems & provide meaningful analysis
  - Reduce false positives & prioritize alerts
- Must locate and contain security breaches
- Should resolve problems automatically

Observations

Details: HotNets'05, MobiSys'06, NSDI '07

- Desktop PC's with good *wired* connectivity are ubiquitous in enterprises

- Outfitting a desktop PC with 802.11 wireless is inexpensive
  - Wireless USB dongles are cheap
    - As low as $6.99 at online retailers
  - PC motherboards are starting to appear with 802.11 radios built-in

Combine to create a dense deployment of wireless sensors

DAIR: Dense Array of Inexpensive Radios

The DAIR Enterprise Wi-Fi Management System

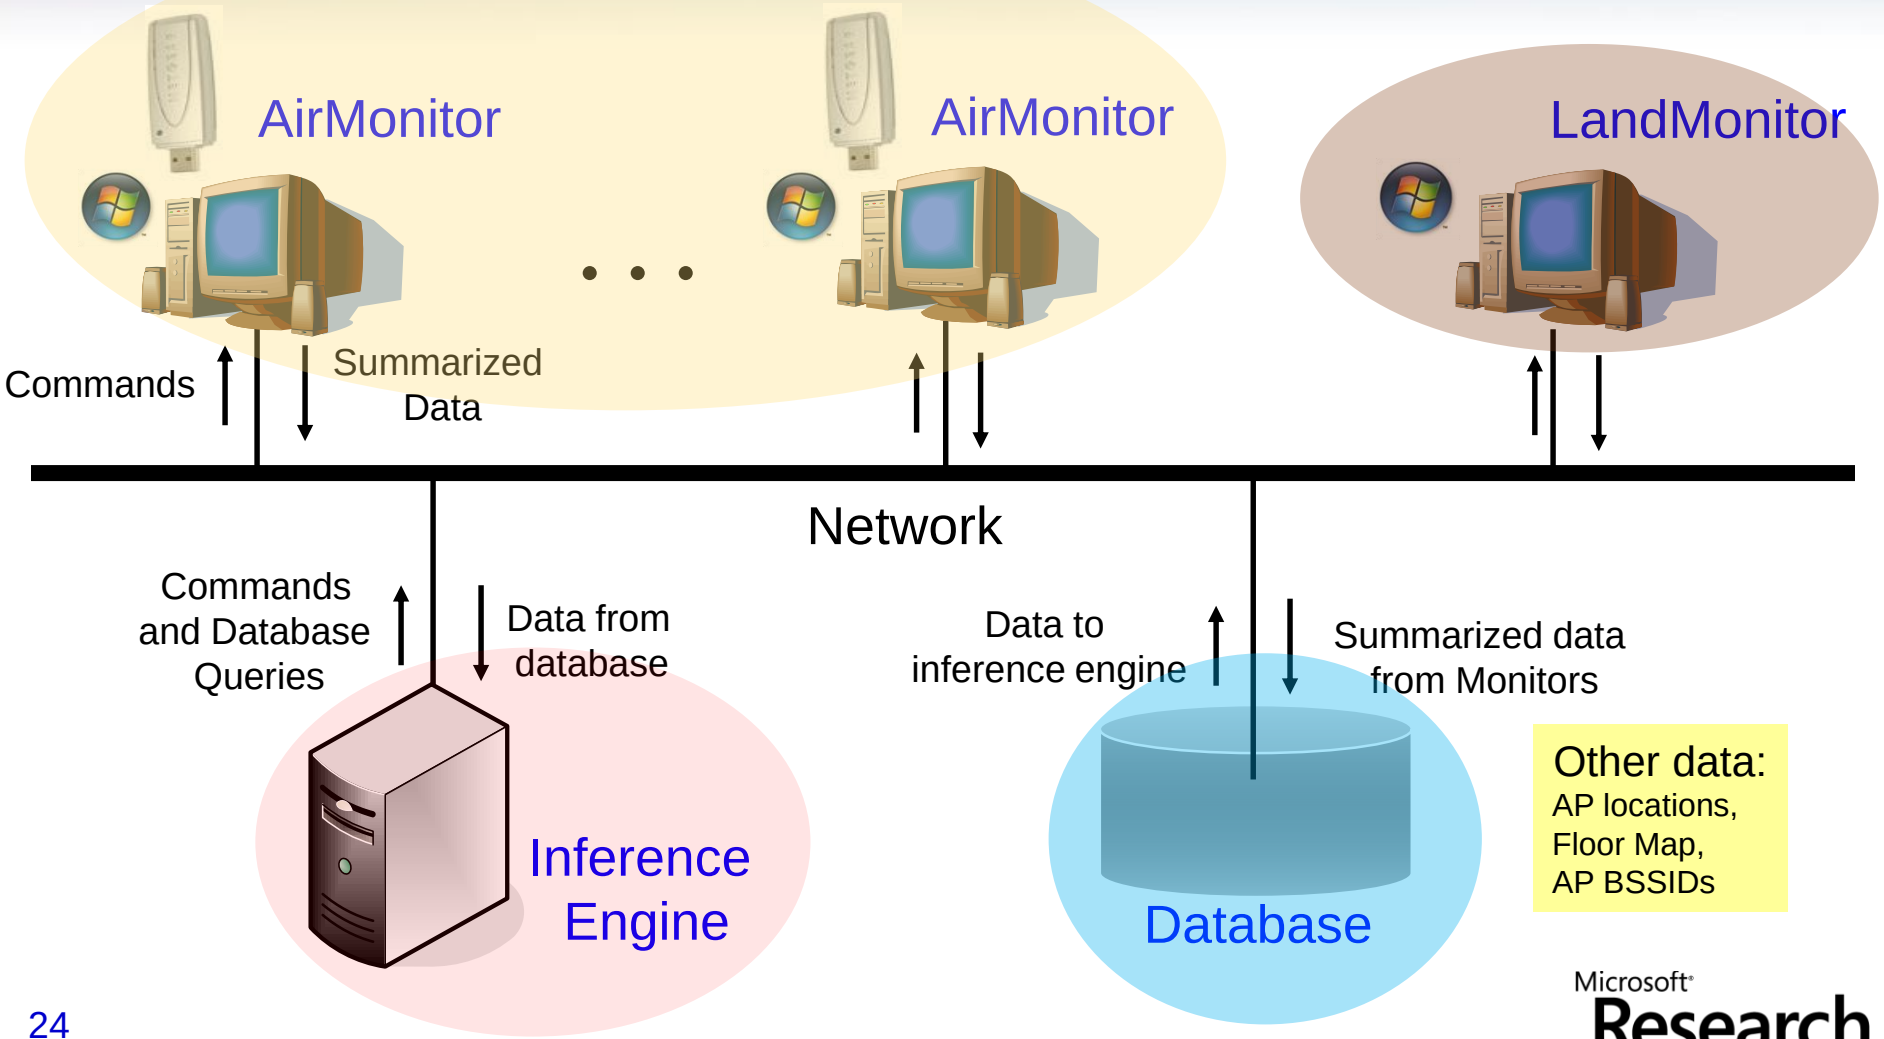

Monitor Software Architecture

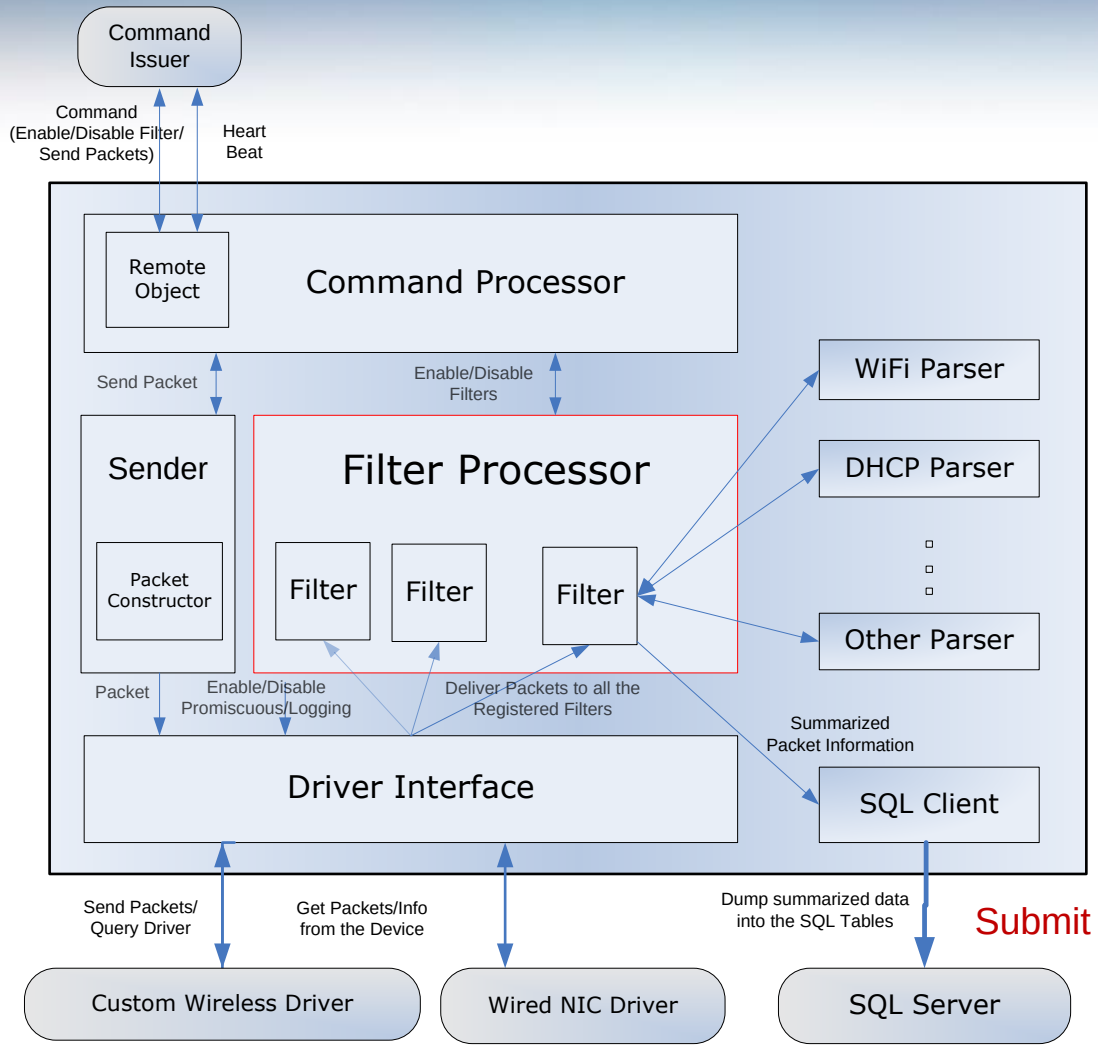

Load on desktops < 2-3%

Network traffic per AirMonitor < 10Kbps

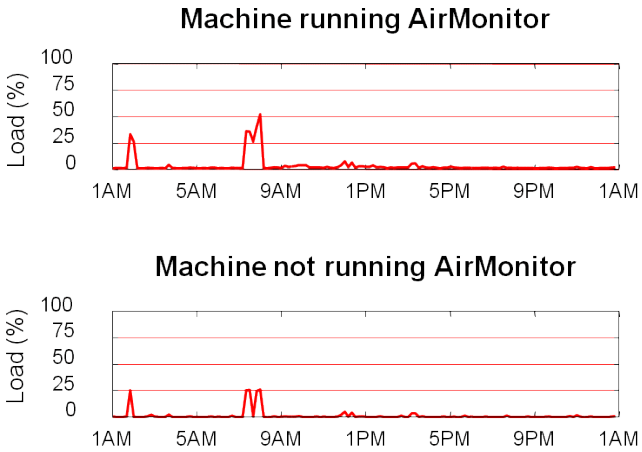

Submit summaries

Sample Research Problems Solved

Details: HotNets'05, MobiSys'06, NSDI '07

Algorithmic Innovations:

- Self-configuring location determination system (DAIR)
- Detecting & attacking rogue wireless nets (DAIR)
- Detecting performance anomalies and RF holes (DAIR)
- Detecting & responding to DoS attacks (DAIR)
- Assigning channel & power; managing handoff (DenseAP)

Systems Innovations:

- Scaling to the size of an enterprise
- Bootstrapping the location system
- Limiting the impact of sensors on office PCs
- Introducing new techniques while remaining backward compatible

Status

60-node system operational for over 8 months, MS-IT & DELL deployment discussions (on-going)

Self-Configuring Indoor Location System

Here's how :

- AirMonitors (AM) automatically determine their position
- AMs collectively profile the RF environment by measuring the signal propagation characteristics between one another
- Inference Engine (IE) uses the RF profiles and signal strength observations at multiple AMs to locate Wi-Fi transmitters

The DAIR system can locate any Wi-Fi transmitter (including non-cooperative ones) to office-level accuracy

AirMonitors Locate Themselves

- Monitor machine activity to determine primary user
- Look up Directory Services (e.g. Active Directory) to determine office number
- Parse office map to determine coordinates of the office
  - Assume AMs to be located at the center of the office
- Improve estimates by verifying & adjusting coordinates by observing which AMs are nearby

Details: INFOCOM '99

RF Propagation Modeling

$$P(d)[dBm] = P(d_o)[dBm] - 10n\log\left(\frac{d}{d_o}\right) - \begin{cases} nW * WAF, & nW < C \\ C * WAF, & nW \geq C \end{cases}$$

MSR RADAR System (1999)

$P(d_o) = 28\ dBm, n = 1.53$

$WAF = 3.1\ dBM, C = 4\ Walls$

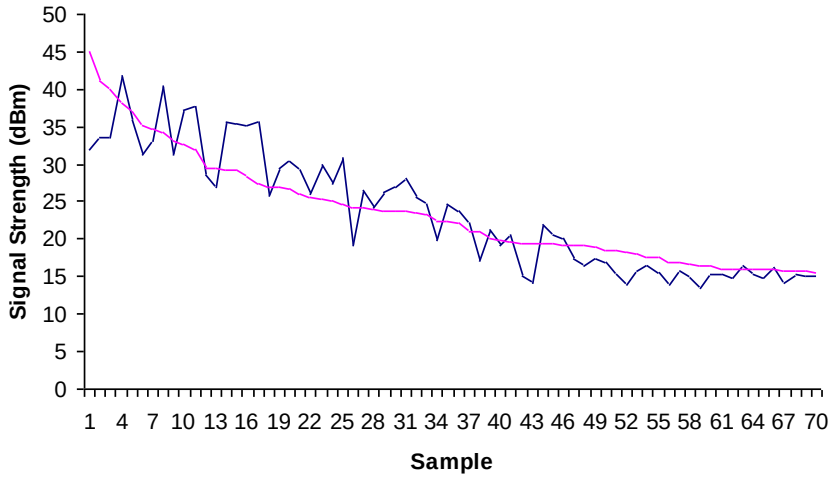

$P(d)[dBm] = m\,d + C$

$m = -1.4, C = 35.7\ dBm$

$P(d)[dBm] = Ae^{\lambda d}$

$A = 60, \lambda = -0.11$

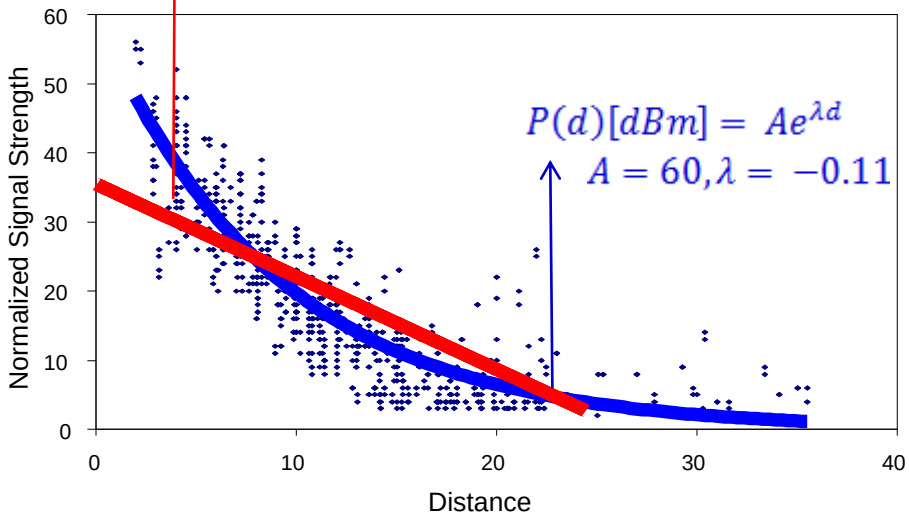

Good News: Don't need sophisticated RF Propagation Models

Each AM determines it's own profile

Microsoft Research

Locating the Wi-Fi Transmitter

Details: NSDI '07

Observed RSSI: 50

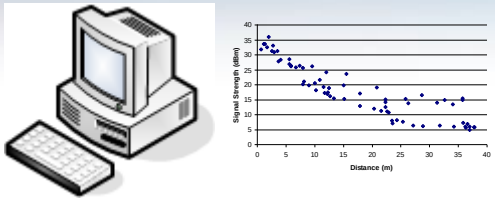

Distance: 3, Estimated RSSI: 54

Distance: 1.3, Estimated RSSI: 51

Observed RSSI: 45

$$P(d)[dBm] = Ae^{\lambda d}$$

Distance: 7.2, Estimated RSSI: 35

Distance: 6.0, Expected RSSI: 41

Observed RSSI: 52

$$P(d)[dBm] = Ke^{\mu d}$$

Distance: 0, Estimated RSSI: 56

Distance: 1.1, Estimated RSSI: 52

Observed RSSI: 44

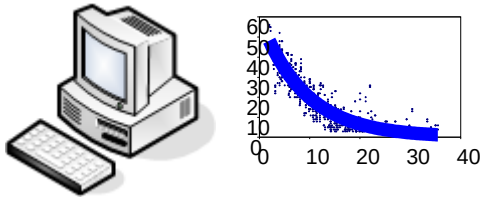

Distance: 6.5, Estimated RSSI: 38

Distance: 6.2, Estimated RSSI: 47

Microsoft Research

Deployment

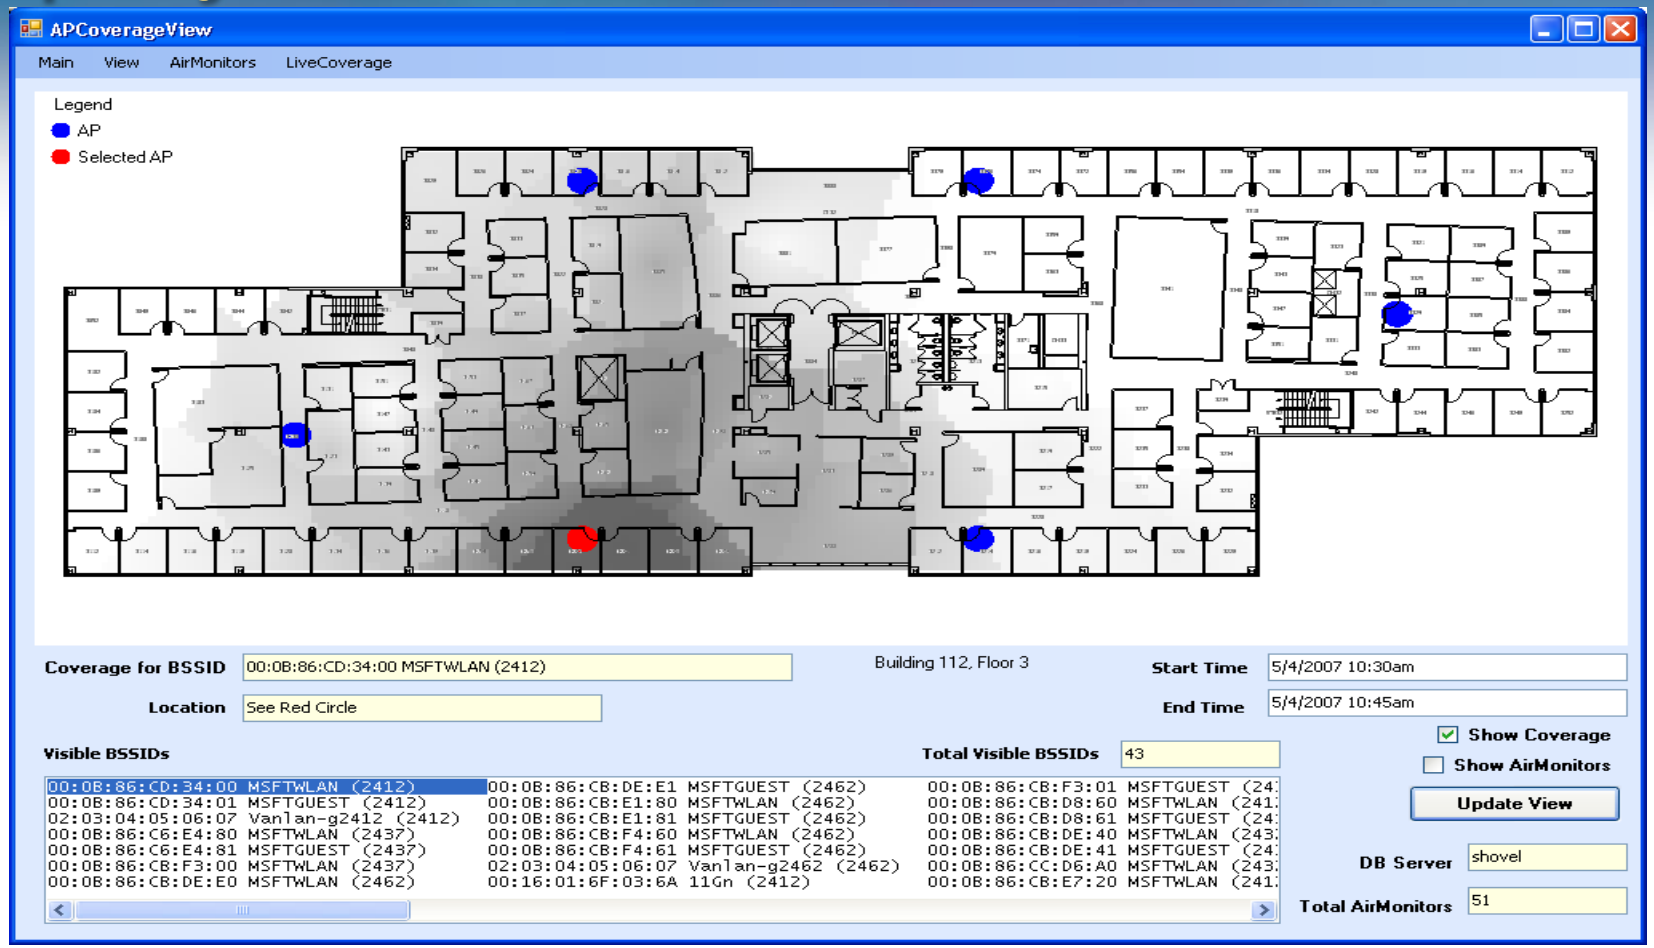

98 meters x 32 meters
150 offices and conference rooms.
Typical office size: 3 meters x 3 meters
Full-height walls. Solid wood doors
59 AirMonitors.

Microsoft Research

DAIR Infrastructure Applications

Performance Management

Isolate performance problems
- Help disconnected clients
- Detect & fix RF Holes
- Detect mis-configuration

Reliability
- Recover from malfunctioning APs
- Compensate for poor association policies

Monitoring
- Site planning: AP placement, frequency / channel selection
- Load balancing

Security Management

Detect rogue wireless nets
- Infrastructure and ad-hoc

Detect DoS attacks
- Spoofing disassociation
- Large NAV values
- Jamming

Contain Attackers
- Attack the attackers

Management

Management

DenseAP project

Access Point Replacement
- Self configuring deployment
- Better spatial reuse

Layer 7 Applications & Services
- Indoor GPS
- Seamless Roaming
- Guest Access

The Wireless Management Ecosystem

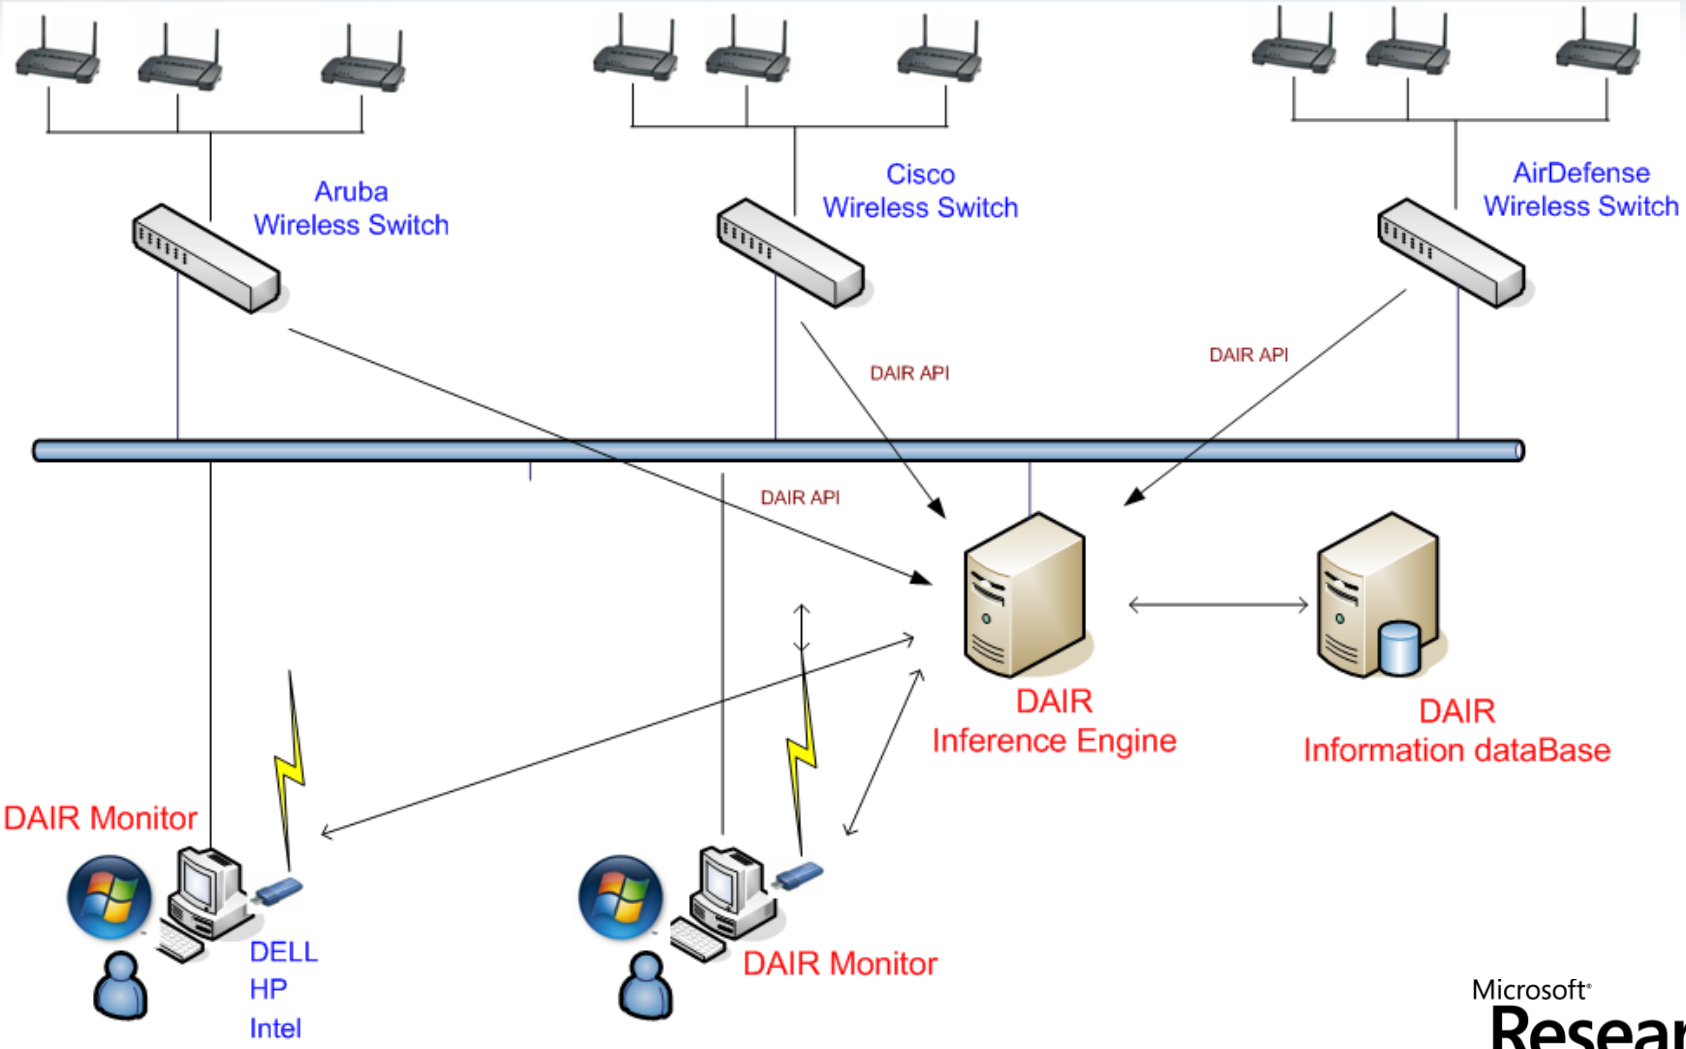

Microsoft Research

Mesh Networking

Managing Meshes

The least well understood area of research

Broadband Connectivity
- Rural & developing areas
- City-wide
- Neighborhoods / Communities
- Wireless Office

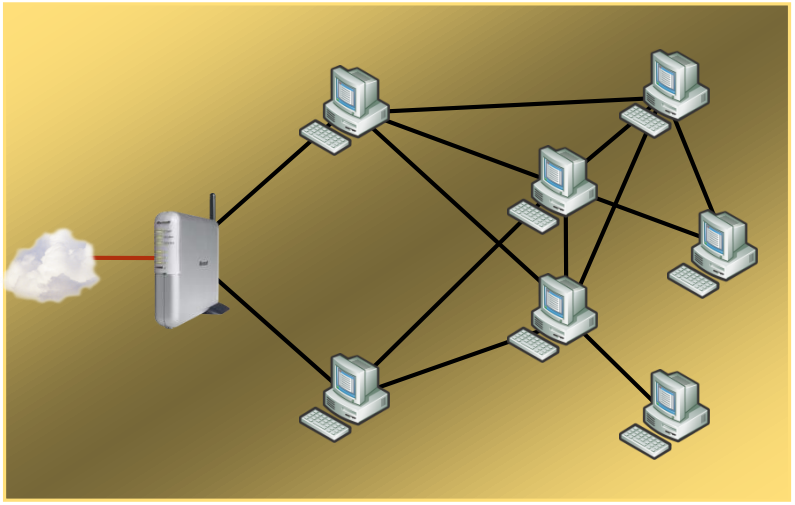

Is this Normal Behavior?

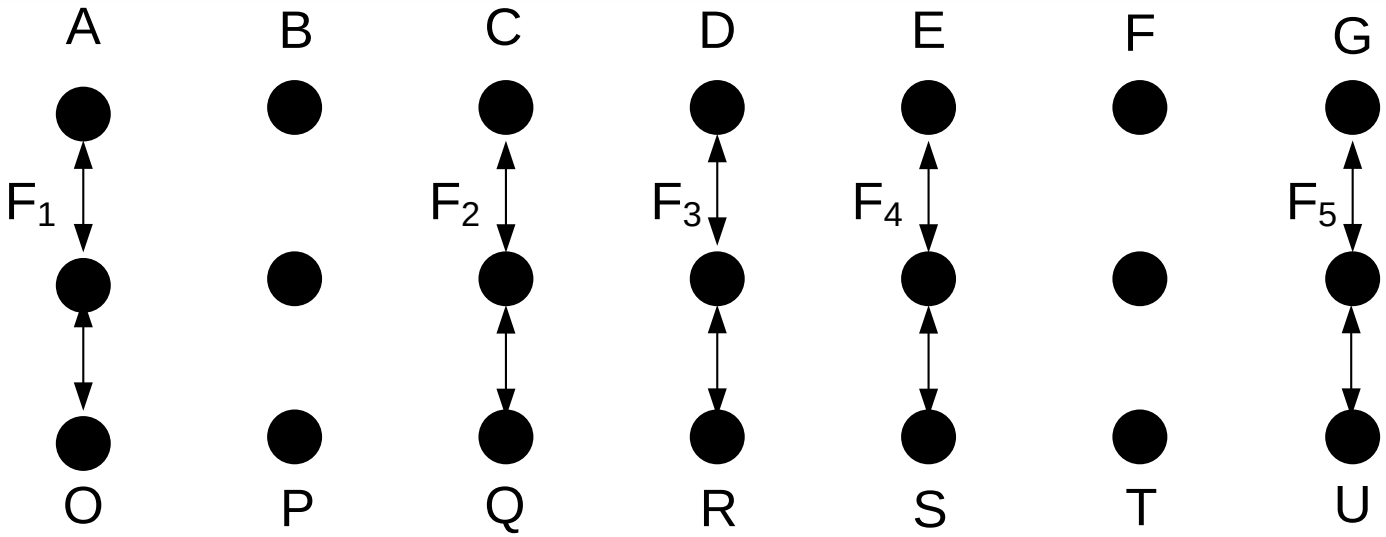

$Flow_1$	$Flow_2$	$Flow_3$	$Flow_4$	$Flow_5$
2.5 Mbps	0.23 Mbps	2.09 Mbps	0.17 Mbps	2.55 Mbps

Control Flow

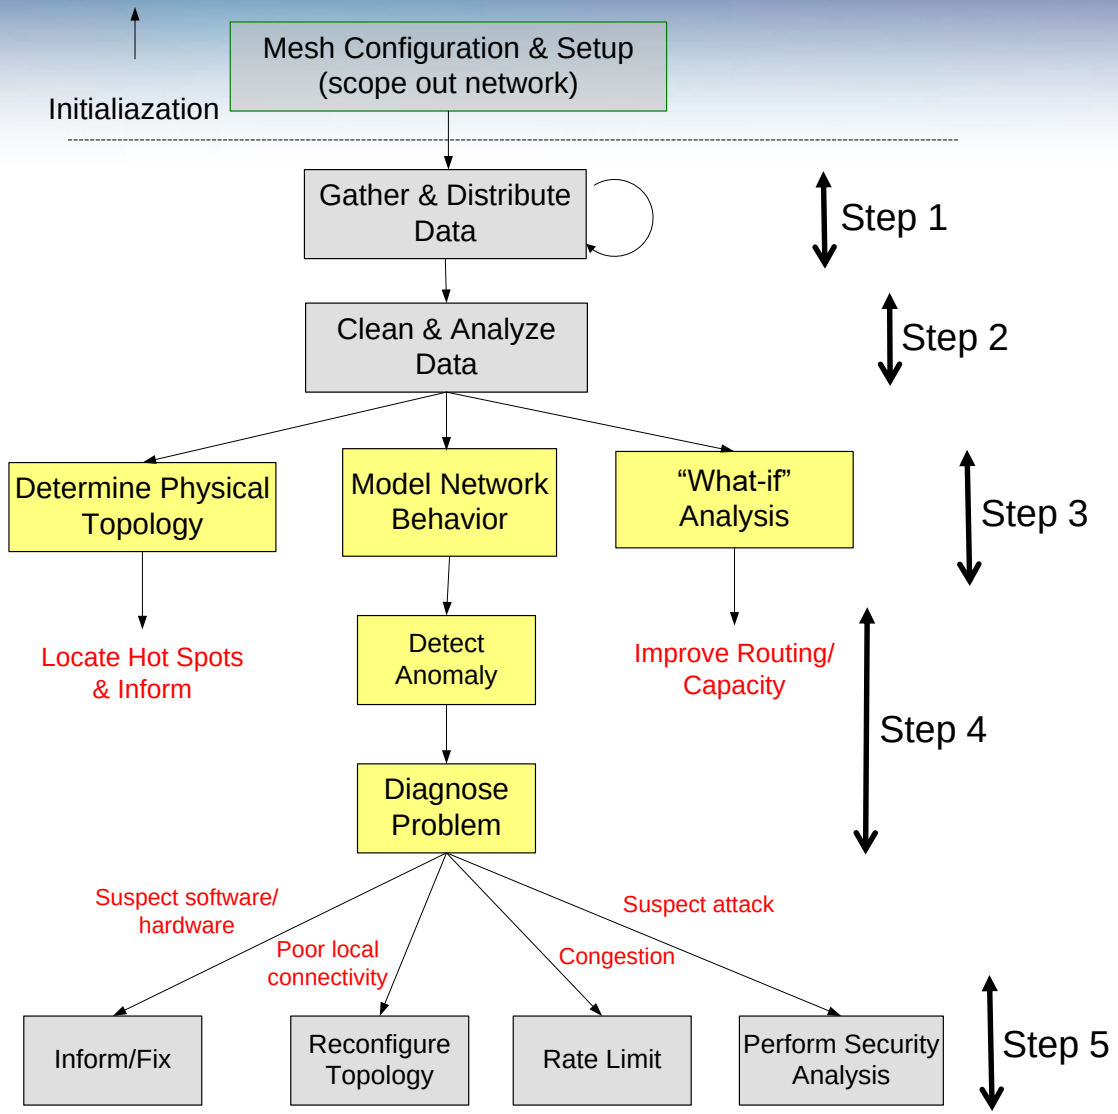

Step 1: Gather & Distribute Data

Monitoring: What should we collect?

- Link Info: Noise level, signal strength, loss rate to direct neighbor (packet retransmission count)
- Connectivity Info: Network topology / connectivity Info (Neighbor Table)
- Traffic Info: Load to direct neighbor
- …

Distribution: Minimize (overhead) bandwidth consumption

- Dynamic scoping
  - Each node takes a local view of the network
  - The coverage of the local view adapts to traffic patterns
- Adaptive monitoring
  - Minimize measurement overhead in normal case
  - Change update period
  - Push and pull
- Delta compression
- Multicast

Step 2: Clean & Analyze Data

Data may not be pristine. Why?

- Liars, malicious users
- Missing data
- Measurement errors

Clean the Data

- Detect Liars
  - Assumption: most nodes are honest
  - Approach:
    - Neighborhood Watch
    - Find the smallest number of lying nodes to explain inconsistency in traffic reports
- Smoothing & Interpolation

Sample Performance
Resiliency against Liars & Lossy Links

Details: CCR '06

Problem
- Identify nodes that report incorrect information (liars)
- Detect lossy links

Assume
- Nodes monitor neighboring traffic, build traffic reports and periodically share info.
- Most nodes provide reliable information

Challenge
- Wireless links are error prone and unstable

Approach
- Find the smallest number of lying nodes to explain inconsistency in traffic reports
- Use the consistent information to estimate link loss rates

Results

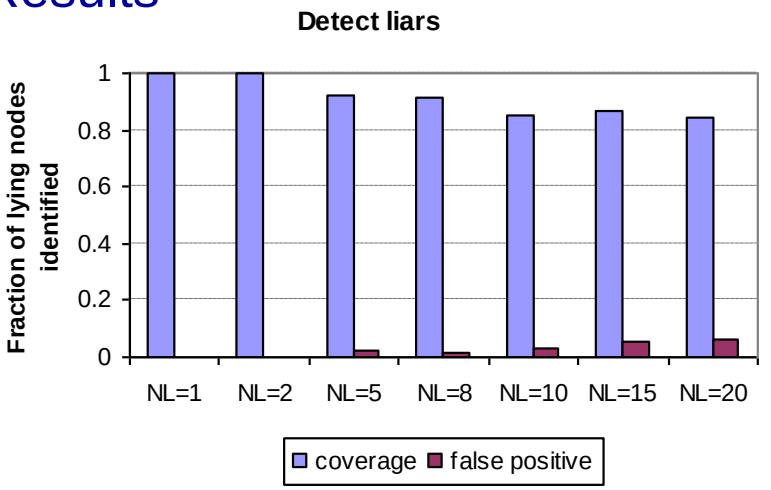

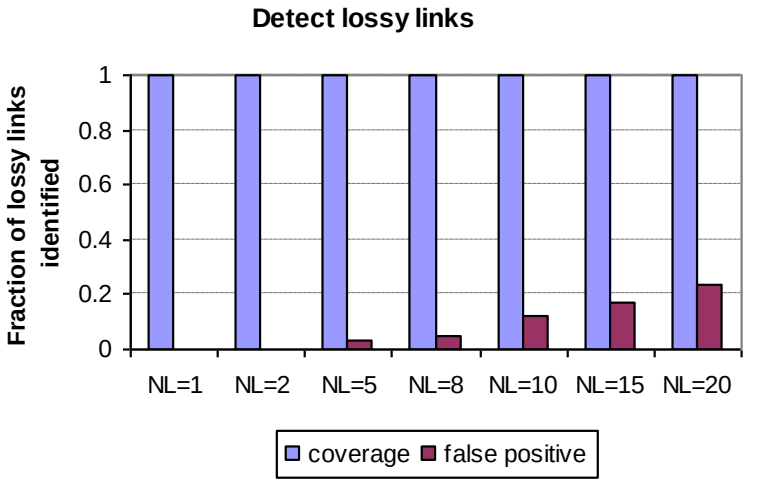

Details: CCR '06

Step 3 & 4:
Model Network & Perform Root Cause Analysis

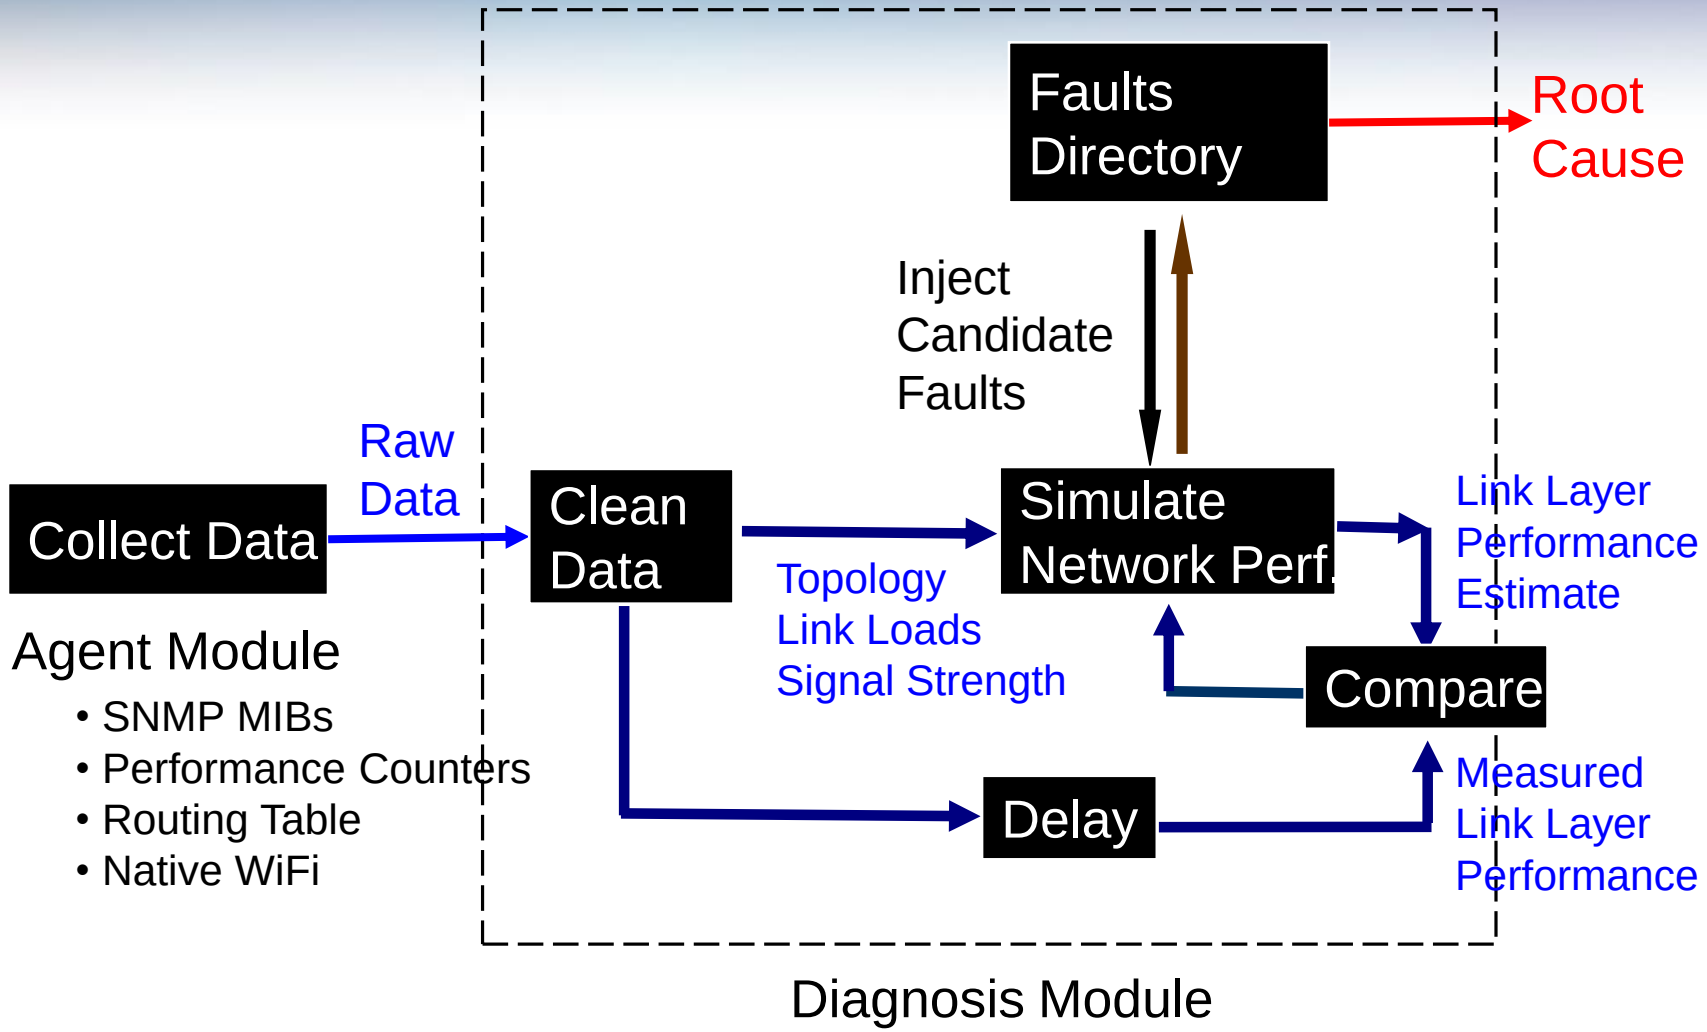

Microsoft Research

Sample Performance

Details: CCR '06

25 node random topology

Number of faults	4	6	8	10	12	14
Coverage	1	1	0.75	0.7	0.92	0.86
False Positive	0	0	0	0	0.25	0.29

Faults detected:

- Random packet dropping
- MAC misbehavior
- External noise

Troubleshooting Framework

Challenges [in Online Simulation based Diagnostics]:

- Accurately reproduce the behavior of the network inside a simulator
- Build a fault diagnosis technique using the simulator as a diagnosis tool

Advantages

- Flexible & customizable for a large class of networks
- Captures complicated interactions

  within the network

  between the network & environment, and

  among multiple faults
- Extensible in its ability to detect new faults
- Allows what-if analysis

Step 5: Mitigation

Responding to troubled spots

- Re-route traffic
- Rate-limit
- Change topology via power control & directional antenna control
- Flag
  - environmental changes & problems
  - Malfunctioning hardware
- Launch DoS attacks against the possible attacker
- etc.

So where does all this leave us……

Think about what's coming?

- Micro-cellular architectures
- Multi-standard, multi-radio devices
- New technologies: WiMax, UWB, .11n, 4G, 60 GHz,…
- Cognitive networking
  - — Reconfigurable adaptive stacks, SDRs, Agile radios
- Data networking in the TV Bands
- Time-sensitive applications
- Sensor Networking

Billions of Devices will have to be Managed

Management & Perfromance is Key!

Wireless networks are complex & difficult to diagnose but diagnostics are critical to wireless deployments

Opportunity to conduct seminal research

- Make networks more deployable in IT-poor markets
- Reduce IT costs in the enterprise
  - Take advantage: infrastructure & end systems owned by same organization

Host-centric approaches show great promise

Tradeoff between gains from management and loss because of overhead

Are we there yet?

Not yet…..

…..but surely getting there

Self-aware, self-healing, easy-to-manage networks

Q/A

http://research.microsoft.com/netres/nethealth/

Microsoft®

Your potential. Our passion.™