18-452/18-750
Wireless Networks and Applications
Lecture 12: MIMO and WiFi Deployments

Peter Steenkiste

Spring Semester 2018

http://www.cs.cmu.edu/~prs/wirelessS18/

Outline

- MIMO and recent WiFi versions
  - » Refresher: spatial diversity
  - » MIMO basics
  - » Single user MIMO: 802.11n
  - » Multi-user MIMO: 802.11ac
  - » Millimeter wave: 802.11ad
- WiFi deployments
  - » Planning
  - » Channel selection
  - » Rate adaptation

How Do We Increase Throughput in Wireless?

- Wired world:

Pull more wires!

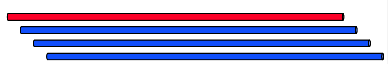

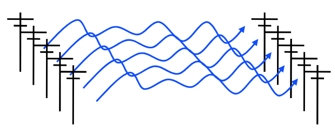

- Wireless world:

How about if we could do the same thing and simply use more antennas?

MIMO
Multiple In Multiple Out

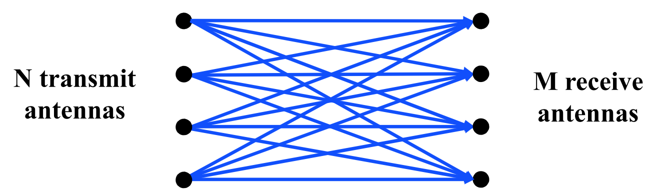

- N x M subchannels that can be used to send multiple data streams simultaneously
- Fading on channels is largely independent
  - » Assuming antennas are separate ½ wavelength or more
- Combines ideas from spatial and time diversity, e.g. 1 x N and N x 1
- Very effective if there is no direct line of sight
  - » Subchannels become more independent

Why So Exciting?

Method	Capacity
SISO	$B \log_2(1 + \rho)$
Diversity (1xN or Nx1)	$B \log_2(1 + \rho N)$
Diversity (NxN)	$B \log_2(1 + \rho N^2)$
Multiplexing	$NB \log_2(1 + \rho)$

802.11 with multiple antennas for dummies, Daniel Halperin, Wenjun Hu, Anmol Sheth, David Wetherall, ACM CCR, Jan 2010

Spatial Diversity

- Use multiple antennas that pick up the signal in slightly different locations
  - » Channels uncorrelated with sufficient antenna separation
- Receiver diversity: i x H x P_R = o

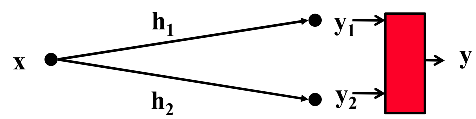

$\vec{y} = \vec{h} * x + \vec{n}$

$y = \vec{h}^* * (h * \vec{x} + \vec{n})$

- Receiver can pick strongest signal: y_1 or y_2
- Or combines the signals: multiply $\vec{y}$ with the complex conjugate h^* of the channel vector h
  - » Can learn h based on training data (Lecture 5)

Other Diversity Options

- Transmit diversity: i x P_T x H = o

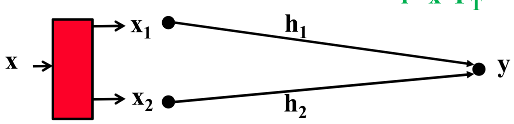

- Combined: i x P_T x H x P_R = o

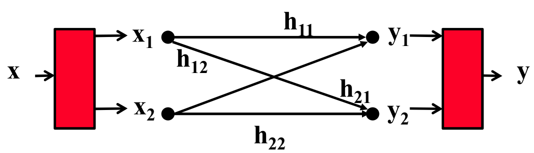

MIMO How Does it Work?

- Transmit and receive multiple data streams
- Coordinate the processing at the transmitter and receiver to overcome channel impairments
  - » Maximize throughput or minimize interference

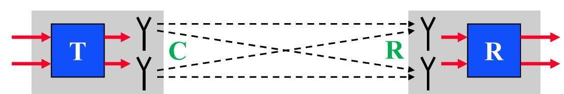

I x P_T x H x P_R = O

Precoding — Channel Matrix — Combining

- Combines previous techniques

Direct-Mapped NxM MIMO Only Receiver Processing (P_T=I)

Effect of transmission: $\underset{M}{\vec{R}} = \underset{MxN}{H} * \underset{N}{\vec{C}} + \underset{M}{\vec{N}}$

Decoding: $\underset{D}{\vec{O}} = \underset{DxM}{P_R} * \underset{M}{\vec{R}}$ $\quad \underset{N}{\vec{C}} = \underset{N}{\vec{I}}$

Results: $\vec{O} = P_R * H * \vec{I} + P_R * \vec{N}$

- How do we pick P_R ? "Inverse" of H: H^{-1}
  - » Equivalent of nulling the interfering signals (zero forcing)
  - » Only possible if the paths are completely independent
- Noise amplification is a concern if H is non-invertible – its determinant will be small

An Example of Space Coding

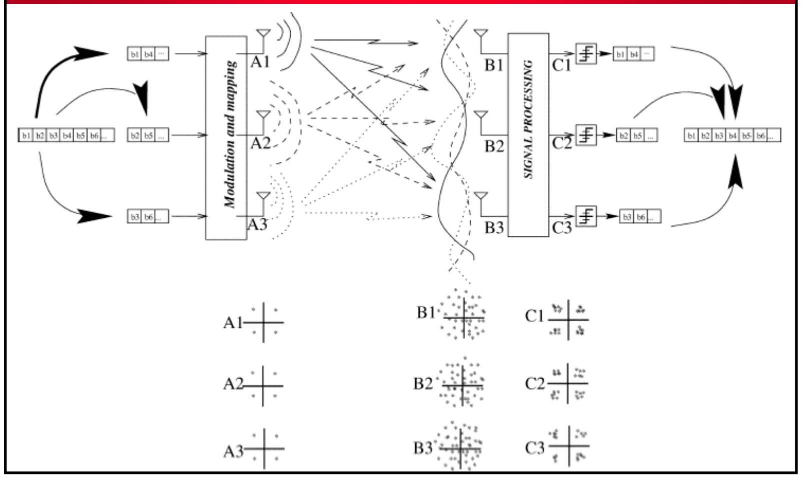

Precoded NxM MIMO

Effect of transmission: $\underset{M}{\vec{R}} = \underset{MxN}{H} * \underset{N}{\vec{C}} + \underset{M}{\vec{N}}$

Coding/decoding: $\underset{D}{\vec{O}} = \underset{DxM}{P_R} * \underset{M}{\vec{R}}$ $\quad \underset{N}{\vec{C}} = \underset{NxD}{P_T} * \underset{D}{\vec{I}}$

Results: $\vec{O} = P_R * H * P_T * \vec{I} + P_R * \vec{N}$

- How do we pick P_R and P_T ?
- Singular value decomposition of H = U * S * V
  - » U and V are unitary matrices – U^H*U = V^H*V = I
  - » S is diagonal matrix

Mechanisms Supported by MIMO

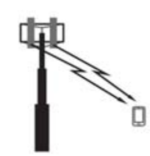

Diversity for improved system performance

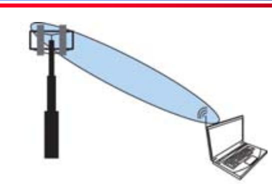

Beam-forming for improved coverage (less cells to cover a given area)

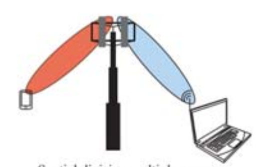

Spatial division multiple access ("MU-MIMO") for improved capacity (more user per cell)

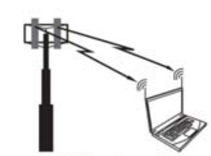

Multilayer transmission ("SU-MIMO") for higher data rates in a given bandwidth

MIMO Discussion

- Need channel matrix H: use training with known signal
- So far we have ignored multi-path
  - Each channel is multiple paths with different properties
  - Becomes even messier!
- MIMO is used in 802.11n
  - Can use two adjacent non-overlapping “WiFi channels”
  - Raises lots of compatibility issues
  - Potential throughputs of 100s of Mbps
- Focus is on maximizing throughput between two nodes
  - Is this always the right goal?

802.11n Overview

- 802.11n extends 802.11 for MIMO
  - Supports up to 4x4 MIMO
  - Preamble that includes high throughput training field
- Standardization was completed in Oct 2009, but, early products had long been available
  - WiFi alliance started certification based on the draft standard in mid-2007
- Supported in both the 2.4 and 5 GHz bands
  - Goal: typical indoor rates of 100-200 Mbps; max 600 Mbps
- Use either 1 or 2 non-overlapping channels
  - Uses either 20 or 40 MHz
  - 40 MHz can create interoperability problems
- Supports frame aggregation to amortize overheads over multiple frames
  - Optimized version of 802.11e

802.11n Backwards Compatibility

- 802.11n can create interoperability problems for existing 802.11 devices (abg)
  - 802.11n does not sense their presence
  - Legacy devices end up deferring and dropping in rate
- Mixes Mode Format protection embeds an n frame in a g or a frame
  - Preamble is structured so legacy systems can decode header, but MIMO can achieve higher speed (training, cod/mod info)
  - Works only for 20 MHz 802.11n use
  - Only deals with interoperability with a and g – still need CTS protection for b
- For 40 MHz 802.11n, we need CTS protection on both the 20 MHz channels – similar to g vs. b
  - Can also use RTS/CTS (at legacy rates)
  - Amortize over multiple transmissions

Interoperability Uses PLC in Three Modes

- Legacy mode: use 802.11a/g OFDM format
  - The L-SIG field contains rate and length information

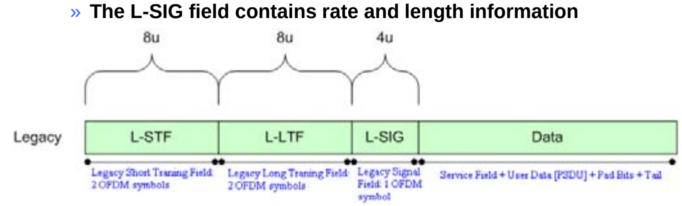

- Mixed mode:
  - Include both an 802.11a/g and 802.11n PLC
  - See next slide for figure
  - 802.11n devices can interpret green field, which includes the L-SIG field (rate and length information)

 http://rfmw.em.keysight.com/wireless/helpfiles/n7617a/mimo_ofdm_signal_structure.htm

Interoperability: High Throughput (HT) Modes

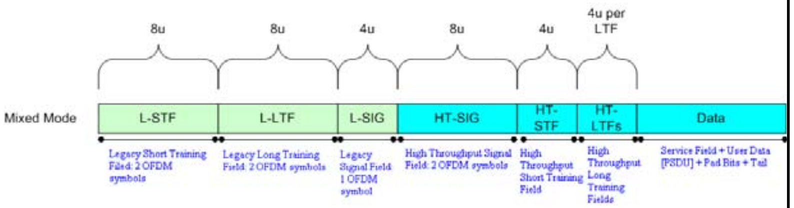

- Green field mode: use 802.11a/g OFDM format

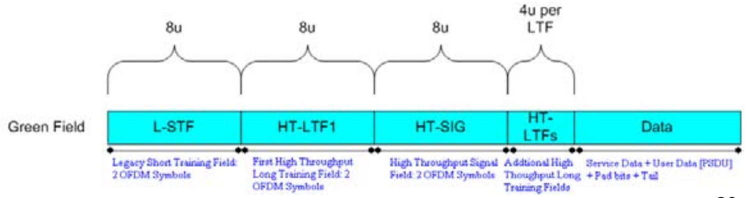

 http://rfmw.em.keysight.com/wireless/helpfiles/n7617a/mimo_ofdm_signal_structure.htm

MIMO in a Network Context

N transmit antennas

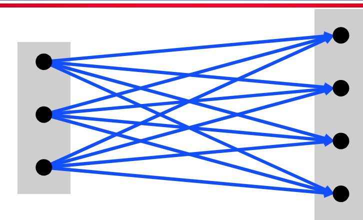

M receive Antennas - 1 receiver

How is this Different?

N transmit antennas

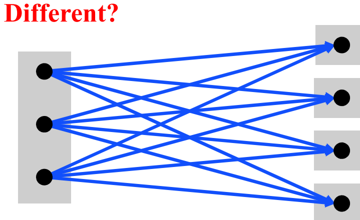

M receive antennas - M receivers

Multi-User MIMO Discussion

- Math is similar to MIMO, except for the receiver processing (P_R)
  - » Receivers do not have access to the signals received by antennas on other nodes
  - » Cannot cancel interference – limits ability to extract useful data
  - » Can only do transmit-side preprocssing
- MU-MIMO versus MIMO is really a tradeoff between TDMA and use of space diversity
  - » Sequential short packets versus parallel long packets

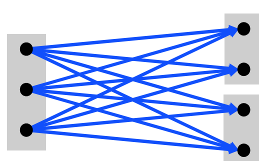

How about This?

N transmit antennas

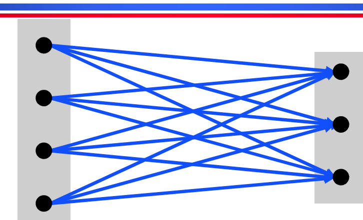

M receive Antennas - 1 receiver

How is this Different?

N transmit antennas

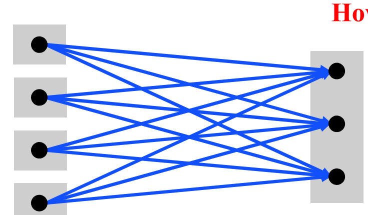

M receive antennas - M receivers

Multi-User MIMO
Up versus Down Link

- Assume an AP with multiple clients
- Uplink: Multiple Access Channel (MAC)
  - » Multiple clients transmit simultaneously to a single base station
  - » Requires fine grain clock coordination among clients on packet transmission – hard problem!
- Downlink: Broadcast Channel (BC)
  - » Base station transmit separate data streams to multiple independent users
  - » Easier to do: closer to traditional models of having each client receive a packet from the base station independently

802.11ac
Multi-user MIMO

- Extends beyond 802.11n
  - » MIMO: up to 8 x 8 channels (vs. 4 x 4)
  - » More bandwidth: up to 160 MHz by bonding up to 8 channels (vs. 40 MHz)
  - » More aggressive signal coding: up to 256 QAM (vs. 64 QAM); both use 5/6 coding rate (data vs. total bits)
  - » Uses RTS-CTS for clear channel assessment
  - » Multi-gigabit rates (depends on configuration)
- Support for multi-user MIMO on the downlink
  - » Can support different frames to multiple clients at the same time
  - » Especially useful for smaller devices, e.g., smartphones
  - » Besides beam forming to target signal to device, requires also nulling to limit interference

Challenges in 802.11ac

- You must have traffic for multiple receivers!
- Channels to the receivers be "orthogonal"

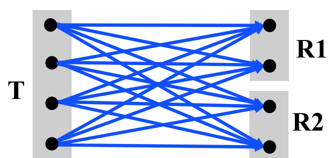

R1: $O_1 = P_{R1} * H_1 * P_T * I + P_{R1} * N$

R2: $O_2 = P_{R2} * H_2 * P_T * I + P_{R2} * N$

  - » The signal that you create with the packet for one destination should have a "null" for the other destination(s)
  - » Important since the other receivers cannot cancel out that signal
- Becomes a scheduling problem: for each "packet" transmission, identify the destinations that have traffic waiting and that are "the most" orthogonal

802.11ad
60 GHz WiFi

- Uses a new physical layer definition specifically for 60 GHz band
  - » Very different signal propagation properties
  - » Does not penetrate walls, but does work with reflections
  - » Shorter distances; up to 7 Gbps
  - » 6 channels of 2.16 GHz
- Compatible with 802.11 in 2.4 / 5 GHz bands
  - » Backwards compatible MAC
  - » E.g., mobile devices can switch between bands
- Has been used for point-point links for a while
  - » APs now available
  - » Combined with other 802.11 versions

Optimizing Communication in 802.11ad

- Transmission range in 60 GHz is limited
- Must use directional antennas to direct energy to the receiver
  - » Increases range and throughput (high signal strength)
  - » Also reduces interference at other nodes!
- Good news: antenna size scales with wave length
  - » Small antennas and narrow beams
- Bad news: how do nodes find each other?
  - » Use iterative algorithm, starting with wider beams

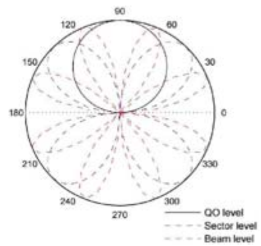

Outline

- WiFi deployments
  - » Planning
  - » Channel selection
  - » Rate adaptation

Infrastructure Deployments Frequency Reuse in Space

- Set of cooperating cells with a base stations must cover a large area
- Cells that reuse frequencies should be as distant as possible to minimize interference and maximize capacity
  - » Hidden and exposed terminals are also a concern

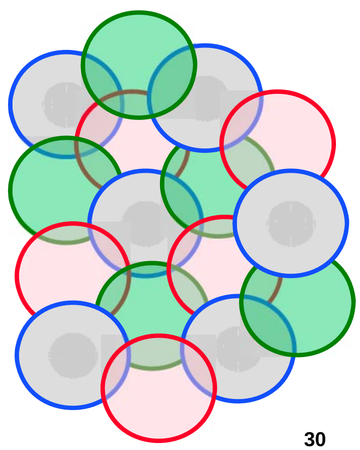

Frequencies are Precious

- 2.4 Ghz: 3 non-overlapping channels
  - » Plus lots of competition: microwaves and other devices
- 5 GHz: 20+ channels, but with constraints
  - » Power constraints, indoor/outdoor, ..
  - » Exact number and rules depend on the country
- 802.11n and ac: bonding of 2-8 channels
- And the world is not flat!

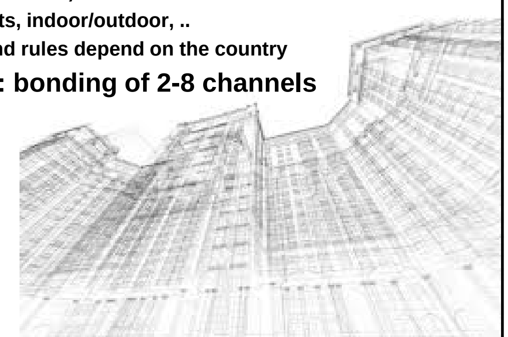

Frequency Planning

- Campus-style WiFi deployments are very carefully planned:
- A lot of measurements to determine where to place the AP
  - » What is the coverage area?
  - » What set of APs has good coverage with few "dead spots"
  - » What level of interference can we expect between cells
  - » What traffic loads can we expect, e.g., auditorium vs office
- Frequencies are very carefully assigned
  - » Can use the above measurements
- Must periodically re-evaluate infrastructure
  - » Furniture is moved, remodeling, …

Centralized Control

- Many WiFi deployments have centralized control
- APs report measurements
  - » Signal strengths, interference from other cells, load, …
- Controller makes adjustments
  - » Changes frequency bands
  - » Adjusts power
  - » Redistributes load
  - » Can switch APs on/off
  - » Very sophisticated!

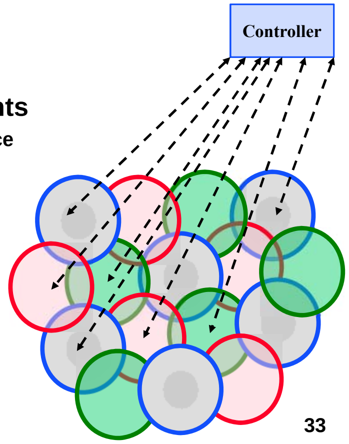

Monitoring the Spectrum

- FCC (in the US) controls spectrum use
  - » Rules for unlicensed spectrum, licenses for other spectrum, what technologies can be used
- … but there is an special clause for campuses
  - » They have significant control over unlicensed spectrum use on the campus
  - » They can even use some "licensed" spectrum if it does not interfere with the license holder
- Network management carefully monitors spectrum use to make sure it is used well
  - » Shut down rogue APs – interference, security
  - » Non-approved equipment - interference
  - » Discourages outdated standards - inefficient

How about Small Networks?

- Most WiFi networks are small and (largely) unmanaged
  - » Home networks, hotspots, …
- Traditional solution: user-chosen frequency of their AP or a factory set default
  - » How well does that work?
- Today, APs pick a channel automatically in a smart way
  - » Monitors how busy channels are or how strong the signals are and then picks the best channel
  - » Can periodically check for better channels

Outline

- WiFi deployments and channel selection
- Rate adaptation
  - » Background
  - » RRAA
  - » Charm

Bit Rate Adaptation

- All modern WiFi standards are multi bit rate
  - » 802.11b has 4 rates, more recent standards have 10s
  - » Vendors can have custom rates!
- Many factors influence packet delivery:
  - » Fast and slow fading: nature depends strongly on the environment, e.g., vehicular versus walking
  - » Interference versus WiFi contention: response to collisions is different
  - » Random packet losses: can confuse "smart" algorithms
  - » Hidden terminals: decreasing the rate increases the chance of collisions
- Transmit rate adaptation: how does the sender pick?

Transmit Rate Selection

- Goal: pick rate that provides best throughput
  - » E.g. SINR 14 dB → 5.5 Mbps
  - » Needs to be adaptive

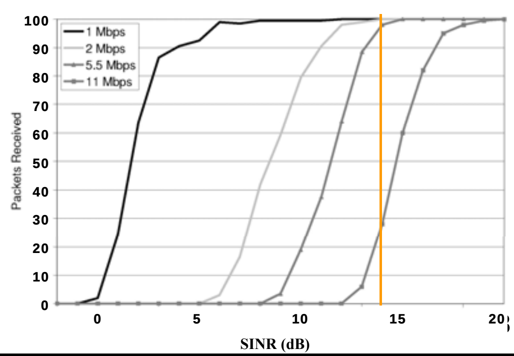

"Static" Channel

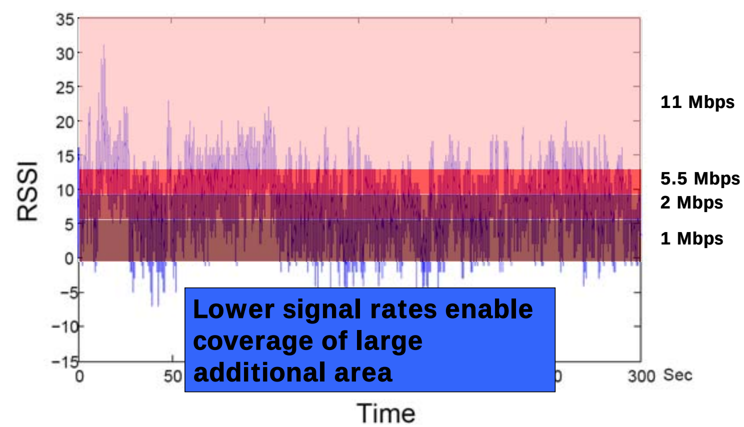

Lower signal rates enable coverage of large additional area

Mobile Channel – Pedestrian

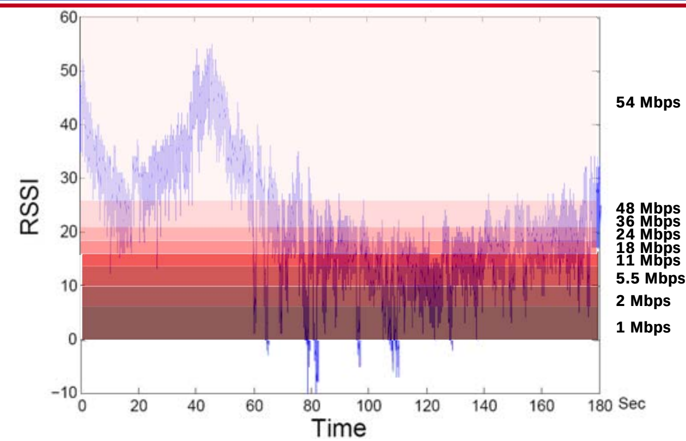

High Level Designs

- "Trial and Error": senders use past packet success or failures to adjust transmit rate
  - » Sequence of x successes: increase rate
  - » Sequence of y failures: reduce rate
  - » Hard to get x and y right
  - » Random losses can confuse the algorithm
- Signal strength: stations use channel state information to pick transmit rate
  - » Use path loss information to calculate "best" rate
  - » Assumes a relationship between PDR and SNR
    - – Need to recover if this fails, e.g., hidden terminals
- Newest class: context sensitive solutions
  - » Adjust algorithm depending on, e.g., degree of mobility, ..

Robust Rate Adaptation Algorithm

- RRAA goals
  - » Maintain a stable rate in the presence of random loss
  - » Responsive to drastic channel changes, e.g., caused by mobility or interference
- Adapt rate based on short term PDR

$$R_{new} = \begin{cases} R^{+} & P > P_{MTL} \\ R_{-} & P < P_{ORT} \end{cases}$$

  - » Thresholds and averaging windows depend on rate
- Selectively enable RTS-CTS

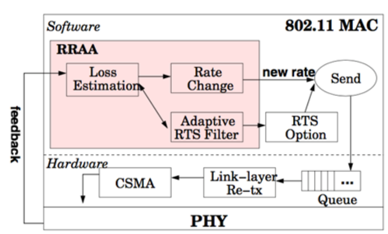

CHARM

- Channel-aware rate selection algorithm
- Transmitter passively determines SINR at receiver by leveraging channel reciprocity
  - » Determines SINR without the overhead of active probing (RTS/CTS)
- Select best transmission rate using rate table
  - » Table is updated (slowly) based on history
  - » Needed to accommodate diversity in hardware and special conditions, e.g., hidden terminals
- Jointly considers problem of transmit antenna selection

SINR: Noise and Interference

$$SINR = \frac{RSS}{Noise + \sum Interference}$$

- Noise
  - » Thermal background radiation
  - » Device inherent
    - – Dominated by low noise amplifier noise figure
  - » ~Constant
- Interference
  - » Mitigated by CSMA/CA
  - » Reported as “noise” by NIC

SINR: RSS

$$RSS = P_{tx} + G_{tx} - PL + G_{rx} \quad (1)$$

$$PL = P_{tx} + G_{tx} + G_{rx} - RSS \quad (2)$$

- By the reciprocity theorem, at a given instant of time
  - » $PL_{A \rightarrow B} = PL_{B \rightarrow A}$
- A overhears packets from B and records RSS (1)
- Node B records P_{tx} and card-reported noise level in beacons and probes, so A has access to them
- A can then calculate path-loss (2) and estimate RSS and SINR at B

CHARM: Channel-aware Rate Selection

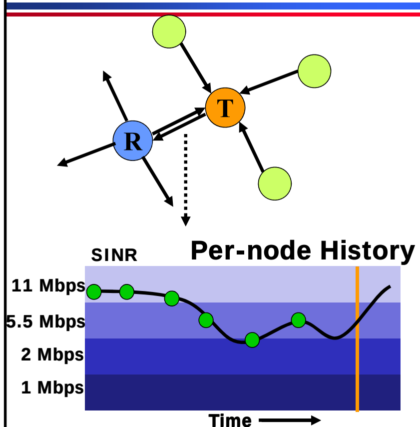

- Leverage reciprocity to obtain path loss
  - » Compute path loss for each host: P_{tx} - RSSI
- On transmit:
  - » Predict path loss based on history
  - » Select rate & antenna
  - » Update rate thresholds