UCL

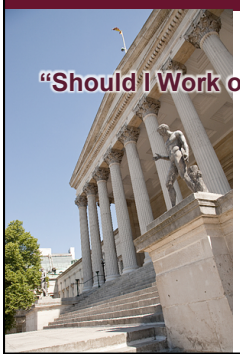

"Should I Work on Wireless Networks?"

ACM S³ Workshop
August 26, 2012
Istanbul, Turkey

Kyle Jamieson
Department of Computer Science
University College London

UCL

A brief disclaimer (and quasi-introduction)

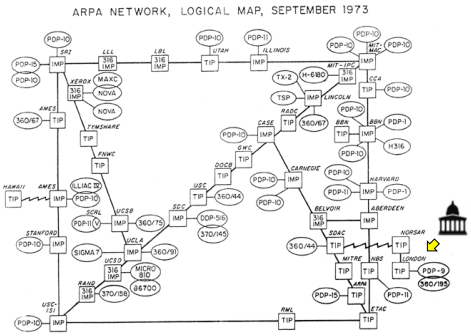

UCL

What this talk is <u>not</u>

- Not a concrete list of open research questions
  - See Craig Partridge's paper in ACM SIGCOMM Computer Communications Review
- Not advice aimed specifically at more experienced graduate students
- Mostly not indisputable facts
- Not "should <u>someone</u> work on wireless?"

UCL

UCL

Today

"As a student, should I work on wireless networks?"

It depends.

- **What makes wireless hard and exciting?**
- What constitutes a good research question in wireless networks?
  - **Heuristics** for choosing good research ideas
- What defines success in wireless networks?

UCL

Some things are well understood…

Q: *What's the capacity of a point-to-point link?*

- Before Shannon:
The only way to make P(error) arbitrarily small is to reduce the rate of communication.

- Shannon:
No! Up to some rate C, coding <u>can</u> make P(error) arbitrary small!

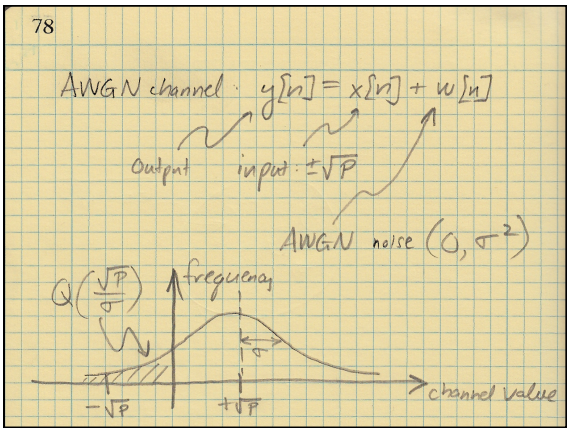

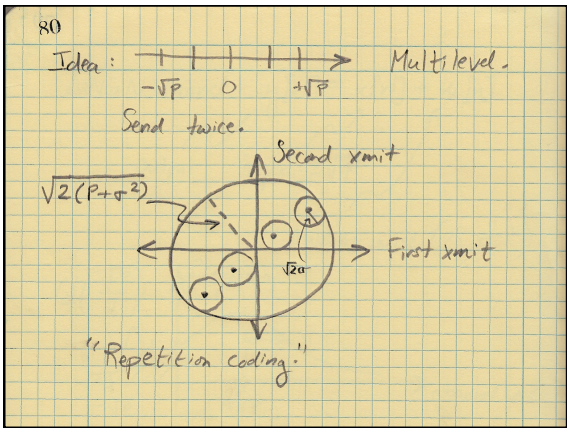

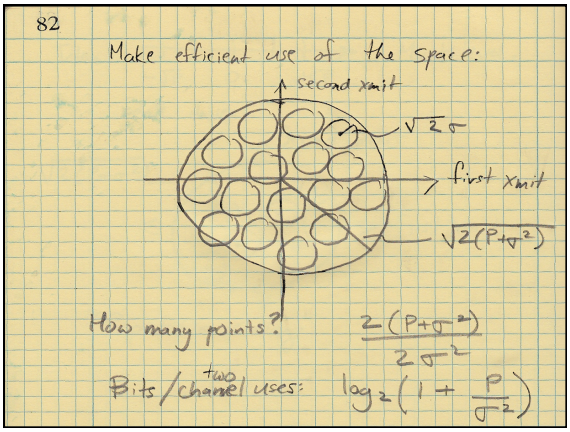

Some things are well understood…

- AWGN channel capacity:

$$C = \frac{1}{2}\log_2\left(1+\frac{P}{\sigma^2}\right) \text{ bits/channel use}$$

- Continuous-time (we didn't derive this one):

$$C = B\log_2\left(1+\text{SNR}\right) \text{ bits/second}$$

…others aren't understood well at all!

Q: What's the capacity of a wireless ***network****?*

A: (information theory) …
A: (CS community) let's build a better medium access protocol!

ALOHA packet radio analysis takes a single-cell perspective

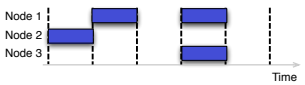

Pr(success)

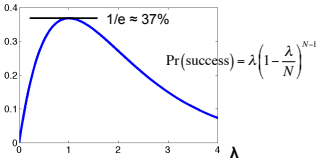

Contrast: Ethernet

- Bob Metcalfe, PhD student at Harvard in early 1970s
  - Working on protocols for the ARPAnet
- Xerox Palo Alto Research Center (PARC), 1973
  - Network the Alto Workstation [Lampson *et al.*]
  - ALOHA packet radio

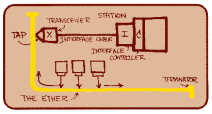

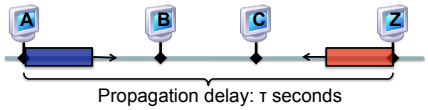

Wireless network system structure

- Optimization across layers
- Optimization across the network, at each layer

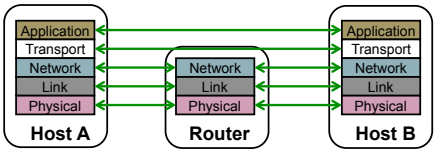

The growing wireless toolbox

Link Layer and above
- MAC protocols
- Routing and handoff

Physical Layer
- Successive Interference Cancellation
- Transmit beamforming
- Receive beamforming
- MIMO spatial multiplexing
- Space-time coding
- Amplify-and-forward

Carrier sense doesn't avoid collisions

Measure goodput during TCP file transfer [Sheth '06]

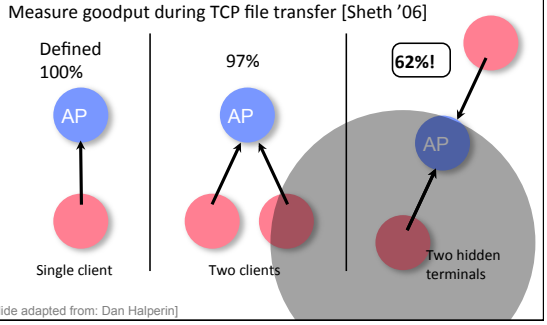

[Slide adapted from: Dan Halperin]

Carrier sense prevents spatial reuse

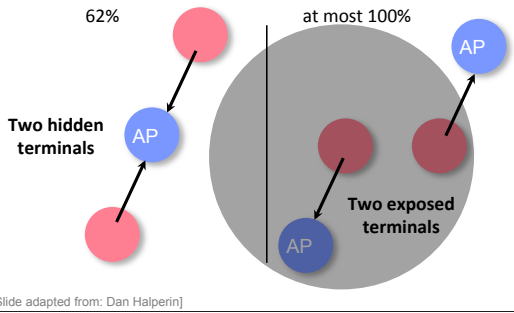

[Slide adapted from: Dan Halperin]

The SINR model

$$\text{SINR} = \frac{\text{Signal}}{\text{Noise} + \sum \text{Interference}}$$

- When relative power of desired signal is large enough, signal received
- Line shows threshold between **reliable** and **lossy** links (> 15% FER)
- Generally receivers require at least 3 dB SINR

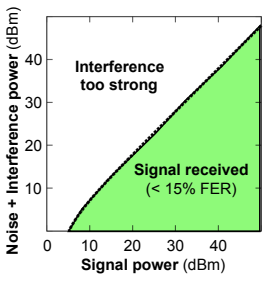

[Slide adapted from: Dan Halperin]

How successive IC (SIC) works

$$\text{Received} = \text{Noise} + \sum \text{Distorted Signals}$$

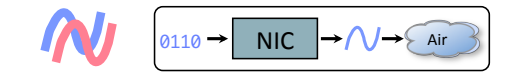

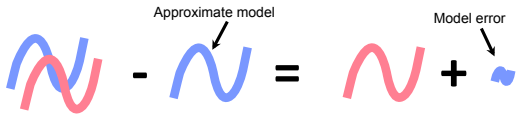

[Slide adapted from: Dan Halperin]

Strong BPSK with weak QPSK

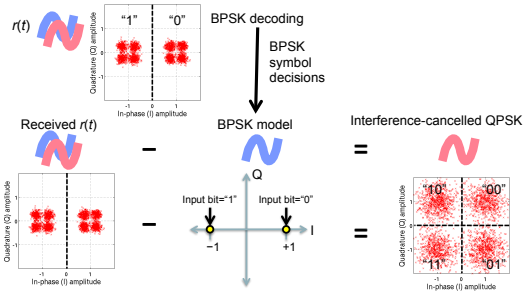

SIC improves reception region

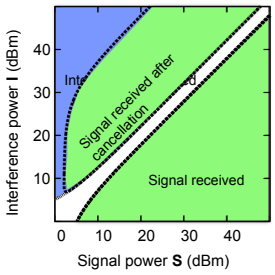

[Slide adapted from: Dan Halperin]

Today

"As a student, should I work on wireless networks?"

It depends.

- ✓ What makes wireless hard and exciting?
- **What constitutes a good research question in wireless networks?**
- What defines success in wireless networks?

What makes a good wireless research problem?

[Partly adopted from Partridge]

1. Worth the attention of multiple people
   - *e.g.*, spatial multiplexing
2. Opens up substantial follow-on efforts
   - Industry, research, or both
3. Likely will have impact
   - Impact may come long after the idea

Choosing a good wireless research problem

How can I choose a good wireless research problem?

Here are some heuristics.

UCL

Heuristic #1: Solve it better

Given a solved problem, ask yourself:

Can I take a fresh look at solving this problem?

Is there anything about the way the researchers solve this problem that I perhaps think can be improved, or even outright disagree with?

UCL

Embrace collisions in RF backscatter networks

- "Efficient and reliable low-power backscatter networks," Wang *et al.*, SIGCOMM 2012

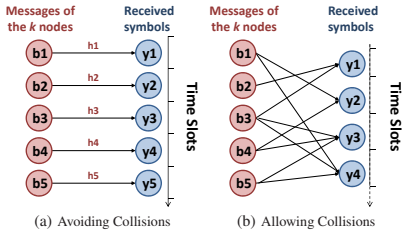

UCL

Indoor location systems today

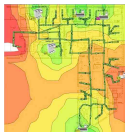

- Map-based approaches: RADAR [Bahl+00]
  - Require calibration to build signal strength *map*
  - Augmented with probabilistic models: Horus [Youssef+05] (60 cm accuracy)
  - Crowdsourcing: EZ [Chintalapudi+10], Zee [Rai+12]
- Radio modeling-based approaches
  - RF propagation model relates signal strength and distance
  - No calibration needed, but accuracy suffers
  - 3 m accuracy [Lim+06], 5.4 m accuracy [Gwon+04]
- Vision-based approaches
  - **Highly accurate** (≈ 20 cm) [Hile+08] but computationally intensive
  - Light conditions aren't always ideal, humans are humans

UCL

Two observations about WiFi access points

1. Ever-increasing number of antennas for MIMO, SDMA

2. High WiFi access point density: usually many nearby

- UCL Department of CS
- Client at a random location sends a packet
- How many APs overhear it?

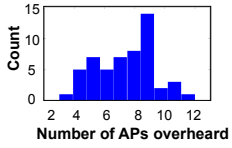

UCL

ArrayTrack: Use AoA for indoor localization

- Client sends a single packet over the air
- Each access point (AP) computes the **physical angles-of-arrival** of a client's transmission: a *pseudospectrum*
- Aggregate pseudospectra at backend server for location

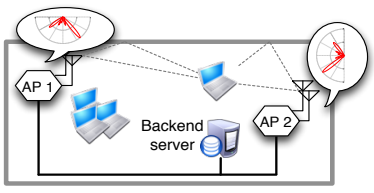

UCL

Phased antenna array: Principle of operation

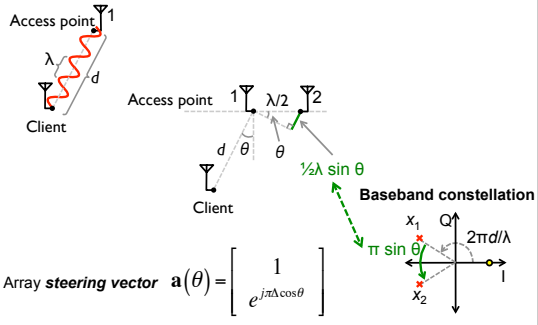

Array ***steering vector*** $\mathbf{a}(\theta) = \begin{bmatrix} 1 \\ e^{j\pi\Delta\cos\theta} \end{bmatrix}$

MUSIC: Geometric interpretation (Three antennas, two signals)

- *MUSIC* algorithm [Schmidt '79] and variants [Shan '85] analyze the *eigenstructure* of $\mathbf{R}_{xx}$

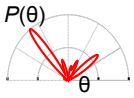

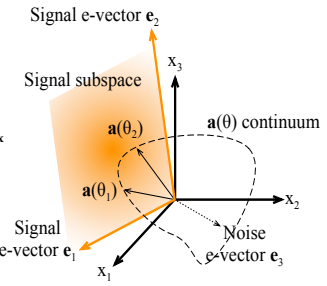

[Adapted from Schmidt, "Multiple Emitter Location and Signal Parameter Estimation"]

Challenge: Multipath propagation

- Multipath is a challenge in **two** distinct ways:

1. MUSIC takes time-averages, gets confused when signals are *coherent*

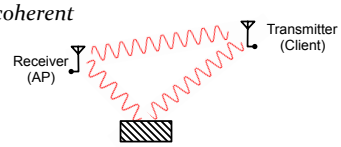

Solution (1): Spatial smoothing

- Well-known technique [Shan+85] to average across spatially diverse groups of antennas:

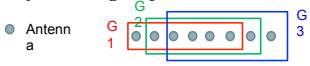

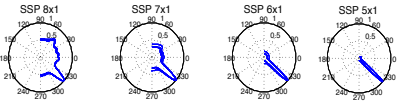

- Tradeoff: Fewer effective number of antennas

Current work: A 16-antenna AP prototype

- Three independent reasons to increase the number of antennas at an AP in the future:
  1. Antenna diversity
  2. MIMO on a single link
  3. SDMA to multiple clients
- Leverage more antennas and more spatial smoothing

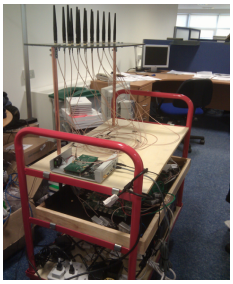

Challenge: Multipath propagation

- Multipath is a challenge in **two** distinct ways:

1. MUSIC takes time-averages, gets confused when signals are coherent

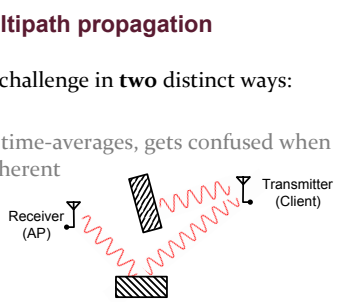

2. Obstacles may block the direct line of sight

Multipath fools naïve approach 18% of the time

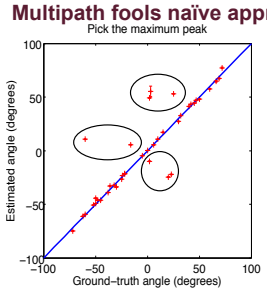

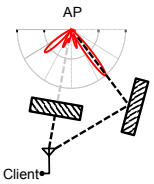

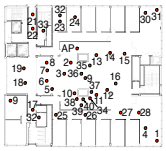

UCL

Combining bearing likelihoods for a location likelihood at the backend server

- **Given** N AP bearing likelihoods $P_1(\theta_1(\mathbf{x}))$, $P_2(\theta_2(\mathbf{x}))$, ..., $P_N(\theta_N(\mathbf{x}))$

1. **Compute** location likelihood $P(\mathbf{x})$ for a given location $\mathbf{x}$:

$$P(\mathbf{x}) = \prod_{l=1}^{N} P_l(\theta_l(\mathbf{x}))$$

$$= \sum_{l=1}^{N} \log P_l(\theta_l(\mathbf{x}))$$

2. **Search** for most likely location with sampling and hill climbing

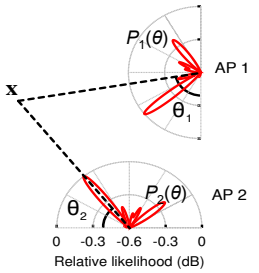

UCL

Sampling and gradient search

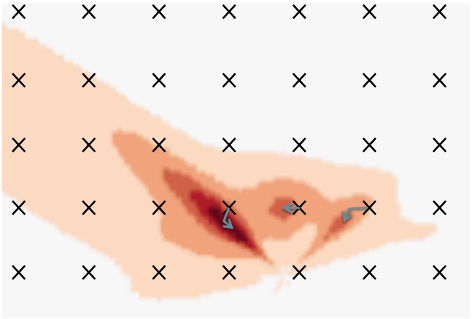

UCL

Additional APs usually improve location estimates

Location heat map cycling through three to six APs

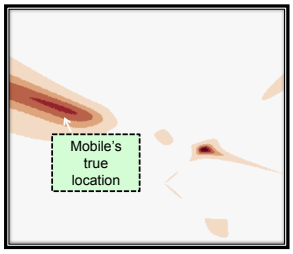

UCL

Pop Quiz: Why could adding an AP make things worse, not better?

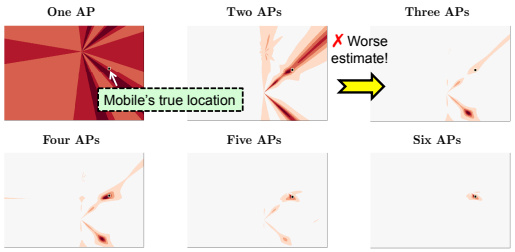

UCL

ArrayTrack achieves high localization accuracy

[ACM HotMobile 2012]

- 33 clients on one floor of an office space in active use
- Ground truth using architectural drawings and laser measure
- Compare with optimal AP subset of *any* size: ArrayTrack is within 2×
  - Reasons to expect even further improvement

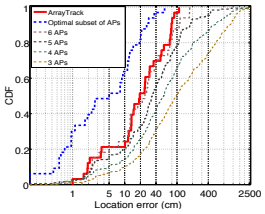

UCL

Heuristic #2: Hack on new hardware

Ask yourself:

What new research questions does new hardware X allow me to investigate?

UCL

Software-defined radio opens up the network to innovation

More on-board processing →

← Faster bus to PC


Ettus Research USRP1


Ettus Research USRP2


Rice WARP v2


Rice WARP v3


MSR Asia Sora

UCL

SoftPHY: Change perspective from packets to symbols

- **Cross-layer** information flow from PHY up
- **Extract use** from high-confidence parts of packets

Network layer
Link layer
SoftPHY interface
Physical layer

- Maintain layered architecture (PHY-independent use)

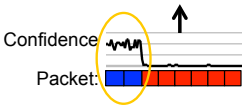

UCL

Partial Packet-ARQ

"Partial Packet Recovery," Jamieson and Balakrishnan, SIGCOMM '07

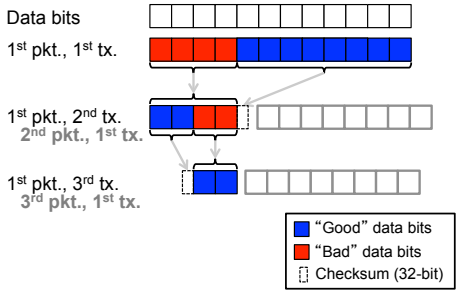

UCL

Fidelity Comtech *Phocus* array access point

- *ca.* 2008
- Commodity Atheros 802.11a/g card
- Eight-element phased-array RF front end

- "DIRC: Increasing Indoor Wireless Capacity Using Directional Antennas," Liu *et al.*, SIGCOMM 2009.

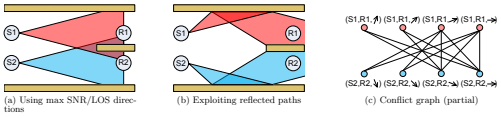

UCL

Wireless Sensor Hints: The Opportunity

- "Improving Wireless Network Performance Using Sensor Hints," Ravindranath *et al.*, NSDI 2011.

- Modern mobile devices have many sensors
- Used by applications
- Ignored by protocols today

Wireless protocols can use hints from sensors to significantly improve performance

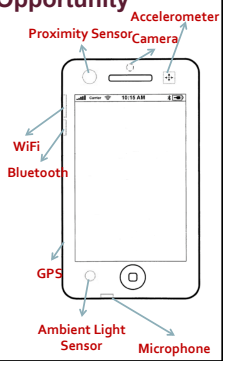

[Slide adapted from: Hari Balakrishnan]

UCL

Sensor Movement Hints

- Has there been movement?
- Heading (direction)
- Speed
- Position

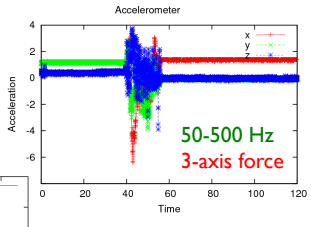

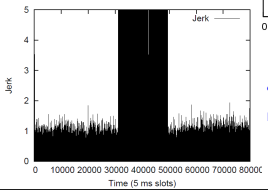

"Jerk" metric detects movement reliably within 10 ms

[Slide adapted from: Hari Balakrishnan]

Architecture

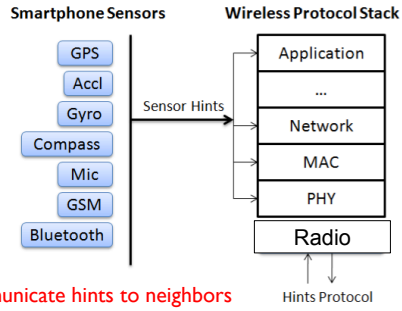

Communicate hints to neighbors

Adapt to neighbor mobility, not just node's own movement

[Slide: Hari Balakrishnan]

Hint-Aware Protocols

- Bit Rate Adaptation
- Topology Maintenance
- Access Point Policies
  - Association
  - Packet scheduling
  - Pruning
- Vehicular network route selection
- And more...

[Slide: Hari Balakrishnan]

Heuristic #3: Take the next step

Given a solved problem, ask yourself:

What's the next step in realizing the solution, and are there any interesting challenges in doing so?

Maranello: Practical Partial Packet Recovery

"Maranello: Practical Partial Packet Recovery for 802.11," Han *et al.*, NSDI '11

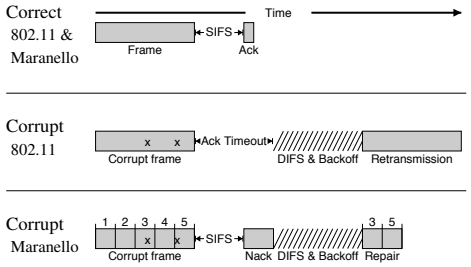

Heuristic #4: Solve a new problem

Ask yourself:

Is there a part of the current paradigm that can be discarded and replaced with something better?

1. Prescience
2. Normal science ("puzzle-solving")
3. Anomalous results lead to "crisis" and "revolution"

Current approach: Modulation adaptation

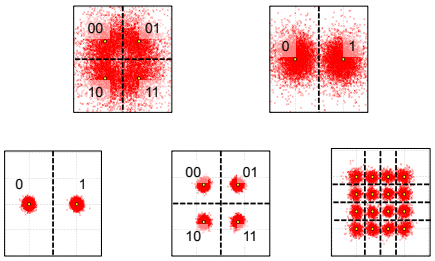

Metric: Signal-to-noise power ratio (SNR)

Current approach: Error control coding adaptation

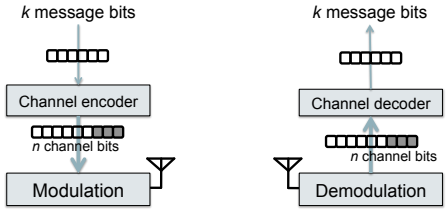

Current approaches adapt a **fixed code rate:** $R = \frac{k}{n}$

Lower code rate ☞ more parity bits ☞ more redundancy

Current approach: "Adapt and discard"

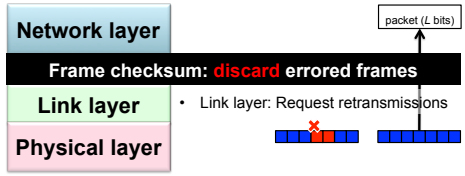

- Physical layer: **Adapt** for a low (~10^{-5}) BER
  - **Adapt** modulation (feedback loop)
  - **Adapt** fixed coding rate (feedback loop)

The wireless channel is dynamic and unpredictable

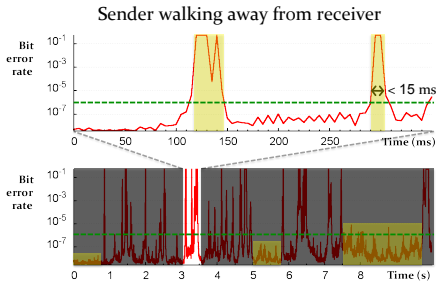

Rateless insight: Avoid committing to a rate *a priori*

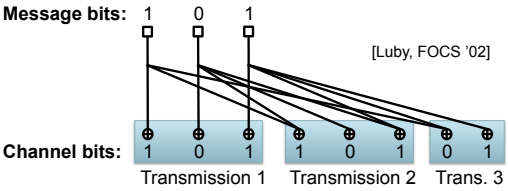

Rate after transmission 1 is sent: 3/3 = **1**
Rate after transmission 2 is sent: 3/6 = **0.5**
Rate after transmission 3 is sent: 3/8 = **0.375**

Receiver: Accumulate, don't discard bits

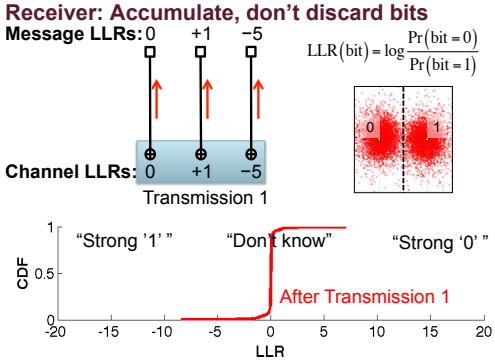

Receiver: Accumulate, don't discard bits

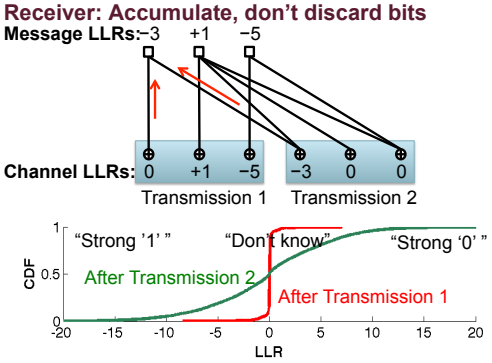

Today

"As a student, should I work on wireless networks?"

It depends.

- ✓ What makes wireless hard and exciting?
- ✓ What constitutes a good research question in wireless networks?
- **What defines success in wireless networks?**

Capacity, capacity, capacity

- But careful, it can be measured in many different ways
  - Network or point-to-point?
  - At many different layers: PHY, link layer, routing
  - Trace-driven or live?
  - Line rate? If not, interaction with the wireless channel?

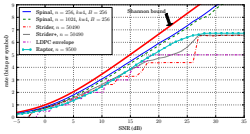

"Spinal Codes," Perry *et al.*, SIGCOMM 2012

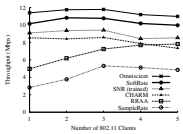

"Cross-Layer Wireless Bit Rate Adaptation," Vutukuru *et al.*, SIGCOMM 2009

Localization

- Interesting problem, not "one size fits all"
- RADAR, infrastructure-dependent approaches
- Moving target: Shift to crowdsourcing and infrastructure-free approaches
- Basic metrics
  - Localization accuracy
  - Latency

In summary

"As a student, should I work on wireless networks?"

It depends.

- ✓ What makes wireless hard and exciting?
- ✓ What constitutes a good research question in wireless networks?
- ✓ What defines success in wireless networks?

The take-away

Great, you've decided to go into wireless networks!

TODO:
1. Solve it better
2. Hack on new hardware
3. Take the next step
4. Solve a new problem