Wireless in the Real World

Principles

- Make every transmission count
  - E.g., reduce the # of collisions
  - E.g., drop packets early, not late
- Control errors
  - Fundamental problem in wless
- Maximize spatial reuse
  - Allow concurrent sends in different places
  - While not goofing up #1 and #2!

Problems

- Today: Deployments are *chaotic*
  - Unplanned: Lots of people deploy APs
    - More planned inside a campus, enterprise, etc.
    - Less planned at Starbucks…
  - Unmanaged
    - Many deployments are "plug-and-go"
      - Becoming increasingly common as 802.11 becomes popular. Not just geeks!
- And it's hard in general. ☺

Making Transmissions Count

- See previous lecture!

Error Control

- Three techniques
  - ARQ (just like in wired networks)
  - FEC (also just like, but used more in wireless)
  - And .. Rate control.
- Remember our Shannon's law discussion
  - Reminder: Capacity = B x log(1 + S/N)
  - Higher bitrates use encodings that are more sensitive to noise
  - If too many errors, can fall back to a lower rate encoding that's more robust to noise.
  - Often called "rate adaptation"

Rate Adaptation

- General idea:
  - Observe channel conditions like SNR (signal-to-noise ratio), bit errors, packet errors
  - Pick a transmission rate that will get best goodput
    - There are channel conditions when reducing the bitrate can greatly increase throughput – e.g., if a ½ decrease in bitrate gets you from 90% loss to 10% loss.

Simple rate adaptation scheme

- Watch packet error rate over window (K packets or T seconds)
- If loss rate > $thresh_{high}$ (or SNR <, etc)
  - Reduce Tx rate
- If loss rate < $thresh_{low}$
  - Increase Tx rate
- Most devices support a discrete set of rates
  - 802.11 – 1, 2, 5.5, 11, etc.

Challenges in rate adaptation

- Channel conditions change over time
  - Loss rates must be measured over a window
- SNR estimates from the hardware are coarse, and don’t always predict loss rate
- May be some overhead (time, transient interruptions, etc.) to changing rates

Error control

- Most fast modulations already include some form of FEC
  - Part of the difference between the rates is how much FEC is used.
- 802.11, etc. also include link-layer retransmissions
  - Relate to end-to-end argument?
  - Compare timescale involved
  - Needed to make 802.11 link layer work within the general requirements of IP ("reasonably low" loss)

Spatial Reuse

- Three knobs we can tune:
  - Scheduling: Who talks when (spatial div)
    - A – B – C – D – E -- F ..
      - A->B, C->D, E-F
      - B->C, D->E
  - Frequency assignment (frequency div)
    - 802.11 has 11 "channels" in the US, but they're not completely independent
      - (draw frequency overlap)
  - Power assignment
    - Many radios can Tx at multiple power levels

Cellular Reuse

- Transmissions decay over distance
  - Spectrum can be reused in different areas
  - Different “LANs”
  - Decay is $1/R^2$ in free space, $1/R^4$ in some situations

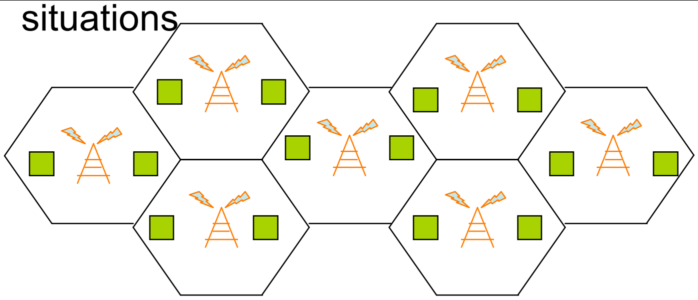

Frequency Allocation

- To have dense coverage
  Must have some overlap
- But this will interfere.
- (Even w/out interference
  if you want 100% coverage)

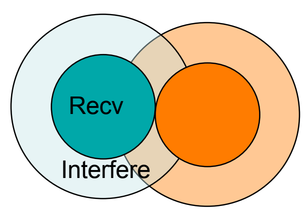

- Answer: Channel allocation for nearby nodes
- Easy way: Cellular deployment. Offline, centralized graph coloring
- Hard way: Ad hoc, distributed, untrusting, …

Ad hoc deployment

- Typically multiple hops between nodes
- Unplanned or semi-planned
- Typical applications:
  - Roofnet
  - Disaster recovery
  - Military
- Even though most wireless deployments are “cellular” systems, they exhibit many of the same challenges of ad hoc…

Power Control

- (diagram)
- Goal: Transmit at minimum necessary power to reach receiver
  - Minimizes interference with other nodes
  - Paper: Can double or more capacity, *if done right.*

Detils of Power Control

- Hard to do per-packet with many NICs
  - Some even might have to re-init (many ms)
- May have to balance power with rate
  - Reasonable goal: lowest power for max rate
  - But finding ths empirically is hard! Many {power, rate} combinations, and not always easy to predict how each will perform
  - Alternate goal: lowest power for max *needed* rate
    - But this interacts with other people because you use more channel time to send the same data. Uh-oh.
    - Nice example of the difficulty of local vs. global optimization

Power control summary

- More power:
  - Higher received signal strength
  - May enable faster rate (more S in S/N)
    - May mean you occupy media for less time
  - Interferes with more people
- Less power
  - Interfere with fewer people
- Less power + less rate
  - Fewer people but for a longer time

Scaling Ad Hoc Networks

- Aggregate impact of far-away nodes
  - Each transmitter raises the “noise” level slightly, even if not enough on its own to degrade the signal enough (S/N…)
- The price of cooperation: In a multi-hop ad hoc network, how much time do you spend forwarding *others* traffic?
- Routing protocol scalability
  - (Next lecture! :-)

Aggregate Noise

- Assume that you can treat concurrent transmissions as noise
  - Example: CDMA spread-spectrum networks do exactly this
- Nodes in a 2d space with constant density p
- Nodes talk to nearest node (multi-hop for far away)
- (This model applies to cooperation, too)
- (diagram)

contd

- Distance to neighbor ~ R_0 = 1/sqrt(p)
- Power level P, attenuation at distance r propto r^{-2} (free space), so signal strength propto r^2
- Total nodes in annulus @ distance *r,* width *dr* from recv:

  2 π r p dr
- Total interference: $\int_0^\infty \frac{2\pi r p dr}{r^2}$

Noise

- Aggregate noise is infinite!
- But the world isn't. Phew. If M nodes total, Rmax node distance is pi R^2 maxp = M
- Solving,integrate from 0 Rmax total signal-to-noise falls off as 1/log M
- Not too bad…

The Price of Cooperation

- In ad hoc, how much of each nodes' capacity is used *for others*?
  - Answer depends strongly on workload.
  - If random senders with random receivers:
    - Path from sender → receiver is length $\sqrt{N}$
      - So every transmission consumes $\frac{\sqrt{N}}{N} = \frac{1}{\sqrt{N}}$ of the network capacity
      - Network has a total capacity of N transmits/time
    - Aggregate network capacity of N nodes scales as sqrt(N)
    - Per-node capacity is $\frac{1}{\sqrt{N}}$

Locality

- Previous model assumed random-random communication
- Locality can help you
  - E.g., geographically dispersed "sinks" to the Internet: Roofnet-style communication
  - E.g., local computation and summary: sensor-network communication
    - Example: Computing the avg, max, min temp
  - "Data" or "content"-centric networking (caching, etc.)

Aside: Flipping Power On Its Head: Power Savings

- Which uses less power?
  - Direct sensor -> base station Tx
    - Total Tx power: distance^2
  - Sensor -> sensor -> sensor -> base station?
    - Total Tx power: n * (distance/n) ^2 =~ d^2 / n
  - Why? Radios are omnidirectional, but only one direction matters. Multi-hop approximates directionality.
- Power savings often makes up for multi-hop capacity
  - These devices are *very* power constrained!
- Reality: Many systems don't use adaptive power control. This is active research, and fun stuff.

Summary

- Make every transmission count
  - MAC protocols from last time, mostly
- Control errors
  - ARQ, FEC, and rate adaptation
- Maximize spatial reuse
  - Scheduling (often via MAC), channel assignment, power adaptation
- Scaling through communication locality
  - e.g., sensor net-style communication