In Defense of Wireless Carrier Sense

Micah Brodsky

Wireless medium is *semi*-shared

- Sometimes networks are largely independent
  - Can transmit concurrently: “spatial reuse” of medium

- Sometimes they are in conflict
  - Throughput will be nearly zero under concurrent transmission; should time-multiplex

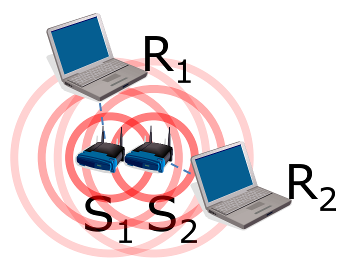

- Someone must make the decision. How?

Solution: Carrier sense?

- Mechanism: Interferer power vs. threshold
  - Defer transmissions when competing packets above threshold
  - Transmit freely when below
  - Used by MACs to answer "Can I talk now?",
- Strikes balance between interference protection and spatial reuse
  - Attempts to use spectrum efficiently while preserving fairness
- Simple – and simple is good!

Reasons to be suspicious…

- Wrong measurement!
  - Power at *receivers* is what matters [Karn ’90]
- Classic example: “hidden terminal”

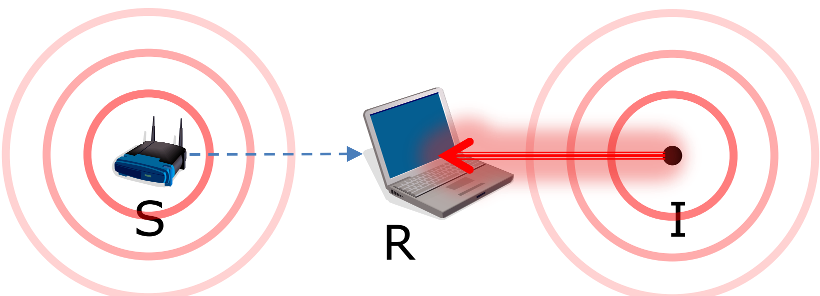

- How can this make sense?

Life's not so simple, either

Desired result: concurrency

Desired result: time-multiplexing

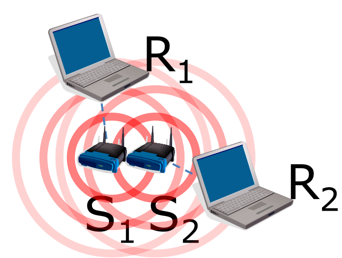

Desired result: ???

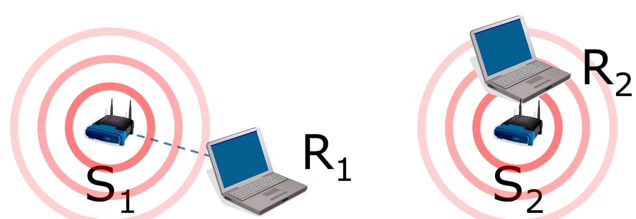

Our question: How well *does* CS work?

- Are collisions and horrible failures the right way to think about carrier sense?
- How common are mistakes? (sub-optimal decisions)
- How much do they cost in throughput?
- How does carrier sense compare to “optimal”?
  - Key metric: Mean expected throughput
  - Also, starvation and similar misbehavior?
- (Also, might things have changed since earlier work?)

Why CS might work: Limiting cases

- "Far" interference:
  - Small distance variation: $\Delta r_1 \approx \Delta r_2$

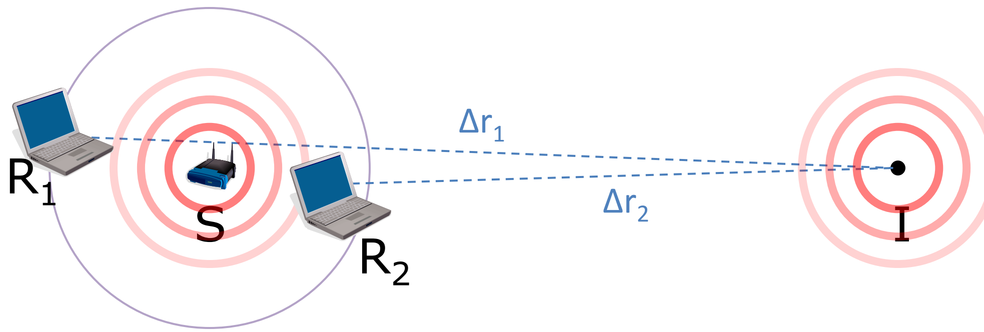

- "Near" interference:
  - Nobody wants concurrency; $SINR_{concurrent}$ <<< $SNR_{multiplexing}$

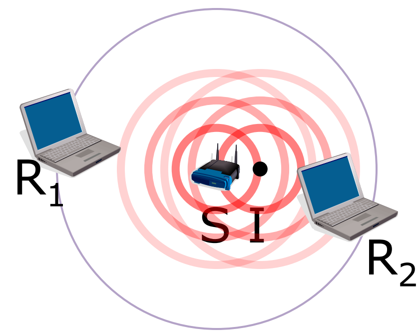

- In both cases, all receivers agree on preferring either multiplexing or concurrency
  - "Agreement" means CS can perform well
- Intermediate distance will be the hard case
- Also, shadows and obstacles?

Let's explore with a simple model

- Simplifications & limitations
  - Only two contending transmitters
  - Transmitters have same power, omni antennas
  - Focus on fundamentals, rather than on a particular implementation
    - No framing, ACKs, slotting, etc.
    - Not modeling capture effects
- Building blocks: Network layout + radio propagation + estimated throughput
- Output: Predictions for average throughput under concurrency, multiplexing, carrier sense, and optimal

Model: layout and averaging

- Place senders at fixed locations
- Assume receivers uniformly distributed within some R_{max}
- Compute mean throughput over both sets of receivers (S1's & S2's)
- Will investigate effect of varying sender-sender distance D, given an R_{max}

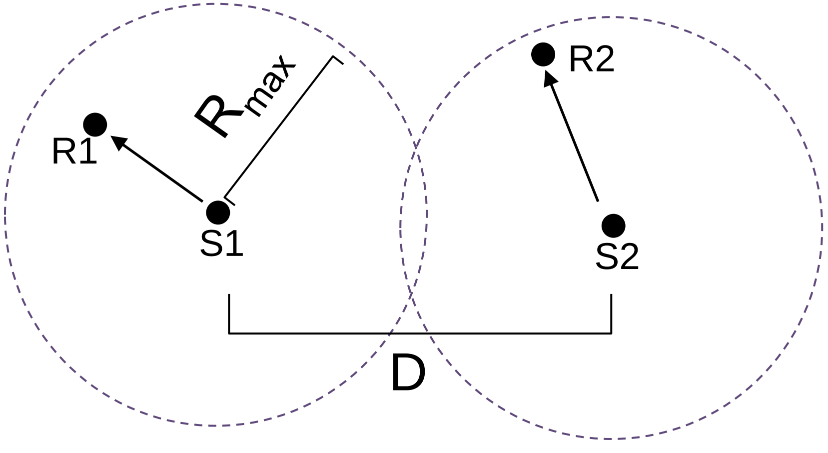

Model: radio propagation

Standard textbook model (e.g. Akaiwa '97):

- Path loss: $r^{-\alpha}$
- Environmental shadowing: $\pm\sigma$ dB
- Multipath fading: Rayleigh variation
  - Wideband channels average this away (mostly)

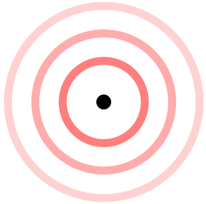

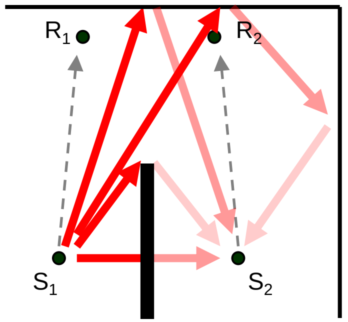

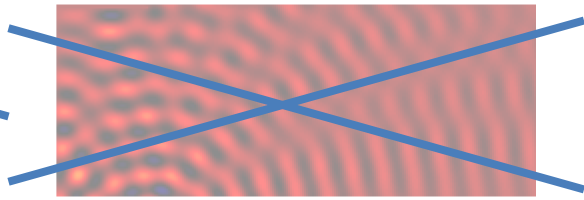

Model: throughput

- Need a way to model throughput as a function of SINR (Signal to Interference + Noise Ratio)
- Adaptive bitrate (ABR) is pervasive nowadays
  - And will turn out to be crucial
- Shannon capacity is a half-decent approximation model for ABR (with nice analytical properties)
  - *Capacity / Bandwidth(Hz) ≈ log(1 + SINR)*

What we’re going to look at

- First, for individual receiver configurations, which choice gives better throughput, concurrency or multiplexing?
- Next, average throughput across the ensemble of different possible receiver configurations
  - Compare CS to concurrency, multiplexing, optimal
- Finally, vary R_{max} (network size) to show that good efficiency holds across the space of possibilities

A first look: individual receivers

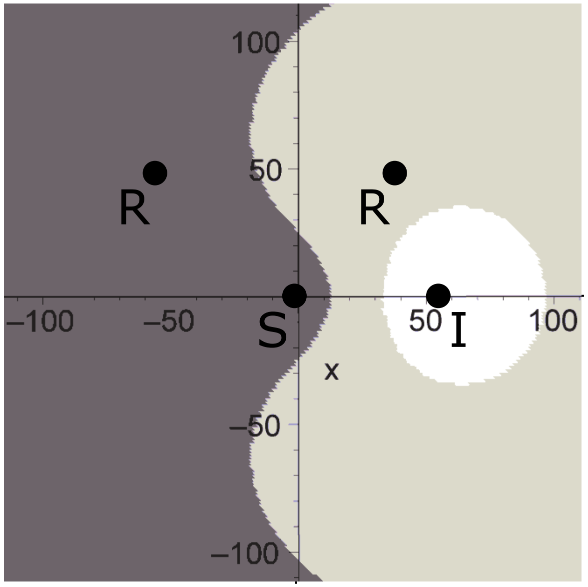

D = 55

- ■ Prefers concurrency
- □ Prefers multiplexing
- □ Starved w/o multiplexing

In detail…

Receiver preference vs. position:

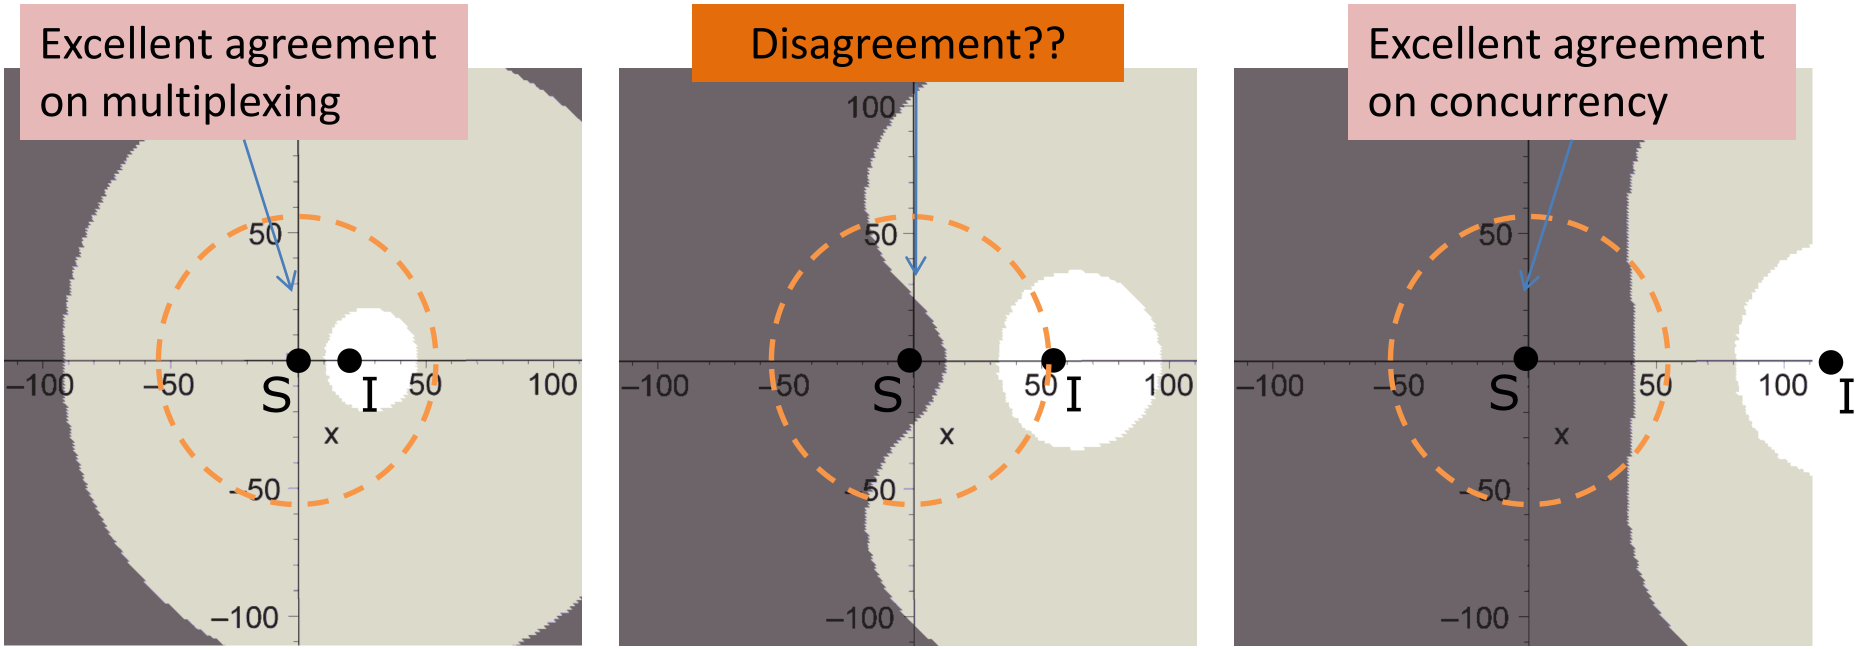

D = 20   D = 55   D = 120

- Prefers concurrency
- Prefers multiplexing
- Starved w/o multiplexing

ABR prevents disaster!

- Intermediate distance can mean poor agreement! But…
- Does "mistaken" concurrency mean near-zero throughput? No. Adapts with lower bitrate.
- Does "mistaken" multiplexing mean 50%-reduced throughput? No. Adapts with higher bitrate.
- "Exposed" and "hidden" terminals are not very useful concepts with ABR

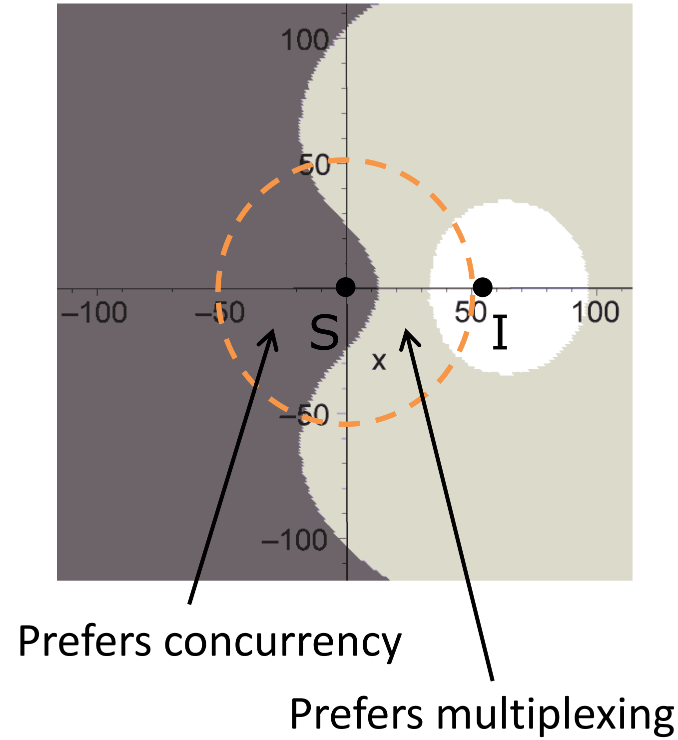

Obstacles aren't fatal

- Most obstacles are not opaque!
- Most configurations have alternate propagation paths
- ±4dB - 12dB variation from path loss is typical
  - (See e.g. COST 231 and other model reviews)
- If shadowing were *much* greater, CS would be no better than random. But it's not.
- (ABR also helps here)

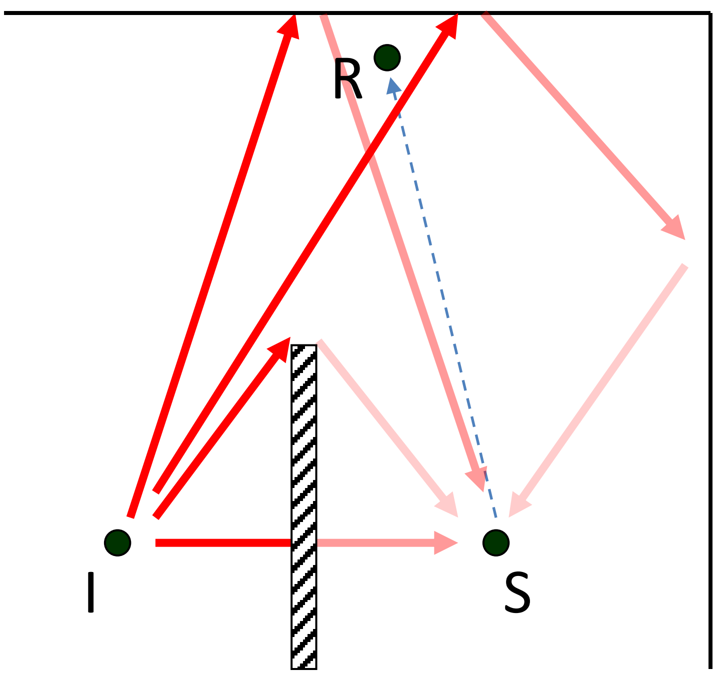

Average throughput: CS works!

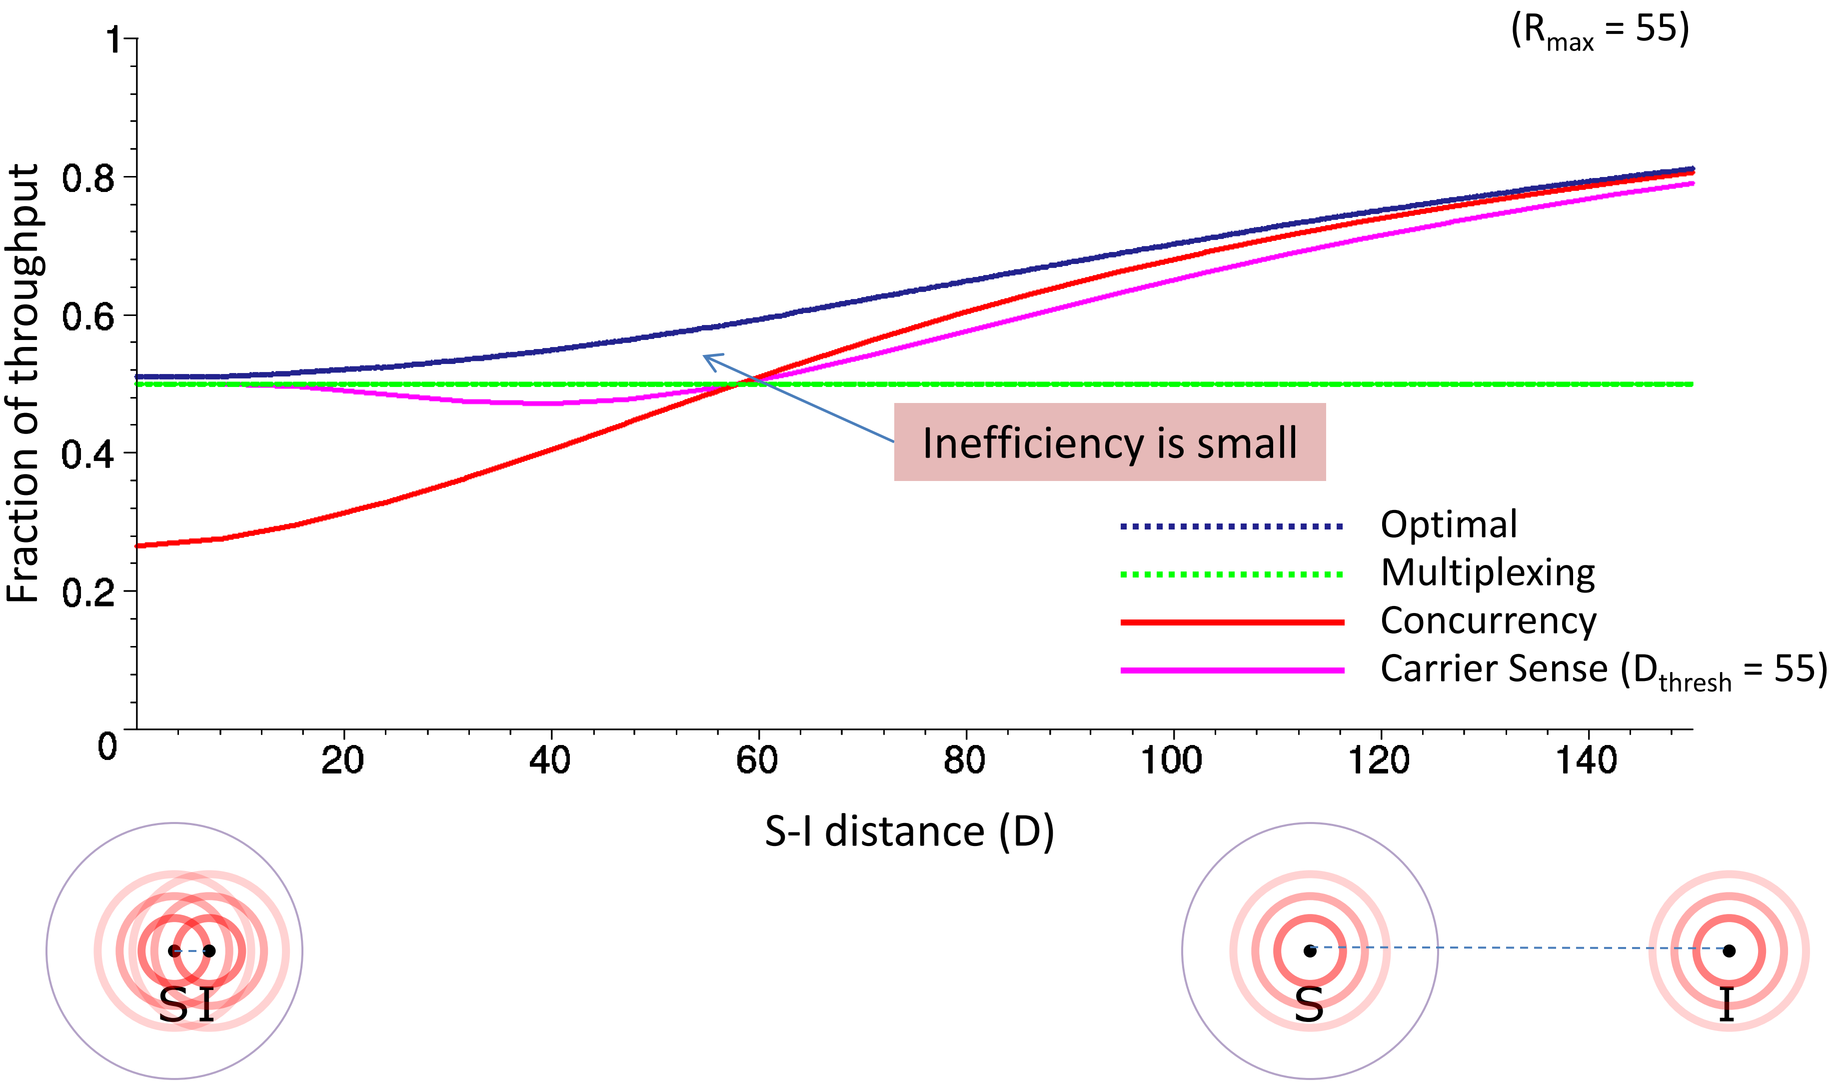

The larger parameter space

- Of course, one example isn't enough
- Need to explore full relevant span of parameters
  - Fortunately, interferer distance and network size capture most of the important features

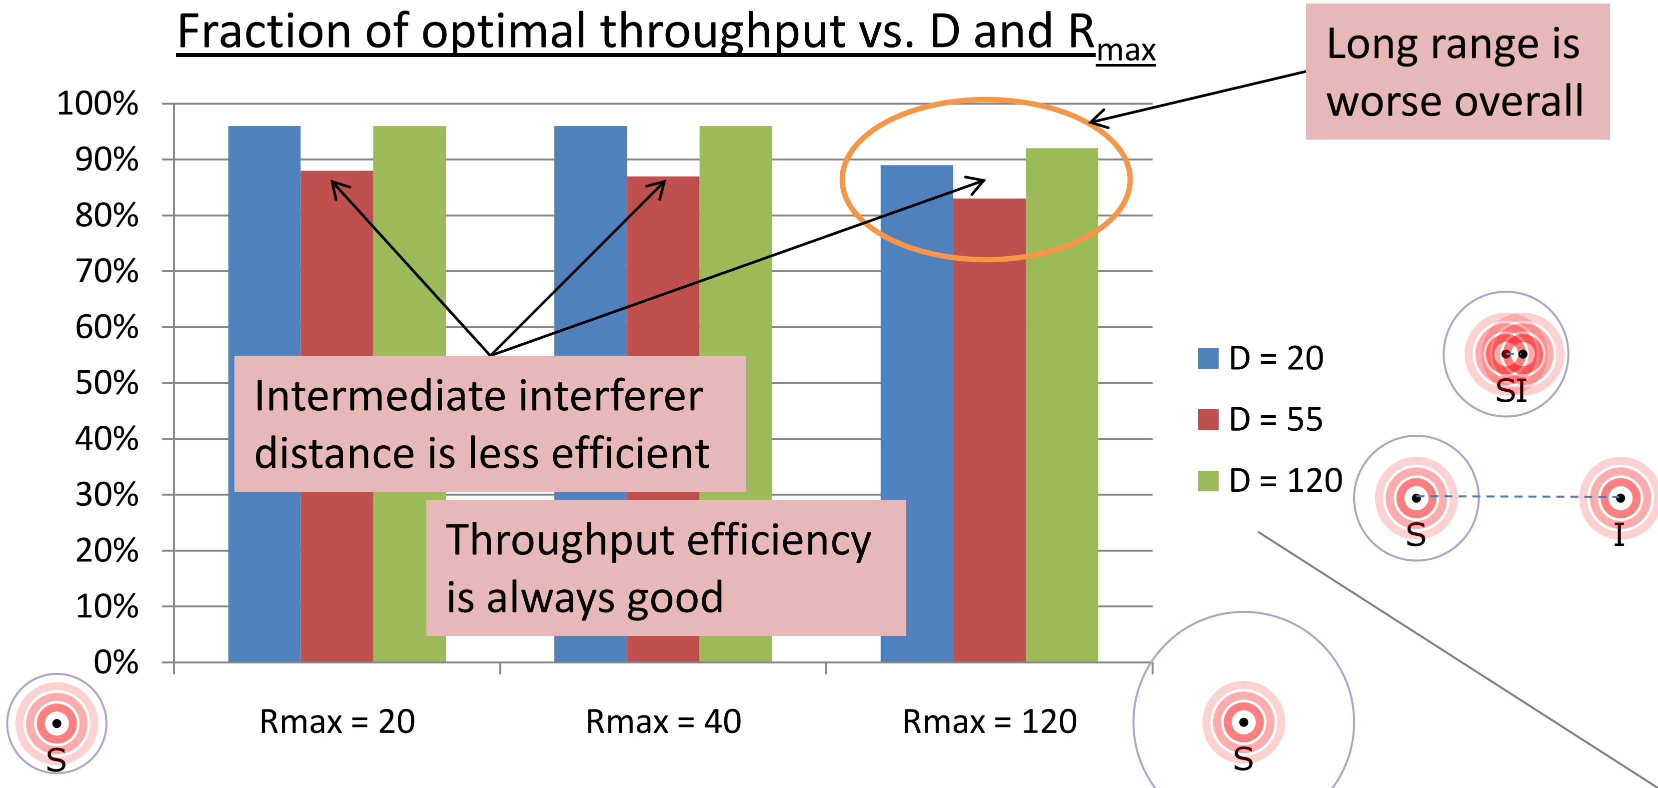

Intuitions summary

- Distant interferers affect receivers uniformly
  - Short range networks switch to multiplexing while interferer still distant
- Nearby interferers don’t – but they’re loud so everybody prefers multiplexing anyway
- So long as most receivers agree, CS performs well
- Rate adaptation smoothes rough edges in between
- Shadowing matters but isn’t big enough to drown out distance

Experiments (brief)

- Experimental hypothesis: We’re not crazy
- Result: We aren’t!
  - Carrier sense mean throughput is close to optimal
  - Short range is excellent
  - Long range is OK
- 802.11a testbed, random pairs of sender-receiver pairs
- Broadcast packets for 15 seconds, try different bitrates, measure throughput under concurrency and multiplexing
- Short range and long range scenarios

One experiment: short range

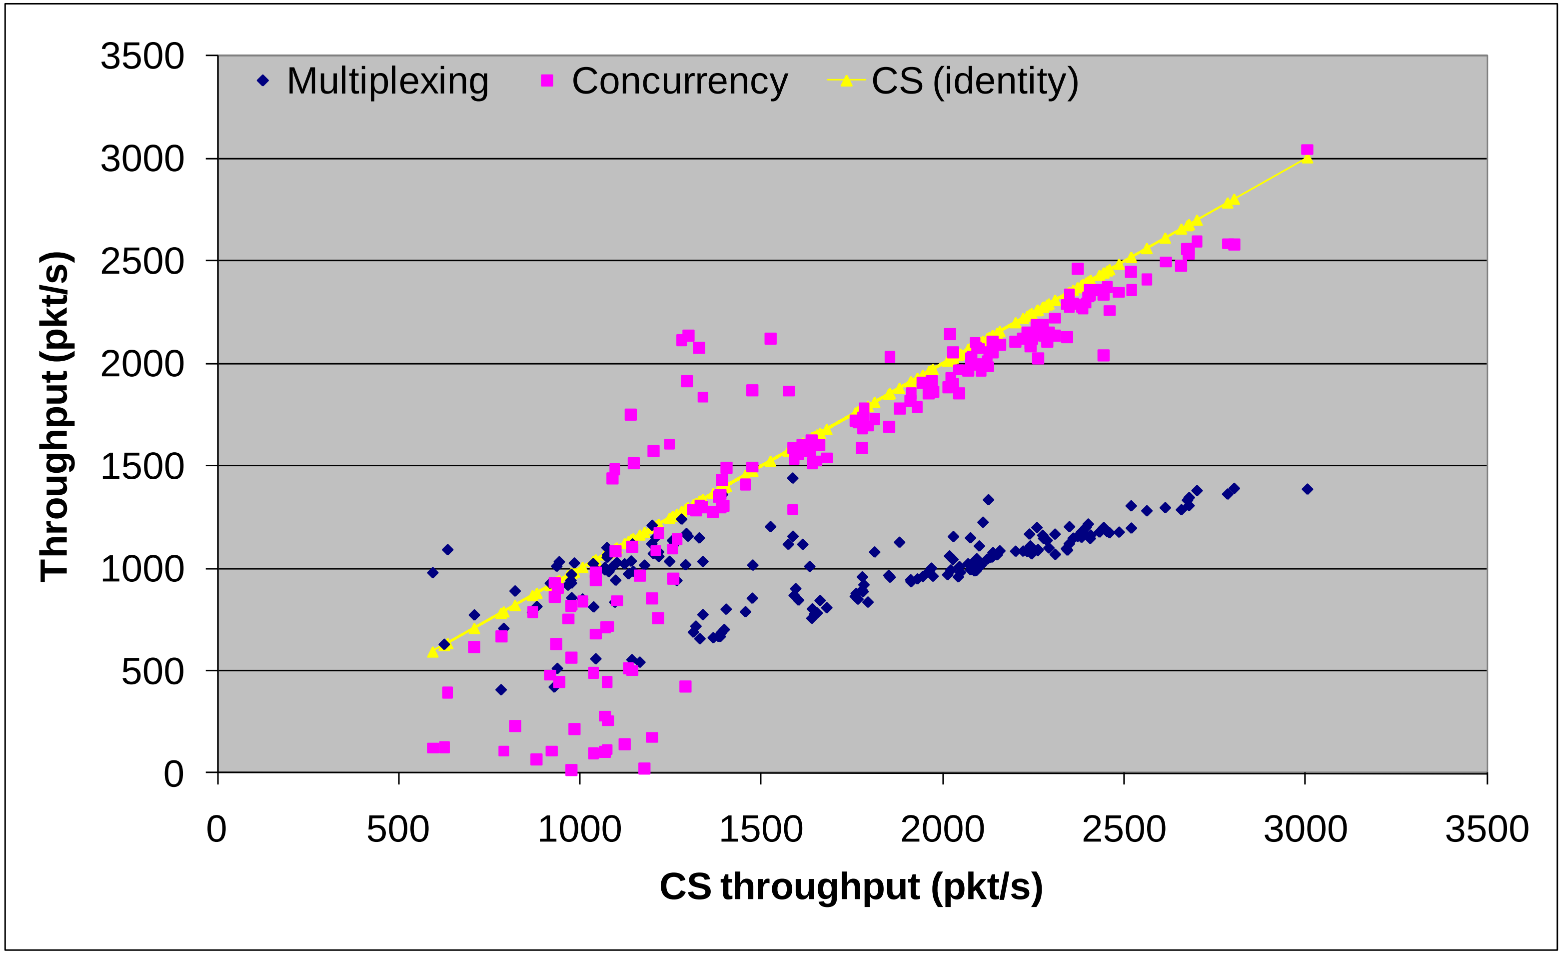

Implications for future research

- Don’t forget bitrate!
  - Much work critical of carrier sense doesn’t consider ABR and so for ABR hardware is pessimistic about CS and optimistic about claimed gains
- Hidden terminals can be a reliability problem but aren’t common and don’t matter much for average performance
  - “Expensive” solutions like RTS/CTS wouldn’t hurt throughput *if* they were only used when needed
- Exposed terminals cost these kinds of networks very little, given ABR
- (Paper argues these three points in more detail)

Conclusions

- Carrier sense *does* work, in a large, important class of networks
  - See paper for discussion of other issues like threshold robustness
- Room for improvement in corner cases, but not much in overall performance
- A fresh look at modeling can help us balance out the idiosyncrasies in experimental wireless work