Multiuser Diversity in Wireless Networks: Smart Scheduling, Dumb Antennas and Epidemic Communication

David Tse
Department of EECS, U.C. Berkeley

August 10, 2001

IMA Wireless Workshop

Communication over Wireless Channels

- Fundamental characteristic of wireless channels: fading.
- A modern view of communication over fading channels is emerging.
- This view has ramifications to the design of the physical and MAC layers as well as the architecture of future wireless data networks.

Multiuser Diversity

Multiuser Diversity

- Downlink scheduling for Qualcomm's HDR (High Data Rate) system. (Tse 99)

Multiuser Diversity

- Downlink scheduling for Qualcomm's HDR (High Data Rate) system. (Tse 99)
- Opportunistic beamforming using dumb antennas (Viswanath, Tse and Laroia 2001)

Multiuser Diversity

- Downlink scheduling for Qualcomm's HDR (High Data Rate) system. (Tse 99)
- Opportunistic beamforming using dumb antennas (Viswanath, Tse and Laroia 2001)
- Multiuser diversity in mobile adhoc networks (Grossglauser and Tse 2001)

Wireless Fading Channels

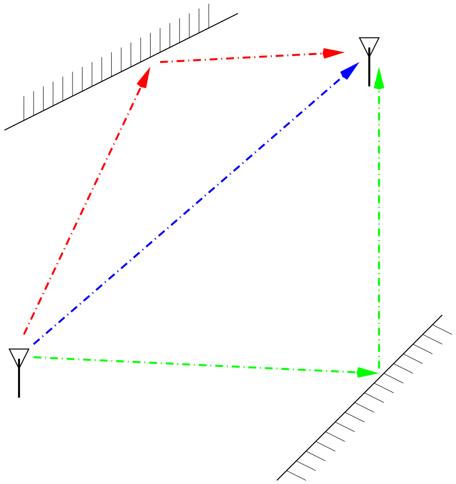

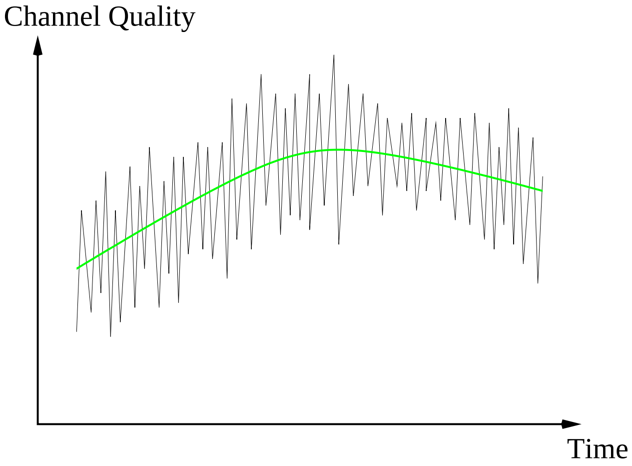

- fast time-scale fading due to constructive and destructive interference between multipaths;
- slow time-scale fading due to shadowing effects and varying distances ($r^{-\alpha}$ path loss).

Qualcomm HDR's DownLink

HDR (1xEV-DO): a wireless data system operating on IS-95 band (1.25 MHz)

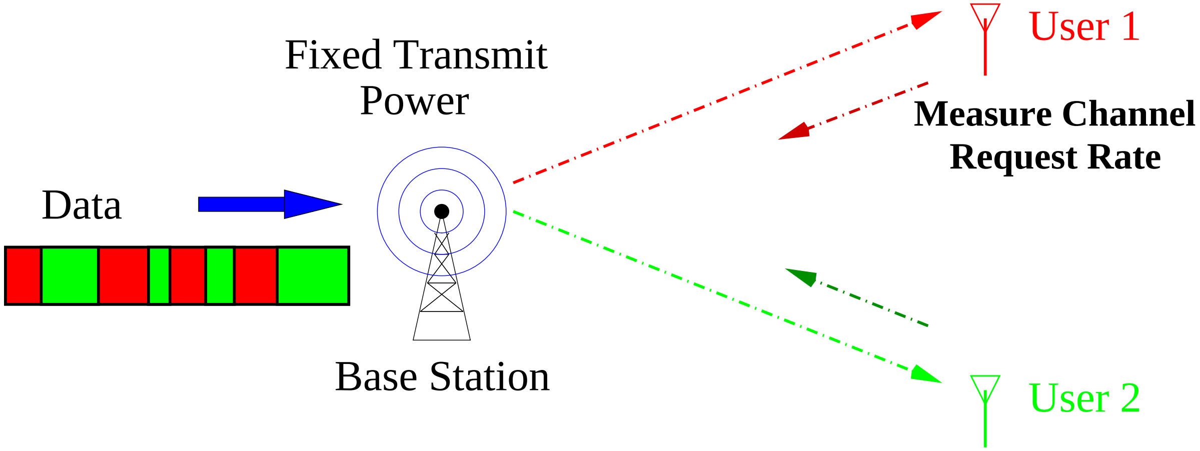

- HDR downlink operates on a time-division basis.
- Scheduler decides which user to serve in each time-slot.

Downlink Multiuser Fading Channel

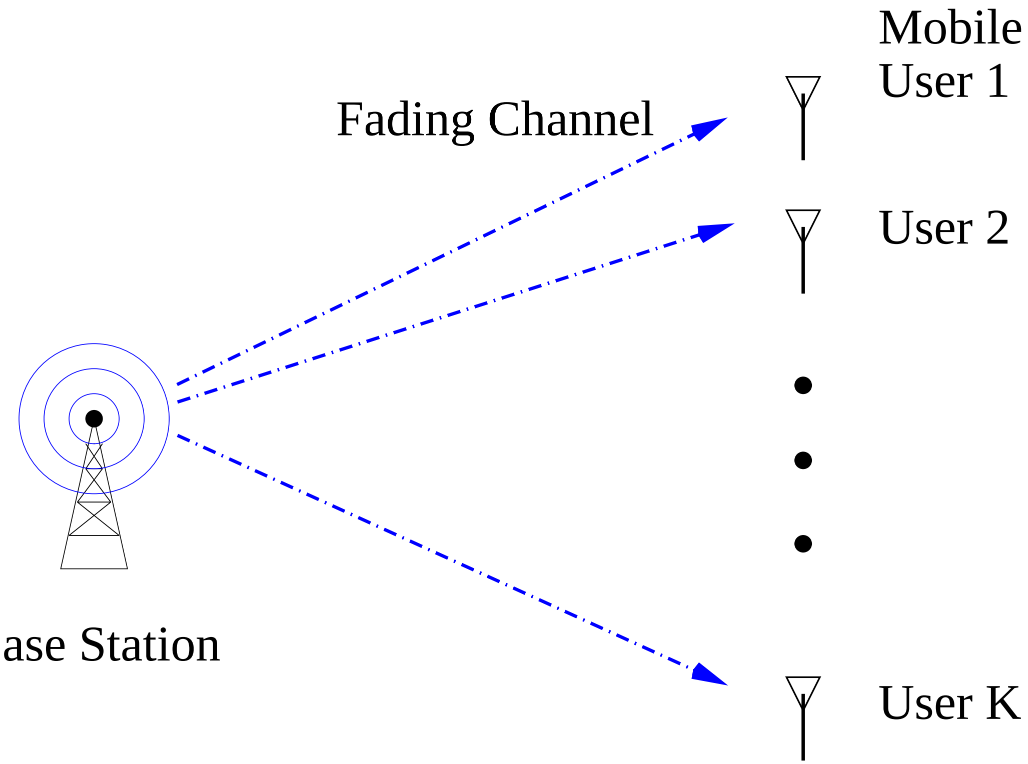

What is the sum capacity with channel state feedback?

Information Theoretic Capacity of Downlink

(Tse 97)

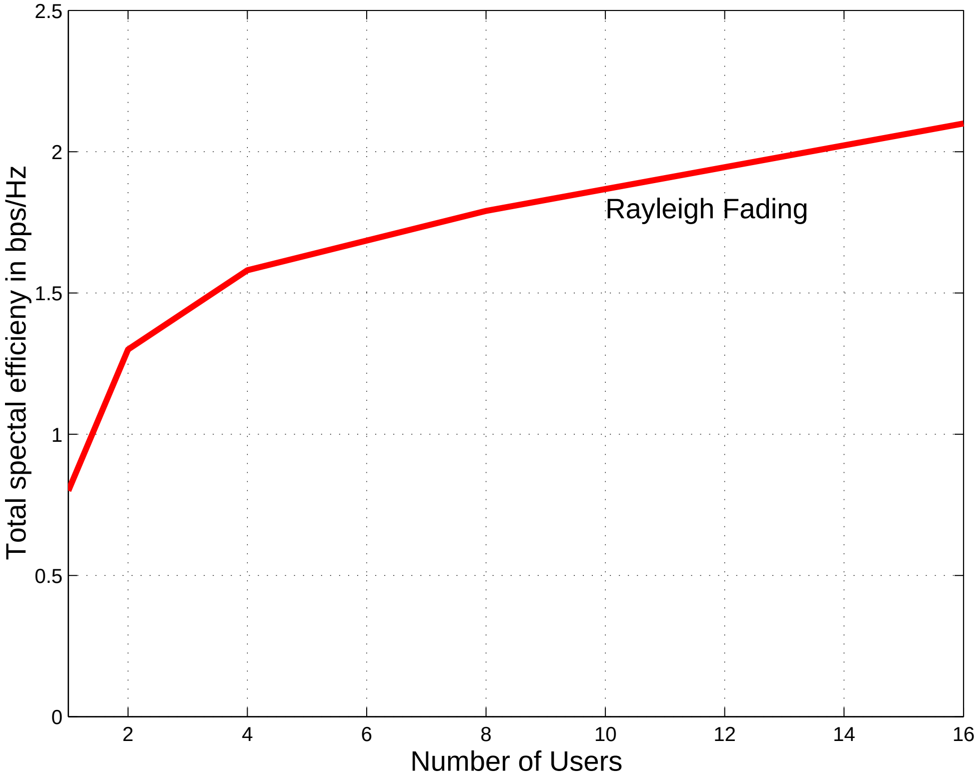

Each user undergoes independent Rayleigh fading with average received signal-to-noise ratio SNR $= 0$dB.

To Fade or Not to Fade?

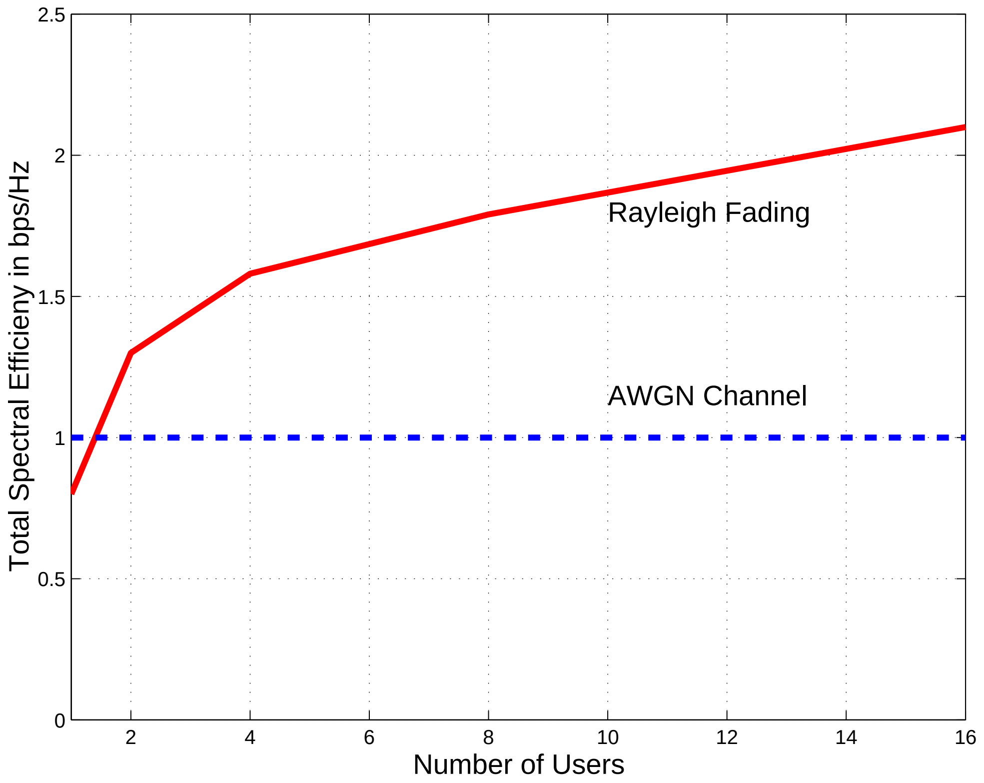

Sum Capacity of fading channel much larger than non-faded channel!

Multiuser Diversity

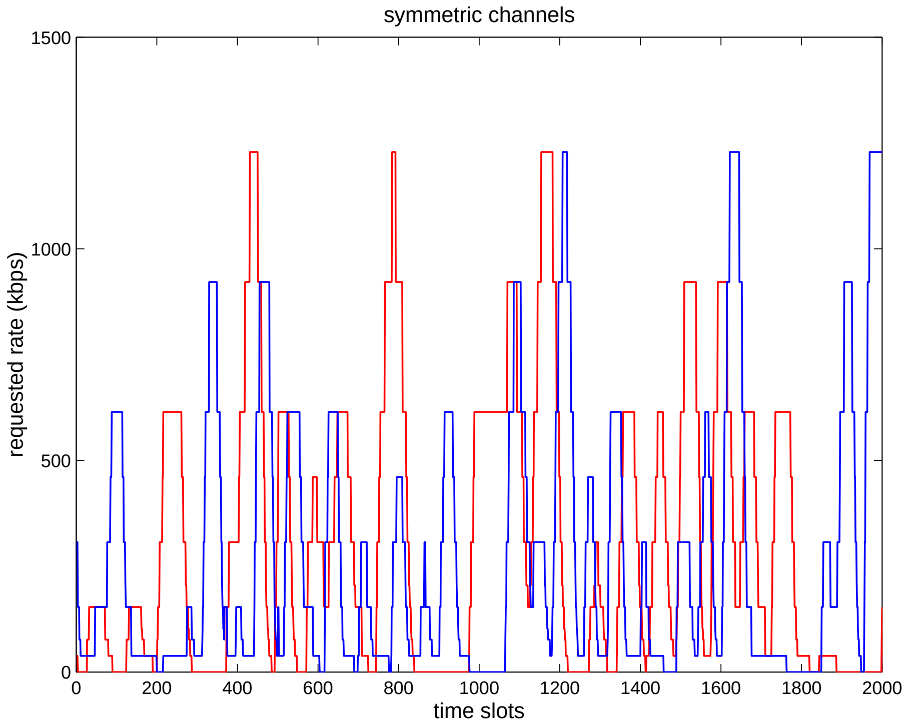

- In a large system with users fading independently, there is likely to be a user with a very good channel at any time.
- Long term total throughput can be maximized by always serving the user with the strongest channel.

(uplink multiuser diversity: Knopp and Humblet 95)

Fairness and Delay

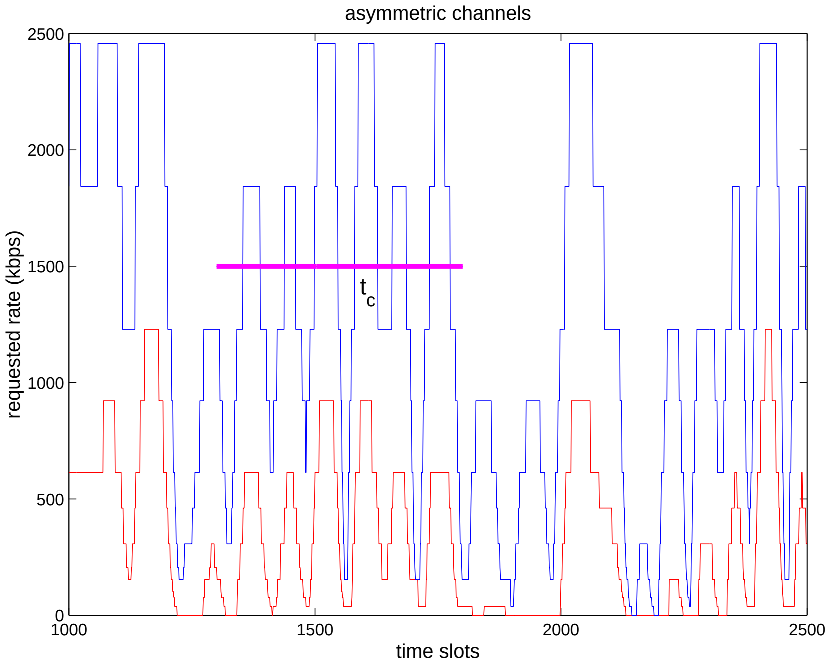

Challenge is to exploit multiuser diversity while sharing the benefits fairly and timely to users with asymmetric channel statistics.

Hitting the Peaks

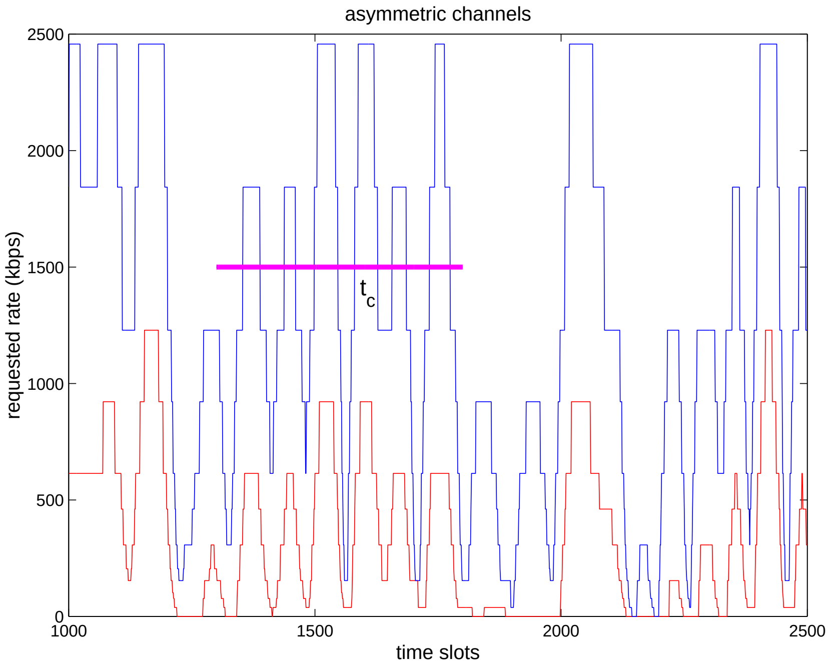

- Want to serve each user when it is near its peak within a latency time-scale t_c.

Hitting the Peaks

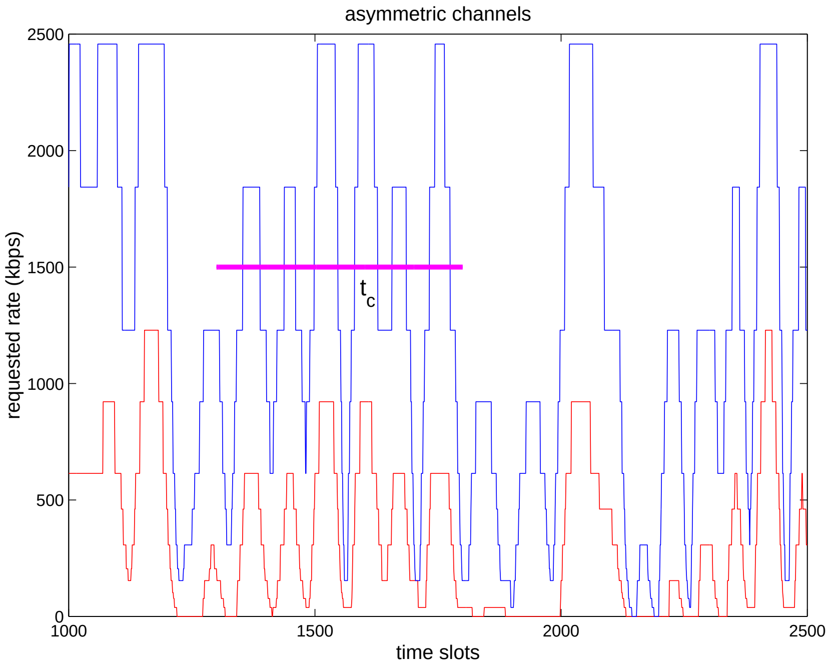

- Want to serve each user when it is near its peak within a latency time-scale t_c.
- In a large system, at any time there is likely to be a user whose channel is near its peak.

Proportional Fair Scheduler

At time slot t, given

1) users' average throughputs $T_1(t), T_2(t), \ldots, T_K(t)$ in a past window.

2) current requested rates $R_1(t), R_2(t), \ldots, R_K(t)$

transmit to the user k^* with the largest

$$\frac{R_k(t)}{T_k(t)}.$$

Average throughputs $T_k(t)$ can be updated by an exponential filter.

Comments

- If users have symmetric channel statistics, this reduces to the greedy policy of transmitting to the mobile with the highest requested rate.
- If channels have different statistics, competition for resource is made fair by normalization
- feedback is built into the metric $R_k(t)/T_k(t)$ to provide a fair bandwidth allocation over the time-scale t_c.

Comparison with Round-Robin Policy

Round-Robin Policy

- Give same number of time slots to all the users in a round-robin fashion, regardless of their channel conditions.

Proportional fair policy:

- Give roughly the same number of time slots to all users, but try to transmit to a user when its channel condition is near its peak.
- Resource fair, but not necessarily performance fair.

Proportional Fairness

Under stationary assumptions, long-term average throughputs $T_1^*, \ldots, T_K^*$ of the scheduler maximizes

$$\sum_k \log T_k$$

among all schedulers.

Throughput of HDR Scheduler: Symmetric Users

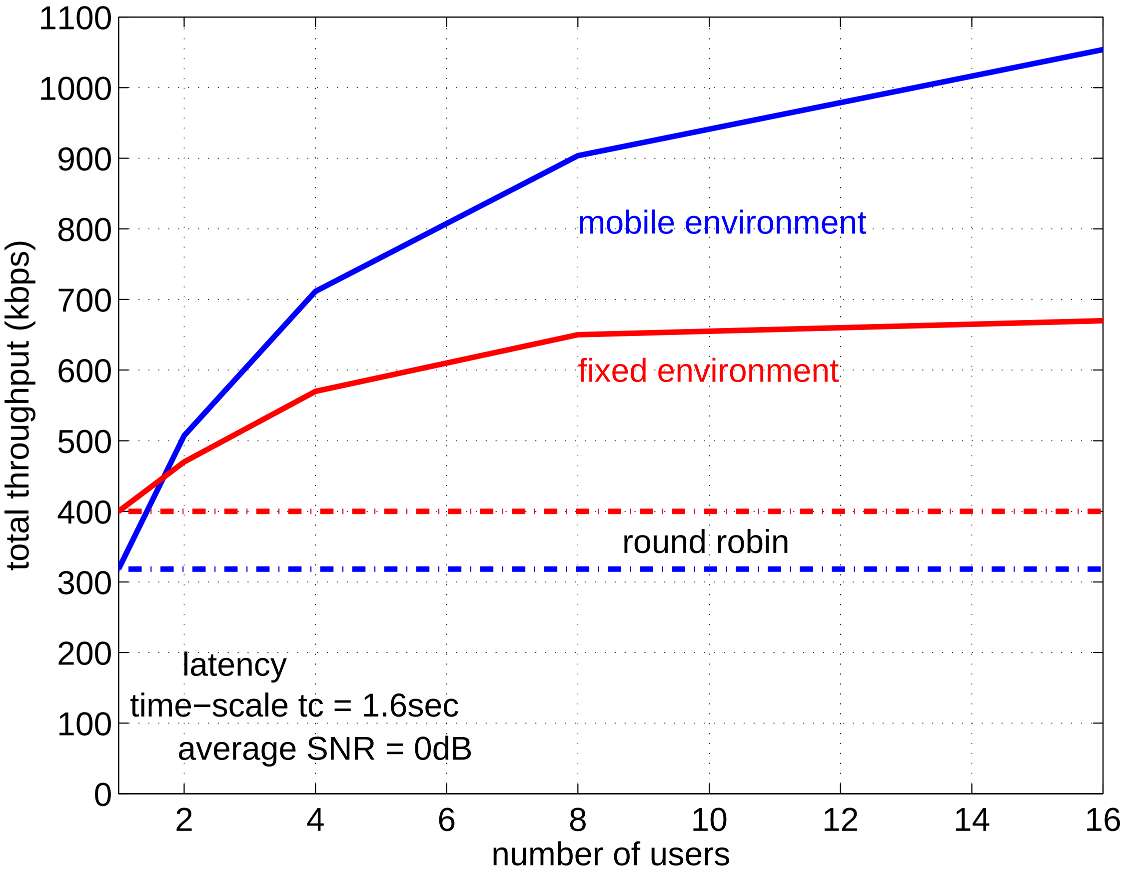

Mobile environment: 3 km/hr, Rayleigh fading

Fixed environment: 2Hz Rician fading with $E_{\text{fixed}}/E_{\text{scattered}} = 5$.

Channel Dynamics

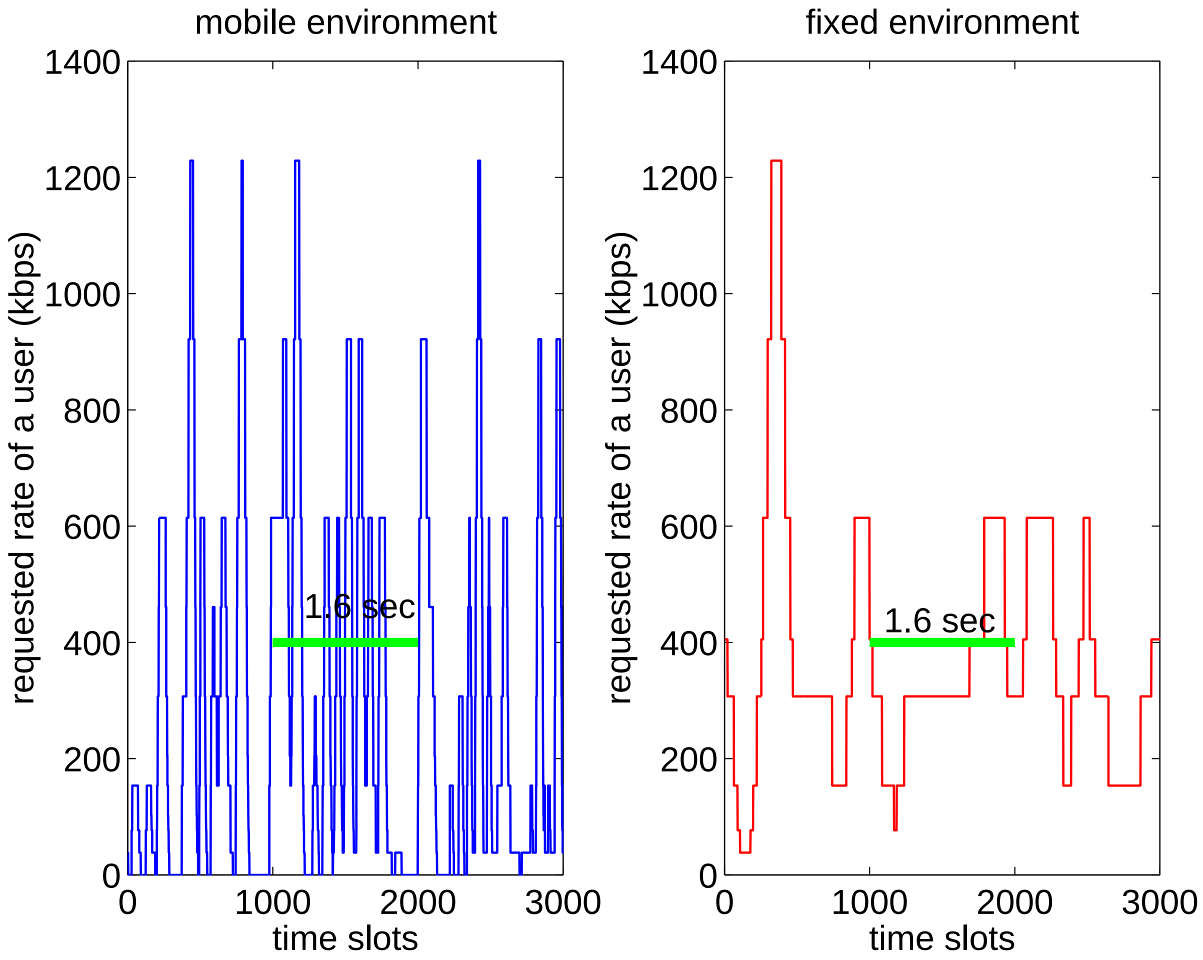

Channel varies faster and has more dynamic range in mobile environments.

Throughput of Scheduler: Cellular Network

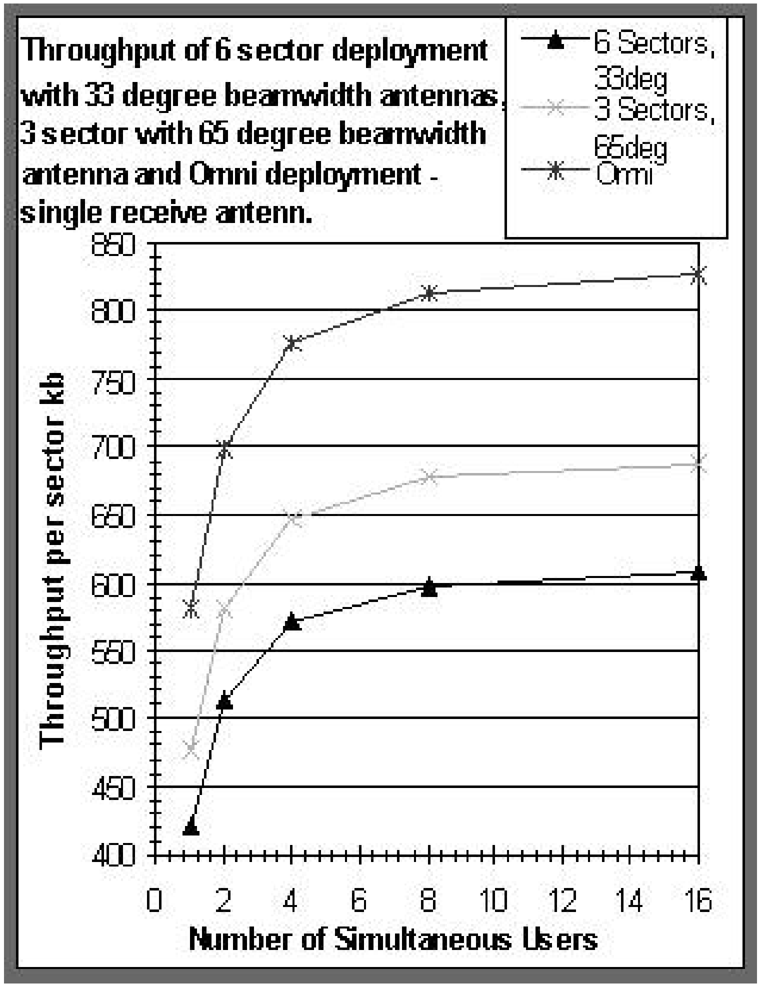

(Jalali, Padovani and Pankaj 2000)

Inducing Randomness

- Scheduling algorithm exploits the nature-given channel fluctuations by hitting the peaks.
- If there are not enough fluctuations, why not induce them artificially?

Dumb Antennas

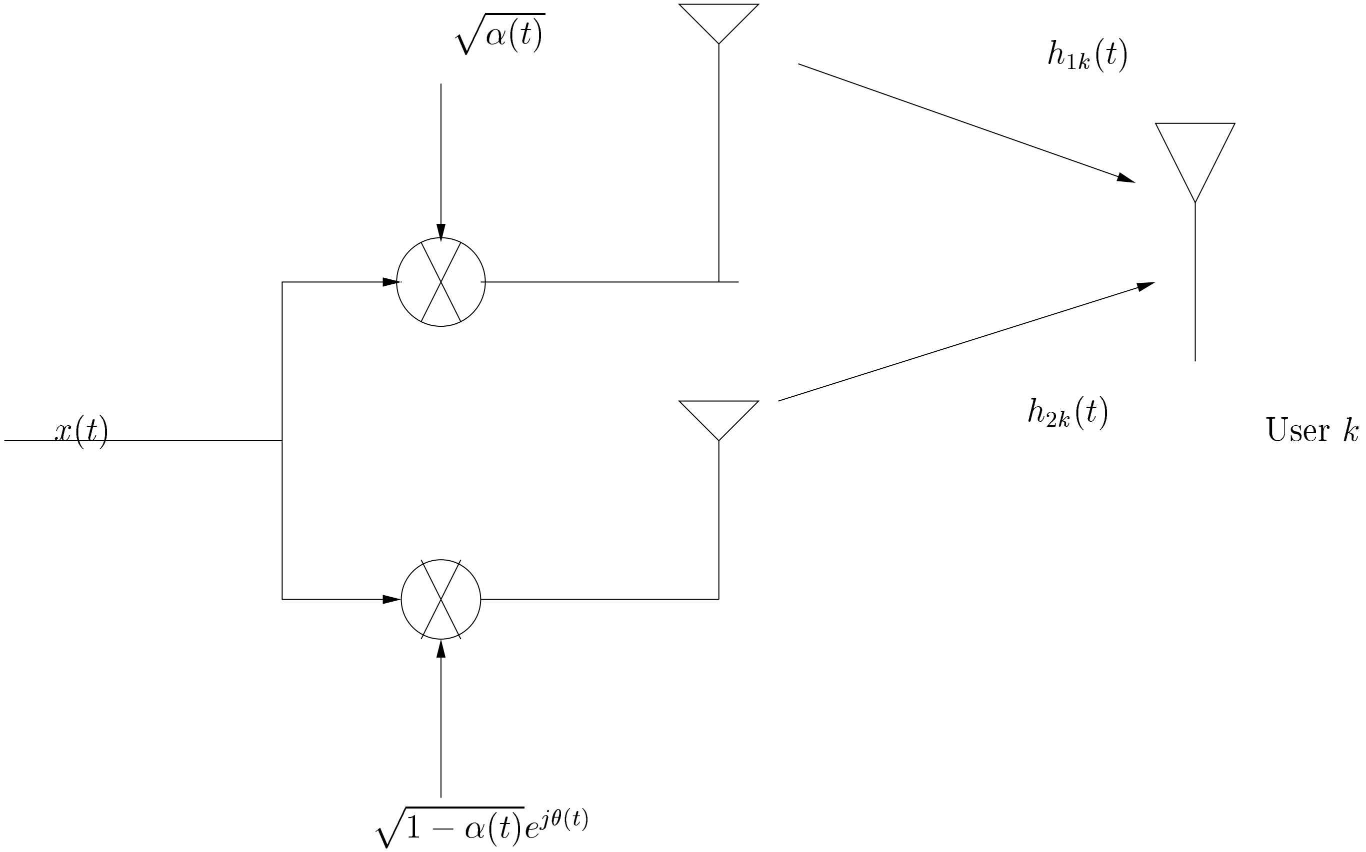

Slow Fading Environment: Before

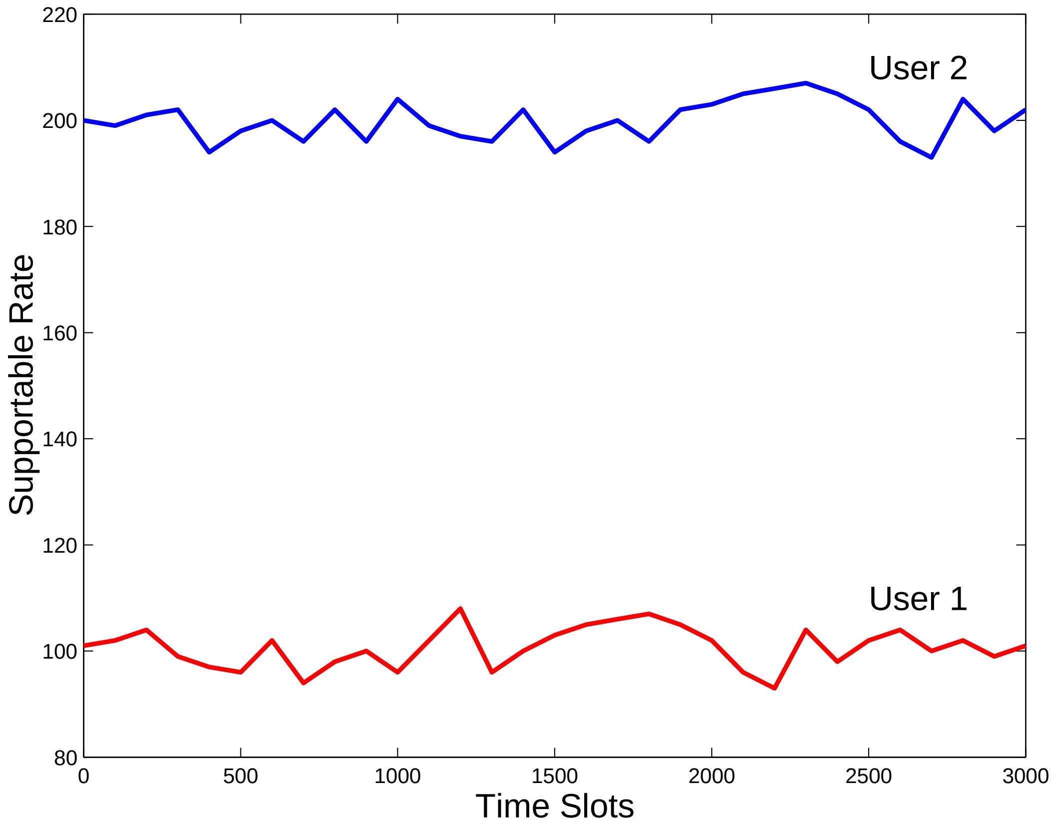

After

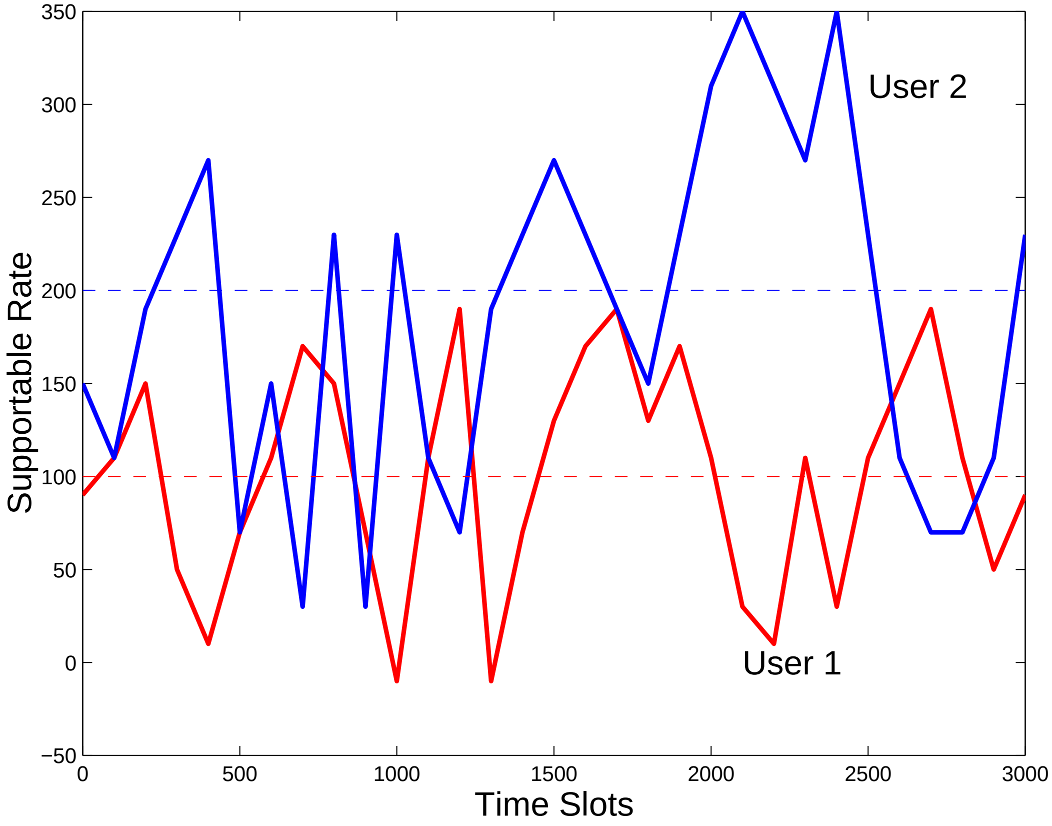

350
300
250
200
150
100
50
0
−50
0
500
1000
1500
2000
2500
3000
Supportable Rate
Time Slots
User 2
User 1

Opportunistic Beamforming: Slow Fading

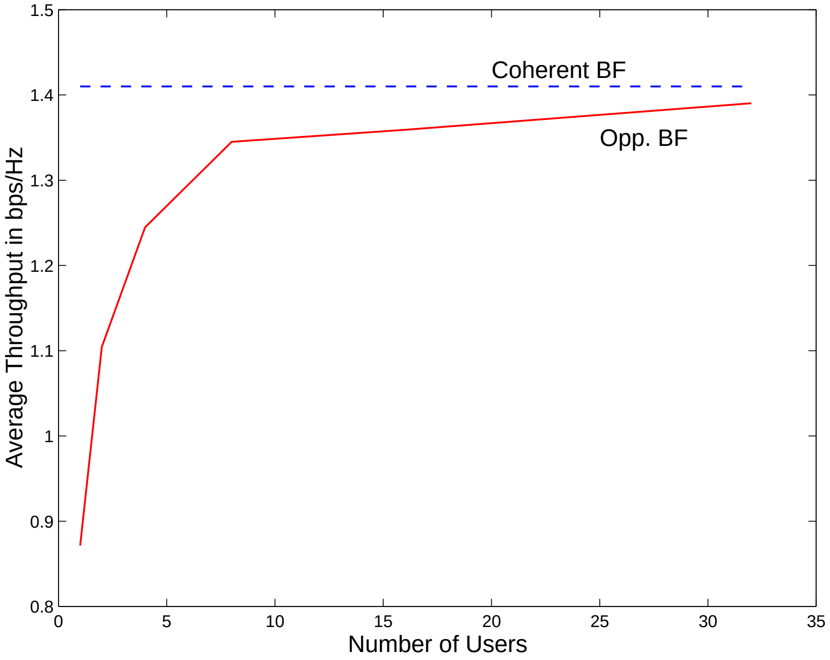

- Dumb antennas randomly sweep out a beam and opportunistically sends data to the user closest to the beam.
- Can approach the performance of true beamforming when there are many users in the systems, but with much less feedback and channel measurements.

Asymptotic Result

Assume that the slow fading states of each user are i.i.d. randomly generated (but fixed for all time).

In a large system of K users, with high probability, the users achieve throughputs

$$T_k = \frac{1}{K} R_k^{\mathrm{bf}}, \qquad k = 1, \ldots K$$

.

where R_k^{bf} is the rate user k gets when it is perfectly beamformed to.

Opportunistic Beamforming: Fast Fading

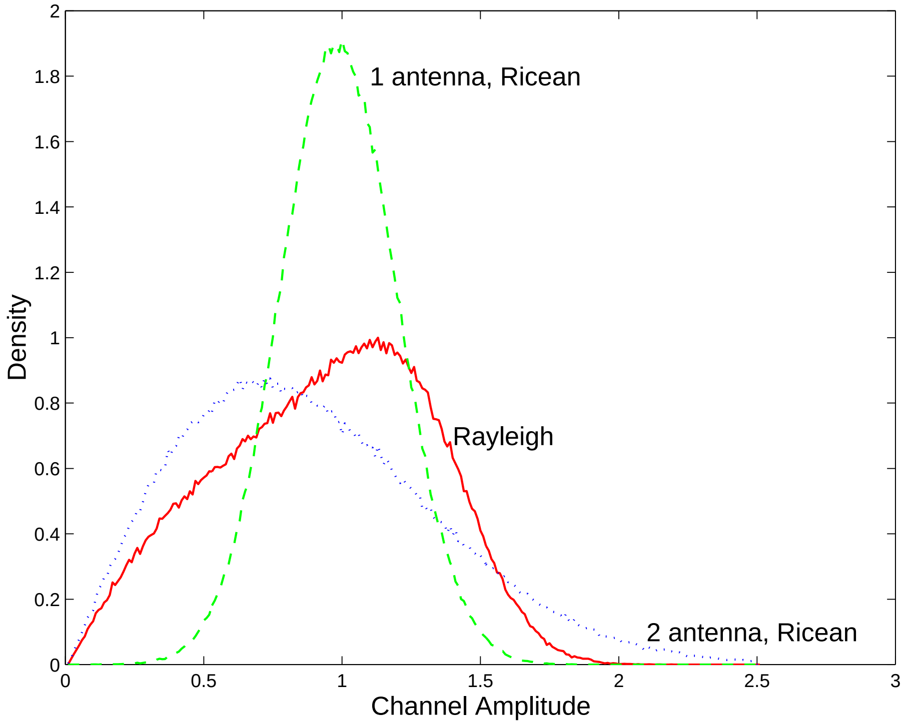

Improves performance in fast fading Rician environments by spreading the fading distribution.

Overall Performance Improvement

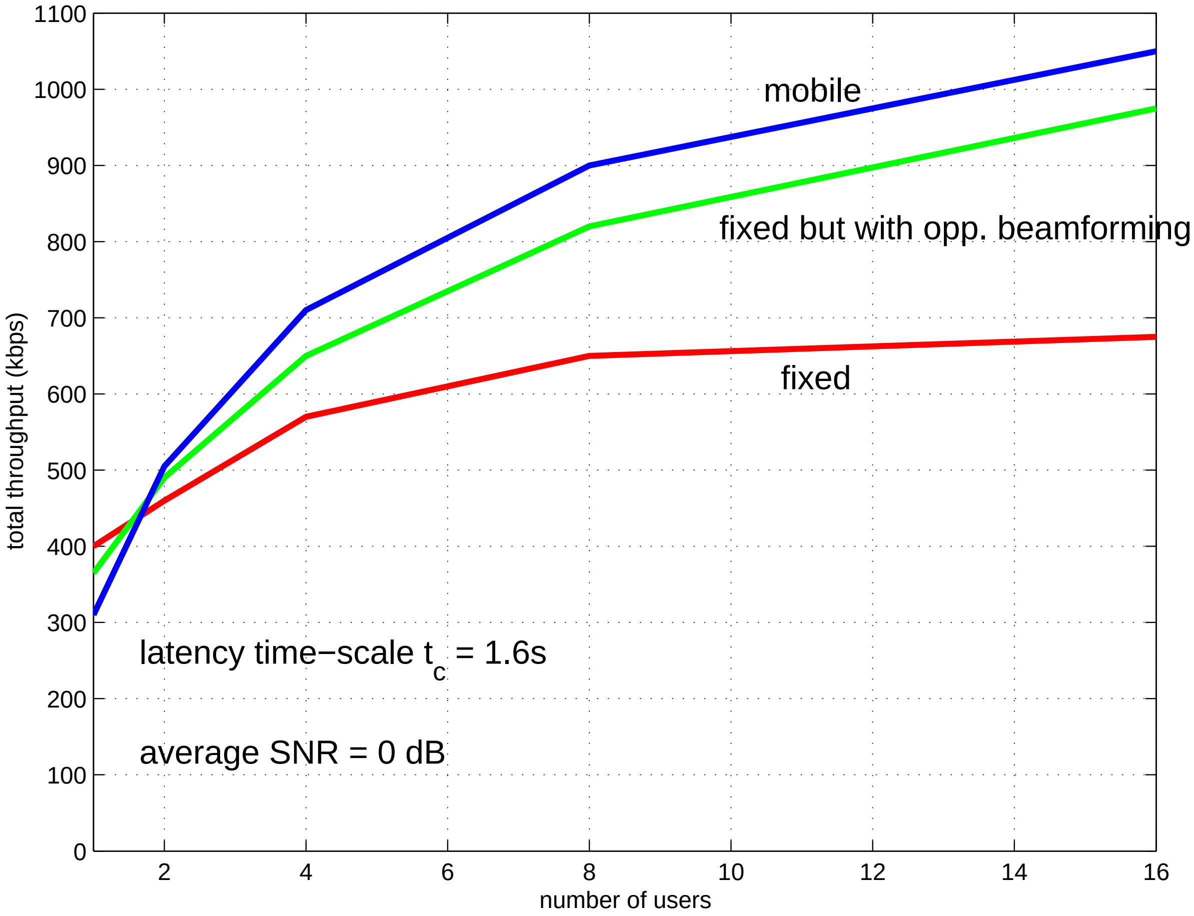

Mobile environment: 3 km/hr, Rayleigh fading

Fixed environment: 2Hz Rician fading with $E_{\text{fixed}}/E_{\text{scattered}} = 5$.

Comparison to Space Time Codes

- No special multi-antenna encoder or decoder nor MIMO channel estimation.
- In fact the mobiles are completely oblivious to the existence of multiple transmit antennas.
- Antennas are truly dumb, but yet can surpass performance of space time codes.

Cellular System: Opportunistic Nulling

- In a cellular systems, users are scheduled when their channel is strong and the interference from adjacent base-stations is weak.
- Dumb antennas provides opportunistic nulling for users in other cells.
- Particularly important in interference-limited systems with no soft handoff.

Traditional CDMA Downlink Design

- orthogonalize users (via spreading codes)

Traditional CDMA Downlink Design

- orthogonalize users (via spreading codes)
- Makes individual point-to-point links reliable by averaging:
  - interleaving
  - multipath combining,
  - soft handoff
  - transmit/receive antenna diversity

Traditional CDMA Downlink Design

- orthogonalize users (via spreading codes)
- Makes individual point-to-point links reliable by averaging:
    - interleaving
    - multipath combining,
    - soft handoff
    - transmit/receive antenna diversity
- Important for voice with very tight latency requirements.

Downlink Design: Modern View

- Shifts from the point-to-point view to a multiuser network view.

Downlink Design: Modern View

- Shifts from the point-to-point view to a multiuser network view.
- Wants large and fast fluctuations of both channel and interference so that we can ride the peaks.

Downlink Design: Modern View

- Shifts from the point-to-point view to a multiuser network view.
- Wants large and fast fluctuations of both channel and interference so that we can ride the peaks.
- Exploits more relaxed latency requirements of data as well as MAC layer packet scheduling mechanisms.

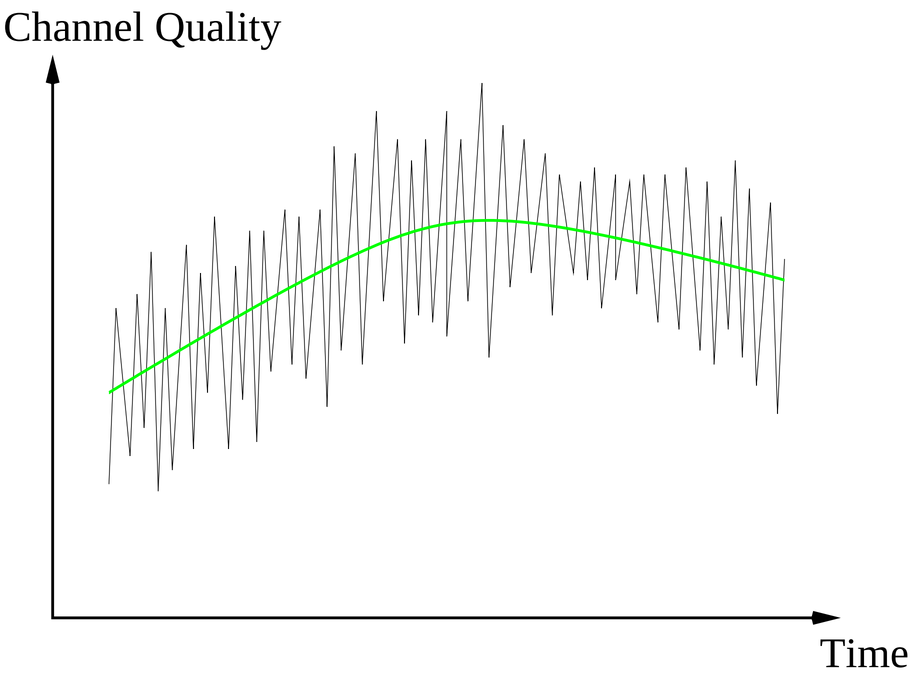

So far we have exploited randomness from fast fading.

How about the slow fading?

Talk Outline

- Downlink scheduling for Qualcomm's HDR (High Data Rate) system. (Tse 99)
- Opportunistic beamforming using dumb antennas (Viswanath, Tse and Laroia 2001)
- Multiuser diversity in mobile adhoc networks (Grossglauser and Tse 2001)

All-Wireless Ad-Hoc Networks

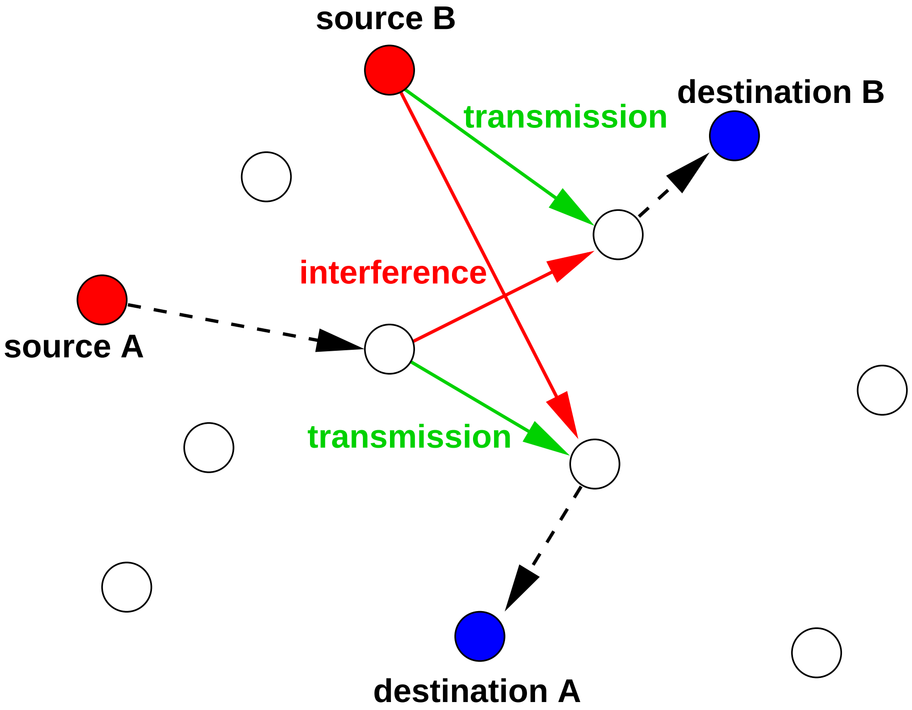

- no wireline infrastructure
- nodes can be relays for other nodes
- packets can go through multiple hops

Scalability of Ad-Hoc Networks

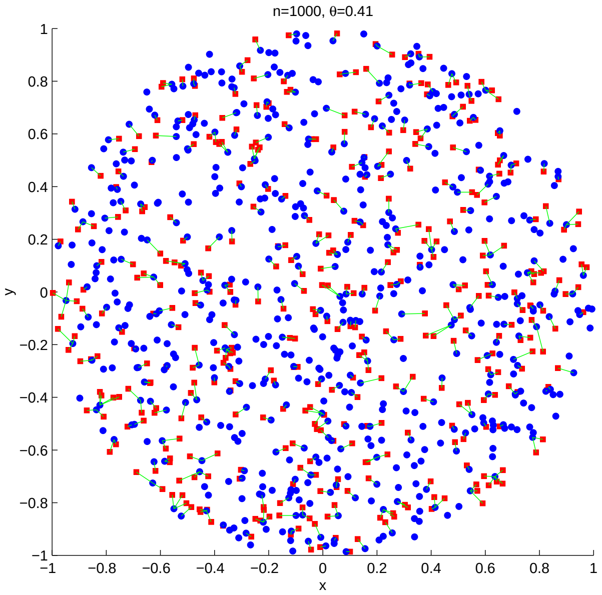

Point-to-point: Suppose each node has a stream of traffic for a particular destination node.

Question: How does the throughput per source-destination pair scale with the number of nodes n per unit area?

Mobile Ad-hoc Network Model

- Mobility model: Nodes move randomly and independently on a disk of unit area.
- Channel model: path loss factor of $r^{-\alpha}$ at distance r, with $\alpha > 2$ (slow fading);
- Communication model: a packet is successfully received if signal-to-interference ratio is greater than a prescribed threshold.

Baseline: Fixed Nodes Case

Suppose nodes are randomly placed but immobile.

Gupta and Kumar:

The best achievable throughput per source-destination pair decreases to zero like $O(\frac{1}{\sqrt{n}})$.

This result holds even allowing optimal and centralized scheduling, power control, relaying and routing.

What’s the Problem?

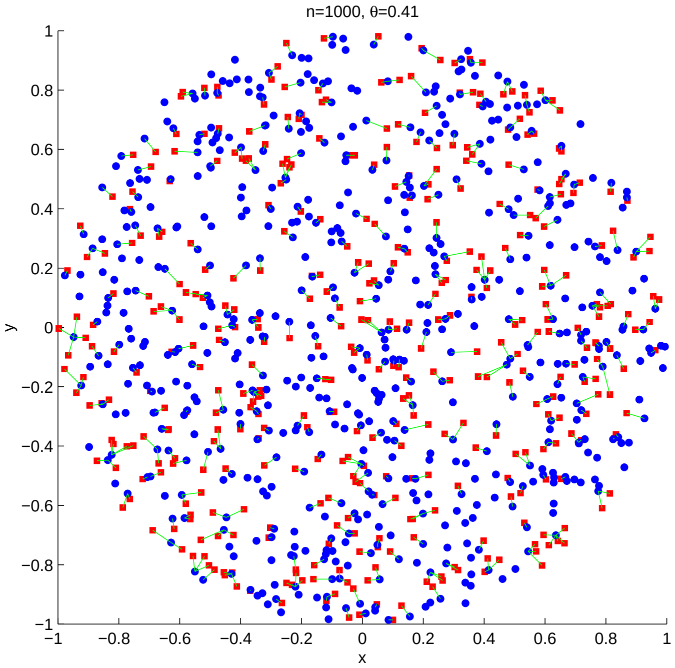

- communication confined to nearest neighbors to allow dense spatial reuse.
- a typical route has order $\sqrt{n}$ relay nodes
- each packet is transmitted of the order $\sqrt{n}$ times.

Mobility Can Help!

Main result: (Grossglauser and Tse 00)

A long-term throughput of $O(1)$ per S-D pair can be achieved when nodes are mobile......

Mobility Can Help!

Main result: (Grossglauser and Tse 00)

A long-term throughput of $O(1)$ per S-D pair can be achieved when nodes are mobile.......

......if one is willing to wait.

Mobility Can Help!

Main result: (Grossglauser and Tse 00)

A long-term throughput of $O(1)$ per S-D pair can be achieved when nodes are mobile.......

......if one is willing to wait.

Throughput is averaged over the time-scale of mobiliy.

Idea #1

Communicate only when the source and destination are nearest neighbors to each other.

Direct Communication Does Not Work

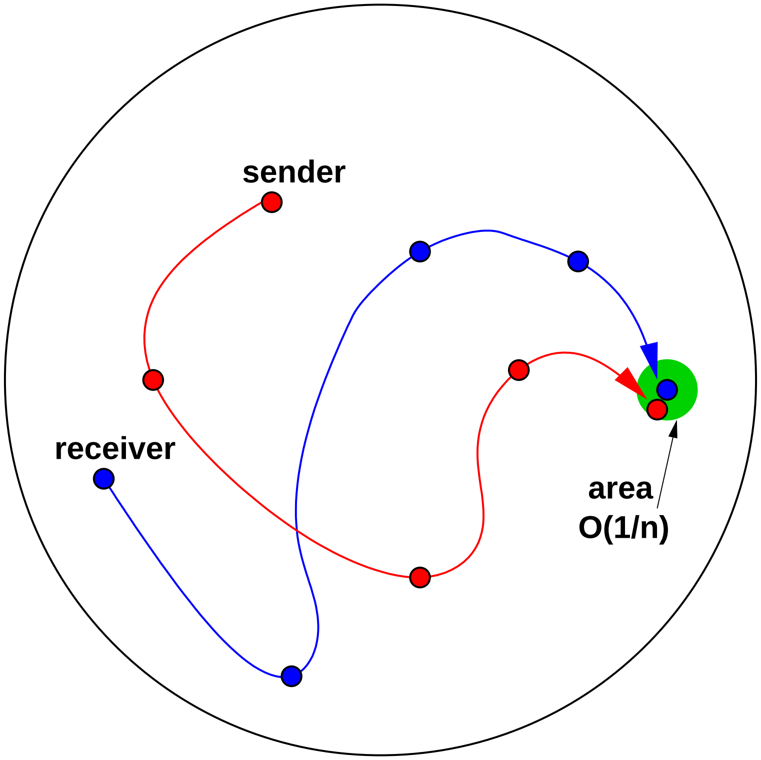

- The source and destination are nearest neighbors only $O(1/n)$ of the time
- In fact, can show S-D throughput is at most $O\left(n^{-\frac{1}{1+\alpha/2}}\right)$ for any policy that does not use relays.

Multiuser Diversity via Relaying

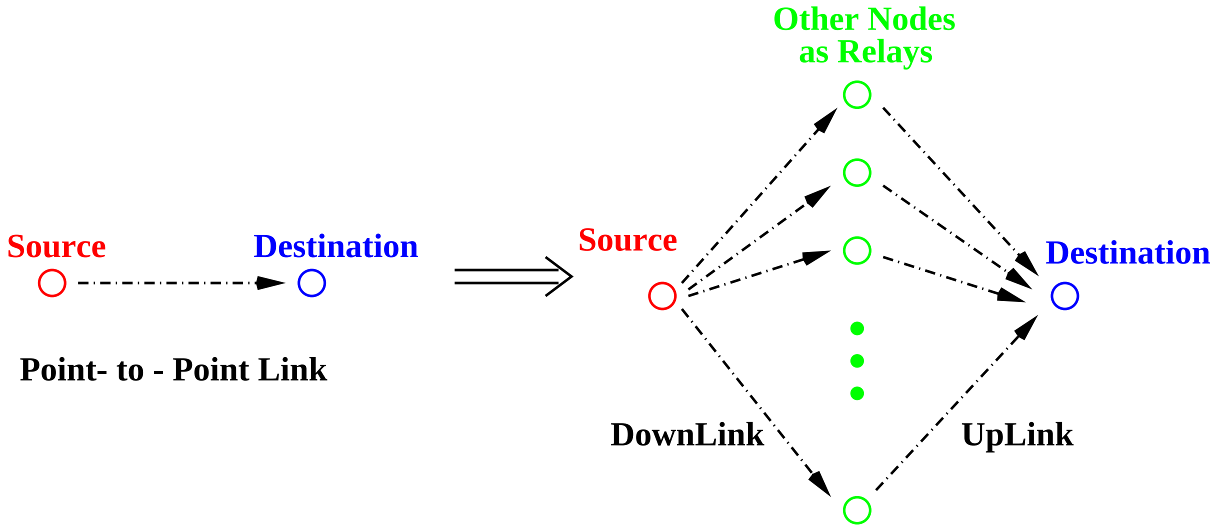

- Multiuser diversity created artificially using all other nodes as relays.
- Channel variation comes from slow rather than fast time-scale fading.

Phase I: Source to Relays

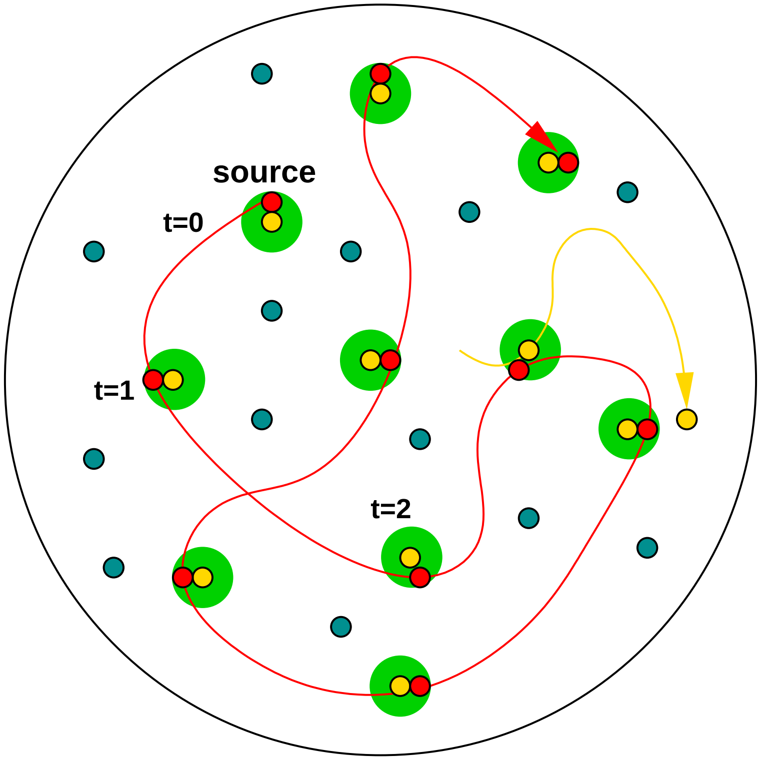

- At each time slot, source relays a packet to nearest neighbor
- Different packets are distributed to different relay nodes.

Phase II: Relays to Destination

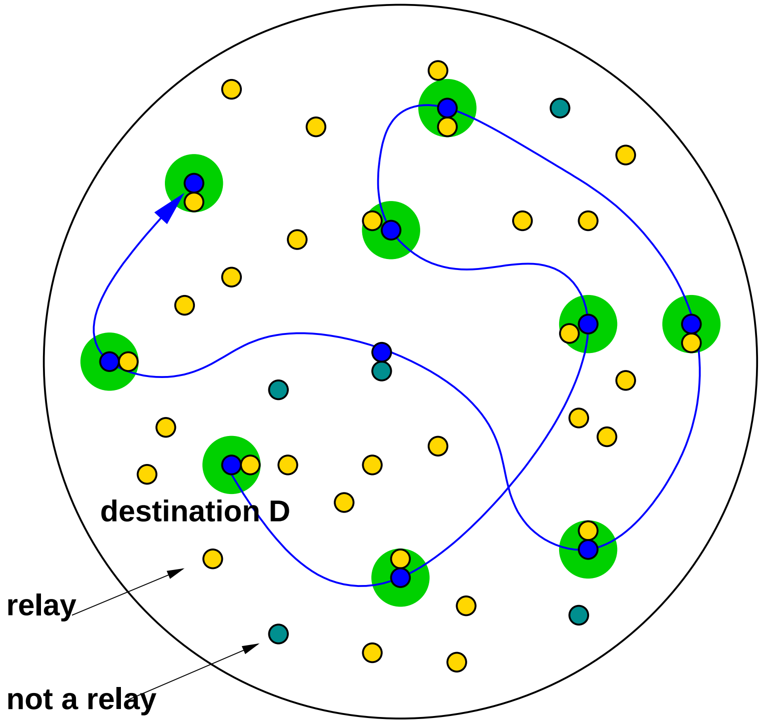

- Steady state: all nodes have packets destined for D
- Each relay node forwards packets to D only when it gets close.

Phase I and II Staggered

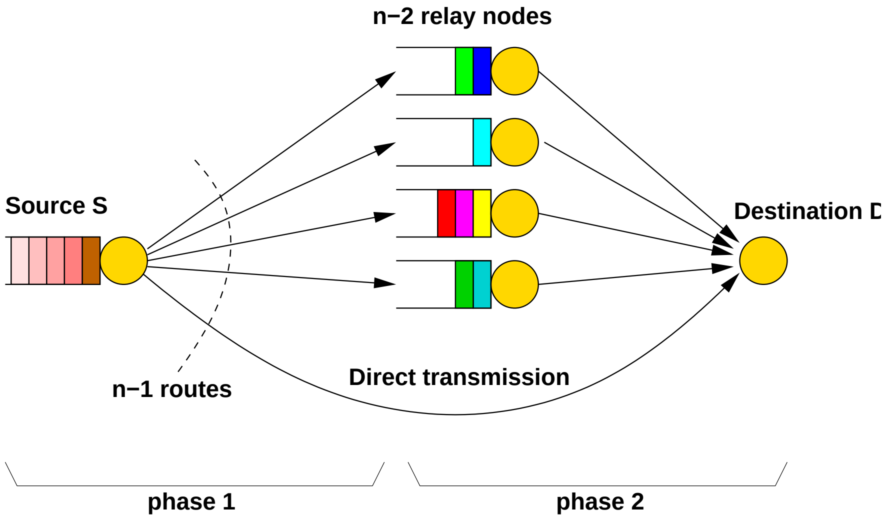

- $O(1)$ throughput from S to D
- Communication is confined to nearest neighbors, but each packet goes through at most two hops.

Network Capacity

- The above discussion pertains to a single source-destination pair.
- We have to show that every S-D pair can follow the same strategy simultaneously.
- $O(n)$ simultaneous nearest neighbor communication is possible, due to power law decay of the received power from a randomly located node.

Discussion

- Throughput is gained at the expense of delays, at the time-scale of significant changes in network topology.
- The result is extreme in the delay-throughput tradeoff, but suggests that mobility should be exploited in designing scalable architectures for delay-tolerant applications.

Conclusions: Three Lessons Learnt

Conclusions: Three Lessons Learnt

- Channel variation can be exploited (even exaggerated) rather than mitigated:
  Hitting the peaks versus averaging.

Conclusions: Three Lessons Learnt

- Channel variation can be exploited (even exaggerated) rather than mitigated:
  Hitting the peaks versus averaging.
- Data versus voice.

Conclusions: Three Lessons Learnt

- Channel variation can be exploited (even exaggerated) rather than mitigated:
  Hitting the peaks versus averaging.
- Data versus voice.
- A system view towards wireless network design.