Packet-Level Diversity From Theory to Practice: An 802.11-based Experimental Investigation

E. Vergetis, E. Pierce, M. Blanco and R. Guérin

University of Pennsylvania
Department of Electrical & Systems Engineering

Mobicom 2006
September 23–26, Los Angeles, CA, USA

Our Starting Point: Diversity

- Diversity: more than one choice in terms of available path/channel
    - A long history of results pointing to its benefits as a means to improve performance

- Some recent developments hinting that simple (blind) deterministic round-robin policies “should” work well most of the times

- Wireless environment is one where trying to take advantage of diversity appears most natural

Our Focus

- Diversity: Lots of ways to leverage it
  - *Open-loop* (blind) system
    - no channel feedback (only general statistics are known, if at all)
    - Pre-determined transmission policy (what packet in a message is sent on what channel)

- Performance: Lots of ways to measure it
  - Metric of interest is *message rate* with an eye to *real-time* applications that require some minimum level of message delivery guarantees
    - No retransmissions
  - Forward error code as another design knob
    - Packet level (N,k) code recovers from loss patterns of up to N-k packets

Outline

- A short primer on open-loop diversity
  - How it works and what it assumes
- Our goals and the experimental setup we used
  - Testing the gap between theory & practice
    - Implementation issues
    - Channel “model”
  - Questions and investigation approach
    - When and/or how much does diversity help/hurt?
- Findings
- Potential Extensions

Open-Loop Diversity – The Theory

- User can choose from *C* channels with “known” statistics
  - Long-term error rate (LTER), expected burst length (EBL)
  - User transmissions do not affect channel statistics
- User distributes packet transmissions across all *C* channels according to some policy
  - Deterministic and probabilistic policies
- User wants to maximize performance
  - Highest possible message (consisting of *k* packets) delivery rate that meets a certain reliability target P_{min}
- Design knobs
  - Transmission policy
    - What set of channels to use and how?
  - Code selection
    - What (N,k) code to choose (smallest N that achieves P_{min})?

Open-Loop Diversity – The Results

- Under certain assumptions
  - Channel independence
  - Stationary, Markovian (Gilbert-Elliot) channels
  - No overhead in switching transmissions from one channel to another

- In scenarios where diversity “helps,” a simple, round-robin policy is close to optimal, and “usually” wont hurt
  - Higher effective message rate (*ER*) and relative insensitivity to errors in channel statistics

Open-Loop Diversity – The Intuition

- Channel diversity is useful because/when
  - It allows breaking-up error bursts
  - It avoids being "stuck" with a bad channel

- Deterministic, round-robin policy works well because
  - It spaces out successive use of a given channel (minimizes the odds of coming back early to an ongoing error burst)

- Distributing packet transmissions across multiple channels yields
  - The ability to use a smaller code length N to satisfy Pmin
  - And/or a higher probability of successful message delivery
  - Most of the gains from diversity are through reducing N

Open-Loop Diversity – The Questions

- How well do the assumptions hold in practice?
  - Independent, stationary channels, with known statistics
  - No impact of user transmissions on channel statistics
  - No channel switching overhead

- What can actually be realized?
  - 802.11b environment
  - Standard end-systems (PCs) without precise control of transmission timings

Experimental Setup

- Two 802.11b Access Points (APs)
  - Intel StarEast board, with one miniPCI NIC each
  - External omni-directional antennas
  - Assigned “non-overlapping” frequency bands
  - Located ~1m from each other
  - Logging of all incoming packets without performance degradation
  - Within reach of other APs interfering in all 11 frequency bands
- Sender
  - Standard laptop with *two* NICs
    - One external PCMCIA NIC, and one internal miniPCI NIC
    - Linux operating system
    - Transmission speed set at 2Mbps
  - Located between 2m and 10m away from the two APs
    - Maintains association with both APs
    - Line-of-sight (LoS) as well as non-LoS (indoor wall) transmissions

Some Other Implementation/Operation Aspects to Consider

- Impact of 802.11 operation
  - RTS/CTS handshake before transmissions
    [**Disabled -** Large RTSThreshold value**]**
  - "Feedback" mechanism: ACK packets
    [**Disabled - ** Broadcast packets **]**
  - Channel access control (contention period)
    - Sensing and exponential backoff
    - Inter-frame spaces (SIFS, DIFS, etc.)
- Processor and OS overhead vs. transmission speed of the NICs
  - Where is the bottleneck and how does it affect transmission timings?

Transmission Timing Scenarios

- Single channel
- Interleaving on one channel
- Perfect timing on two channels
- Bandwidth limited on two channels
- Processor limited on two channels
- Interleaving on two bandwidth limited channels

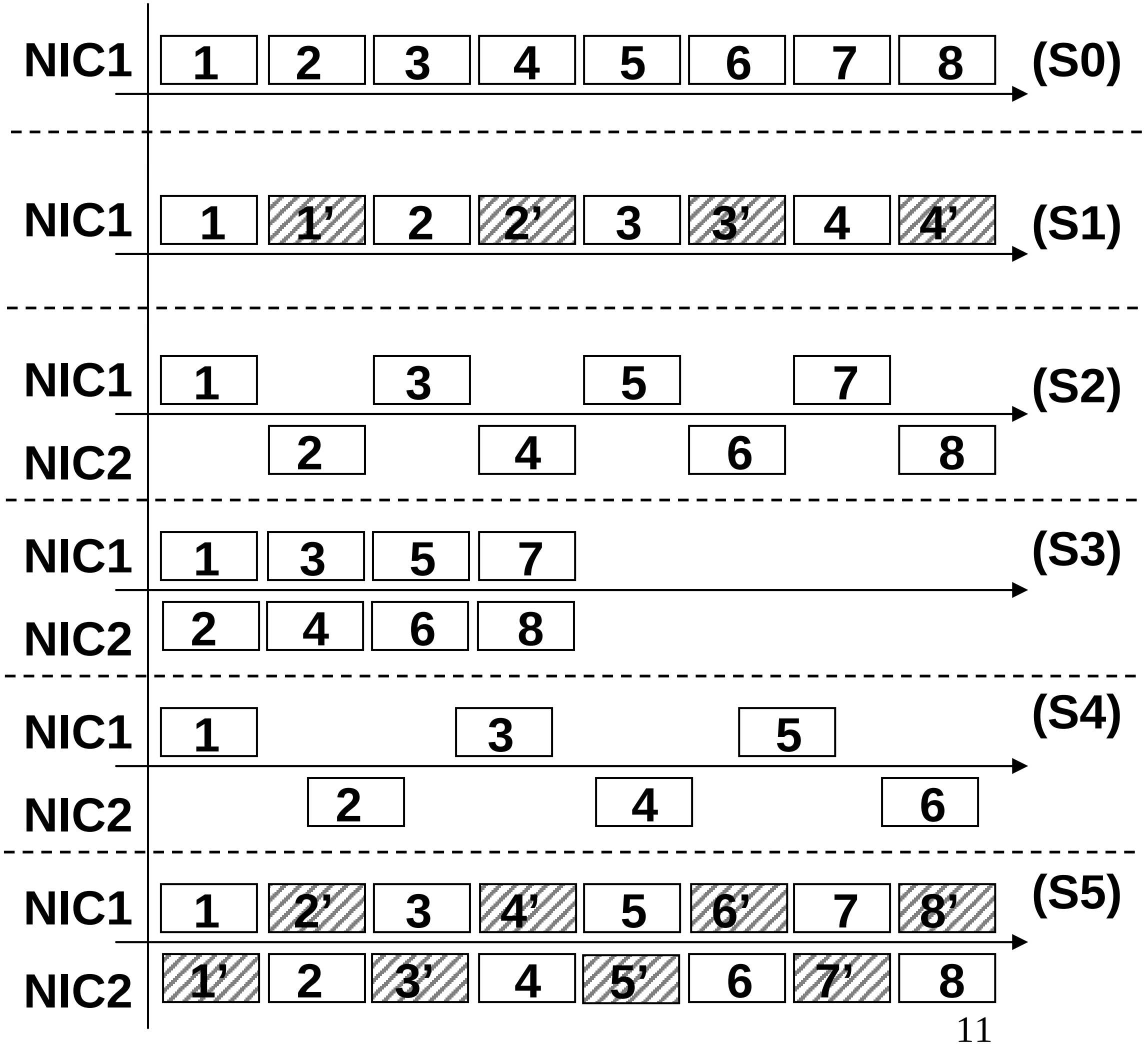

Experimental Approach

- Generate extensive sets of traces
  - "Continuous" transmissions on both NICs
    - 1,000 bytes packets
    - Traces of received packets recorded at each AP
  - Vary
    - Sender location
    - Time-of-day
    - Selection of (non-overlapping) frequency bands
  - Additional configuration to "test" for channel correlation
    - Interferer transmitting in "intermediate" band
- Post-processing of traces to test performance under different configurations/policies
  - Vary coding overhead (N), message size (k), target performance ($P_{\min}$)
  - Explore impact of channel combination, inter-leaving, transmission policies (sticky policies to overcome switching overhead)

But First, What Does an 802.11 Channel Look Like?

- Answer: It's all over the place…
- There is no "average" 802.11 channel
  - Stationary G-E model not particularly accurate
  - Significant time-of-day and location dependent variations
- Across 10 minute intervals, channel characteristics fluctuate widely
  - LTER can range from 0.01% to 70%
  - EBL varies between 1 and 40 packets
    - Actual error bursts were between 1 and several hundred packets
- Similar observations made by others

- Question: What remains of the theoretical "findings" on the benefits of diversity?

Exploring What Remains of the Benefits of Diversity

- A three-prong approach

1. Known channel characteristics are available to identify (and use) “optimal” diversity code
   - Assesses impact of 802.11 channel fluctuations and effect of end-system behavior
2. Unknown channel characteristics – Short-term
   - Explore benefits, if any, of diversity for different levels of coding overhead over short (10mins) “adaptation” periods
3. Unknown channel characteristics – Long-term
   - Evaluate advantages of systematic use of diversity versus single-channel transmissions

- Consider various configurations
  - Channel combinations, sticky policies, etc.

Benefits With Known Channels
Two "Average" Channels

- Channel characteristics:
  - $LTER_1$~ 11%, EBL_1~ 11 pkts
  - $LTER_2$ ~ 10%, EBL_2~ 5 pkts
- Benefits can be substantial *IF*
  - Performance target P_{min} is a constraint
- Interleaving does not seem to help much
  - Conjecture: Operation of the 802.11 protocol itself already creates small "gaps" between packets

k = 10 packets

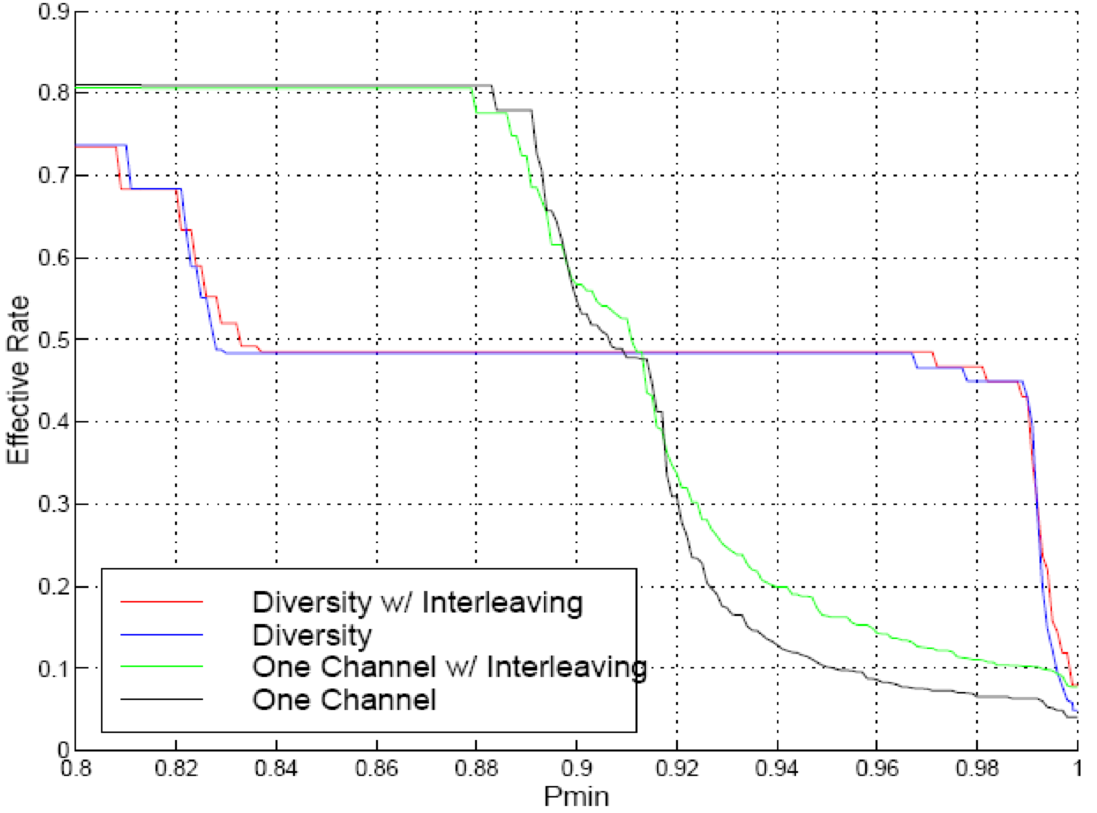

Effective Rate (*ER*) is relative to maximum transmission rate

Things Look Even Better When One Channel Is Bad

- Channels:
  - $LTER_1 \sim 11\%$
  - $EBL_1 \sim 5$ pkts

  - $LTER_2 \sim 66\%$
  - $EBL_2 \sim 22$ pkts

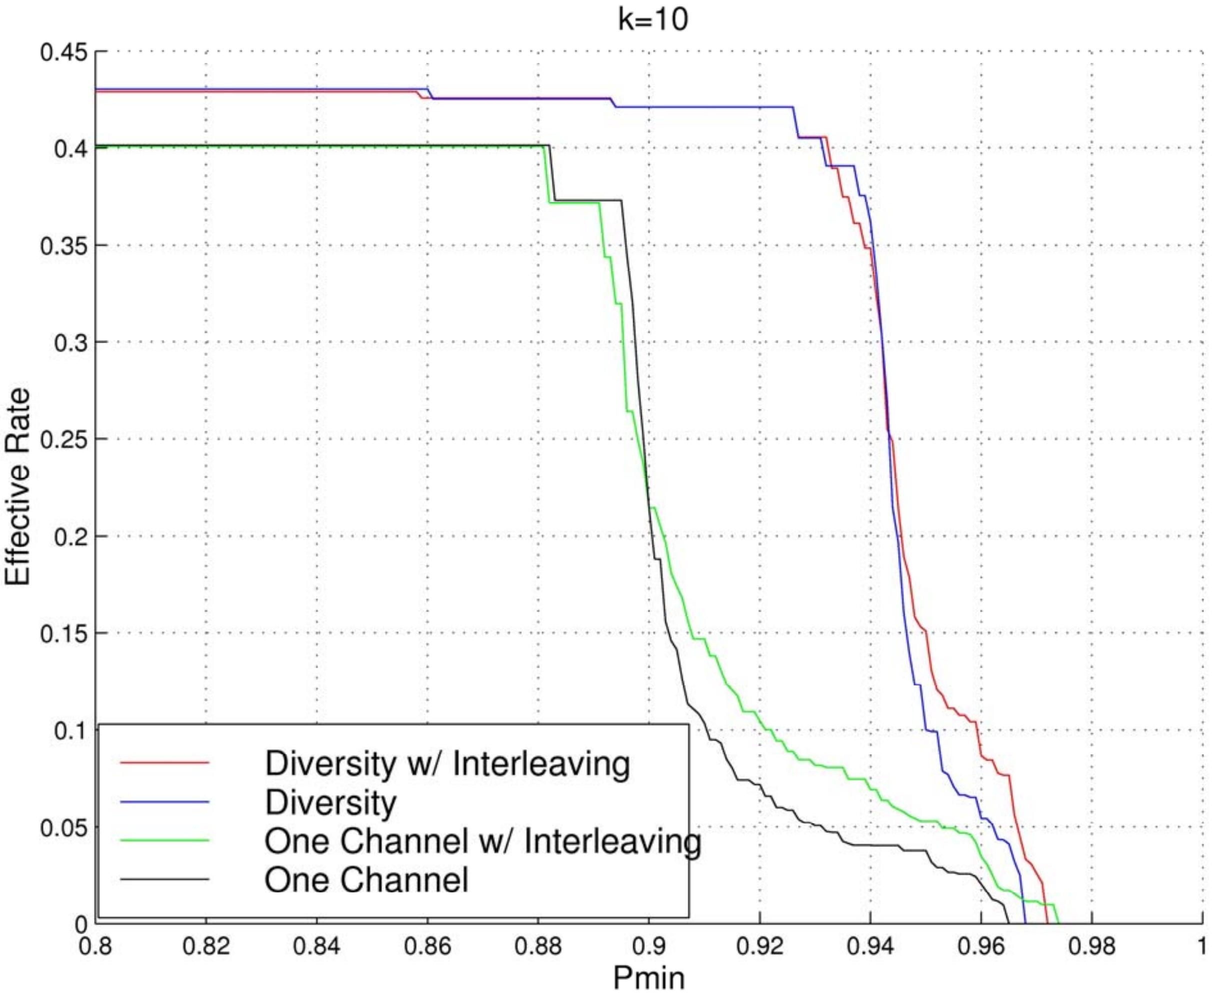

And Vice-Versa When One Channel Is Good

- Channels:
  - $LTER_1$ ~ 11%
  - EBL_1 ~ 10 pkts

  - $LTER_2$ ~ 4%
  - EBL_2 ~ 1 pkt

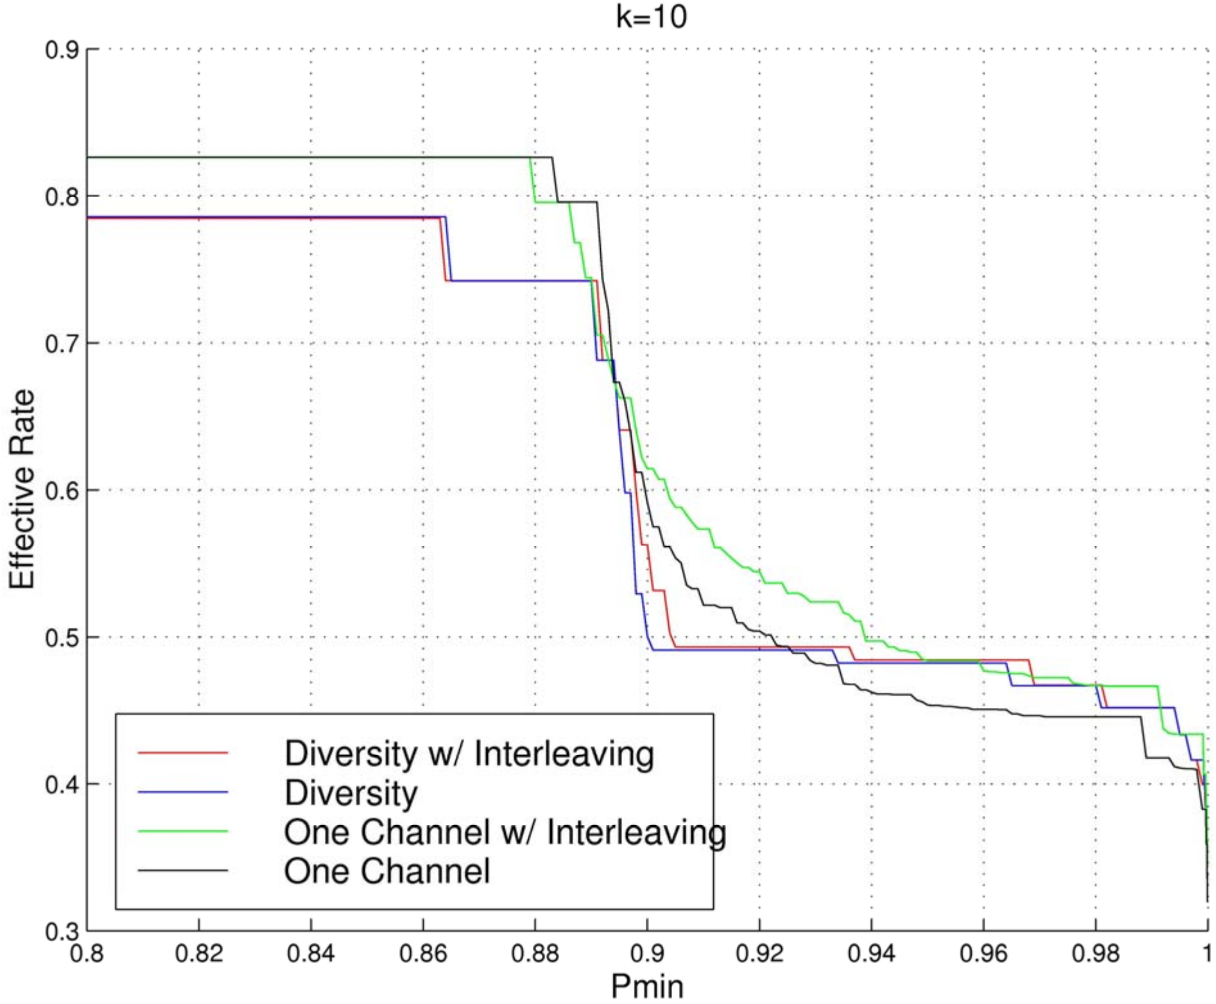

Taking Stock of Where We Stand

- Assuming that
  - We “know” the channel statistics, and
  - Our performance requirements stress channel quality
- Diversity seems to still be potentially helpful in spite of
  - Simplistic (not always optimal) transmission policy
  - Highly variable channel characteristics
  - Lack of precise transmission timing

- So lets now drop our assumptions
  - We know nothing about channel statistics, but
  - We are willing to pay some coding overhead (insurance premium…)

Unknown Channels
One Average, One Bad

- Channels:
  - $LTER_1$ ~ 11%
  - EBL_1 ~ 5 pkts

  - $LTER_2$ ~ 66%
  - EBL_2 ~ 22 pkts
- *Qualitatively* similar results as with known channels, but *quantitatively* quite different
  - Improvements now limited to increasing probability of successful message delivery

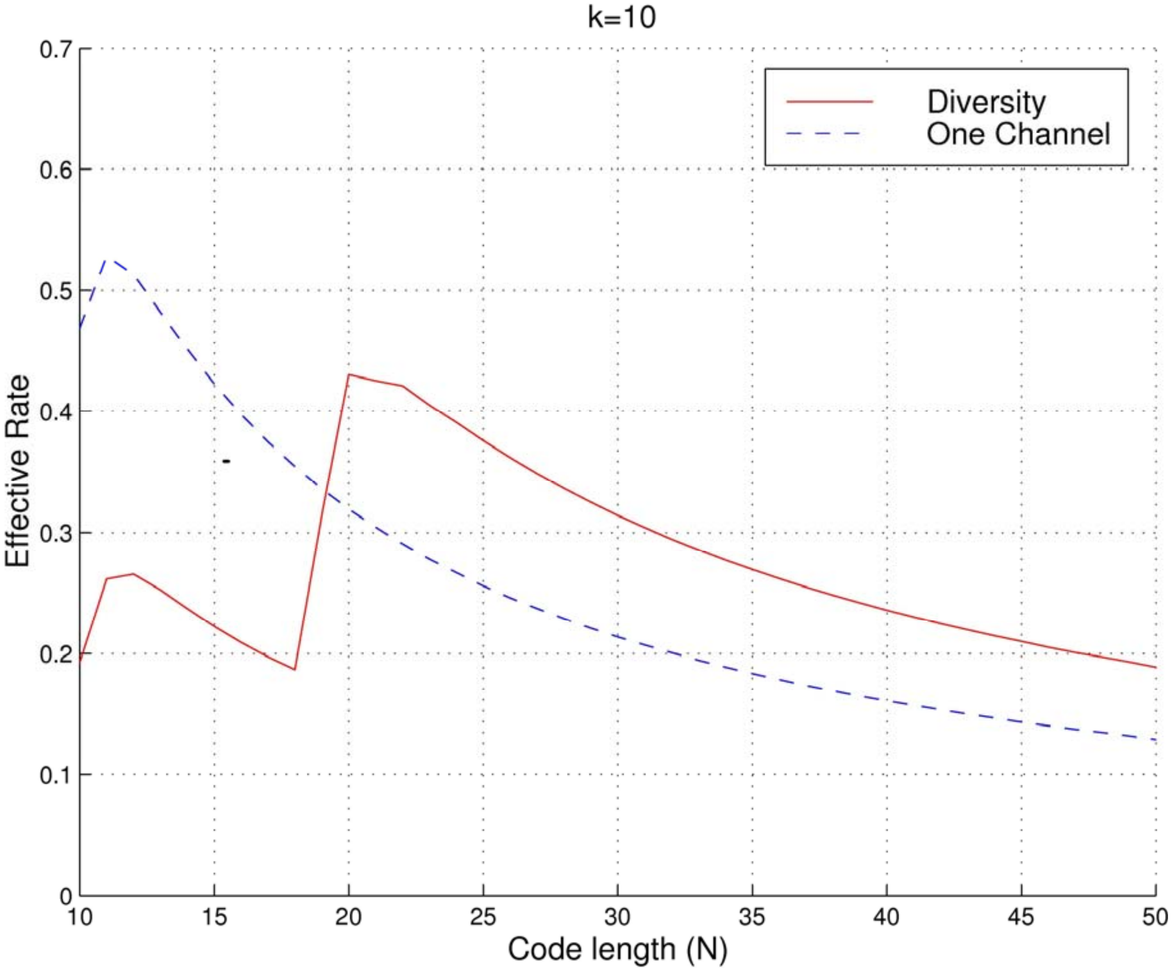

What Can We Conclude So Far?

- In spite of the 802.11 channel fluctuations and the impact of end-systems behaviors, the benefits of diversity (mostly) remain
  - Especially so for reasonably stringent performance requirements
  - Even for unknown channels, diversity rarely “hurts”
  - And the benefits are most visible when one of the channels is BAD

- So, how often are 802.11 channels bad, and can diversity help survive bad channel periods?
  - Monitoring the benefits of diversity over “extended” periods

Diversity As a Performance Stabilizer

- Two channels:
  - $LTER_1$ ~ 11.4%
  - EBL_1 ~ 10 pkts
  - $LTER_2$ ~ 29.2%
  - EBL_2 ~ 11 pkts
- *ER* measured over a 200 messages sliding window
  - Mean value improves by 6%/30%
  - Variance decreases by 60%/90%

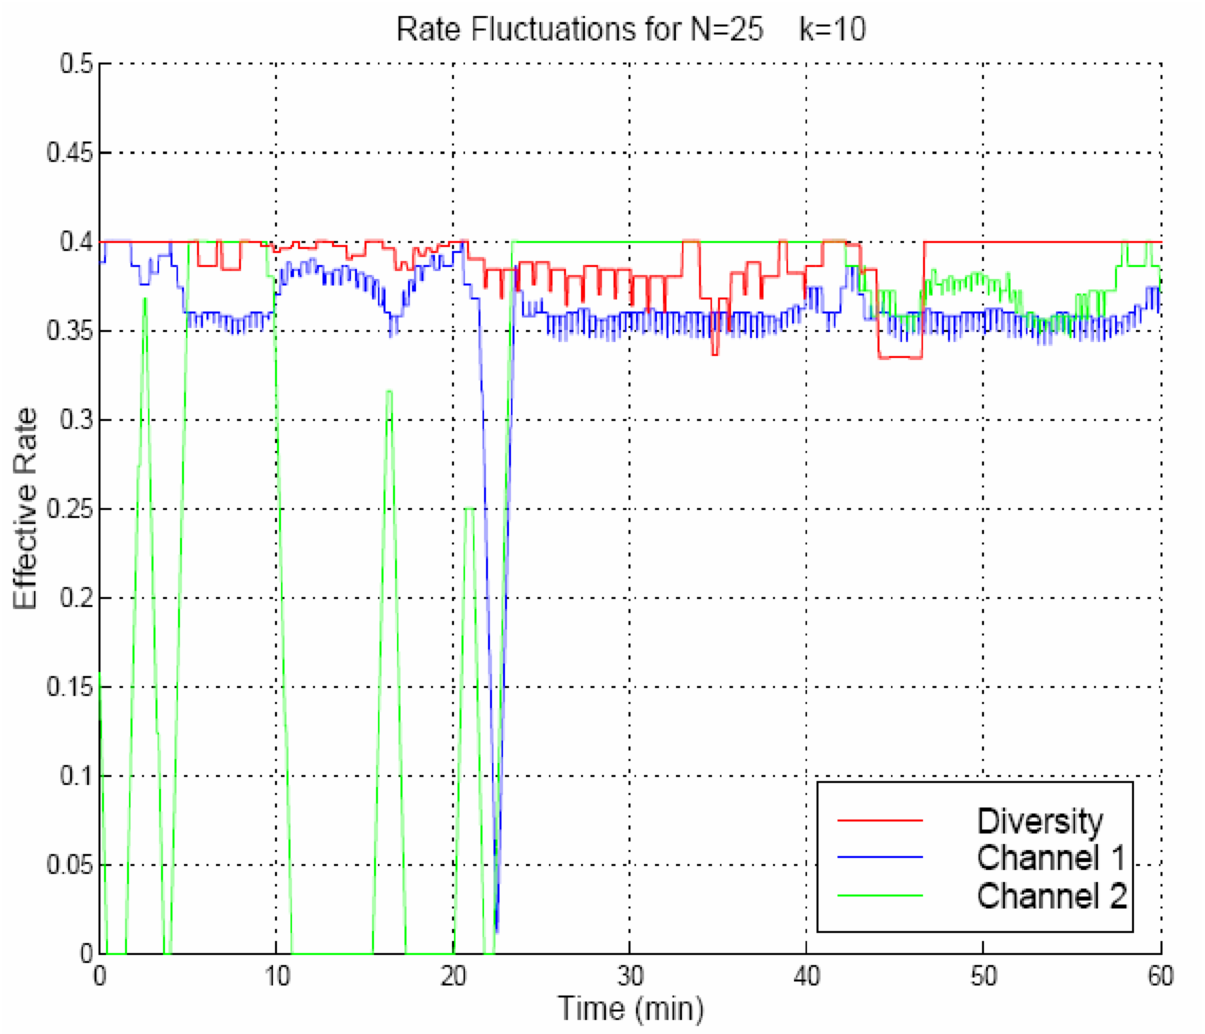

Emulating Sender Configurations That Involve Only One NIC

- Motivation
  - Don’t want twice the hardware
  - Frequency agility is “possible”
    - Use a single NIC to tune to different frequency bands
  - But, today’s channel switching times are high
    - Currently about 25-30 msec, which corresponds to ~23-27 packets
- Approach: “Sticky” transmission policies
  - Channel selection applies to block of packets
  - But, there is a trade-off
    - Large blocks minimize overhead, but diminish burst evasion capability

Sticky Policies

- Basic conclusions are as expected
  - Sticky policies can help realize the “right” trade-off and achieve some of the benefits of diversity in spite of switching overhead
- But
  - Benefits drop-off fast unless overhead is of the order of a packet transmission time or less
- And
  - We have not accounted for any overhead related to APs association if required

Last But Not Least, What About Channel Correlation?

- Correlated channels all but eliminate the benefits of diversity…
- Our investigation suggests that non-overlapping 802.11 frequency bands are reasonably uncorrelated
  - Correlation coefficient between 0 and 0.1
    - Similar to observations by others
  - Similar findings in the presence of a man-in-the-middle interferer
- A few potential reasons for this.
  - Non-overlapping frequency bands are sufficiently far apart
  - As pointed out by others, multipath fading appears to be the dominant source of errors

Summary

- Our initial investigation indicates that
  - Given our performance metric of *message* rate and a willingness to tolerate a non-negligible coding overhead, then
  - Diversity is a reasonable “insurance policy” against the wide range of fluctuations that 802.11 channels experience
- In practice, two NICs are required to take advantage of it
  - Sticky policies offer a possible alternative with a single NIC, but
    - Lower switching delays than what is currently feasible are still needed
    - And a number of open issues remain

Extensions & Open Issues

- Consider other performance metrics than message rate
  - Adaptive applications such as TCP
- We assumed that user transmissions did not affect channel errors – but what if everyone uses diversity?
  - Impact of diversity on channel collisions in 802.11 (and others) systems
- We focused solely on open-loop policies
  - Some (simple) feedback information might provide meaningful improvements
  - What feedback for what improvement?

- And then there are quite a few things to take care of to truly build a system that *hides* all the details of diversity from the users