Packet-Level Diversity
Basic Results and Initial Experiments*

R. Guérin
University of Pennsylvania

* joint work with S. Sarkar and E. Vergetis

Broad Problem Setting

- An increasing number of transmission options
  - Wireline and wireless
- Basic question
  - Can we take advantage of that diversity to improve communication "performance"?
- Numerous approaches and associated challenges
  - Physical (MIMO, OFDM, etc.), MAC (channel allocation), Network (path diversity) layer solutions
  - Closed-loop vs. open-loop

Partial List of Related Works

- Overview
  - Diggavi et al. (2004)
- Open-loop
  - Golubchik et al. (2002) and Abdouni et al. (2005)
  - Tsirigos & Haas (2004)
- Closed-loop
  - Chandra et al. (2004), Miu et al. (2005)
- Diversity codes and physical layer
  - Rabin (1989), Ayanoglu et al. (1993), Biersack (1993), Shacham & McKenney (1990)
  - Laneman et al. (2003, 2004), Apostolopoulos (2001), Miu et al. (2003), Mao et al. (2003), Nguyen & Zakhor (2003), etc.

Our Focus

- Open-loop, packet-level solutions (more on this in a minute)
  - Portable across channel "types"
  - Minimal added complexity
- Goals
  - Better understanding of when and what performance benefits are achievable, and *how* to achieve them
  - Experimental validation
    - 802.11 setting where channels are available frequency bands

The (Initial) System Model

- One user, C channels
  - User wants to transmit messages (data blocks) of size k packets with a target probability of successful delivery $P_{\min}$
  - Channel statistics are "known"
    - Long-term error rate (LTER), expected burst length (EBL)
    - User transmissions do not affect channel statistics
    - Channels are independent
- Diversity coding is used to ensure reliable transmissions
  - (N,k) code can recover from any combination of $i \leq N-k$ lost packets
- User distributes packet transmissions across all C channels according to some policy
  - Deterministic and probabilistic policies
- Performance is measured through the *Effective Rate* (*ER*) realized by the user
  - Number of *messages* successfully delivered per unit of time (unit of time is packet transmission time)

Summarizing Models & Notation

- Channels characteristics – Gilbert-Elliot model
  - Long-term error rate (*LTER*); Expected burst length (*EBL*)
  - Two-state Markov chain alternates between Good and Bad states
    - Packets are lost when channel is in Bad state
- Transmission model
  - Fixed size message consisting of k packets
  - Messages sent using (N,k) diversity code, $N \geq k$ is variable
  - Policy A determines which packet is sent over which channel
    - Decisions at the packet granularity
  - $P_{\min}$: Target probability of successful **message** transmission
  - $P^{A}_{succ}(N,k)$: Probability of successful message transmission given a code of length N, and policy A
- Performance model
  - Effective Rate under policy A and code length N $$ER_A(N,k) = \frac{k \cdot P^{A}_{succ}(N,k)}{N}$$

Transmission Policies

- Probabilistic policies
  - Before each packet transmission select channel i, $1 \leq i \leq C$, with probability p_i
  - $\underline{p} = [p_1\ p_2 \ldots p_C]$ characterizes the channel selection policy across C channels
- Deterministic policies
  - For N-packet messages, pre-determine the channel c_i that packet i , $1 \leq i \leq N$, is to be sent on
  - Schedule $S=[c_1, c_2, \ldots, c_N]$ specifies transmission policy for all N packets of a message

Measuring Performance

- Let A and B denote two arbitrary transmission policies
- The **relative gain** in ER is given by

$$G_{ER}(A,B) = \frac{ER_B(N_B, k_B) - ER_A(N_A, k_A)}{ER_A(N_A, k_A)}$$

- Basic performance metric for diversity:
  - Policy A corresponds to no diversity, i.e., user transmits on only one channel
  - Policy B is "a" policy for distributing packet transmissions across the available channels

Why Should We Expect Better Performance?

- Spreading transmissions across multiple channels allows us to
  - Avoid being stuck with a really bad channel
  - Decrease the effective length of error bursts, which can facilitate recovery (fewer lost packets)

- Potential for improvements arises from
  - A higher probability of successful message transmission
  - The ability to satisfy P_{min} with a smaller code length N

In most settings reducing the code length N needed to satisfy P_{min} is what yields the biggest improvement

Identifying Optimal Policies

Probabilistic Policies

- Calculate $P^{A}_{succ}(N,k)$ given the channel characteristics
  - Recursive solution
  - 4-state Markov Chain for two *independent* GE channels
  - For C independent channels, you end-up with a Markov chain with 2^{C} states
- Search through all policies to find optimal selection

Deterministic Policies

- Deterministic schedule allows each channel to be viewed independently
  - Compute statistics of the associated embedded Markov chains (one for each channel)
- Total number of errors is sum of independent random variables (number of errors when using each channel)
  - Use convolution to compute overall probability of success
- Search through all policies to find optimal selection

Computational Challenges

- Computing optimal policies is feasible but complex
- No intuitive and/or computationally tractable "closed-form" solutions
  - Caused in part by discrete nature of problem
  - No consistent behavior of optimal policy
    - Identical channels need not be used equally
    - Bernoulli channels not always preferred over burstier channels
    - More channels does not always improve performance
- Need some "heuristic" that enables us to easily determine whether or not to use diversity and how

Methodology and Results Summary

- **Step 1**: Run through a broad range of channel combinations to identify if/when meaningful performance improvements are attainable
  - Most scenarios that yield "meaningful" ($\geq$ 30%) improvements involve distributing transmissions roughly equally across channels
- **Step 2**: What sets of channels are likely to result in optimal policies that distribute transmissions equally across channels?
  - Concept of *equivalent* channels
  - Necessary but not sufficient condition
- **Step 3**: Simple heuristic to quickly identify if using diversity pays off
  - Necessary and sufficient condition for meaningful gains

Step 1

- Broad range of channel pairs
  - Initial focus on 2-channel scenarios
  - *LTER* ranges from 1% to 9%
  - *EBL* ranges from 1.01 to 20 packets
- Max. benefits when channels are used roughly *equally*

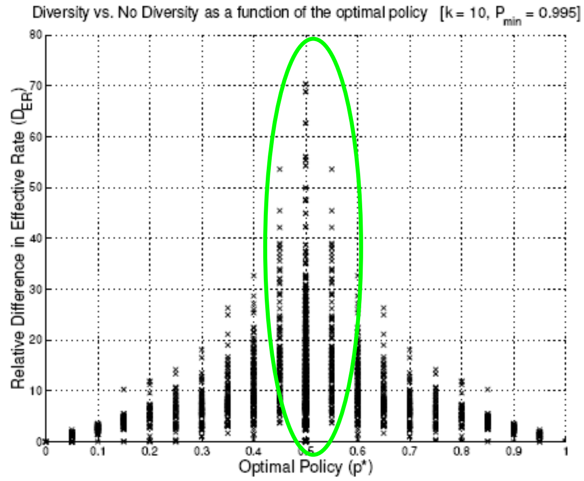

Equal channel use $\Rightarrow$ Focus on deterministic policies

 13

Deterministic vs. Probabilistic Policies

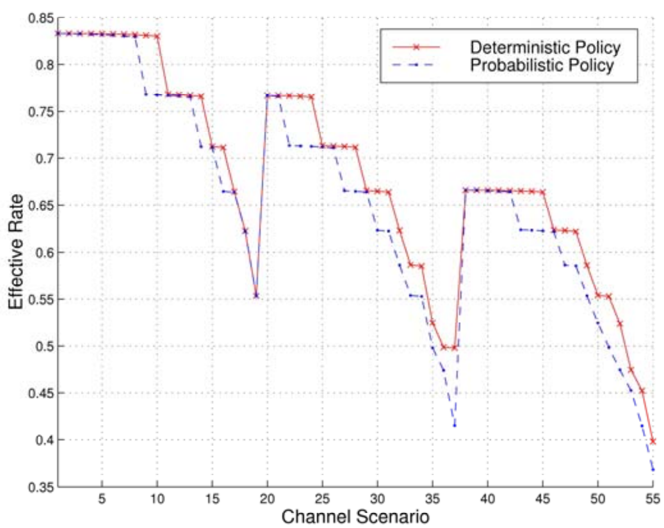

 14

Step 2

- When are channels used equally?
  - This holds for identical channels in most (but not all) settings
    - Any other scenarios?
- Three possible perspectives
  - Used equally under the optimal diversity policy
  - Identical individual performance (*ER*)
  - Produce maximum improvement under the optimal diversity policy
- Interestingly all three perspectives are nearly identical, although not entirely
  - Simple "test" for identifying channels that when paired would be used equally

"Equivalent" Channels – (1)

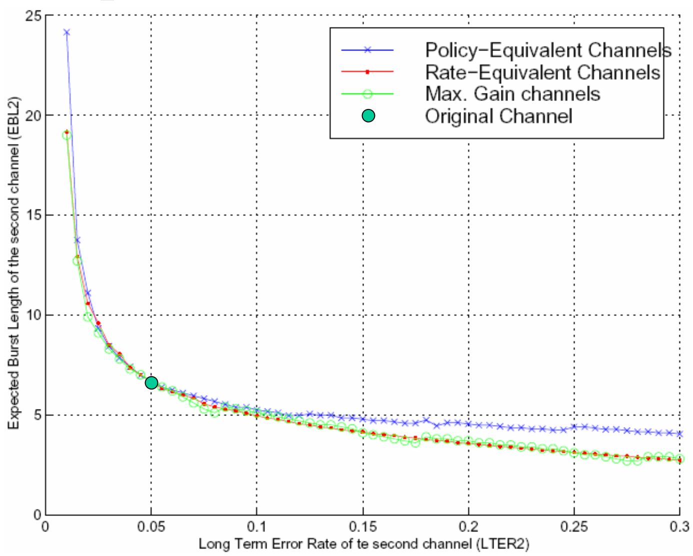

“Equivalent” Channels – (2)

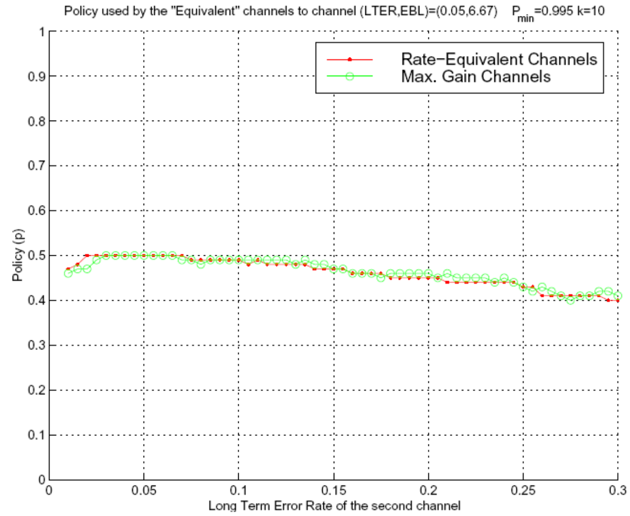

- Optimal policy remains close to 0.5 for “equivalent” channels

The Price of Uniformity

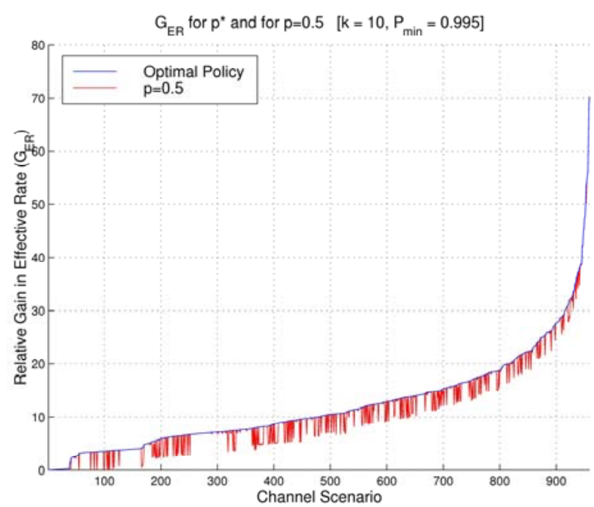

Step 3 - Simple Heuristic

- Given C channels
  - Identify subsets of ~ equivalent channels
  - Group equivalent channels into "equivalence classes"
- For channels in the same equivalence class of size n
  - Compute ER_i, for each channel i, $1 \le i \le n$
- Compute ER_{equal} achieved by cycling equally through all n channels
  - If $\dfrac{ER_{\text{equal}} - \max_{1\le i\le n}(ER_i)}{\max_{1\le i\le n}(ER_i)} \ge \vartheta$ use all n channels
  - Else use best single channel

Heuristic's Loss in Performance Gain (L)

Average	1.28%
Standard Deviation	2.71%
Minimum	0%
Maximum	15.72%
Median	0.03%
L less 5%	81.62% of cases
L less 10%	98.77% of cases
L less 15%	99.04% of cases
L less 20%	100% of cases

Understanding Better When Diversity Helps

- Three parameters of interest:

1. Channel characteristics, i.e., *EBL* and *LTER*
2. Performance target P_{min}
3. Number of channels available

Impact of Channel Characteristics

- 25% performance improvement when combining identical channels with these characteristics
- The higher *LTER*, the smaller the *EBL* needed to achieve a given level of improvement

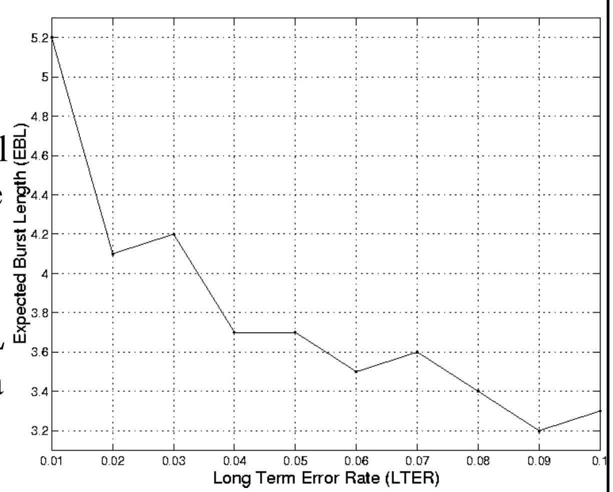

Sensitivity to P_{min}

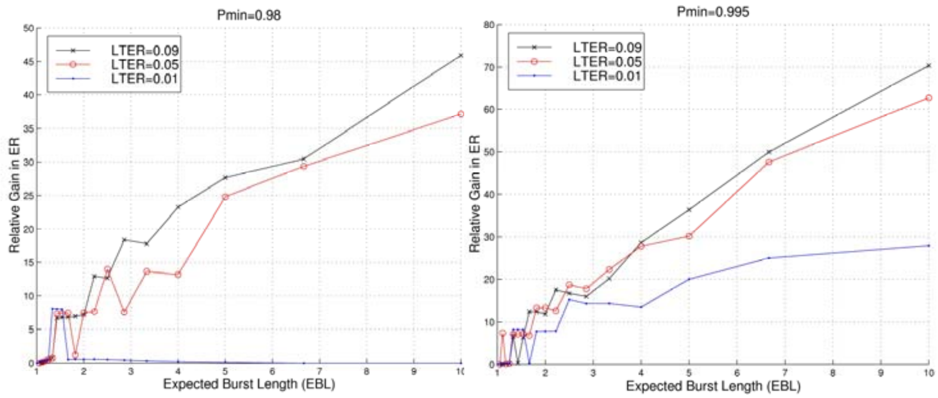

- Potential for improvement increases as
  - P_{min} gets tighter (up to a point – more on this next)
  - *EBL* and *LTER* increase

23

Sensitivity to P_{min} – Bursty Channels

- As P_{min} gets large and N increases correspondingly, the benefits of diversity start decreasing
  - Relative burst size goes down

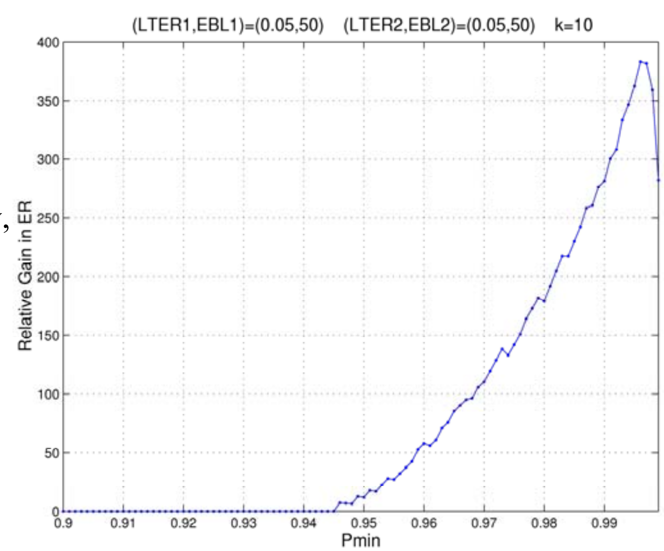

Impact of Number of Channels

- GSM channel scenario
  - Benefits tapers off after 2 or 3 channels
  - Non-monotonous behavior because of discrete nature of transitions

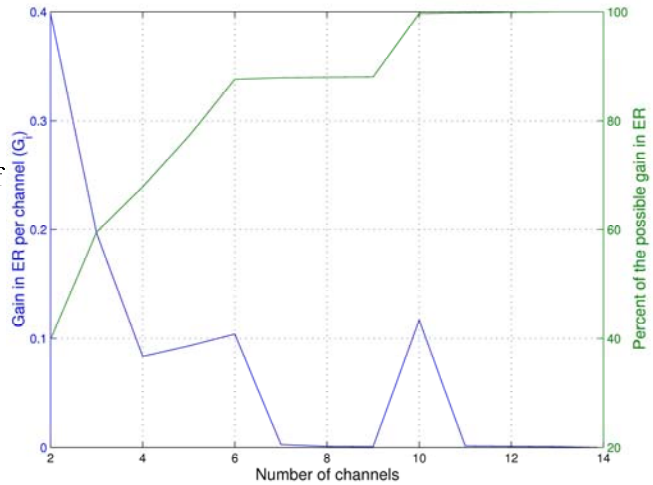

Performance vs. Robustness?

- Two concerns
  1. Are we optimizing ourselves into a corner?
     - Quality of channels can change over time
     - Measurements might be inaccurate
  2. Can I trade-off performance improvements for robustness against channel degradations
- Explore sensitivity to
  - Changes in channel parameters (*EBL* and *LTER*)
  - Changes in distribution of duration of error bursts
    - Impact of the GE channel model
- Investigate relationship between performance improvements and ability to maintain P_{min} over degraded channels

Impact of Channel Degradations

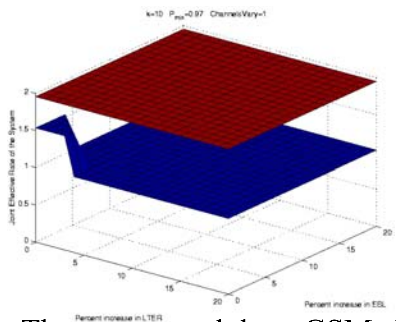

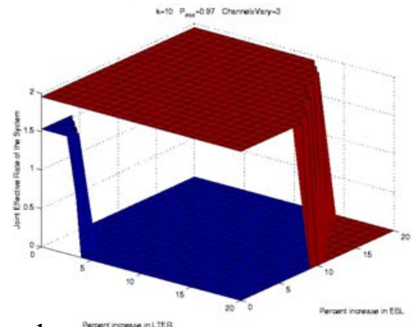

- Three users and three GSM channels
  - Two scenarios: (1) each user is assigned one channel; (2) all three users (optimally) share the three channels
  - Both *EBL* and *LTER* are progressively made worse
    - First on only one channel (left), then on all three channels (right)
- Use of diversity helps improve **both** performance and robustness
  - There some loss of “isolation” in the single bad channel case, but it happens quite late (≥ 40% in both *EBL* and *LTER*)

Trading-Off Performance for Robustness

- System scenario
  - *EBL* and *LTER* are made progressively worse on three channels
  - We vary the code length *N* that the diversity system uses
  - A larger *N* makes the system more robust to errors, but lessens the potential performance improvement under “normal” conditions
- We assess the trade-off between the two

System	D_{ER} Compared to a no diversity system	Percent increase in ***both*** LTER and EBL so that $P_{\min}$ is not satisfied
No diversity (*N*=19)	0%	2%
Diversity (*N*=15)	27.6%	16%
Diversity (*N*=16)	20.7%	37%
Diversity (*N*=17)	14.2%	63%
Diversity (*N*=18)	8.2%	92%
Diversity (*N*=19)	2.7%	> 100%

Impact of Changes in Channel Statistics

Variance Multiplier	No Diversity		Channel Diversity					
	N = 19		N = 15		N = 16		N = 19	
	ER	P_{succ}	*ER*	P_{succ}	*ER*	P_{succ}	*ER*	P_{succ}
Original	1.534	0.971	1.956	0.978	1.850	0.987	1.574	0.997
x 0.25	1.555	0.985	1.947	0.973	1.840	0.982	1.574	0.997
x 0.5	1.547	0.980	1.942	0.971	1.837	0.980	1.568	0.993
x 1	1.538	0.974	0	0.968	1.83	0.976	1.562	0.989
x 2	0	0.963	0	0.962	0	0.968	1.552	0.986
x 4	0	0.961	0	0.949	0	0.957	1.538	0.974
x 8	0	0.953	0	0.941	0	0.949	0	0.966

- We use three users and three GSM channels with $P_{\text{min}} = 0.97$
  - The variance of the error burst periods is varied from 0.25 to 8 times that of the GSM channel using a Gamma distribution (non-Markovian)
- Again diversity allows trading-off performance for robustness

What's Left of All This in Practice?

- Experimental 802.11 setup
  - Two access points assigned non-overlapping frequency bands (Intel Stareast boards)
  - One user with two radio cards (standard laptop)
  - Focus on (blind) equal channel use policies
- Generic issues
  - 802.11 channel model?
  - Controlling the timing of packet transmissions

802.11 Channel "Model"

- No meaningful "model"
- Over 10mins intervals
  - LTER varies from 0.01% to 70%
  - EBL ranges from 1 packet to 40 packets
    - Actual sizes of error bursts went from 1 packet to several hundred packets
  - Statistics were far from being well approximated by a Gilbert-Elliot model
  - Significant time-of-day and location dependent variations
- Consistent with conclusions from earlier measurements

Timing of Packet Transmissions

- S0: Single channel
- S1: Single channel with inter-leaving
- S2: Perfect diversity timing
- S3: Bandwidth constrained diversity timing
- S4: Processor constrained diversity timing
- S5: Bandwidth constrained diversity timing with inter-leaving

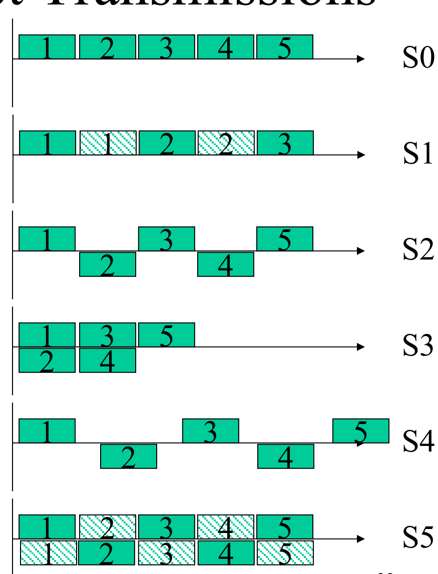

Methodology

1. Assume first that channel characteristics are known and use “optimal” diversity code
   – Goal is to assess if benefits remain in spite of the 802.11 channel characteristics and limited control of transmission timings
2. Assume next that channel characteristics are unknown and assume that the user is willing to incur a certain level of coding overhead
   – When and how much gain does diversity retain
3. Explore long-term stability benefits of diversity
   – Overcome wild fluctuations in 802.11 channels

Known Channel Characteristics

- Multiple channel combinations
  – One average channel and one bad channel
  – Two average channels
  – One average channel and one good channel
- Diversity and no-diversity
  – Inter-leaving and no inter-leaving

One Average Channel – One Bad Channel

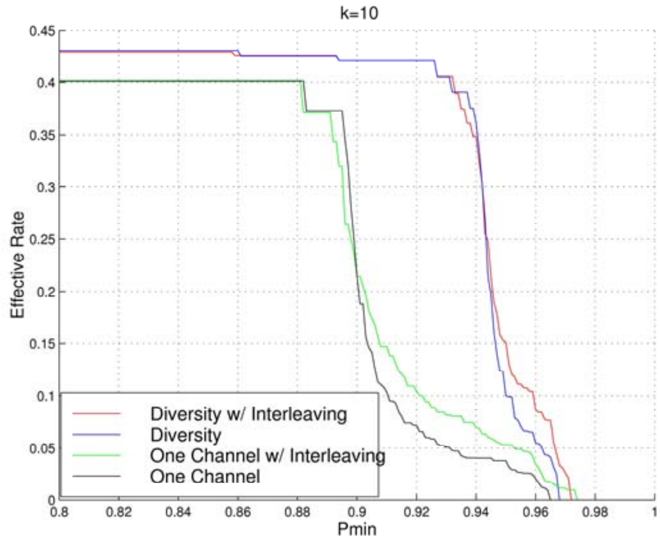

Two Average Channels

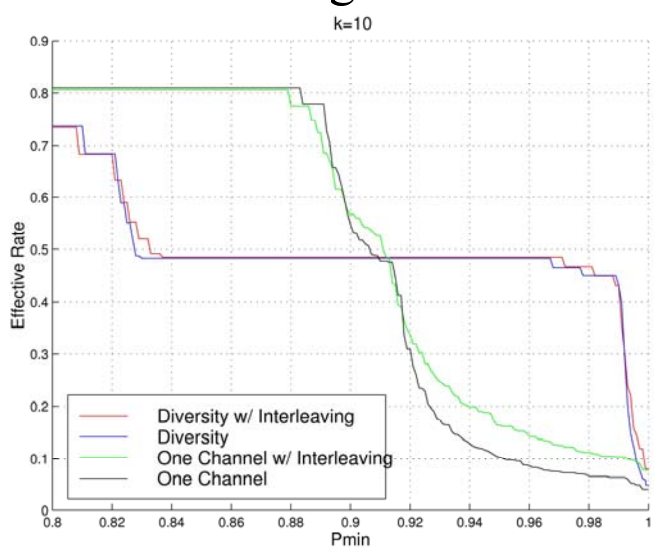

One Average Channel – One Good Channel

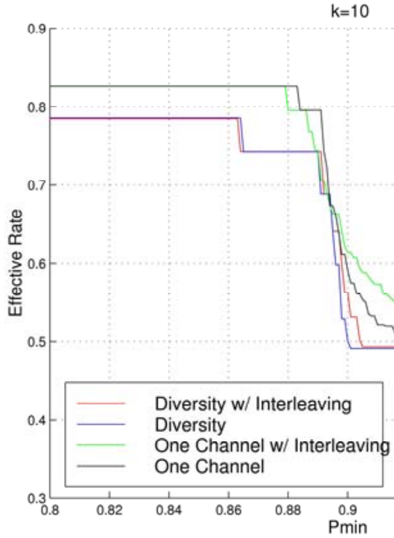

Unknown Channel Characteristics

- Similar channel combinations
  - One average channel and one bad channel
  - Two average channels
  - One average channel and one good channel
- Diversity and no-diversity
  - Varying levels of coding overhead

One Average Channel – One Bad Channel

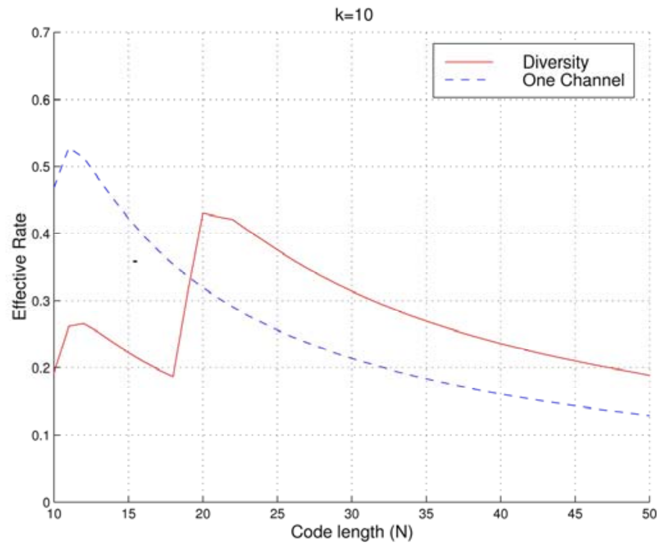

Two Average Channels

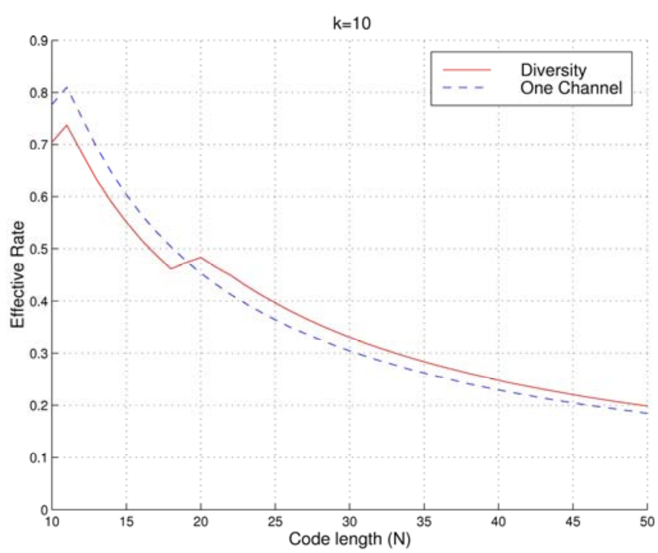

One Average Channel – One Good Channel

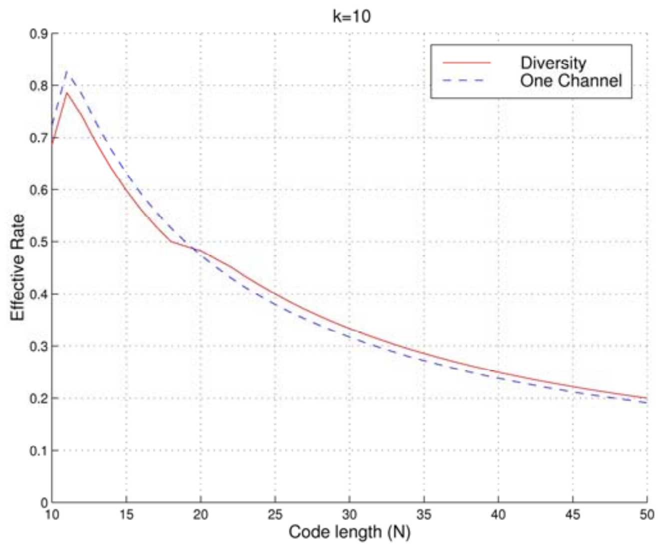

41

Taking Stock of Previous Findings

- Benefits remain in spite of 802.11 channel "characteristics" **IF** channel characteristics are known
- For unknown channels, meaningful benefits remain **ONLY** in the presence of bad channels
- So, how common are scenarios involving one bad channel?

42

Diversity as a Performance Stabilizer

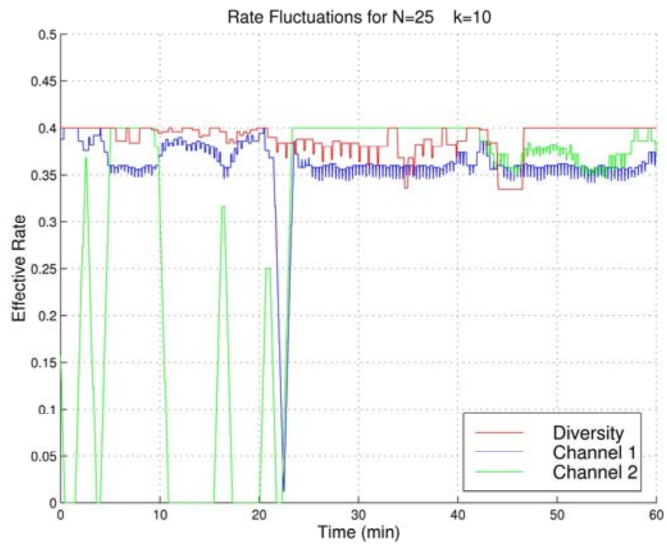

Summary

- Diversity can help even with very little information
  - Meaningful improvements when channel characteristics are "known"
  - Potential for better/more stable performance in highly variable environments, e.g., 802.11, where channel characteristics are hard to predict
- Implementation cost of open-loop policies is minimal, which makes them attractive for wide-spread deployment

Some Ongoing/Future Work

- Impact of channel "stickiness" and packet size
  - Make transmission decisions for a block of *b* packets or make packets larger
  - Reduces the channel switching overhead
  - But, also reduces the ability to avoid bursts
- More general channel models
  - Bit vs. packet level impairments (impact of time-scale)
  - More complex channel statistics, e.g., an 8-state Markov Chain
  - Correlated channels
- What if everybody starts doing it?
  - Access point association scenario
    - Users register with multiple access points (to implement transmission diversity)
    - More users per access point ⇒ greater potential for collision, but
    - More access points per user ⇒ lesser load per user on a given access point

References

1. E. Vergetis, R. Guerin, and S. Sarkar, *"Improving Performance Through Channel Diversity in the Presence of Bursty Losses."* In Proceedings of ITC'19, Beijing, China, August 2005.
2. E. Vergetis, R. Guerin, and S. Sarkar, *"Realizing the Benefits of User-Level Channel Diversity."* ACM SIGCOMM Computer Communication Review, Vol. 35, No. 4, October 2005