



September 2006

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

Evangelos Vergetis

University of Pennsylvania, vergetis@alumni.upenn.edu

Eric Pierce

University of Pennsylvania, pierceer@seas.upenn.edu

Marc Blanco

University of Pennsylvania, marc.blanco@alumni.upenn.edu

Roch A. Guérin

University of Pennsylvania, guerin@acm.org

Follow this and additional works at: http://repository.upenn.edu/ease_papers

Recommended Citation

Evangelos Vergetis, Eric Pierce, Marc Blanco, and Roch A. Guérin, "Packet-Level Diversity - From Theory to Practice: An 802.11-based Experimental Investigation", . September 2006.

Postprint version. Copyright ACM, 2006. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version will be published in ACM MOBICOM 2006, Los Angeles, CA, Sept. 24-29, 2006

Publisher URL: <http://www.sigmobile.org/mobicom/2006/>

This paper is posted at ScholarlyCommons. http://repository.upenn.edu/ease_papers/194

For more information, please contact repository@pobox.upenn.edu.

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

Abstract

Packet-level diversity, or the ability to distribute packet transmissions over multiple, diverse channels, offers a number of benefits in improving communication performance and robustness to variations in channel quality. There have been a number of works that have analyzed and quantified those benefits, and developed transmission policies to realize them. However, translating those potential benefits into practice still faces numerous challenges. These range from uncertainty regarding the adequacy of the channel models on which the development of policies was predicated, to the many implementation constraints one faces when attempting to realize the precise transmission schedules that these policies mandate. This work is an initial step in assessing how much of the benefit that packet-level diversity promises actually remains once confronted with the many practical challenges we have just alluded to. Our investigation is carried out in the context of an 802.11 testbed, where diversity is realized through the different frequency bands available for transmissions between hosts and access points. We use our testbed to evaluate the impact of various parameters, including transmission policies, channel characteristics, channel correlation, and various end-system constraints that affect our ability to precisely control the timing of transmissions. Our investigation reveals that in spite of the many gaps that exist between theory and practice, packet-level diversity still provides a simple solution to improving transmission performance and robustness across a broad range of configurations.

Keywords

Wireless, path diversity, 802.11, experiments

Comments

Postprint version. Copyright ACM, 2006. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version will be published in ACM MOBICOM 2006, Los Angeles, CA, Sept. 24-29, 2006

Publisher URL: <http://www.sigmobile.org/mobicom/2006/>

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

Evangelos Vergetis, Eric Pierce, Marc Blanco, and Roch Guérin^{*}

Dept. Elec. & Sys. Eng., University of Pennsylvania
Philadelphia, PA 19104, USA

{vergetis@alumni, pierceer@seas, marc.blanco@alumni, guerin@ee}.upenn.edu

ABSTRACT

Packet-level diversity, or distributing packet transmissions over multiple, diverse channels, offers benefits in improving communication performance and robustness to channel variations. Previous works have analyzed and quantified those benefits, and developed transmission policies to realize them. However, translating those benefits into practice still faces numerous challenges from uncertainty in the adequacy of the channel models used to develop policies, to implementation difficulties in realizing the precise transmission schedules they mandate. This work is a first step in assessing what remains of those benefits once confronted with such practical challenges. Our investigation is carried out over an 802.11 testbed, where diversity is realized through the different frequency bands available for transmissions between hosts and access points. We use the testbed to evaluate the impact of transmission policies, channel characteristics, channel correlation, and various end-system constraints that affect our ability to precisely control transmissions timing. Our investigation reveals that in spite of the many gaps that exist between theory and practice, packet-level diversity still provides a simple solution to improve transmission performance and robustness across a broad range of configurations.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication

General Terms

Experimentation, Reliability

Keywords

Channel Diversity, 802.11, Open-Loop, Robustness

^{*}This work was supported in part by NSF grants ITR-0085930, CNS-0435306, and NCR-0238340.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiCom '06, September 23–26, 2006, Los Angeles, California, USA.
Copyright 2006 ACM 1-59593-286-0/06/0009 ...\$5.00.

1. INTRODUCTION

The recent past has seen a rise in the transmission options available to end-users¹. This trend is present in both wired, e.g., multi-homing and overlays, and wireless, e.g., Bluetooth, 3GPP, multiple WiFi bands, WiMax, etc., settings. This diversity offers opportunities when it comes to improving transmission performance, because in most cases not all channels experience degradations at the same time. As a result, it is possible to improve performance by “intelligently” distributing transmissions across channels.

In Section 6, we review recent proposals that explore approaches for taking advantage of channel diversity. This paper focuses on solutions that share two key characteristics, namely, channel selection decisions are at the *packet-level* and *open-loop*. Such solutions are attractive because of their “portability” and relatively low overhead. Because of their open-loop nature, they only require cursory knowledge of channel characteristics, and their reliance on packet-level decisions lets them operate across many different physical layers. These advantages notwithstanding, a natural question is how such simple policies compare to more sophisticated schemes, e.g., closed-loop solutions that sense the different channels in an attempt to always transmit on the best one.

Clearly, access to more (channel) information can only improve performance, but it comes at the cost of added overhead and greater dependency on the characteristics of individual physical layers (sensing mechanisms often need to be adapted to each channel type). In addition, a number of recent works that focused on packet-level, open-loop solutions, [1, 8, 17, 19, 20], have demonstrated that they indeed have the potential to offer meaningful performance improvements. This makes them attractive candidates for providing simple solutions capable of delivering benefits across a broad range of environments. In particular, [20] identified numerous channel combinations, where a simple round-robin transmission policy cycling through the available channels afforded throughput increases ranging from 25% to 80%.

This is the starting point for this paper, which focuses on assessing what remains of these benefits once implementation constraints and actual channel conditions are factored in. For example, when different channels are accessible through distinct Network Interface Cards (NICs), precise control of transmission timings, something that is important to even the simplest round-robin, open-loop transmission policy, may not always be feasible. Understanding the im-

¹In this paper, “user” and “sender” are synonymous.

pact this may have on performance, and whether it can be mitigated while preserving implementation simplicity is of interest. Conversely, when all channels are accessed through a single frequency-agile NIC, understanding the effect of the channel switching overhead is also of importance. Furthermore, even when reasonably accurate channel models exist, significant differences between these models and actual channel characteristics are not uncommon. For instance, as noted by various authors, there are no “typical” characteristics for an 802.11 channel, and error rates can fluctuate between 0.1% and 70%, e.g., [22]. Assessing how differences between reality and models affect the performance benefits of diversity is another of our goals.

An exhaustive investigation of all these issues is clearly beyond the scope of a single paper. In this paper, we concentrate on evaluating the benefits of diversity in an 802.11b setting, where senders can distribute their packet transmissions across multiple frequency bands. Our focus is primarily on real-time applications that need a minimum guarantee of successful message delivery, and for which retransmissions are undesirable, e.g., because of latency. In that context, we are interested in several performance metrics. The first assumes that the user is targeting a certain level of transmission reliability, i.e., probability of successful message delivery, and we want to identify whether and by how much can diversity help reduce the overhead required to achieve this level of reliability. Conversely, another metric of interest is whether, given a certain level of overhead, diversity can help an application improve how often it meets its reliability target. We explore these issues in settings that involve multiple NICs, one per channel, and where a major challenge is controlling the timing of transmissions across channels. We also consider environments where a single NIC is used, and where the overhead of switching from one channel to another affects the potential benefits of diversity. Our general findings are that while system constraints and discrepancies between channel models and actual channel characteristics do affect the benefits of diversity, its use remains largely beneficial and can help improve transmission performance and robustness across a wide range of channel conditions.

The rest of the paper is organized as follows. Section 2 describes the system model, transmission policy, and performance metrics we assume. Section 3 discusses various implementation issues and their possible solutions. Section 4 gives a detailed description of our experimental setup. Section 5 presents the results of our experiments and comments on their implications. Section 6 reviews related works, while Section 7 concludes the paper.

2. SYSTEM MODEL

In this section, we introduce system parameters of interest including the simple transmission policy we rely on, and define the metrics we use to evaluate and compare approaches.

We consider an 802.11 system consisting of a sender within reach of several access points operating over distinct frequency bands. The sender can associate with more than one access point as well as determine on a packet-by-packet basis which channel (to which access point) to transmit an individual packet on. In all our experiments, transmissions on different channels are realized by means of separate NICs, with each NIC associated with a given access point. However, we also “emulate” scenarios where all transmissions would take place over a single NIC whose frequency could

be tuned to different bands. In such configurations, the overhead (switching time) involved in tuning the transmitter from one band to another and possibly associating with a new access point are likely to affect the benefits achievable from diversity, and exploring this is also of interest.

2.1 Performance Metrics

As far as performance is concerned, our metric is the *message* transmission rate, where a message corresponds to an application data unit that maps into k network/link layer packets. Because of the possibility of packet losses, redundancy is added by the sender to ensure a probability of successful message delivery greater than a target value P_{\min} . For each message, redundancy is in the form of additional packets using an (N, k) diversity code [3, 4, 13, 16] that guarantees successful message delivery if at least k out of N packets are correctly received. Diversity codes are attractive as they offer simple (packet-level) implementations and reasonable performance. The sender chooses a code length N based on channel characteristics and its target P_{\min} .

We now formally define the message transmission rate or *Effective Rate* (ER), through which we quantify the benefits of the open-loop, packet-level diversity solutions we investigate in this paper. The ER of a given (N, k) code under transmission policy S is the amount of user information successfully delivered per unit of time, where the “unit of time,” is the time required to send one packet. A user message consists of k packets, so that with an (N, k) code, k units of information are sent in N units of time. A message being successfully delivered with probability $P_{succ}^S(N, k)$ under policy S , the corresponding ER is defined as

$$ER_S(N, k) = \frac{k}{N} \cdot P_{succ}^S(N, k). \quad (1)$$

Given channel models and the specification of the transmission policy S , i.e., which channel to use for each packet of a message, it is possible, e.g. see [1, 8, 17, 19, 20], to explicitly compute $P_{succ}^S(N, k)$ for any (N, k) code. However, because one of our goals is to assess the impact of potential discrepancies between “modeled” and real channels, rather than rely on such expressions, we use traces collected in our experiments to compute actual values for $P_{succ}^S(N, k)$ for different code lengths N , with and without diversity.

Note that the “Effective Rate” metric can be easily converted to “throughput,” measured in Mbps, simply by multiplying the maximum channel transmission rate T_s by ER . For example, if $T_s = 7 \text{ Mbps}^2$ and $ER=0.4$, the throughput of the system is 2.8 Mbps. In the rest of this paper we use ER as our performance metric, and use the terms “Effective Rate” and “throughput” interchangeably.

There are two options for comparing the performance of diversity to what is achievable with only one channel. The first compares diversity to the performance achievable using the best single channel. This is, however, somewhat unrealistic, as it requires the sender to know in advance the statistics of all channels. A fairer comparison is to assume that the sender has no advance knowledge of channel statistics, and picks (or get assigned to) one of the available

²In 802.11b, possible transmission speeds are 1, 2, 5.5, and 11 Mbps. However, because of protocol overhead, the maximum feasible throughput T_s is smaller. For example, when an 11 Mbps transmission speed is chosen, the maximum value of T_s is approximately 7 Mbps in practice.

channels at random. In order to take this random channel choice into account, we therefore calculate the performance that can be achieved by using each channel individually, and then use the average of those as the “single channel performance.” In particular, in Sections 5.2 and 5.3 we compare diversity to the single-channel case by averaging the ER obtained from using only channel 1 or only channel 2. In other words, if the throughputs over channels 1 and 2 are ER_{ch1} and ER_{ch2} , respectively, then we define $ER_{one\ channel}$ as

$$ER_{one\ channel} = \frac{ER_{ch1} + ER_{ch2}}{2}.$$

Next, we briefly review several packet-level policies that have been proposed to take advantage of diversity, and motivate the simple policy we consider in this paper.

2.2 Transmission Policy

Transmission policies determine the channel on which to transmit a packet, and can be deterministic or probabilistic. Deterministic policies follow a set schedule that determines ahead of time the channel on which each packet transmission is to take place. Probabilistic policies randomly select a channel for each packet transmission according to a pre-defined probability vector. Both types of policies are relatively simple to implement, even if as discussed in [20], probabilistic policies extend more readily to transmission schedules that do not use all channels equally. The goal of any policy is to maximize performance, i.e., ER , as a function of the (N, k) code in use and assuming some knowledge of channel characteristics, e.g., long-term error rates (LTER) and other relevant statistics³ such as the expected burst length (EBL) or average number of consecutive lost/corrupted packets. A number of earlier works, e.g., [17, 19], describe how to compute optimal policies that maximize ER for certain types of channel models. However, they also show that in most instances where diversity yields “meaningful” improvements⁴, a simple policy that deterministically schedules packet transmissions in a round-robin fashion across the available channels performs close to optimally. This is of practical significance, as even if optimal policies are reasonably robust to variations in channel characteristics [20], they are nevertheless dependent on the underlying channel model and statistics used to derive them. Given the wide range of channel behaviors that have been reported for 802.11 channels, e.g., see [22] and our own results of Section 5.1, the availability of a single, simple policy that seems to perform reasonably well across a wide range of channel scenarios is obviously desirable. For the rest of this paper we therefore focus on a round-robin deterministic policy that simply cycles through the available channels. Determining whether its use indeed delivers some benefits in practice is one of the goals of this paper.

Another dimension of any transmission policy, including the simple round-robin policy we focus on, is the granularity at which it makes decisions. In particular, should new channels be selected for each packet transmission, or should channel selection decisions extend to blocks of consecutive packets? We refer to this dimension as the “stickiness” of the

³Recall that open-loop policies assume that there is no real-time knowledge of the channel state.

⁴For sake of argument, we deem the benefits of diversity to be meaningful if they afford an improvement of more than 25% in Effective Rate.

policy. In practice, sticky policies may be needed to amortize over multiple packets the latency of switching from one channel to another. As mentioned earlier, the senders used in our experiments are equipped with multiple (two) NICs so that switching from one channel to another can be realized simply by selecting a different NIC, which incurs little or no overhead. However, as described in Section 5.5 it is also possible to rely on this setup to emulate an environment involving a single tunable NIC, and explore how different levels of switching overhead (latency) affect the performance of diversity schemes. Hence, we also consider policies of various levels of stickiness to explore the trade-off that exists between leveraging diversity by distributing packet transmissions across as many channels as possible, and the cost associated with channel switching operations.

3. IMPLEMENTATION ISSUES

Before presenting the specific results of our investigation, we step back and comment on several generic implementation hurdles one faces when designing a system capable of realizing the benefits of channel diversity. The main aspect of interest and one that is likely to be present and affect the behavior of any system, is controlling the timing of transmissions that diversity policies implicitly assume. Specifically, the benefits of diversity mostly arise because of its ability to effectively “break up” extended periods of error bursts. This is best understood by means of an example. Consider a scenario involving a sender using a $(20, 15)$ code, i.e., a 15-packet long message that can be recovered as long as at most 5 out of the 20 transmitted packets are corrupted. Two channels are available for transmission, one of which is experiencing an error burst lasting for the equivalent of 10 packet transmissions. If the sender properly sequences its packet transmissions to successively proceed on one channel, then on the other, and so on, it will then experience only 5 instead of 10 packet losses, and the message will be successfully recovered by the intended recipient. This success is, however, predicated on the proper serialization of packet transmissions by the sender, namely, the transmission of one packet only begins *after* the transmission of the previous packet has completed. Failure to do so, e.g., if packet transmissions proceed in parallel on both channels, can all but eliminate the benefits. The question that then arises is how difficult such timing control may be in the context of general purpose end-systems that are our primary target for open-loop, packet-level diversity solutions.

For example, for a system with two NICs where packet transmissions alternate between the two corresponding channels, the transmission pattern assumed by diversity policies is as shown in S2 of Figure 1, which also illustrates the range of possible patterns that may arise as a result of end-systems behavior. Specifically, packets are built and scheduled for transmission by the Operating System (OS), and although the OS properly alternates writing successive packets to each NIC, this does not necessarily result in the consecutive transmissions of those packets by each NIC. Instead, both NICs will often be transmitting packets simultaneously. Barring the addition of special purpose timing control, the extent to which this is the case depends on the relative difference between the time it takes the NIC to transmit a packet and the time it takes the OS to write it to the NIC. When the latter is much shorter, then all the packets in a message end-up being queued to the transmit buffers of the two NICs,

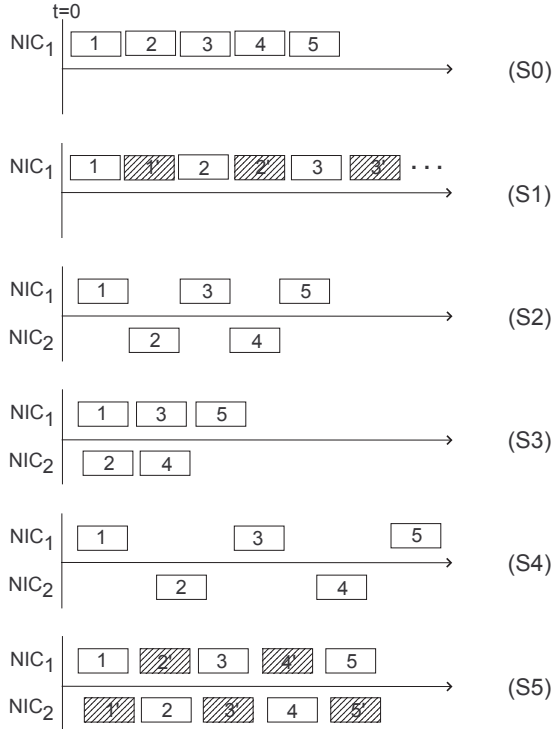


Figure 1: Possible transmission patterns.
S0: Only one channel is used.
S1: Interleaving over only one channel.
S2: “Ideal” transmissions over two channels.
S3: Two-channel “bandwidth-limited” system.
S4: Two-channel “processor-limited” system.
S5: Interleaving in two-channel “bandwidth-limited” system.

so that packet transmissions proceed essentially in parallel. We call such systems “bandwidth-limited” and S3 of Figure 1 illustrates the corresponding transmission pattern. On the other hand, if the time it takes the OS to write a packet to the NIC is larger than the time it takes the NIC to transmit it, then consecutive packets are transmitted one after the other on each channel with the possible addition of gaps between them, whose duration depends on how much slower than the link the processor is. We call such systems “processor-limited” and S4 of Figure 1 shows the resulting transmission pattern⁵. It should be noted that in the context of an 802.11b transmission system, most end-systems are likely to belong to the “bandwidth-limited” category, i.e., the bottleneck is the link rather than the OS/CPU, at least for reasonable packet sizes.

There are a number of possible options for realizing the desired timing of packet transmissions, i.e., transmission pattern S2 of Figure 1. For example, the sender of a bandwidth-limited system can insert appropriate pauses between transmissions of consecutive packets. However, a pausing mech-

⁵Note that a similar behavior is observed if the sender is equipped with only one NIC and uses, for example, Microsoft’s MultiNet software [5] in order to switch between channels. Then the “gaps” of system S4 can be viewed as the time it takes for the NIC to switch to a different channel.

anism implies non-trivial changes to the underlying system. Moreover, the accuracy required to insert pauses of a duration equal to the time it takes to transmit one packet (of the order of hundreds of μ s or one ms) is likely to exceed the capabilities of software solutions and therefore require additional hardware support. Another possibility is for the sender to *interleave* packets from multiple messages (N -blocks). In other words, a sender whose diversity policy calls for distributing packet transmissions of a message across b channels could use b distinct messages and rotate the assignment of their packets across the available channels. This option is illustrated in S1 and S5 of Figure 1. Scenario S1 corresponds to the case of $b = 1$, i.e., no diversity, where the packet interleaving is aimed at emulating some of the benefits of diversity by breaking-up the error bursts that a message could experience through the spreading of its packets over time, i.e., temporal versus spatial diversity. Scenario S5 corresponds to the base diversity case of $b = 2$, where packets of a message are distributed over 2 channels. The major drawback of interleaving, besides the OS or application modifications it requires, is the added latency it introduces. Specifically, packet transmissions at the sender are delayed by the time it takes to generate b messages. Additionally, the receiver now needs to be able to perform more complex message reassembly.

In order to allow for a consistent comparison of the throughput, i.e., ER , achieved across different configurations, we use the following convention. When only one NIC is available, we simply rely on Eq. (1) to compute ER , whether or not diversity is used. When using diversity over one NIC, Eq. (1) is accurate under the assumption that there is no overhead involved in tuning the NIC from one channel to the other. In cases where there is a non-zero switching delay SD , we compute ER by multiplying the result of Eq. (1) by $\frac{s}{s+SD}$, where s denotes the number of packets between channel switching decisions and SD is in units of packet transmission times. The situation is slightly different for configurations involving multiple NICs. For example, in “bandwidth-limited” systems with two NICs, packet transmissions can proceed continuously on both channels, and therefore presumably achieve a rate twice as high as what would be feasible with a single NIC. For consistency, the ER value we report in those cases, is either the value as computed from Eq. (1) if no diversity is used, i.e., transmissions actually proceed only on one NIC, or it is the value obtained from Eq. (1) divided by the number of available NICs (typically 2) when diversity is used. This convention was chosen to facilitate comparisons across scenarios. As we shall see, it often actually penalizes diversity whose total throughput across both NICs often exceed the sum of the throughputs realized by using each NIC individually (see Section 5.5 for an example).

In addition to the relation between processor and link speeds, many other factors affect the timing of packet transmissions. For example, the time it takes a NIC to transmit a packet is typically proportional to the packet size. However, packet size often has little or no impact on the time it takes the OS to deliver a packet to the NIC. Furthermore, the transfer of a packet from the OS to a NIC is itself a complex process that involves several additional components, each with the potential to affect transmission timings. The speed of the hardware bus connecting the system memory (where packets reside) to the NIC, as well as the type of memory transfer supported further contribute to differ-

ences across systems. For example, an Ethernet interface directly on the motherboard will in general receive packets much faster than a miniPCI interface on a daughter board connected to the motherboard via a PMC PCI interface⁶.

Another aspect to consider when implementing and experimenting with diversity, is the impact of link layer mechanisms. In an 802.11b system, there are two mechanisms of relevance. The first is the use of ACKs that trigger link layer retransmissions. Because of our focus on real-time applications and open-loop policies, we disabled the use of acknowledgments⁷. The second relevant mechanism is the RTS/CTS hand-shake used to control channel access. This mechanism can interfere with the scheduling of packet transmissions and add substantial latency. As a result, we configured our end-systems to operate with RTS/CTS disabled (by setting a large enough packet size threshold). Finally, although our focus on packet-level diversity is in part motivated by our desire to develop solutions that are not tightly coupled to the underlying physical layers, it is nevertheless important to understand any physical layer requirement to indeed ensure adequate channel diversity. One such aspect is the need for transmitting antennas to be physically separated by approximately six inches⁸ in order for the power levels not to create interferences between channels. The setup we describe next satisfies this requirement, and antenna separation was never found to be an issue in our experiments.

4. EXPERIMENTAL SETUP

The main goal of our experimental setup was to reproduce an environment as close as possible to what one could expect to encounter in practice if one were to use diversity across multiple 802.11b channels. This is reflected in our choice of systems and configurations. Our sender was based on a standard laptop equipped with two wireless cards, while the access points/receivers it could connect to consisted of two Intel StarEast boards (see below for more details). Our channel environment is representative of an indoor “office” environment, both in terms of propagation and through the presence of other users and access points distributed across the building. Specifically, our sender was placed about ten meters away from the access points (in an adjacent room of the same building), and was not in their line-of-sight. In addition, more than ten other Access Points operating in various channels and creating different levels of interference in all 11 frequency bands where in the vicinity of our setup.

4.1 Sender

Our sender is a Dell laptop (Intel Pentium III 933 MHz, with 384 MBytes of memory) running Linux 2.6.8, and equipped with two wireless interfaces. One is an internal 3Com miniPCI card and the other a Lucent PCMCIA card using the Orinoco driver. Each NIC is set to a different frequency, i.e., a different channel, and a simple user-level program was used to alternate packet transmissions between the two NICs.

The “bandwidth limited” nature of our sender, which results in packet transmissions taking place simultaneously

⁶See the following section for illustrative examples from the experimental setup we used.

⁷This can be easily accomplished without any system modifications, by declaring all packets as broadcast packets.

⁸The actual desired separation depends on the types of antennas used and the power levels of the NICs, but [15] mentions that six inches should be enough for most setups.

over both channels, has implications for the choice of frequencies that can be used on each NIC. Because the two access points associated with each NIC will obviously receive *both* transmitted signals, whether intended for one or the other, it is important that they use non-overlapping frequency bands⁹ to avoid self-interferences caused by the sender’s simultaneous transmissions on the two channels. Ideally, this means that the two channels should be separated by at least 5 frequency bands. In other words, if one of the NICs is, for example, on frequency band 3, then the second NIC should be on frequency bands 8, 9, 10, or 11.

4.2 Access Points

We briefly outline the structure of our access points. We used StarEast stackable systems, which include a baseboard equipped with an Intel IXP425 network processor, and an adapter daughter board with two miniPCI interfaces. The baseboard also includes 133 MHz/256 MBytes of on-board SDRAM, and 32 MBytes of on-board Intel StrataFlash[®] memory, and provides two fast Ethernet ports, one UART, and two mirror PMC PCI interfaces to connect the daughter board¹⁰. The boards use the uClinux 2.4.24 operating system. We used Senao NL-2511MP Plus miniPCI cards, and attached external omni-directional antennas (YSC-RE05T) to achieve satisfactory signal strength. The boards were tested to ensure that the packet “logging” they were required to perform did not degrade their performance, and hence affect the results of the experiments. In other words, both boards were capable of sinking and logging all the packets they were receiving, so that corrupted or lost packets were indeed the result of link/physical layer errors, e.g., because of multipath fading, collisions or insufficient SNR.

5. RESULTS

This section introduces the set of experiments performed over our testbed, motivates the aspects of diversity we explored, and articulates the conclusions we believe those experiments allow us to draw regarding the benefits of diversity when used in a realistic setting. All the results presented in this section correspond to multiple pairs of 10-minute traces collected simultaneously over two different channels. We used 1,000 byte packets, and the sequence numbers of correctly received packets were logged at each access point. Those logs were then combined and post-processed to calculate performance for different combinations of transmission policies and system parameters: (i) using only one channel or both channels, (ii) using different code lengths, (iii) using interleaving or not, and (iv) emulating the behavior of various “sticky” transmission policies. For example, to investigate interleaving, consecutive packets are interpreted as coming from different rather than the same N -block. Similarly, to emulate “sticky” policies, s consecutive packets are

⁹Although the 802.11b standard specifies 11 distinct frequency bands, most are overlapping. Specifically, while there is no precise definition of the “width” of a frequency band, the standard specifies the center frequency of each, along with a spectral mask which requires that the signal be attenuated by at least 30 dB from its peak energy at 11 MHz from the center frequency, and at least 50 dB at 22 MHz from the center frequency. Hence, signals from channels more than 5 frequency “bands” apart should be sufficiently attenuated to minimally interfere with each other.

¹⁰For more details, see Flexcomm’s website.

processed from one log file, followed by s consecutive packets from the other, and so on¹¹. Finally, multiple consecutive 10-minute traces were also combined to create traces that allow us to look at the “average” performance of diversity over longer time periods. We expand on this in Section 5.4.

A first aspect in designing our experiments is to ensure adequate coverage of “all” possible channel conditions. As we shall see and as pointed out by others, the characteristics of 802.11b channels can vary widely. As a result, we conducted experiments across many combinations of frequency bands and over extended periods of time and at different hours of the day. In addition, given our goal of assessing the performance benefits of using “minimalist” diversity solutions, we proceed in steps of decreasing knowledge at the sender regarding channel conditions in support of diversity.

In a first step to verify whether the benefits predicted by the theory remain over real 802.11b channels, we assume that the sender has some knowledge of the channel characteristics, i.e., the sender has reasonable estimates for the values of LTER and EBL prevailing on both channels. Such information could possibly be acquired, at some cost, from gathering over time the statistics of successful packet transmissions. Under such an assumption, the sender is then capable, albeit under some (unrealistic) channel models, e.g., using a simple Gilbert-Elliott model, of determining the coding overhead it requires to meet its performance target (P_{\min}). Our second set of experiments eliminates the requirement that the sender be aware of channel characteristics, which may indeed be a more realistic assumption given the wide fluctuations observed on 802.11b channels. In this setting, we consider the case of a sender willing to pay a certain price in terms of coding overhead, i.e., choose a code length $N \geq k$ given its original message length of k , and we then explore if and when diversity still yields performance benefits, i.e., higher ER values. Our third set of experiments expands on this by further assessing the benefits of diversity by tracking the evolution of the throughput it yields over an extended period of time (hours), and comparing it to the performance that would have been achieved by a user who would have selected a single channel for its transmissions. The results show that the simple diversity scheme we investigate not only increases ER in most cases, but is also successful in significantly reducing its variations over time.

The first three sets of experiments use two NICs, and there is, therefore, no channel switching overhead. In the next set of experiments, we emulate the transmission patterns that would have been generated using a single frequency-agile NIC. In that setting, we investigate the impact of different channel switching latencies, and explore the use of sticky transmission policies to compromise between diversity and switching overhead. We finally comment on the effect of correlated errors between channels, and their potential impact on the benefits of diversity.

In qualitative terms, our results establish that in most cases diversity improves performance, be it in terms of higher throughput, or by reducing throughput variations over time. Additionally, blindly using diversity even when it is unlikely to offer benefits, rarely hurts performance. In other words, it appears that using the simple diversity scheme we investigate is “safe” across different types of channel conditions.

¹¹Recall that s denotes the number of packets between channel switching decisions.

5.1 Channel Characteristics

Before moving on to results that are specific to the use of diversity, we first comment on the general channel characteristics we encountered during our experiments. The key message is that error patterns fluctuate widely depending on several parameters, e.g., time of the day, choice of channel, location of the measurements, distance between the sender and the receiver, behavior of other users, etc. In other words, unlike other technologies, e.g., GSM-based systems, where the quality of the channel is more or less “stable,” one cannot make statements of the form “the *average* 802.11b channel has an LTER of $X\%$ and an EBL of Y packets.” As a result of these large fluctuations and because channel diversity operates within a finite time-horizon, i.e., it “sees” channels at the time granularity of N consecutive packet transmissions, it is unlikely that theoretical performance results predicated on simple channel models can accurately predict the benefits of diversity over 802.11b channels.

Grouping the data collected during our experiments in 10 minutes intervals, we observe LTERs between 0.01% and 70%, and EBLs ranging from 1 packet (essentially a Bernoulli error process) to 40 packets. The actual sizes of error bursts varied from 1 packet to several hundreds of packets. Similar observations have been made by other authors (see for example Willig et al. [22] for extensive 802.11 channel measurements), and we conclude that systems have to be designed with this wide range of characteristics in mind. This calls for either designing diversity systems that account for the “worst case” channel conditions, or for systems that via some feedback mechanism periodically update their estimates of channel statistics. For reasons of simplicity articulated earlier, and to assess what might be doable using a “bare-bone” diversity scheme, we proceed with our evaluation of the simple open-loop transmission policy that rotates packet transmissions across channels. Our findings will show that in spite of the expected “gaps” between theory and practice, the benefits of even this simplest of diversity schemes remain non-negligible.

5.2 Benefits of diversity under “known” channel characteristics

As mentioned earlier, in this first set of experiments we assume known channel characteristics, so that the sender can compute the best code (the one that maximizes ER given a target P_{\min}) for both diversity and no-diversity scenarios. We compare the performance of two systems that do not use diversity, and two systems that do. The first system uses just one channel (S0 of Figure 1). The second one uses one channel but interleaves the packets of two N -blocks on that channel (S1 of Figure 1). The third system uses diversity across two channels (S3 of Figure 1), and the fourth system not only uses diversity across two channels but also interleaves the packets of two N -blocks (S5 of Figure 1).

We present results that cover a fairly broad range of channel characteristics and combinations, e.g., channels whose error rates vary from approximately 4% to 66%, by using three combinations of channel pairs: one “average” and one “bad” channel; two “average” channels; one “average” and one “good” channel¹². We also investigated cases with two “good” channels, i.e., error rates of less than 4%, and for reasons that will become apparent later, the benefit of diversity

¹²See the captions of Figures 2-4 for exact channel statistics.

in those cases was fairly small. The left hand side (LHS) of Figures 2-4 shows the throughput achieved by the diversity and no-diversity systems as a function of P_{\min} . Note that for the no-diversity systems, the throughput values reported in the figures correspond to the *average* of the throughput realized over each individual channel. As mentioned earlier, the motivation for using the average across the two channels, is to provide for a fairer comparison that accounts for the fact that in practice users will pick (be assigned to) one of the two channels at random. In addition, given the varying nature of 802.11b channels, it is also likely that over time any one channel will experience the conditions we select for the two channels for the duration of our experiment.

We first notice that interleaving does not appear to have a major impact, i.e., while there are some differences in performance between systems S0 and S1 and systems S3 and S5, they are probably not sufficient to warrant the added latency and message reassembly that interleaving requires. Additionally, while in some cases interleaving over one channel (system S1) offers advantages over the vanilla system S0, its performance is in most cases still far from the performance that can be achieved via diversity (with or without interleaving). One of the potential reasons for this finding is the channel sensing and back-off mechanism of 802.11, which will in some cases add spaces between successive packet transmissions of the same N -block. Further increasing this space via the use of interleaving does not improve the performance substantially. We therefore conclude that even if it attempts to emulate the break-up of error bursts that diversity affords, interleaving over one channel alone cannot yield meaningful performance improvements. As a result, we do not pursue further the use of interleaving in the rest of the experiments. Instead, we focus on quantifying the benefits of diversity over transmissions using just one channel.

In general, the LHS of Figures 2 to 4 demonstrate that diversity is beneficial and can even yield substantial improvements, at least when the target probability of success P_{\min} is commensurate with the quality of the channels, i.e., neither too lax (greater than 90%, which is what most real-time applications typically call for) nor overly aggressive (so as to anyhow require an unacceptably high coding overhead).

In the case where the two available channels consist of an “average” and a “bad” channel (Figure 2), we see that diversity consistently outperforms the single channel configurations. In particular, when $0.90 \leq P_{\min} \leq 0.95$ the improvement is dramatic, e.g., for $P_{\min} = 0.94$ diversity delivers an improvement of more than 400% in ER . For the case where the two available channels are “average” (Figure 3), we see that there exists a transition point, i.e., a value of P_{\min} ($P_{\min} \approx 0.9$) before which diversity can be hurtful. The intuition is that for low P_{\min} a high rate code can be used, i.e., $N \approx k$. Because such high rate codes can recover from only one or at most two lost packets, it is preferable for the sender to stay on one channel and hope that no error burst will occur. However, as mentioned before, scenarios corresponding to such relatively high message loss rates are of limited interest given our focus on real-time applications. The situation reverses when performance requirements become more stringent, i.e., $P_{\min} > 0.9$, where the benefits of diversity can again be as high as 330%. Finally, in the case of one “average” and one “good” channel, the LHS of Figure 4 shows a similar behavior as in Figure 3, albeit with less pronounced differences between the diversity and no diver-

sity scenarios. This is because as the quality of the channels improves, the required coding overhead becomes smaller and so does the potential for any rate improvement.

5.3 Benefits of diversity under unknown channel characteristics

Having established that the benefits of diversity predicted by the theory are still largely present, at least when the user has some rough knowledge of the characteristics of the available 802.11b channels, we proceed next to remove this assumption. Specifically, we consider the more practical setting where because of the widely varying nature of 802.11b channels, the user does not attempt to optimize the selection of a code for a particular set of channel conditions. Instead, it selects a level of coding overhead it is willing to tolerate as a performance “safeguard.” Our goal is then to quantify how such “blind” selection affects the benefits available from diversity. We explore this aspect for different levels of coding overhead deemed acceptable by the user.

The right hand sides (RHS) of Figures 2-4 plot the throughput achieved across different scenarios as a function of the coding overhead that the user selects. The general conclusions one can draw are similar to those based on the LHS of the figures, with one anticipated difference given the use of a fixed code, independent of actual channel characteristics. The similarities are that the use of diversity is again beneficial in most scenarios of practical interest, namely, scenarios that result in acceptable performance, i.e., a code length $N \geq 20$, which is the value required in most cases to deliver a message loss probability below 10%. The magnitude of those benefits is, however, lower than in the LHS of the figures, which included improvements in excess of 400%. Instead, improvements now range from 2% to about 40%, and meaningful improvements are present only for relatively poor channels, e.g., one “bad” channel. This is not unexpected since the fact that the code length is now *fixed* means that the only possible dimension for throughput increase is to improve the probability of successful message delivery P_{succ} . When channels are reasonably good and P_{succ} is close to 1, there is little room for improvements. In summary, blind diversity appears to still deliver benefits when it comes to overcoming poor performance on individual channels, and those benefits do not come at the cost of reduced performance in other configurations. At least not when focusing on configurations that are capable of delivering acceptable performance to applications, i.e., $P_{succ} \geq 0.9$. We explore this aspect further in our next set of experiments that track performance over an extended period of time, during which the quality of individual channels exhibits a wide range of fluctuations.

5.4 Diversity as a performance “stabilizer”

In this section, we explore the benefits that our “blind” diversity policy affords in terms of its ability to mitigate the impact of the high variability of 802.11b channels. Specifically, we select two distinct 802.11b channels that we monitor over a one hour period during which their quality displays substantial variations. Channel 1 is average (LTER=11.4%), except for a few short periods of poor performance. Channel 2 is worse than channel 1 (LTER=29.2%), in that it experiences greater variability, including several extended periods of poor performance. We compare the ER achieved when using both channels according to our blind diversity





Figure 2: Channel Scenario 1: One “average” channel and one “bad” channel. $LT\epsilon R_1=11.34\%$, $EBL_1=5.4$, $LT\epsilon R_2=65.76\%$, $EBL_2=22.1$.





Figure 3: Channel Scenario 2: Two “average” channels. $LT\epsilon R_1=11.39\%$, $EBL_1=11.1$, $LT\epsilon R_2=9.83\%$, $EBL_2=5.0$.


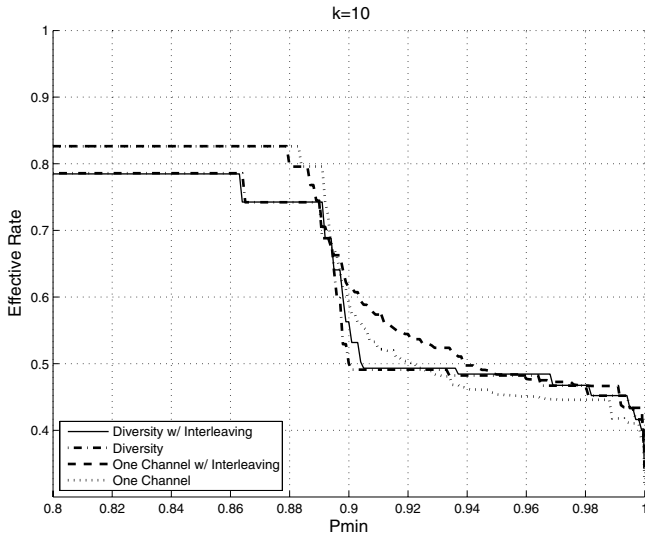


Figure 4: Channel Scenario 3: One “average” channel and one “good” channel. $LT\epsilon R_1=11.27\%$, $EBL_1=10.9$, $LT\epsilon R_2=4.17\%$, $EBL_2=1.1$.

policy, to the ER values realized using either of the two channels alone. In order to capture the evolution of ER over time, we compute its value over a sliding window of the last 200 N -blocks that were transmitted. As in the previous set of experiments, the coding overhead selected by the user is fixed and remains constant for the entire duration of the experiment, i.e., we have $N = 25$ and $k = 10$. Note that this corresponds to a relatively large coding overhead, but we have found this to be typically required to deliver a reasonable probability of success, i.e., $P_{\min} > 0.9$, in spite of the wide range of fluctuations that 802.11b channels experience.

The results are shown in Figure 5 and confirm earlier intuition. In spite of the occasional poor quality of channel 2, diversity yields not only a higher overall¹³ ER , but also more stable performance throughout the experiment. For example, consider a user that decided to only use channel 2. During the first 20 minutes, its throughput would actually be zero or close to zero for fairly long periods of time. This is because the channel is bad enough that a (25,10) code consistently fails to successfully deliver messages (i.e., in every block of 25 packets, at least 15 are in error). In the next 20 minutes, the quality of channel 2 improves dramatically, and as a result it now yields the best performance when used alone. However, its advantage over the diversity solution is minimal, and certainly nowhere near sufficient to make up for the lost throughput of the first 20 minutes. Similar, albeit less pronounced observations hold for channel 1. Note that even better performance could clearly be achieved if the user was somehow capable of always selecting the best channel for its transmissions. Such a dynamic path switching mechanism is certainly possible, e.g., see [12], but as discussed earlier, it involves added cost to monitor the quality of individual channels. Furthermore, our focus in this paper is on investigating the benefits achievable from diversity through a simple open-loop approach.

We performed similar experiments for various other code lengths, and a few representative results are summarized in Table 1. It presents statistics for the average and standard deviation of ER when using diversity, channel 1 alone, and channel 2 alone¹⁴. The table shows that diversity consistently outperforms a single channel system. While increases in ER for the particular codes we chose are not as substantial as the ones reported earlier, the standard deviation of ER is significantly reduced, approximately by 60-90%, when diversity is used. This translates into much more stable performance, which is desirable for most real-time applications.

5.5 Diversity with a single frequency-agile NIC and “sticky” transmission policies

As mentioned earlier, one motivation for conducting experiments using multiple NICs is the channel (frequency) switching overhead when only a single NIC is available. In particular, although device level switching latency in an IEEE 802.11 NIC is approximately $80\mu\text{s}$ [9], there are many other factors that contribute to increasing this value, such as synchronization requirements possibly including association handshakes with the new access point. Recently, Chandra et al. [5] reported switching latencies of the order of 25-30 ms,

¹³Note that without the previously mentioned scaling of the ER of the diversity scheme (division by 2), its value (twice that reported in Figure 5) would often significantly exceed the sum of the rates achieved over channels 1 and 2.

¹⁴The column $N=25$ gives statistics for the traces of Figure 5.

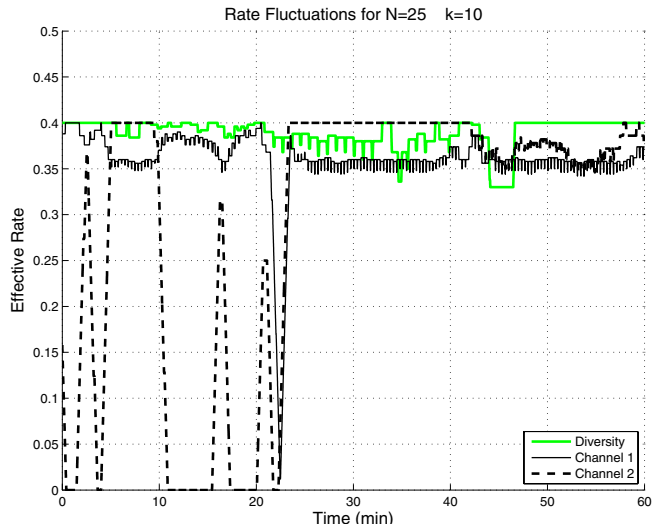


Figure 5: ER of diversity and no-diversity. A (25,10) code is used over a period of one hour.

	N=22	N=25	N=30
Average ER_{div}	0.437	0.381	0.326
Std. Dev. of ER_{div}	0.017	0.015	0.011
Average ER_{ch1}	0.408	0.360	0.300
Std. Dev. of ER_{ch1}	0.046	0.039	0.031
Average ER_{ch2}	0.336	0.296	0.247
Std. Dev. of ER_{ch2}	0.175	0.153	0.126
Increase in Average ER : diversity vs. channel 1	7.1%	5.8%	8.7%
Reduction in $\text{Var}(ER)$: diversity vs. channel 1	63.0%	61.5%	64.5%
Increase in Average ER : diversity vs. channel 2	30.1%	28.7%	32.0%
Reduction in $\text{Var}(ER)$: diversity vs. channel 2	90.3%	90.2%	91.3%

Table 1: Long term performance of diversity vs. no diversity for different code lengths.

which is clearly an overhead that cannot be incurred after every packet¹⁵. One option to enable the use of a single frequency-agile NIC is to rely on a policy that only switches channels every $s \geq 1$ packets. The larger the value of s the less the overhead, but on the other hand a large s also limits the ability of diversity to break-up error bursts (as a matter of fact, a value of $s \geq N$ is equivalent to no diversity). In order to investigate this trade-off, we perform a number of experiments using different values for s , namely, $s = 1, 4, 7, 10$, as well as two channel switching delays SD of one and five packet transmission times¹⁶.

As in Section 5.2, we first consider a scenario where the sender has some knowledge of channel characteristics so that

¹⁵The transmission time of an 1,500 byte packet at 11 Mbps is about 1.1 ms, which means that channel switching time is roughly equivalent to 25 packet transmissions. Hence a policy that transmits successive packets in different channels suffers an unacceptable throughput reduction.

¹⁶This was motivated by the fact that although Chandra et al. report switching times of about 25 packet transmission times, they expect substantial reductions in the near future.


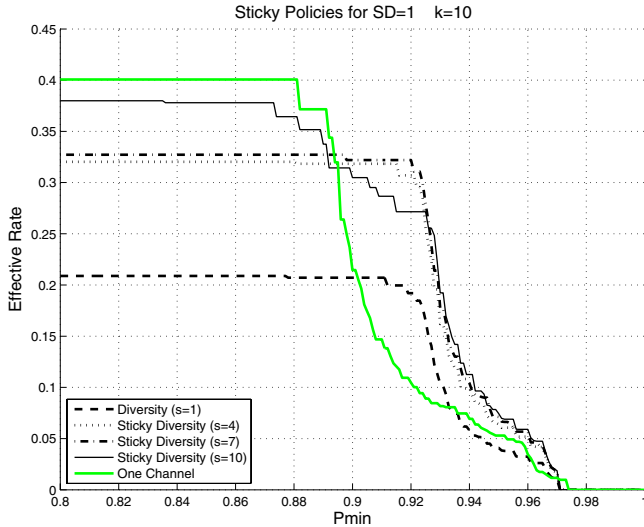


Figure 6: *ER* of “sticky” diversity systems (same channel scenario as in Figure 2).

it can select the best possible code length N given its target P_{\min} . Figure 6 shows the performance achieved by diversity for the two previously mentioned values of SD (LHS: $SD=1$; RHS: $SD=5$), and with policies of different stickiness (values of s). The figure shows that diversity can still realize some of the benefits of Figure 2, but their magnitude is clearly affected by the switching overhead and a non-negligible stickiness is required. To better assess the impact of SD , we consider next the same channel scenario, but now fix the code length N and vary SD . As before, we evaluate *ER* for different stickiness values. For reference purposes, the user’s performance target is set to $P_{\min} = 0.92$, and in each scenario the code length is selected to meet this target assuming zero switching delay. Figure 7 illustrates the intuitive result that when switching overhead is low a “sticky” policy with a small s should be used, while as the overhead grows a larger s performs better as it amortizes switching cost over a longer block of packets. More interestingly, the figure quantifies the impact of channel switching delays on the improvement available from diversity. It shows that even if with the appropriate level of stickiness “some” improvements are present for SD up to about 15 packets, meaningful benefits really call for a sub-ms level channel switching delay.

5.6 Correlated Channel Error Processes

An important issue we have not discussed so far is how correlation in the error processes across channels affects the benefits of diversity. Intuitively, correlated channels should yield smaller benefits, because if error bursts tend to occur on both channels at about the same time, switching channel will not “break up” those bursts. At one extreme, with perfectly correlated channels, diversity offers no benefits.

In an 802.11 system, channels that are more than 5 frequency bands apart are thought to be “orthogonal.” However this does not necessarily mean that their error processes are uncorrelated. For example, consider a diversity system that uses channels 2 and 7. If another user is transmitting over channel 4, these transmissions can create correlated errors on both channels 2 and 7, since they have the same source of interference. The traces we collected during our experiments and the results we presented up to now, might therefore correspond to channels that exhibit some level of

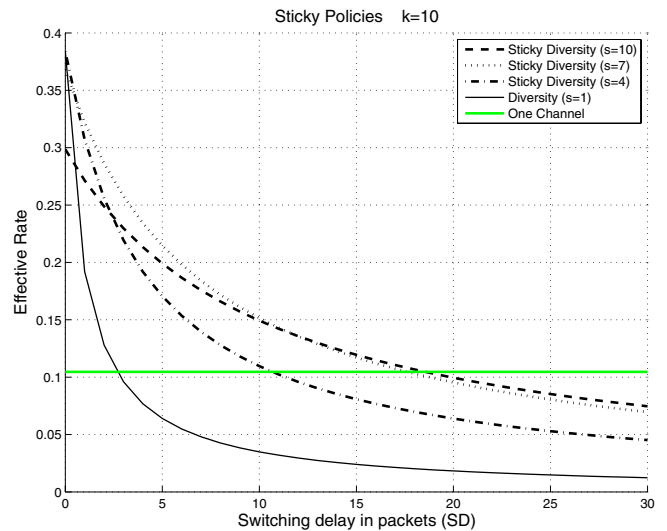


Figure 7: *ER* of “sticky” diversity systems as a function of channel switching time (in packet transmission times) - same channel scenario as in Figure 2.

correlation. To better understand the extent to which this was the case, we calculated correlation coefficients¹⁷ for all pairs of traces, and found them to be consistently very small, i.e., between 0 and 0.1. This lack of significant correlation between 802.11 channels is consistent with what others have reported [2]. A possible explanation offered in [2] is that the dominant cause for losses is multipath fading, and not necessarily interferences from users transmitting in adjacent

¹⁷Denoting the trace for channel 1 as x and the trace for channel 2 as y , the correlation coefficient ρ is defined as

$$\rho = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2} \sqrt{\sum_i (y_i - \bar{y})^2}},$$

where x_i and y_i take the value 1 if packet i is correctly received and 0 if packet i is in error. \bar{x} and \bar{y} correspond to the averages of each trace. ρ takes values in $[-1, 1]$, and when ρ is close to 0, the two traces are considered “uncorrelated.”

channels¹⁸. Willig et al. [22] also make the observation that “multipath fading instead of noise is the dominant source of errors.” Additionally, in order to examine if the performance of diversity would have been different had the channels been truly uncorrelated, we “time-shifted” by 5 minutes one trace from each pair of channels, and processed the traces again. As expected given the small correlation coefficients we had observed, we did not see any noticeable difference between the “correlated” and the “uncorrelated” results.

It is, however, conceivable that in environments busier than ours, e.g., in public libraries, cafés, etc., higher correlation levels might exist even if careful frequency planning of access points within transmission range of each other can limit these effects. Such planning may, however, not always be possible, especially because multiple independent service providers often deploy their networks within the same small geographic area (e.g., different companies occupying different floors of the same building, or a café located close to a university campus) and without much or any coordination. In order to emulate such a scenario, we used an interference generator to create interferences on channel 4, while our diversity system was using channels 1 and 6 for which we then collected packet traces. Even with the interference generator located close to the access points (receivers) we did not observe much correlation between channels 1 and 6. As mentioned earlier, we suspect that this is because multipath is the dominant source of errors, especially when the interfering user is 2-3 frequency bands apart from each of the channels used for diversity.

We therefore conclude that channel correlation is unlikely to be an issue when using diversity in 802.11-based systems, especially when the access points accessible to the user use frequency bands that are sufficiently “far” apart.

6. RELATED LITERATURE

Diversity has received significant recent attention, but its experimental validation has been more limited. In this section, we review some relevant works targeting experimental work in an 802.11 environment, and in particular those aimed at assessing the benefits of channel diversity. We also identify several recent works focused on exploring the advantages of diversity and developing policies to exploit them.

Robinson et al. [15] experiment with a multi-radio mesh network and identify several issues that arise in the deployment of such networks (e.g., interference, antenna separation, etc.). Aguayo et al. [2] deployed Roofnet, an outdoor 802.11b mesh network and present their findings regarding the quality of the links between the nodes. Raychaudhuri et al. [14] present ORBIT, an indoor testbed developed for evaluating wireless protocols. Wu et al. [23] use ORBIT to experiment with rate control and frequency selection in 802.11, and comment on the use of “channel diversity.” However, their use of the term corresponds to neighboring pairs of nodes using different frequencies, which differs from the concept of diversity used in this paper. Vaidya et al. [18] develop a testbed for wireless experiments and comment on the implementation issues that arise in creating realistic environments. Karrer et al. [10] discuss the challenges in building a scalable and widely deployed wireless network, and argue in favor of a multi-hop wireless backbone that uses

directional antennas. Draves et al. [7] discuss routing in multi-radio wireless mesh networks, and develop a 23-node indoor testbed in which all nodes are equipped with two radios. Finally, several other wireless testbeds have been developed [11, 21, 24, 25], but none of them are used to investigate issues related to channel diversity.

More specific to channel diversity, Bahl et al. [5] develop a software-based approach that enables the use of multiple channels using only one 802.11 NIC. As mentioned earlier, the channel switching time is relatively large (around 25-30 ms), which as shown in Section 5.5 makes the use of packet-level diversity solutions somewhat impractical. Miu et al. [12] also present a system that uses path diversity, and evaluate it over experimental testbed. Although the motivation of their work is similar to ours, namely, improve loss resiliency in wireless networks, their approach is fairly different. They use multiple paths to transmit multiple copies of the same frame which get combined at the receiver end to recover from errors, and they also implement a low-overhead retransmission scheme. The approach we investigate does not change the number of packets being transmitted and simply distributes them across multiple channels. In addition, we do not consider the use of retransmissions.

In terms of theoretical investigations of path diversity, [6] provides a comprehensive overview of various relevant issues, from the physical to the network layer. The papers by Golubchik et al. [8] and Abdouni et al. [1], and the paper by Tsirigos and Haas [17], together with our work [19, 20] formulated various approaches to evaluate the performance of diversity for different channel models and policies. However, none of these works focused on experimental investigations.

7. CONCLUSION

In this work, we investigated the potential benefits of simple open-loop, packet-level diversity solutions in a realistic setting by evaluating their performance in an 802.11b testbed. We developed an understanding of many of the implementation issues that are relevant to building such a system, and probed its performance across a broad range of channel scenarios and system configurations.

Our main finding is that in spite of major gaps between theory and practice, substantial advantages can still be attained. In particular, even though the wide fluctuations in quality of 802.11b channels make it impossible to select “optimal” codes tuned for best performance, a simple diversity policy that evenly distributes packet transmissions across channels can still yield meaningful improvements across most code selections. This is especially the case in configurations where “reasonable” performance is actually feasible. Those improvements are not so much in terms of higher transmission rates, although diversity commonly delivers higher effective rates, but more because of the greater *stability* it affords in helping ride out performance variations in individual 802.11b channels.

Our investigation also revealed that unlike what the theory assumes, the benefits of diversity are not overly dependent on the precise timing of packet transmissions. Furthermore, although access to multiple NICs clearly simplifies implementation, diversity solutions could also be realized with a single frequency-agile NIC, *provided* its switching overhead was not too high (lower than what is feasible today) *and* channel switching decisions were “sticky.” Last but not least, we confirmed that as long as they are not too close to

¹⁸We note, however, that [2] used directional antennas, which can limit the level of correlation among channels.

each other, different 802.11b channels exhibit only limited correlation. This is of practical significance, since diversity is obviously of little or no use when channels are correlated.

There are many natural extensions to the initial investigation of this paper. In this paper, the focus was on exploring potential advantages in terms of the maximum rate that real-time applications could sustain, but extending this to adaptive applications such as TCP is an obvious next step. Another natural extension is to involve multiple senders all using diversity. Our current experiments did capture interactions with other users, as we operated in an environment with many other access points besides our own. However, they did not systematically investigate how users all sharing access points through diversity would interact. Finally, another area we are interested in is the introduction of some minimal feedback about channel conditions, to better guide diversity policies and the choice of codes.

8. REFERENCES

- [1] B. Abdouni, W. Cheng, A. L. Chow, L. Golubchik, W.-J. Lee, and J. C. Lui. Multi-path streaming: Optimization and evaluation. In *Proc. of MMCN'05*, San Jose, CA, January 2005.
- [2] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *Proc. of SIGCOMM'04*, Portland, OR, August 2004.
- [3] E. Ayanoglu, C.-L. I, R. Gitlin, and J. Mazo. Diversity coding for transparent self-healing and fault-tolerant communication networks. *IEEE Transactions on Communications*, 41(11), November 1993.
- [4] E. Biersack. Performance evaluation of forward error correction in an ATM environment. *IEEE J. Select. Areas Commun.*, 11(4):631–640, May 1993.
- [5] R. Chandra, P. Bahl, and P. Bahl. MultiNet: Connecting to multiple IEEE 802.11 networks using a single wireless card. In *Proc. of INFOCOM'04*, Hong Kong, China, March 2004.
- [6] S. Diggavi, N. Al-Dhahir, A. Stamoulis, and A. Calderbank. Great expectations: The value of spatial diversity in wireless networks. *Proceedings of the IEEE*, 92(2), February 2004.
- [7] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *Proc. of MobiCom'04*, Philadelphia, PA, September/October 2004.
- [8] L. Golubchik, J. C. Lui, T. Tung, A. L. Chow, W.-J. Lee, G. Franceschinis, and C. Anglano. Multi-path continuous media streaming: What are the benefits? *Performance Evaluation*, 39, September 2002.
- [9] F. Herzel, G. Fischer, and H. Gustat. An integrated CMOS RF synthesizer for 802.11a wireless LAN. *IEEE Journal of Solid-State Circuits*, 38(10), October 2003.
- [10] R. Karrer, A. Sabharwal, and E. Knightly. Enabling large-scale wireless broadband: The case for TAPs. In *Proc. of HotNets'03*, Cambridge, MA, November 2003.
- [11] H. Lundgren, D. Lundberg, J. Nielsen, E. Nordström, and C. Tschudin. A large-scale testbed for reproducible ad hoc protocol evaluations. In *Proc. of WCNC'02*, Orlando, FL, March 2002.
- [12] A. Miu, H. Balakrishnan, and C. E. Koksal. Improving loss resilience with multi-radio diversity in wireless networks. In *Proc. of MobiCom'05*, Cologne, Germany, August/September 2005.
- [13] M. Rabin. Efficient dispersal of information for security, load balancing, and fault tolerance. *Journal of the Association for Computing Machinery*, 36(2):335–348, April 1989.
- [14] D. Raychaudhri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremono, R. Siracusa, H. Liu, and M. Singh. Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols. In *Proc. of WCNC'05*, New Orleans, LA, March 2005.
- [15] J. Robinson, K. Papagiannaki, C. Diot, X. Guo, and L. Krishnamurthy. Experimenting with a multi-radio mesh networking testbed. In *Proc. of the 1st WiNMee'05*, Trentino, Italy, April 2005.
- [16] N. Shacham and P. McKenney. Packet recovery in high-speed networks using coding and buffer management. In *Proc. INFOCOM'90*, San Francisco, CA, April 1990.
- [17] A. Tsirigos and Z. Haas. Analysis of multipath routing-Part I: The effect on the packet delivery ratio. *IEEE Transactions on Wireless Communications*, 3(1), January 2004.
- [18] N. H. Vaidya, J. Bernhard, V. V. Veeravalli, P. R. Kumar, and R. K. Iyer. Illinois wireless wind tunnel: A testbed for experimental evaluation of wireless networks. In *Proc. of the E-WIND workshop (at SIGCOMM'05)*, Philadelphia, PA, August 2005.
- [19] E. Vergetis, R. Guérin, and S. Sarkar. Improving performance through channel diversity in the presence of bursty losses. In *Proc. of the 19th International Teletraffic Congress (ITC)*, Beijing, China, August/September 2005.
- [20] E. Vergetis, R. Guérin, and S. Sarkar. Realizing the benefits of user-level channel diversity. *ACM Computer Communication Review*, 35(5), October 2005.
- [21] B. White, J. Lepreau, and S. Guruprasad. Lowering the barrier to wireless and mobile experimentation. In *Proc. of HotNets-I*, Princeton, NJ, October 2002.
- [22] A. Willig, M. Kubisch, C. Hoene, and A. Wolisz. Measurements of a wireless link in an industrial environment using an IEEE 802.11-compliant physical layer. *IEEE Transactions on Industrial Electronics*, 49(6), December 2002.
- [23] Z. Wu, S. Ganu, I. Seskar, and D. Raychaudhri. Experimental investigation of PHY layer rate control and frequency selection in 802.11-based ad-hoc networks. In *Proc. of the E-WIND workshop (at SIGCOMM'05)*, Philadelphia, PA, August 2005.
- [24] M. D. Yarvis, W. S. Conner, L. Krishnamurthy, and A. Mainwaring. Real-world experiences with an interactive ad hoc sensor network. In *Proc. of ICPPW'02*, Vancouver, Canada, August 2002.
- [25] Y. Zhang and W. Li. An integrated environment for testing mobile ad-hoc networks. In *Proc. of MobiHoc'02*, Lausanne, Switzerland, June 2002.