October 2005

Realizing the Benefits of User-Level Channel Diversity

Evangelos Vergetis

University of Pennsylvania, vergetis@seas.upenn.edu

Roch A. Guérin

University of Pennsylvania, guerin@acm.org

Saswati Sarkar

University of Pennsylvania, swati@seas.upenn.edu

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

Recommended Citation

Evangelos Vergetis, Roch A. Guérin, and Saswati Sarkar, "Realizing the Benefits of User-Level Channel Diversity", . October 2005.

Postprint version. Copyright ACM, 2005. 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 was published in *ACM SIGCOMM Computer Communication Review*, Volume 35, Issue 5, October 2005, pages 17-28.

Publisher URL: <http://doi.acm.org/10.1145/1096536.1096539>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/ease_papers/138
For more information, please contact repository@pobox.upenn.edu.

Realizing the Benefits of User-Level Channel Diversity

Abstract

Channel or path diversity is known to improve performance in physical layer designs, channel access strategies, path switching mechanisms, etc. In this paper, we focus on "user-level" mechanisms that operate simply by distributing packet transmissions across multiple channels. We seek to understand when, why, and to what extent this can be of benefit, and equally important, whether these benefits can be realized with as little of an added cost as possible. In that context, our main contribution is not so much in identifying optimal policies for leveraging channel diversity, but in introducing the concept of channel "equivalence" and demonstrating that channel diversity yields substantial benefits mostly when channels are approximately equivalent. We build on this finding to investigate the robustness of these improvements against errors in the characterization of the available channels or changes in their characteristics. We also explore the sensitivity of the results as the number of available channels varies. The findings of the paper demonstrate that by allowing packet transmissions from multiple users to intelligently share channels, it is possible to improve overall performance and robustness through simple and portable user-level mechanisms.

Keywords

Channel Diversity, Robustness, Cross-Layer, Network

Comments

Postprint version. Copyright ACM, 2005. 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 was published in *ACM SIGCOMM Computer Communication Review*, Volume 35, Issue 5, October 2005, pages 17-28. Publisher URL: <http://doi.acm.org/10.1145/1096536.1096539>

Realizing the Benefits of User-Level Channel Diversity

Evangelos Vergetis
University of Pennsylvania
200 South 33rd Street
Philadelphia, PA 19104, USA
vergetis@seas.upenn.edu

Roch Guérin
University of Pennsylvania
200 South 33rd Street
Philadelphia, PA 19104, USA
guerin@ee.upenn.edu

Saswati Sarkar
University of Pennsylvania
200 South 33rd Street
Philadelphia, PA 19104, USA
swati@ee.upenn.edu

ABSTRACT

Channel or path diversity is known to improve performance in physical layer designs, channel access strategies, path switching mechanisms, etc. In this paper, we focus on “user-level” mechanisms that operate simply by distributing packet transmissions across multiple channels. We seek to understand when, why, and to what extent this can be of benefit, and equally important, whether these benefits can be realized with as little of an added cost as possible. In that context, our main contribution is not so much in identifying optimal policies for leveraging channel diversity, but in introducing the concept of channel “equivalence” and demonstrating that channel diversity yields substantial benefits mostly when channels are approximately equivalent. We build on this finding to investigate the robustness of these improvements against errors in the characterization of the available channels or changes in their characteristics. We also explore the sensitivity of the results as the number of available channels varies. The findings of the paper demonstrate that by allowing packet transmissions from multiple users to intelligently share channels, it is possible to improve overall performance and robustness through simple and portable user-level mechanisms.

Categories and Subject Descriptors

C.2.5 [Computer-Communication Networks]: Access Schemes

General Terms

Reliability, Performance

Keywords

Channel Diversity, Robustness, Cross-layer Designs, Open-Loop Control

1. INTRODUCTION

IP networks nowadays need to provide minimum service guarantees to a range of applications. Such guarantees can be provided by (a) altering the network architecture and protocols, or (b) by utilizing the availability of multiple (diverse) network paths (or channels¹) and the resiliency in the applications. Multiple channels are naturally available in

This work was supported in part by NSF grants ITR-0085930, ANI-0106984, CNS 0435306, and NCR-0238340.

¹Throughout the paper we interchangeably use the terms path and channel.

several networks. For example, in wireless networks multiple distinct frequency “bands” are available for transmission, e.g., in Wi-Fi systems, cellular systems, or sensor networks. We therefore investigate the latter approach as it eliminates the need to change the network, and allows it to remain simple. Specifically, we consider the combination of increased application resiliency through the addition of coding, and the leveraging of channel diversity by making transmission decisions based on channel characteristics. In particular, we focus on packet level transmission decisions, as they are reasonably easy to implement across a broad range of diverse network technologies.

The first question, however, is whether diversity and resiliency schemes introduce significant implementation complexity, or require significant modifications to existing systems. We need to determine what information a user needs to know about the network and the complexity involved in acquiring and using such information. If a user knew the instantaneous states of the channels, it could clearly optimize performance by selecting the best channel. But channel characteristics fluctuate over time, often rapidly, so that a user can only assess the quality of all the available channels by repeatedly probing them before each transmission². The bandwidth and energy consumed transmitting probe packets in different channels can be significant when the system has multiple channels. Furthermore, probe packets sent on a channel can disrupt transmissions from other users on that channel. We therefore consider scenarios where the users know only the statistical characteristics, and not the instantaneous qualities of the channels. The transmission decisions are essentially open-loop and do not depend on feedback information about the actual channel states.

Although the use of “open-loop” mechanisms greatly simplifies the task of implementing a system capable of realizing the benefits of channel diversity, there are still many questions that need to be addressed to determine if it is worth the added complexity. In [26], we demonstrated that it was possible to identify simple open-loop channel diversity policies that offered substantial performance improvements across a broad range of channel configurations. In this paper, we concentrate on several aspects that we believe are key to better assess whether channel diversity can provide a

²A user can estimate the quality of a channel from the feedback obtained after previous packet transmissions in the channel. However, when the number of available channels becomes large, the last transmission in a given channel is increasingly likely to have been a long time back, and therefore the associated feedback will not provide a reliable estimate of the current channel quality.

meaningful solution to improving transmission performance in practice. First, we try to “isolate” the system characteristics that affect the gains due to diversity. Towards that end, we determine the combinations of channel characteristics, the number of available channels, and the minimum desirable performance requirements in which channel diversity yields substantial gains. We demonstrate that such benefits are significant when channels have similar characteristics, or each channel delivers similar performance when used alone. We denote these channels as “equivalent channels,” and show that diversity strategies that use these equivalent channels for equal amounts of time perform close to optimal. As we shall discuss, this offers a number of advantages when it comes to identifying good diversity policies and in limiting their complexity. In addition, we also demonstrate that the bulk of the gains realizable from diversity are typically achievable with only a small number of channels. This further limits the associated complexity, be it in terms of coordination among users or in the number of channels that an individual user needs to manage.

From a practical standpoint, the paper shows that significant benefits from diversity can be achieved via systems with relatively low implementation complexity, namely by using only a few number of channels together with open-loop transmission policies that use all channels for the same amount of time.

In addition to demonstrating that channel diversity can offer substantial performance improvements at a limited cost, we also explore the extent to which it increases robustness to variations in channel characteristics. Specifically, a natural concern when optimizing performance is whether such optimizations render the system more vulnerable to minor variations in operating conditions. In other words, are the performance improvements of channel diversity heavily dependent on the accuracy of the channel models used, and do those improvements cease to exist, or even worse, turn into degradations as the channel parameters vary slightly? We demonstrate that this is typically not the case, and that, furthermore, channel diversity provides a mechanism for trading-off performance improvements for increased robustness.

The remainder of the paper is organized as follows. In Section 2 we briefly review related works. Section 3 introduces the channel model we rely on, the types of transmission policies we consider, and the metric we use to measure performance improvements. Section 4 explores the question of when channel diversity is most beneficial and demonstrates that its biggest gains are in scenarios where channels are used approximately equally. It also introduces the concept of “equivalent” channels to facilitate this exploration. Section 5 presents our findings regarding the evolution of the performance improvements that channel diversity affords as a function of the number of available channels. Section 6 is concerned with how channel diversity affects the relationship between performance and robustness, and describes how it is possible to trade one for the other. Section 7 concludes the paper.

2. RELATED LITERATURE

The use of diversity for single-user or multi-user communication has been the topic of much recent attention, and [8] provides a comprehensive overview of the relevant issues, from the physical to the network layer. Laneman et al. [15]

compare application-layer and physical-layer approaches to diversity. Wang et al. [27, 28] and Zimmermann et al. [29] design protocols that take advantage of path diversity in wireless ad hoc networks. Laneman et al. [16], Pradhan et al. [20] and Puri et al. [21] analyze diversity from an information theory point of view. Akella et al. [2] explore the effects of diversity from a multi-homing point of view, and show that significant benefits can be achieved by using three different Internet service providers. Gummadi et al. [12] carried out a similar investigation in the context of overlay networks, while Tao et al. [23] looked at the same problem in both multi-homing and overlay settings.

Path diversity has often been used in conjunction with techniques that add resiliency to combat packet losses, e.g., coding. The use of coding to overcome the impairments associated with lossy channels is obviously not new, but a number of schemes were devised to specifically take advantage of path diversity, e.g., [5] and [22]. Diversity offers the opportunity to mitigate the impact of periods (bursts) of significant losses on one path, something that is particularly important to real-time applications that often cannot tolerate retransmission delays. Many recent proposals using coding together with path diversity have, therefore, been motivated by video applications [4, 7, 17, 18, 19]. A key difference between those works and this paper is that they typically target the design of codes that maximize application resiliency to packet losses, or the identification of the best set of paths over which to send packets. In contrast, we assume a *given code and set of paths/channels* and focus on the performance gains that can be realized by intelligently utilizing the available paths.

From that standpoint, the papers by Golubchik et al. [11] and Abdouni et al. [1] come closest to our work. They follow a similar approach for coding the data and coordinating transmissions at the packet level, and for modeling the error process of the available channels. However, they do not explicitly consider the transmission policies we use in this work, and more importantly do not focus on developing an understanding of when and to what extent channel diversity can be of benefit. In addition, the performance metrics they target are also different and in particular they do not investigate the trade-off between performance improvements and robustness that we also consider.

The paper by Tsigas and Haas [25] is another work closely related to ours. It considers a wireless setting where multiple channels are available for transmissions, and as in our previous work [26] it seeks to identify the optimal strategy for maximizing throughput when diversity coding is used. A major difference is in the channel model used, which is an “on-off” model where each channel is assumed to be either fully available or entirely unavailable for the duration of transmissions. In contrast, we rely on more general channel models, which as we discuss later, allow us to capture a significantly broader range of channel scenarios. In addition, the aspect of what channel combinations are most conducive to allowing significant improvements, be it in terms of performance or robustness, is also not addressed in the paper.

3. SYSTEM MODEL

We consider a system where senders transmit messages consisting of k packets to a common access point. Messages are encoded into blocks of $N \geq k$ packets to ensure a proba-

bility of successful message delivery greater than or equal to $P_{\min}$. An (N, k) code guarantees that a message is successfully delivered if at least k out of N packets are correctly received [5]. A sender chooses a code length N and a transmission policy, i.e., which channel to use when, based on the channel characteristics and its target value for $P_{\min}$.

We assume that in any transmission slot a node transmits one packet in one of the available channels, and that senders are synchronized so that no two senders transmit in the same channel at the same time. In other words, there are no packet collisions³. The complexity of achieving this synchronization obviously varies with the number of users and channels shared, as well as the type of transmission policy, and we discuss this issue further later in the paper.

In the rest of this section, we first introduce our channel model (Section 3.1), and then define the classes of policies (Section 3.2) and the metrics (Section 3.3) that we use to evaluate the performance of path diversity.

3.1 Channel Model

Channels are modeled by Markov chains that transition between states of different loss rates according to some transition probabilities. Our analysis is general enough to allow arbitrary chains, but for the sake of computational efficiency, we focus on the Gilbert-Elliott (G-E) model [10, 9]. This is a well known model that captures the bursty nature of channels, and is simple enough to allow for a computationally efficient analysis. While more accurate models for, say, the GSM channel are available [13, 14], the G-E model is sufficient for our purposes of investigating the benefits of path diversity within a broad range of channels. For example, in the context of evaluating the benefits of path diversity, the same model was used by Golubchik et al. [11], by Abdouni et al. [1], by El Al et al. [3], while an even simpler model was used by Tsigros and Haas [25]. We make the assumption that channels are independent from each other (see [24] for a discussion on path diversity over correlated channels). We are currently developing an experimental testbed that will allow us to consider the case of correlated channels, but this is beyond the scope of this paper.

According to the G-E model, each channel has two states, good and bad (G and B), and at the end of each packet transmission, it transitions to the bad state with a known probability (P_e and P_b , respectively). In the good (bad) state, the probability of correct transmission is $P_G = 1$ ($P_B = 0$)⁴. Senders do not know the current or previous states of the channels, and only know their long term statistics, i.e., their Long Term Error Rate (LTER) and Expected Burst Length (EBL)⁵. In other words, there is no “active” monitoring of channel quality.

Note that when channels are Bernoulli, i.e., the loss pro-

cess is random ($P_e = P_b$), then the optimal transmission strategy trivially sends all packets on the channel with the lowest long term error rate. This clearly maximizes the probability of correctly receiving at least k out of every N packets, and hence the probability of successful message delivery. Therefore, we concentrate on scenarios where the loss process on at least one channel is bursty ($P_b > P_e$), which is common in both wireline and wireless settings [6, 13, 14]. See Figure 1 for a schematic of the channel model.

Figure 1: The Gilbert-Elliott channel model.

3.2 Transmission policies

The first set of policies that we consider is the class of probabilistic policies. When C channels are available, any strategy in this class is fully identified by its policy vector $\underline{p} = [p_1 \ p_2 \ \dots \ p_C]$, where $p_1 + p_2 + \dots + p_C = 1$. Prior to each packet transmission, policy $\underline{p}$ selects channel 1 with probability p_1 , channel 2 with probability p_2 , and so on. This allows using each channel in different proportions. An optimal policy $\underline{p}^*$ corresponds then to a probability vector that maximizes performance given the characteristics of the available channels. Performance can be computed using a recursive algorithm we previously developed for computing the probability of success of an (N, k) code under an arbitrary probabilistic channel selection policy $\underline{p}$. The algorithm uses an aggregate Markov chain to represent the joint state of the channels with transition probabilities that also incorporate the impact of the transmission policy that determines which channel is used in each slot. We then use a recursion that calculates the probability of having exactly m errors in n packets (denoted as $P(m, n)$), based on $P(m, n-1)$ and on $P(m-1, n-1)$. This allows us to finally compute the probability of k or more errors out of N transmissions. See [26] for a complete description of the algorithm.

Another important class of transmission policies consists of those that deterministically send each packet according to a predefined schedule. Clearly, there are C^N such policies, where C is the number of available channels and N is the length of the frame to be sent. More formally, given C and N , a policy in this class can be specified by a vector $\underline{d} = [c_1 \ c_2 \ \dots \ c_N]$, where each of the c_i 's denotes the channel over which packet i will be sent, and therefore takes values in $\{1, 2, \dots, C\}$. The performance of such deterministic policies can also be computed relatively easily, as it involves a simple recursion (much simpler and more scalable than the recursion needed for probabilistic policies) and the use of convolution, since the deterministic schedule essentially allows us to consider the channels as decoupled and focus on the sum of errors across them.

³The impact of collisions can be taken into account assuming a given number of independent senders sharing multiple channels. However, we will not expand on this further in this paper.

⁴The analysis can be extended to the case where the channels are modeled by a Markov chain with more than two states, and/or the case where $0 \leq P_B \leq P_G \leq 1$.

⁵Note that LTER and EBL can be converted to P_e and P_b , and vice versa. When $P_G = 0$ and $P_B = 1$, it is easy to show that $\text{LTER} = \frac{P_e}{1-P_b+P_e}$ and $\text{EBL} = \frac{1}{1-P_b}$. Equivalently, $P_e = \frac{\text{LTER}}{(1-\text{LTER})\text{EBL}}$ and $P_b = \frac{\text{EBL}-1}{\text{EBL}}$. The derivations are omitted due to space limitation.

3.3 Performance metrics

We now formally define the metric we use to evaluate performance. We define the *Effective Rate* of an (N, k) code under policy S , as the amount of information successfully delivered per unit of time. Our “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 is successfully delivered with probability $P_{succ}^S(N, k)$. The Effective Rate (ER) of an (N, k) code under policy S is defined as

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

As mentioned in Section 3.2, $P_{succ}^S(N, k)$ can be calculated using simple recursive algorithms for both probabilistic and deterministic transmission policies.

The relative difference in Effective Rate between transmission policy S with code (N_S, k) and policy T with code (N_T, k) can then be computed as

$$D_{ER}(S, T) = \frac{ER_S(N_S, k) - ER_T(N_T, k)}{ER_T(N_T, k)}.$$

4. WHEN IS DIVERSITY BENEFICIAL?

In this section, we attempt to develop a better understanding of *when* and *why* channel diversity might be of benefit. Specifically, while there are many evidences that distributing packets across multiple channels can improve performance, this comes at a cost in terms of added complexity. As a result, even if we contemplate using only simple policies, it is important to determine when using diversity is worth that added cost. This calls for both a better sense of the intrinsic cost of path diversity, as well as for an assessment of the magnitude of the improvements it affords, and in particular under what circumstances these are large.

4.1 The Cost of Diversity

The basic cost of diversity is in being able to control the selection of which channel to use for each packet transmission. This obviously depends on the physical characteristics of the system through which the channels can be accessed. For example, a frequency agile transmitter whose frequency hopping pattern is controlled directly within the communication sub-system would make the use of diversity mostly transparent (free). This is not so when channels are associated with different interfaces, so that controlling which interface is used for transmitting a packet may require modifying the standard packet forwarding routines of the end-user operating system. This aspect of the cost of diversity is nevertheless relatively generic, as these are issues that have been dealt with in many other settings, e.g., controlling frequency agile transmitters is common in a number of communication schemes, and distributing packet transmissions across multiple interfaces is implemented in many load-balancing schemes. Another challenge is the previously mentioned need for coordination between senders in order to avoid collisions when one or more users transmit packets in the same channel. This can again be addressed using standard techniques, and furthermore is a requirement that could be relaxed at the cost of a decrease in the achievable performance gains. Without underestimating these aspects, we focus on a few issues that are specific to channel diversity.

A major concern is the complexity associated with the policies that control packet transmissions, and the coordination of which policy which user should use. We dispense with the cost of identifying good policies, as although it can be high and grow rapidly with the number of channels, e.g., see [26], it is a one-time cost that can be amortized over time. Instead, we focus on the impact that the policies have on user behavior. Specifically, when policies amount to distributing packet transmissions evenly across channels, their cost is relatively minor, and they can be easily implemented using either probabilistic or deterministic approaches. The situation is quite different when policies impose *uneven* use of the channels. This affects the complexity of the policies themselves, i.e., although uneven channel use can be easily accomplished with probabilistic policies this is not so for deterministic policies, but, more importantly, calls for additional coordination among users in order to ensure fairness.

For example, consider the case of two users, U_1 and U_2 , and two channels, c_1 and c_2 . Assume c_1 is better than c_2 and that the optimal (probabilistic) policy, i.e., the policy that maximizes the total transmission rate across both users, calls for one user, say U_1 , to use channel c_1 with probability $p_1 > 0.5$ and channel c_2 with probability $p_2 = 1 - p_1 < 0.5$, while the other user (U_2) relies on the complementary policy, i.e., uses channel c_1 with probability $1 - p_1$ and channel c_2 with probability p_1 . Because the two channels are not of the same quality, user U_1 will likely get a higher ER than user U_2 . In order to remedy this potential unfairness, it is necessary for the two users to regularly “switch roles” and, therefore, policies. Specifically, if user U_1 initially uses code (N_1, k) with the policy $[p_1, 1 - p_1]$, while user U_2 uses code (N_2, k) with the policy $[1 - p_1, p_1]$, then the two users will switch policies every $q \cdot \text{LCM}(N_1, N_2)$ ⁶ time units. Clearly, this role switching procedure adds complexity to the operation of the two users.

Implicit in the above discussion was the notion that users sharing channels but using different policies, as would be the case when channels are not used evenly, might also rely on different codes, i.e., code (N_1, k) for user U_1 and code (N_2, k) for user U_2 in the previous example. This may be required in order for both users to meet their minimum performance requirement $P_{\min}$. This then implies that in the case of uneven channel use, users need to not only switch policies but also codes, which introduces further complexity.

To summarize, taking advantage of channel diversity can come at the cost of significant additional complexity. This complexity increases significantly when the policies needed to take advantage of channel diversity require uneven use of the available channels, e.g., because users may need to switch policies and codes at regular intervals. The complexity also increases with the number of channels being shared, as coordination between users becomes harder. Given this potential cost, it is imperative to develop a better understanding of the performance gains that channel diversity can yield, and in particular whether meaningful gains are achievable in scenarios that have lesser complexity, e.g., when a small number of channels are used evenly across users. Investigating this issue is the topic of the next section, which establishes that fortunately many of the scenarios where channel diversity can provide significant benefits are also among those with lower complexity.

⁶ $\text{LCM}(\alpha, \beta)$ denotes the least common multiple of α and β , and q is any integer greater than 0.

4.2 Large Gain Scenarios

In exploring which channel diversity scenarios can yield significant benefits, we focus on the simplest possible setting, namely, two users and two channels. Our investigation is based on assessing the performance gains achievable (using an optimal policy) across a semi-exhaustive list of possible channel combinations with varying LTER and EBL values. Our main finding is that the scenarios where channel diversity yields the most benefits are scenarios where all channels are used for approximately the same fraction of time.

The system parameters we use for our investigation assume a message size of $k = 10$ blocks and a required probability of success $P_{\min} = 0.995$. These are reasonably typical values, and as shown in [26] a more stringent $P_{\min}$ will typically increase the potential for improvements achievable through channel diversity. We consider 55 different combinations of channels with LTER between 1% and 9%, and with EBL between 1.01 and 20 packets. More specifically, there are 3 groups of channels, with LTER equal to 1% for the first group, 5% for the second group, and 9% for the third group. In each group, there are channels with different EBLs, ranging from 1.01 packets to 20 packets. Note that these channels cover a reasonably broad range of realistic channels, and include a number of typical wireless channels, e.g., the G-E model of [14] for the GSM channel. For each combination, we compute the total effective rate of the system, i.e., the sum of the effective rates of each user, using the policy that maximizes this value and allowing the users to use different codes if warranted. In this investigation, we restrict ourselves to probabilistic policies, as they are more readily capable of handling uneven channel uses.

Figure 2 shows the performance (ER) improvements achieved by diversity and the corresponding optimal probabilistic policy p^* ⁷ for all channel combinations. The figure shows that diversity yields the highest benefits when $p^* \approx 0.5$, i.e., the two channels are used approximately equally. For example, when the relative difference in ER between the diversity and the no diversity systems is larger than 35%, then the optimal policy is within the range $[0.45, 0.55]$, and in about 63% of the scenarios where D_{ER} was larger than 35%, the optimal policy was given by $p^* = 0.5$. These findings are summarized in the following observation:

Observation 1: Scenarios for which diversity yields significant improvements are such that the optimal policy uses all channels approximately equally.

This is certainly fortuitous in light of the earlier discussion on complexity. In the remainder of this section, we explore further under what conditions channel diversity uses channels approximately equally. Intuitively this should be the case in most cases where channels are identical, but it is of interest to understand if this also arises in other situations.

4.3 “Equivalent” channels

For the purpose of exploring when the optimal diversity policy amounts to using channels approximately equally, we introduce two definitions of channel “equivalence.” The first definition considers a particular channel and classifies as “equivalent,” any channel for which when used together with the first channel, the optimal (probabilistic) policy is 0.5

⁷For the case of two channels, the optimal policy is fully specified by a single number $p^* \in [0, 1]$. More formally, for users U_1 and U_2 , $\underline{p}_1^* = [p^*, (1 - p^*)]$ and $\underline{p}_2^* = [(1 - p^*), p^*]$.

Figure 2: Relative Difference in Effective Rate between a diversity and a no diversity system, as a function of the optimal policy.

(“Policy-equivalent”). The second “equivalence” class contains all channels that individually deliver the same performance as the original channel (“Rate-equivalent”). Clearly, the original channel belongs to the second equivalence class, i.e., it delivers the same performance as itself, and across all configurations we experimented with, it was also found to belong to the first equivalence class, i.e., across all pairs of identical channels the optimal policy consisted of using them equally. Our interest is now to not only investigate the range of channels that belong in each equivalence class, but also to determine the extent to which these different classes include the same set of channels. This would facilitate the search for channels combinations that can easily deliver good performance improvements when used jointly. In particular, identifying a “Rate-equivalent” channel is relatively easy, since it only involves evaluating their performance in isolation. Furthermore, if the optimal policy with “Rate-equivalent” channels turns out to be using each channel equally, this would have the added benefit of lower complexity. Another goal of our investigation, is to assess the magnitude of the improvements achievable across the different equivalent channels and how these compare to the maximum feasible. Again, we would ideally like to be able to gauge the benefits of channel diversity using “Rate-equivalent” channels and a policy that uses all channels equally. We explore these issues in the rest of this section.

Our investigation is carried out using the GSM channel as the reference channel. We then determine for this channel the composition of the above two equivalence classes, as well as the evolution of the corresponding improvements in ER for channels in each class. Specifically, given a GSM channel, we consider using it together with a second channel whose LTER varies across a broad range of values. For each LTER value, we identify the EBL value associated with the corresponding channel in the “Rate-equivalent” and “Policy-equivalent” classes, as well as the channel that for this LTER value yields the maximum possible D_{ER} . We call the latter channels the “Max. Gain channels.” Results are reported

Figure 3: “Equivalent” and “Max. Gain” channels. The first channel is the GSM channel.

in the left-hand-side (lhs) plot of Figure 3, which establishes that while the three types of channels are not identical, they are nevertheless very close to each other. When combined with the results reported in the right-hand-side (rhs) of Figure 3, we conclude that there exists a large number of channels whose optimal diversity policy when used with a GSM channel, involves using both channels roughly equally. In addition, identifying these channels can be done relatively easily based on the definition of “Rate-Equivalence,” since all three channel types are close to each other.

The potential inaccuracy in evaluating the *ER* improvement achievable through diversity by relying on “Rate-equivalent” channels and an equal channel use policy is illustrated in Figure 4. The lhs of the figure plots the *ER* improvement achievable for all three channel classes. It shows that in spite of some differences, all three offer roughly the same level of improvements. The rhs of the figure specializes this further by assuming that the policy used is the equal channel use policy, and plots the absolute loss in *ER* compared to the optimal policy for “Rate-equivalent” channels. The results again show that the differences are relatively minor. It should also be noted that not all “equivalent” channels yield the same improvement in *ER*. In particular, as channels with increasing LTER are being considered, the benefits available through channel diversity decrease. This is because a smaller EBL is needed in order to make the channels “equivalent” in spite of the larger LTER, and a smaller EBL makes for a more “Bernoulli-like” channel for which diversity does not help as much.

In summary, the results of this section show that the majority of scenarios where diversity can significantly improve performance involve combinations of channels where the optimal policy is to use channels equally. This is clearly of advantage given our earlier discussion regarding the potential complexity of implementing channel diversity policies that do not use all channels equally. In addition, the concept of “Rate-equivalent” channels provides a relatively simple way for identifying channel combinations where the optimal policy consists of using both channels equally. In the next section, we investigate another aspect that also affects the

complexity of implementing channel diversity, namely, the number of channels being shared among users. The issue of sensitivity of our results to both the choice of policy, i.e., the choice of an equal channel use policy instead of the optimal policy, and to variations in channel characteristics is the topic of Section 6.

5. IMPACT OF THE NUMBER OF CHANNELS

In this section, we focus on the evolution of the benefits of diversity as the number of channels grows large. From a system design point of view, this investigation will provide an answer to the question of whether one should allow all senders to use all available channels⁸, or whether partitioning senders in smaller groups and allocating a subset of the available channels to each group yields a better trade-off between performance and complexity.

Given the results of Section 4, we assume “equivalent” channels, and for simplicity concentrate on the case of identical channels for which the optimal policy uses all available channels equally. In such a setting, it is possible to rely on either probabilistic or deterministic policies, and for our investigation we will select a simple round-robin policy that cycles across available channels.

More formally, we define the round-robin deterministic policy D as the policy that transmits packet z_1 over channel c_1 , packet z_2 over channel c_2 , and so on. If the number C of available channels is greater than or equal to the number of packets (i.e., if $C \geq N$), then all N packets are transmitted over distinct channels⁹. If the number of packets is greater than the number of available channels (i.e., if $C < N$), the first C out of N packets are transmitted on channels c_1 to c_C , then packet z_{C+1} is transmitted on channel c_1 , packet z_{C+2} is transmitted on channel c_2 , and so on. This guar-

⁸Intuitively, spreading transmissions across more channels should benefit performance, but comes at the cost of greater complexity.

⁹Note that when this deterministic policy is used, there is no need for more than N channels, since for $C = N$ all packets are guaranteed to be transmitted over distinct channels.

Figure 4: Impact of using “Rate-equivalent” channels and an equal channel use policy on ER . The first channel is the GSM channel.

antees that channels will be reused after exactly C slots, therefore maximizing the number of slots between successive transmissions over the same channel. Intuitively, this ability to maximize the time separating successive transmissions over the same channel should improve performance in comparison to that of a probabilistic policy that also uses channel equally. This intuition is confirmed in Figure 5 for the case of two channels.

Using these results, we now investigate the impact of the number of available channels on the performance of channel diversity. We show that only a small number of channels are needed in order to take advantage of most of the benefits of path diversity.

Figure 5: Deterministic vs. Probabilistic Policies - Two Identical Channels

Although it need not be true in general, for all cases of multiple identical channels that we investigated, the maximum achievable ER was achieved when all N packets are

transmitted over distinct channels, i.e., when $C \geq N$, and that performance steadily increased as C increased up to N . Note that when $C = N$, the Effective Rate is the same as it would have been by using a single random (Bernoulli) channel with the same LTER as the C available channels. Given that we consider identical channels and a deterministic round-robin policy, the total system throughput is simply C times that of one user¹⁰ since they all achieve the same ER value. Our goal is to evaluate different system configurations ranging from an “optimal” setting where each sender uses all C channels, to one where senders are grouped in clusters of g ¹¹, and each cluster uses only g of the available C channels. Our goal is to quantify the evolution of the increase in ER that users experience as the number of channels available to them increases. As mentioned earlier, we noticed that in all cases that we investigated, the more channels (up to N) the better the performance, but the higher the coordination cost.

In order to quantify this behavior, let ER_i be the Effective Rate achieved when i channels are available ($1 \leq i \leq C$). Then, ER_N denotes the Effective Rate that can be achieved when N channels are available. Although it need not be true in general, as we noticed earlier, this is the maximum possible Effective Rate, since all N packets are transmitted over distinct independent channels. We denote the increase in Effective Rate from using q channels to using w channels ($q < w$) as $ER_{q \rightarrow w} = ER_w - ER_q$. Then, the total potential increase in Effective Rate is denoted as $ER_{1 \rightarrow N} = ER_N - ER_1$. The relative gain in Effective Rate by adding the i^{th} channel ($i \in \{1, \dots, N\}$) can now be quantified as

$$G_i = \frac{ER_{(i-1) \rightarrow i}}{ER_{1 \rightarrow N}}.$$

For $i > N$, G_i is trivially equal to zero.

For the case of identical GSM channels, Figure 6 plots the percentage of possible increase in Effective Rate as a function of the number of available channels, as well as G_i as a

¹⁰We assume a system where C users are sharing C channels.

¹¹ g is an integer that denotes the cluster size.

Figure 6: The benefits of using additional channels. Left axis: Contribution of adding each channel. Right axis: Percent of the possible gain in Effective Rate as a function of the available GSM channels.

function of i . From this figure we notice that we need 13 identical GSM channels in order to achieve the maximum possible benefits that path diversity can offer¹². However, we see that by using two channels, the users can take advantage of 40% of the potential benefit of path diversity. By using three channels they can take advantage of 60% of the potential benefit, and by using four channels they can take advantage of 68% of the potential benefit. Note that there are two factors that affect the improvement in Effective Rate (see Equation (1)). The first factor is an increase in the probability of success of a particular code. The second factor is the possibility of using a smaller code (i.e., reducing the code length N) in order to achieve the minimum $P_{\min}$ requirement. This is the reason why the relative gain per channel (G_i) is not monotonically decreasing. For example, adding the seventh, eighth or ninth channels in the system investigated in Figure 6, only slightly improves the probability of success of the same code. But adding the tenth channel finally allows for a smaller code to be used, and therefore the “jump” in the absolute Effective Rate of the system and in the value of G_{10} .

The important observation from Figure 6 is that most of the benefits (about 60%) can be realized with only three channels. From a complexity standpoint, this suggests grouping senders in groups of three, and allocating three channels to each group. Such a choice strikes a balance between implementation complexity (many users having to coordinate with each other, users having to switch between a large number of channels, etc.) and performance. Given this finding, our investigation of the sensitivity/robustness of diversity schemes in the next section will focus on a configuration with three channels and three users.

¹² As mentioned earlier, the performance of the diversity system is limited by the performance that one would achieve by using one Bernoulli channel with the same LTER as the available bursty channels. The code that achieves $P_{\min}$ over such a Bernoulli channel is $N = 13$. Therefore, adding more than 13 channels will result in no increase in the Effective Rate of the system.

6. ROBUSTNESS OF PATH DIVERSITY

In this section, we explore the sensitivity of the performance gains of channel diversity to deviations from the optimal policy and/or errors in the channel characteristics assumed in selecting a diversity policy. As we saw in the rhs portion of Figure 3, using a policy that uses all channels equally instead of the exact optimal policy did not result in significant differences (at least in scenarios of most interest, i.e., when using “equivalent” channels for which the potential gain from diversity is largest). It is nevertheless important to understand the extent to which this robustness holds as channel conditions deviate more and more from their original assumed values. In general, it is of interest to understand the broader trade-off that exists between improving performance and the sensitivity of these improvements to changes in channel parameters. In particular, one might opt for a smaller improvement if this improvement can be preserved across a broad range of channel conditions. Alternatively, one might decide to entirely forfeit all performance improvements, and instead focus on ensuring a base level of performance across the widest possible range of channel combinations. Exploring these issues is the focus of this section.

Specifically, we consider two different types of variations in channel parameters. The first is variations in LTER and EBL. The second corresponds to changes in the distribution of the burst lengths. We keep the expected length of the bursts the same, but change their variance. We therefore quantify the sensitivity of the system performance to higher order moments. In a sense, this is a measure of the sensitivity of the system performance to the G-E channel model that we use throughout this paper. Across all experiments, the policy used is kept as the equal channel use policy, irrespective of the changes in channel parameters. We focus on both the ER value achieved for different levels and types of channel variations, as well as the policy’s ability to maintain performance above the target value of $P_{\min}$.

6.1 Variations in LTER and in EBL

We start by examining the robustness of the system when there are variations in the long term error characteristics of the channels. In particular, we consider the case of three senders and three available channels, and let the error characteristics of the channels (LTER and EBL) increase by 20% (in steps of 1%). First, we consider the case where the characteristics of only one channel are varying, and then we consider the case where the characteristics of all three channels are varying. In all cases, the senders use the (N, k) code that maximizes ER given the original characteristics of the channels, i.e., the code with the smallest value of N that satisfies the $P_{\min}$ requirement under either diversity or no diversity. As a result, when the channels get worse, this code will eventually fail to satisfy the $P_{\min}$ requirement. When that happens, the Effective Rate of the users is set equal to zero.

In the case where only one channel is changing, only one of the senders of a system that does not use path diversity (i.e., a system where one channel is dedicated to each sender) will notice a difference in performance. In a system that implements path diversity (i.e., a system where all senders use all three channels) all senders will notice some difference in performance. However, as shown in Figure 7, the dif-

Figure 7: Joint Effective Rate of all three senders as a function of the percent change in the long term error rate (LTER) and the expected burst length (EBL) of the available channels. Originally, three identical GSM channels are available. The required performance is $P_{\min} = 0.97$.

lhs: The characteristics of only one channel are changing.

rhs: The characteristics of all three channels are changing.

Table 1: Performance vs. Robustness trade-off associated with the use of different diversity systems. Same scenario as in the rhs of Figure 7.

System	D_{ER} compared to a no diversity system	Percent increase in <i>both</i> 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%

ference in the joint Effective Rate¹³ of all three senders is usually smaller when path diversity is used. Assuming three GSM channels, Figure 7 plots the Effective Rate of the two systems (path diversity and no path diversity) for the case where one channel is changing (lhs) and for the case where all three channels are changing (rhs).

From the lhs of Figure 7 (only one channel is changing), we note that the joint ER of the system remains mostly “stable” when diversity is used. This is not the case without diversity, as the user assigned to the degraded channel quickly experiences a drop in performance below its target $P_{\min}$, hence a corresponding ER value of 0. However, while this scenario clearly highlights the benefits of diversity in making the overall system less sensitive to degradations on one channel, it should be pointed out that the sharing of channels can also result in the performance of *all* users being affected by one bad channel. In particular, it is indeed possible for the joint ER value to drop down to 0 when diversity is used, while the worst case drop is 1/3rd without diversity (only the user assigned to the bad channel is affected). Such an extreme scenario only occurs under very severe channel

degradations, i.e., a 40% increase in *both* LTER and EBL, which is outside the range displayed in the figure. The situation is somewhat different when degradations take place on all three channels (rhs of Figure 7), as both the diversity and the no diversity system can now result in a joint ER of 0 when the channels become bad enough. However, the advantage of diversity in preserving performance remains and is even stronger, as such a precipitous drop takes place only after *both* LTER and EBL have experienced a 16% increase, while a mere 3 – 4% increase in *either* LTER or EBL is sufficient to shut-down the no diversity system.

When looking at the above results a natural question is whether the improvement in ER that diversity offers in the absence of channel degradations can be traded-off to further improve robustness. The answer is clearly yes, as forfeiting improvements in ER typically means using a larger value for N than is necessary to guarantee $P_{\min}$ under normal conditions, and this in turn translates into a stronger code capable of absorbing harsher channel conditions. However, what is less clear is the magnitude of the improvements in robustness that are achievable when giving up a certain amount of improvement in ER . In other words, if the code length achieving the maximum value of ER when diversity is used

¹³Recall that this is defined as the sum of the Effective Rates of all users in the system.

is N^* , what is the added robustness that is achievable if the senders choose a code length of $N^*+1, N^*+2, \dots, N'$, where N' is the code used by the no diversity system. Investigating this aspect in the context of the scenario of the rhs of Figure 7 (for which $N' = 19$ and $N^* = 15$) is the topic of Table 1. Specifically, Table 1 shows that the optimal diversity system ($N = 15$) achieves an improvement in ER of 27.6% over the no diversity system, and is robust up to an increase of about 16% in both LTER and EBL (as opposed to 2% for the no diversity system). In contrast, a diversity system that trades most of the improvements in ER for greater robustness, i.e., uses a code length of, say, $N = 18$, achieves a performance improvement of only 8.2%, but is robust even to a 92% increase in both LTER and EBL. This highlights the fact that based on application needs, channel diversity can be used to either improve or preserve performance, or both.

Note that due to the discrete nature of the selection of a code length N , the relative increase in LTER and/or EBL that is needed to degrade performance down to $ER = 0$ will vary. In particular, it depends on the original characteristics of the channels and the value of $P_{\min}$. Suppose for example that $P_{\min}=0.97$, and consider two cases (i.e., two sets of channels with different characteristics). In the first case, a $N = 14$ code achieves $P_{succ}=0.9699$, and a $N = 15$ code achieves $P_{succ}=0.9785$. In the second case, a $N = 17$ code achieves $P_{succ}=0.9612$, and a $N = 18$ code achieves $P_{succ}=0.9701$. Notice that in the first case, the $N = 15$ code performs much better than the required $P_{\min}$. Therefore, it will take a relatively large degradation in channel characteristics for the system's ER to drop to zero. In the second case, however, the $N = 18$ code is just barely capable of achieving the desired $P_{\min}$ value. As a result, even a relatively small change in the channel characteristics may be sufficient to drop the system's ER to 0. This points to the fact that in reality, with channel characteristics being more often than not only approximate estimates, seeking the maximum improvement that diversity can afford need not be desirable or even wise. Selecting a code length that is 1 larger than the optimal value might be advisable in practice, e.g., in Table 1, using a code length of $N = 16$ instead of the minimum possible value of $N = 15$ more than doubles the range of performance degradations that the system can tolerate, while still providing meaningful improvements in ER .

6.2 Variations in the Distribution of the Burst Lengths

In this section, we extend the investigation of the previous section, and focus on the impact of our simple G-E channel model. Recall from Section 3.1, that since the G-E channel model corresponds to a two-state Markov chain, the length of error bursts is geometrically distributed, which will often be a relatively loose approximation of the actual burst length distribution. It is, therefore, important to evaluate the impact that different, non-geometric distributions can have. We focus on the case where differences in burst length distributions are primarily reflected in different standard deviations. We choose the gamma distribution to carry out this investigation, as it allows us to change the variance of the burst lengths while keeping the same mean. Since the error process is not any more described by a Markov chain, we rely on simulations to quantify the difference in performance.

We compare the performance of a system that does not use path diversity to the performance of three versions of a path diversity system. The system that does not use path diversity dedicates one channel to each sender, and uses the code length N' that results in the highest possible performance. The first path diversity system is optimized for performance, and therefore uses the smallest possible code length N^* that satisfies $P_{\min}$. The second path diversity system trades some of the performance gains for robustness. It therefore uses a code that is one packet larger than the optimal code (i.e., it uses a code length of $N^* + 1$). The third system trades all of the performance gains of path diversity in order to provide for maximum robustness against changes in the channel parameters, namely, in the variance of the burst lengths. In doing so, it uses the same code length N' that is used by the system that does not use path diversity. In all scenarios, there are three senders and the reference channel is again the GSM channel. The variance of burst lengths is then changed while keeping all other channel parameters (LTER and EBL) constant, and we examine the sensitivity of system performance to these changes.

Table 2 presents the results for the four systems. The burst length variance of the GSM channel ($Var(GSM)$) is 37.78, and we consider channels whose burst length variance varies from $\frac{1}{4} Var(GSM)$ to $8 \cdot Var(GSM)$. Note that when the variance multiplier is equal to 1, the channels have the same long term characteristics as the original GSM channels, but the distribution of the burst lengths is different (gamma instead of geometric). Under the original GSM channel, the system that does not use path diversity uses a (19, 10) code in order to achieve the highest possible performance, and the path diversity system that is designed for maximum performance uses a (15, 10) code. The required performance target is $P_{\min} = 0.97$, and the statistics are averaged over 10,000 N -blocks.

From the results of Table 2, we first note that a system that is optimized for maximum performance (regardless of whether or not path diversity is used) is fairly sensitive to changes in the variance of the burst lengths. However, diversity provides a powerful solution for improving robustness to changes in channel characteristics without sacrificing performance. When using a diversity system that uses the same code length ($N' = 19$) as the no diversity system, and therefore achieves the same performance (actually slightly better because of a better P_{succ}), we are able to maintain the target $P_{\min}$ even when the variance of burst lengths increases by a factor 4. This again demonstrates that diversity is a versatile mechanism when it comes to improving or preserving performance, even in environments where channel characteristics are variable or only known approximately.

7. CONCLUSION

This paper explored the potential benefits of “user-level” channel diversity solutions that instead of dedicating one channel to each user, let all users distribute their packet transmissions across all channels. The paper identified that the channel configurations for which channel diversity affords meaningful benefits are fortunately also those where it can be implemented at a relatively low cost, i.e., by using a simple round-robin policy to distribute packets across channels. The paper also introduced the concept of “equivalent” channels, and demonstrated that there exist a broad range of channels that can be considered equivalent and that the

Table 2: System performance as a function of the variance of the burst lengths. Three senders and three GSM channels.

Variance multiplier	No path diversity		Path diversity systems					
	$N = 19$		$N = 15$		$N = 16$		$N = 19$	
	ER	P_{succ}	ER	P_{succ}	ER	P_{succ}	ER	P_{succ}
Original (Markov)	1.534	0.971	1.956	0.978	1.850	0.987	1.574	0.997
0.25	1.555	0.985	1.947	0.973	1.840	0.982	1.574	0.997
0.50	1.547	0.980	1.942	0.971	1.837	0.980	1.568	0.993
1	1.538	0.974	0	0.968	1.830	0.976	1.562	0.989
2	0	0.963	0	0.962	0	0.968	1.552	0.986
4	0	0.961	0	0.949	0	0.957	1.538	0.974
8	0	0.953	0	0.941	0	0.949	0	0.966

exact definition of equivalence is relatively robust, at least in terms of characteristics that are meaningful in the context of diversity. The paper also showed that reaping the benefits of diversity can typically be accomplished with a relatively small number of channels, which is also significant in the context of keeping complexity low. Finally, the paper explored the aspect of sensitivity to changes in channel conditions or errors in their estimated characteristics. It showed that diversity solutions are not only relatively insensitive to slight inaccuracies in estimating channel parameters, but that they also provide a powerful solution to improve robustness against a broad range of channel degradations.

We are currently in the process of deploying an experimental IEEE 802.11 testbed in order to evaluate the extent to which the “conceptual” benefits of channel diversity identified in this paper can be realized in a practical system. In doing so, we are in the process of investigating several issues, including an assessment of the limitations of the G-E channel model, the impact of correlated channels, and the issue of the same channel delivering different performance to different users. We are also exploring the influence of selecting different packet sizes (and assessing the impact of the packet header overhead) as well as making transmissions decisions at a coarser grain than on a per-packet basis, i.e., transmitting a group of packets on one channel. In a sense, using larger packets or making transmission decisions at a coarser grain are expected to have a similar impact in terms of the constraint they introduce regarding sensitivity to bursts given an underlying physical channel model. Finally, we are extending our investigation to systems that do not require coordination between users, i.e., where there is the possibility of collisions, and have also been considering the use of more complex solutions that involve channel monitoring and feedback.

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 12th Annual Multimedia Computing and Networking (MMCN)*, San Jose, CA, January 2005.
- [2] A. Akella, J. Pang, B. Maggs, S. Seshan, and A. Shaikh. A comparison of overlay routing and multihoming route control. In *Proc. SIGCOMM '04*, Portland, OR, August 2004.
- [3] A. A. E. Al, T. Saadawi, and M. Lee. Improving interactive video in ad-hoc networks using path diversity. In *Proc. IEEE International Conference on Mobile Ad-hoc and Sensor Systems (MASS)*, Ft. Lauderdale, FL, October 2004.
- [4] J. Apostolopoulos. Reliable video communication over lossy packet networks using multiple state encoding and path diversity. *Visual Communications and Image Processing (VCIP)*, pages 392–409, January 2001.
- [5] 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.
- [6] J. Bolot. End-to-end packet delay and loss behavior in the Internet. In *Proc. ACM SIGCOMM'93*, San Francisco, CA, September 1993.
- [7] J. Chakareski, S. Han, and B. Girod. Layered coding vs. multiple descriptions for video streaming over multiple paths. In *Proc. ACM Int'l. Conf. on Multimedia*, Berkeley, CA, November 2003.
- [8] 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.
- [9] E. Elliott. Estimates of error rates for codes on burst-noise channels. *Bell System Technical Journal*, 42(5), September 1963.
- [10] E. Gilbert. Capacity of a burst-noise channel. *Bell System Technical Journal*, 39(5), September 1960.
- [11] 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 Journal*, 39, September 2002.
- [12] K. Gummadi, H. Madhyastha, S. D. Gribble, H. M. Levy, and D. J. Wetherall. Improving the reliability of Internet paths with one-hop source routing. In *Proc. of 6th Symposium on Operating Systems Design & Implementation (OSDI)*, San Francisco, CA, December 2004.
- [13] P. Ji, B. Liu, D. Towsley, Z. Ge, and J. Kurose. Modeling frame-level errors in GSM wireless channels. *Performance Evaluation (Internet Performance Symposium - IPS 2002)*, 55(1-2), January 2004.
- [14] A. Konrad, B. Zhao, A. Joseph, and R. Ludwig. A Markov-based channel model algorithm for wireless networks. *Wireless Networks*, 9(3), May 2003.
- [15] J. Laneman, E. Martinian, G. Wornell, J. Apostolopoulos, and S. Wee. Comparing application- and physical-layer approaches to diversity on wireless channels. In *Proc. of IEEE International*

- Conference on Communications (ICC)*, Anchorage, AK, May 2003.
- [16] J. Laneman, D. Tse, and G. Wornell. Cooperative diversity in wireless networks: Efficient protocols and outage behavior. *IEEE Trans. Inform. Theory*, 50(12), December 2004.
 - [17] S. Mao, S. Lin, S. Panwar, Y. Wang, and E. Celebi. Video transport over ad hoc networks: Multistream coding with multipath transport. *IEEE J. Select. Areas Commun., Special Issue on Recent Advances in Wireless Multimedia*, 21(10):1721–1737, December 2003.
 - [18] A. Miu, J. Apostolopoulos, W. Tan, and M. Trott. Low-latency wireless video over 802.11 networks using path diversity. In *Proc. IEEE Int'l. Conf. on Multimedia and Expo (ICME)*, Baltimore, MD, July 2003.
 - [19] T. Nguyen and A. Zakhor. Path diversity with forward error correction (PDF) system for packet switched networks. In *Proc. IEEE INFOCOM 2003*, San Francisco, CA, April 2003.
 - [20] S. Pradhan, R. Puri, and K. Ramchandran. n-Channel symmetric multiple descriptions – Part I: (n,k) source-channel erasure codes. *IEEE Transactions on Information Theory*, 50(1), January 2004.
 - [21] R. Puri, S. Pradhan, and K. Ramchandran. n-Channel symmetric multiple descriptions – Part II: An achievable rate-distortion region. *IEEE Transactions on Information Theory*, 51(4), April 2005.
 - [22] 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.
 - [23] S. Tao, K. Xu, Y. Xu, T. Fei, L. Gao, R. Guérin, J. Kurose, D. Towsley, and Z.-L. Zhang. Exploring the performance benefits of end-to-end path switching. In *Proc. ICNP 2004*, Berlin, Germany, October 2004.
 - [24] A. Tsirigos and Z. Haas. Analysis of multipath routing, Part 2: Mitigation of the effects of frequently changing network topologies. *IEEE Transactions on Wireless Communications*, 3(2), March 2004.
 - [25] 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.
 - [26] 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.
 - [27] J. Wang, H. Zhai, and Y. Fang. Opportunistic packet scheduling and media access control for wireless LANs and multi-hop ad hoc networks. In *Proc. of the IEEE Wireless Communications and Networking Conference (WCNC)*, Atlanta, GA, March 2004.
 - [28] J. Wang, H. Zhai, W. Liu, and Y. Fang. Reliable and efficient packet forwarding by utilizing path diversity in wireless ad hoc networks. In *Proc. Milcom '04*, Monterey, CA, October 2004.
 - [29] E. Zimmermann, P. Herhold, and G. Fettweis. A novel protocol for cooperative diversity in wireless networks. In *Proc. of the Fifth European Wireless Conference*, Barcelona, Spain, February 2004.