

Improving Performance Through Channel Diversity in the Presence of Bursty Losses

Evangelos Vergetis, Roch Guérin and Saswati Sarkar*

University of Pennsylvania, Department of Electrical and Systems Engineering
200 South 33rd Street, Philadelphia, PA 19104-6390, USA
{vergetis@seas, guerin@ee, swati@ee}.upenn.edu

Abstract. As more applications migrate to IP networks, ensuring a consistent level of service is increasingly important. One option is for the network to offer service guarantees. Another is to leverage the path diversity that the Internet intrinsically offers. Our focus is on understanding if and when one can indeed take advantage of multiple disjoint paths to improve performance. We consider an environment where loss patterns are bursty and where coding is used to provide robustness against packet losses. We assume that only long-term loss statistics are known about each path, and we seek to identify the best strategy for sending packets over the available paths. Our contributions are two-fold. First we demonstrate that even with minimal knowledge of channel characteristics and using simple transmission policies, path diversity can help significantly improve performance. Second, we derive an efficient method for identifying optimal policies, and more importantly characterize when having access to multiple paths can be of benefit.

Keywords: Channel Diversity, Cross-layer designs, Reliability and fault tolerance.

1 Introduction

IP networks are nowadays carrying many diverse applications, including multimedia applications with requirements for minimum service guarantees. Service guarantees are, however, not “native” to IP networks, and while there have been proposals to introduce Quality-of-Service (QoS) in the network, another alternative is to make applications more resilient to variations in network performance. This has the advantage of keeping complexity out of the network and allows custom tailoring for each application, e.g., by adjusting the level of resiliency. Resiliency can be provided by letting the application adapt to changes in network performance, and by adding redundancy to the data it transmits. Redundancy is typically in the form of coding, that lets applications recover from a certain number of network losses. In this paper, we investigate whether such resiliency can be further enhanced by also taking advantage of diversity in the transmission options, e.g., multiples paths or channels¹ available from the network. Intuitively, diversity can be of benefit by helping avoid extended periods of poor performance on any given channel, which would translate into a long burst of errors (lost packets).

Our starting point consists of applications that in order to achieve some resiliency to losses, use packet level encoding so as to be able to recover data whenever at least k packets out of a frame of N transmitted packets are correctly received. A number of such codes have been proposed for use in packet networks, e.g., information dispersal [1] and diversity coding [2], that are relatively simple to implement. Depending on the quality of the network and the “strength” of the code, the sender achieves a certain “effective” transmission rate, i.e., successfully delivers to the receiver a number of frames² per unit of time. Typically, the sender targets a frame loss probability P_{min} that corresponds to a minimum desired level of application quality, and selects its code, i.e., the values of k and N , so as to meet this target while maximizing the code rate $\frac{k}{N}$. When there is only one path, the sender transmits all its packets on the

* This work was supported in part by NSF grants ITR-0085930, ANI-0106984, ANI-9902943, ANI-9906855 and CNS 0435306, and by grant NCR-0238340.

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

² A frame corresponds to the data encoded in a frame of N packets, i.e., that can be recovered whenever at least k out of N packets are received.

same path, but in the presence of multiple paths, the sender has the option of sending different packets over different ones.

As a result, we consider a setting with multiple (independent) channels between a sender and a receiver, and explore the extent to which this diversity can be leveraged, even when the information available about each channel is minimal. Specifically, we assume that only general channel statistics such as average loss rate and loss burstiness are known, and that neither sender nor receiver are capable of sensing or predicting channel states. In other words, we limit ourselves to static channel information, which has the advantage of both simplicity and portability. Simplicity because it allows the exploitation of diversity through minor modifications to traditional transmission systems, i.e., the ability to switch successive packet transmissions from one channel to another. This can be implemented through small device driver level modifications to allow the distribution of packet transmissions across multiple network interface cards (NICs). Alternatively, a single, frequency-agile NIC could also be used³. More importantly, unlike many physical layer diversity solutions that rely on the precise knowledge of channel characteristics, e.g., OFDM or MIMO systems, the resulting solutions are portable in that they can operate across a broad range of channel types, and can also be easily adjusted to accommodate updates in channel estimates.

The type of static channel characteristics we assume known are commonly available for both wireline and wireless channels, and similarly access to multiple such transmission options (channels) is itself not unusual. For example, in wireline networks the gateways connecting different sites often have multiple connections between them, either through multi-homing or through overlay paths. These paths are monitored to ensure compliance with service level agreements, and these measurements allow the construction of simple channel models, e.g., [4, 5]. Similarly, in wireless networks multiple channels associated with different frequency bands, different physical antennas, or antenna polarization are also not uncommon. For many if not all of these channels, models are available, e.g., [6, 7], that provide the type of information we assume available for channel selection decisions.

Our goal in this paper is, therefore, to identify if, when, and how access to multiple choices can be of benefit, when selecting a channel for packet transmission. Note that implicit here is the assumption that sending successive packets over distinct channels does not in itself negatively impact performance. In particular, it is expected that out-of-order packet transmissions can be readily handled by the receiver. This ability is commonly available in most applications, as the Internet itself does not guarantee in-order delivery, and simple resequencing buffers have been shown to be quite effective [8, 9], even in the context of mobile ad hoc networks [10], where path changes are frequent.

In order to gain some insight into the potential benefits that channel diversity offers, we focus on a family of transmission strategies that are both simple to investigate and implement. Specifically, rather than consistently transmitting all its packets on the same channel, the sender *randomly* selects a channel prior to each packet transmission. Random policies do not require much additional state at the sender (they are fully specified by the transmission probabilities associated with each channel), and do not change the number of packet transmissions at the sender. This ensures that any benefit derived from (randomly) using multiple channels is solely because of diversity, and not because of additional transmissions and therefore a greater power or bandwidth consumption. In addition to focusing on random policies, we also initially limit our investigation to the case when only two channels are available. This is in part because this represents the most common environment, and also because, as we shall see, the bulk of the benefits achievable through diversity are often realized with only two channels.

In the context of this investigation, our contributions are as follows. First and foremost, we demonstrate that by taking advantage of path diversity and with minimum changes to the transmission methods they currently use, applications can achieve substantial performance improvements in the amount of information they successfully transfer per unit of time. The magnitude of those improvements obviously varies with the relative quality of the available paths, the desired target probability of successful transmissions, and the frame size, but there are numerous configurations where an application substantially improve its effective transmission rate. For example, in the case of two identical GSM channels [11] and a target loss probability of 0.5%, randomly alternating transmissions between the two channels can result in an effective rate increase of 47%. Another contribution of the paper is the development of a computationally efficient technique for identifying the optimal policy, which for a given code and desired probability of successful transmission yields the highest possible transmission rate. Finally, we broadly characterize the set of configurations, i.e., channel characteristics, target probability of successful transmissions, frame size, etc., for which path diversity can be expected to yield meaningful performance improvements.

³ The device level switching latency in an IEEE 802.11a NIC is approximately $80\mu s$ [3]. Clearly this does not account for all the system level overhead, but points to the potential feasibility of packet-level channel switching decisions.

The remainder of the paper is organized as follows. In Section 2 we briefly review related works. Section 3 formally introduces a model for the system we consider. Section 4 defines metrics for comparing different policies, and Section 5 presents an efficient scheme for computing the probability of success for different (random) transmission strategies. In Section 6 we quantify the benefits available from path diversity as well as characterize the channel scenarios where path diversity is indeed beneficial. Finally, Section 7 summarizes our results and outlines several extensions we are currently working on.

2 Related Literature

The use of diversity has been the topic of much recent attention, and [12] provides a comprehensive overview of the relevant issues, from the physical to the network layer. Laneman et al. [13] compare application-layer and physical-layer approaches to diversity. Wang et al. [14] and Zimmermann et al. [15] design protocols that take advantage of path diversity in wireless ad hoc networks. Laneman et al. [16], Pradhan et al. [17] and Puri et al. [18] analyze diversity from an information theory point of view.

Path diversity has been traditionally motivated by real-time applications and has been used in combination 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, and several previously proposed schemes, e.g., [1] and [2], were also motivated by the availability of multiple paths. Diversity offers the opportunity to mitigate the impact of periods of significant losses on one the paths, 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 [19–22]. 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. [23] and Abdouni et al. [24] come closest to our work. They follow a similar approach for coding the data and for modeling the error process of the available channels, but do not explicitly consider the probabilistic policies that we use in this work. Moreover, they focus on different metrics when quantifying the benefits of path diversity.

Tsirigos and Haas [25] also share our goal of identifying an optimal transmission strategy based on knowing only long term channel characteristics. The paper considers a wireless setting where multiple channels are available for transmission, and where fading causes channels to be in either a “bad” or a “good” state. When a channel is in a bad state, *all* the packets transmitted on that channel are lost. A channel has the same state for the entire duration of each frame. This represents a simpler channel model than the one we use in this paper, as we allow for general statistical loss patterns, i.e., different average error rates and expected burst lengths. We also allow the state of the channel to change on a per packet basis, not on a per frame basis. Given that frames can potentially be relatively large (tens or hundreds of packets) and that channel conditions can change relatively quickly, our model can adequately capture the behavior of such a system. Our channel model does translate into a more complex analysis, but more importantly, it extends the applicability of the results to a broader range of channels, including wireline channels and many typical wireless channels, e.g., the GSM channel of [11], that do not fit a purely “on-off” model.

3 System Model

In this work we consider a system comprised of a sender with many available transmission paths to a particular destination. Each of these paths or channels is modeled as a Markov chain that transitions between states associated with different loss rates. Our analysis is general enough to allow for an arbitrary Markov model for the channels, but for the sake of computational efficiency, we focus on the Gilbert-Elliott model [26, 27]. This is a well known model that captures the bursty nature of channels, while at the same time being simple enough to allow for a computationally efficient analysis of the system performance. While more accurate models for, say, the GSM channel are available [7, 11], the G-E model is sufficient for our purpose of demonstrating the benefits of path diversity within a broad range of channels. 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. See

⁴ Our analysis can be extended to the case where $0 \leq P_B \leq P_G \leq 1$.

Figure 1 for a schematic of the channel. A node transmits each packet in only one channel, but may select different channels for successive packets. It does not know the current or previous states of the channels, but it knows the long term statistics, P_G , P_B , P_e and P_b . Note that when the channels are Bernoulli channels, i.e., the loss process has no memory ($P_e = P_b$), then the optimal strategy is to send all packets on the channel with the lowest loss probability. Thus, we concentrate on scenarios where the loss process on at least one channel is bursty, which is common in both wireline and wireless settings [4, 7, 11].

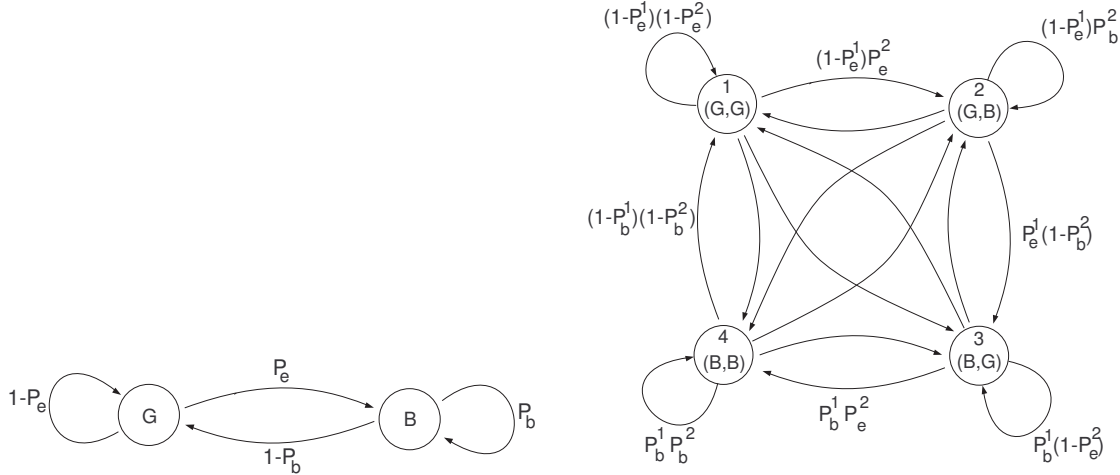


Fig. 1.

Left: The simplified Gilbert-Elliott channel model.

Right: The 4-state Markov Chain corresponding to the case of two *mutually independent* Gilbert-Elliott channels. A few representative transitions are labeled.

Frames are transmitted in blocks of N packets and have been encoded using (N, k) coding such that a receiver can decode the original information if and only if it correctly receives at least k packets; otherwise all N received packets are discarded. For any given frame length k , the higher the value of N , the more the coding redundancy and therefore the lower the fraction of “information” transmitted per frame. On the other hand, the lower the value of N , the smaller the probability that at least k packets are delivered successfully in each frame and the frame is decoded correctly. The quantity $\frac{k}{N}$ is defined as the *rate* of the code. For a given frame of length k , we seek to quantify the system trade-offs associated with different values of N and transmission strategies. Specifically, we are interested in identifying the *optimum* channel selection strategy for any given (N, k) code, which is the channel selection strategy that maximizes the probability that k or more packets out of N are received correctly.

For any given N , a node can select the transmission channel of each packet in several different ways. For example, it can transmit all packets over one channel, or distribute transmissions among multiple channels. In the latter case, the node may use a deterministic strategy, e.g., transmit the first i_1 packets in channel 1, the next i_2 packets in channel 2, etc., or it can randomize its channel selection decisions. In this paper, because of its simplicity we focus on the latter approach. Furthermore and as mentioned earlier, we also limit ourselves to a base scenario involving only two channels. Whenever a packet is ready to be transmitted, channel 1 is selected with probability p and channel 2 with probability $1 - p$, so that any strategy in this class is fully specified by its parameter $p \in [0, 1]$. Note that this class of strategies includes the basic strategies of transmitting all N packets on the same channel, i.e., $p = 0$ or $p = 1$.

The goal is then to develop a computational framework that will let us characterize as a function of the code length N , and the channel parameters P_e and P_b , the optimal p^* that maximizes the sender’s effective rate. Conversely, this framework can then also be used to identify an optimal code, i.e., the parameter N , that for channels parameters P_e and P_b , and for given frame of length k and associated target probability of successful frame transmissions P_{min} , maximizes the sender’s effective rate.

4 Performance Metrics

We define a metric that captures the trade-off attained by different transmission policies and codes, and use it to introduce our definition of an optimal policy.

The *Effective Rate* (ER) of a code under a given channel selection policy A , a quantity that captures the useful data throughput, is defined as the useful information that is transmitted per unit of time. By “unit of time,” we refer to the time required to send one packet. Since each frame in an (N, k) code consists of N packets, the time required to send a frame is equal to N . The amount of useful information in a frame is k , and this information is correctly transmitted with probability $P_{succ}^A(N, k)$. Therefore, the Effective Rate of an (N, k) code under some arbitrary channel selection policy A can be written as

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

We consider the case where for a given frame length k , any code length N can be used that meets the target success probability P_{min} . We define the optimal transmission policy p^* as the policy that maximizes the Effective Rate over all N . Let N^* denote the value of N that together with the optimal channel selection policy p^* attains the maximum possible Effective Rate. The relative gain in Effective Rate of the optimal combination consisting of code (N^*, k) and policy p^* over policy A (used in conjunction with the code (N_A, k) that maximizes the Effective Rate of that policy A) is defined as the ratio

$$G_{ER}(A) = \frac{ER_{opt}(N^*, k) - ER_A(N_A, k)}{ER_A(N_A, k)}. \quad (2)$$

5 Computing the Optimal Policy

We discuss how to compute the optimal channel selection strategy p^* and the optimal code parameter N^* . Note that the Effective Rate can be computed for any channel selection policy if the probability of successful transmission of each frame under the policy is known. We therefore develop a recursive scheme to compute this probability for an arbitrary channel selection policy.

A system with 2 channels can be modeled by a 4-state Markov chain with states (G, G) (state 1), (G, B) (state 2), (B, G) (state 3) and (B, B) (state 4). See Figure 1 for a schematic of the Markov chain. The first (second) component of each state of the Markov Chain represents the state of the first (second) channel. Let π_i be the stationary probability of the chain being in state i . Let $P_{ij}^A(m, n)$ be the conditional probability of having m errors in a sequence of n transmissions under channel selection strategy A given that the initial state was i and the ending state was j . Note that $P_{ij}^A(m, n)$ depends on the 4-state Markov chain and the channel selection strategy. We now write an expression for $P_{succ}^A(N, k)$ as follows.

$$P_{succ}^A(N, k) = \sum_{m=0}^{N-k} \sum_{i=1}^4 \sum_{j=1}^4 \pi_i P_{ij}^A(m, N). \quad (3)$$

The first sum is over all possible numbers of errors that can be corrected by the (N, k) code. The next two sums are over all possible starting and ending states of the system. Since the system has two channels and each channel can be in one of two possible states (G or B), there are $2^2 = 4$ such states.

Note that we assumed that the system is in steady state at the start of each frame transmission or block of N packets. We justify this assumption as follows. Consider sampling the state of the Markov Chain after transmitting every frame of N packets. Let $s_0, s_1, \dots, s_j, \dots$ be the states observed at the sample epochs. Let the indicator variable I_i^j be equal to 1 if the Markov Chain is in state j at the i^{th} sample, and equal to 0 otherwise. Thus, the fraction of times the system is in state j at the start of transmission of a frame is $\lim_{i \rightarrow \infty} \frac{I_1^j + I_2^j + \dots + I_i^j}{i}$ which equals π_j by ergodicity ([28], Chapter 4.3).

We now present an algorithm for recursively computing $P_{ij}^A(m, n)$ for any given channel selection strategy A . For all $n = 0, 1, \dots$ and $m = 0, 1, \dots, n$, and for all $i, j \in \{1, 2, 3, 4\}$,

$$P_{ij}^A(m, n) = \sum_{k=1}^4 P_{ik}^A(m, n-1) \cdot P_{kj} \cdot P^A\{\text{no error} \mid \text{state is } j\} + \sum_{k=1}^4 P_{ik}^A(m-1, n-1) \cdot P_{kj} \cdot P^A\{\text{error} \mid \text{state is } j\}, \quad (4)$$

where $P^A\{\text{no error} \mid \text{state is } j\}$ is the conditional probability that a packet is successfully transmitted under policy A given that the state is j , and $P^A\{\text{error} \mid \text{state is } j\}$ is the conditional probability that a packet is unsuccessfully transmitted under policy A given that the state is j . P_{kj} denotes the transition probability from state k to state j . The initial conditions for the recursions are given by $P_{ij}^A(m, n) = 0$ for all $m > n$ and $m < 0$. When $m = n = 0$, the initial conditions are $P_{ij}^A(0, 0) = 1$ if $i = j$, and $P_{ij}^A(0, 0) = 0$ otherwise. Thus, $P_{ij}^A(m, n)$ depend on the transition probabilities as well as on $P^A\{\text{no error} \mid \text{state is } j\}$ and $P^A\{\text{error} \mid \text{state is } j\}$, which depend on the channel selection policy A .

From Equations (3) and (4), we see that $P_{succ}^A(N, k)$ can be computed recursively with a computational complexity of order $O(N^3)$. Our main motivation for following such an approach is computational tractability. In particular, while one could derive a closed-form expression for $P_{succ}^A(N, k)$ by extending results available for the single Gilbert-Elliott channel case [29], we believe that such an approach would both prove computationally inefficient and fail to provide insight into the impact of individual parameters. As a matter of fact, a similar situation motivated the development of recursive expressions in the single Gilbert-Elliott channel case [29, 30], as the available closed-form expressions were numerically unstable and so complex that they provided little intuition for the role of individual channel parameters.

Using the approach we outlined, the probability of successful frame transmission and, therefore, the Effective Rate of any channel selection policy can be computed for all channel and code combinations. Hence, the optimal strategy can be identified. Extending the above recursion to more than two channels is readily accomplished, but computational complexity increases rapidly with the number of channels.

6 Evaluating the Benefits of Path Diversity

As mentioned before, path diversity arises in many different scenarios (multi-homing, overlay paths, different frequency bands, different physical antennas, antenna polarization, etc.), and it involves different control overhead in each. Assessing when and if taking advantage of this diversity is beneficial calls, therefore, for the ability to compare achievable performance improvements and implementation costs in each scenario. The framework we have just introduced will allow a system designer to make such an assessment, and in this section we demonstrate through numerical evaluations how it is possible to articulate broad guidelines for when path diversity can be effective.

We first show that in many cases path diversity can substantially improve performance. We make this assessment by evaluating the relative gain in effective rate when using two channels instead of only one, i.e., $p \in \{0, 1\}$, and characterize the channel scenarios in which this gain is significant. We show that adding a second channel, even a relatively bad one, can result in significant performance improvements, which we quantify for different combinations of system parameters.

Specifically, we consider the case of two mutually independent channels with transition probabilities P_e^1, P_b^1 and P_e^2, P_b^2 (Figure 1), and different long term error rates and expected burst lengths. In particular, for the initial investigation of quantifying the benefits of diversity, we consider 55 different channels with long term error rates between 1% and 9%. As far as burst characteristics are concerned, we consider channels ranging from Bernoulli (i.e., not bursty), to channels where the expected burst length is 10 (when $P_b = 0.9$). This combination of 55 channels allows us to cover a reasonably broad range of realistic channel characteristics in order to gain a better understanding of when path diversity can be beneficial. In addition, the 55 channels also include a number of specific cases of practical importance, such as the GSM channel that has parameters $P_e = 0.0079$ and $P_b = 0.8509$ [7], which result in a long term error rate of 5% and an average burst length of 6.67 packets. Because of its importance, we will highlight this channel throughout our investigation. Later on we will consider channels with long term error rate up to 30% and arbitrary expected burst lengths.

We assume that the frame length k , and the minimum required probability of success P_{min} are derived from application requirements. The smaller the frame length, k , the smaller the coding delay, and the larger P_{min} , the more reliable the transmissions. The optimal path diversity solution corresponds to the channel selection policy p^* and code length N^* that maximize the effective rate under these constraints. We consider different values of k and P_{min} .

First note that when no coding is used and the only requirement is to maximize the maximum number of packets successfully transmitted, then the optimal channel selection strategy is to select the channel with the smallest long-term error rate independent of the expected burst lengths. Thus, channel diversity need not be used in such cases. However, as previously discussed, coding is necessary in many practical cases where a minimum successful frame transmission probability P_{min} is required. We now seek to assess whether channel diversity is beneficial when coding is used ($N > k$).

As a starting point, we consider specific values of k and P_{min} : $k = 10$, $P_{min} = 0.995$ or 0.980 . If only one channel is to be used, the sender calculates the maximum Effective Rate achievable over either of the two channels by optimally choosing the value of N to use over each channel. Then, the sender uses the channel and the corresponding code that result in the highest Effective Rate. In the case of path diversity, the sender optimally picks a code and channel selection policy so as to maximize the Effective Rate when using *both* channels. For the case of two identical channels (using three groups within the 55 channels described before that correspond to a Long-Term Error Rate (LTER) of 1%, 5% and 9% and different Expected Burst Length (EBL) values), Figure 2 plots the relative gain in Effective Rate (G_{ER}) as a function of the EBL and the LTER for two different values of P_{min} ⁵. We notice that as a general trend, when the EBL and/or the LTER increase, then the relative gain in ER also increases, i.e., path diversity becomes more useful⁶. Given that the relative gain can be as high as 70% (for the case of two GSM channels and $P_{min} = 0.995$, the gain was 47.60%), or barely noticeable, it is critical to identify the channel scenarios in which diversity does indeed help.




Fig. 2. Relative Gain in ER as a function of EBL and LTER for two different values of P_{min} .

We therefore now seek to explain *when* and *why* channel diversity is beneficial. Clearly, the benefits will depend on the long term error rates and the average burst lengths of the available channels. For example, if the long term error rate of one channel is significantly less than that of the other, then intuitively the optimal channel selection strategy will always select the former. Thus, a natural first step to investigate the benefits of channel diversity with two statistically identical channels. For two GSM channels and $P_{min} = 0.995$, for $k = 10$ and $k = 25$ the relative differences in ER over the benchmark policies with $p = 0$ and $p = 1$ are 47.5% and 24.41% respectively. Clearly, these benefits are significant. We now seek to understand why channel diversity is beneficial in this case. The availability of an additional channel allows the packet transmissions to “avoid” error bursts when they occur. In fact, using the previously mentioned 55 channels, in all cases where two identical channels were available, the optimal policy always distributed packet transmissions equally across both channels, i.e., $p^* = \frac{1}{2}$. The intuition behind this behavior is that by transmitting on both channels, the policy allows the sender to “escape” part of the error bursts. Moreover, when channels are identical, the sender is equally likely to hit a burst on either channel, so that there is no benefit in favoring one over the other.

The above results suggest that channel diversity increases the Effective Rate because it helps avoid error bursts. This would seem to imply that channel diversity would be more helpful with bursty channels, and less so when channels are not bursty. We now systematically examine whether this is indeed the case. We ask the question of what is the smallest burstiness (in terms of the EBL) that is needed, in order for

⁵ When $P_{min} = 0.98$, no coding is needed over the channels that have LTER of 0.01. Therefore, we only report the results for channels that have LTER of 0.05 and 0.09.

⁶ Note here that the relative gain is not always increasing as LTER or EBL increase. This is due to the discrete nature of the codes that are used. For example, a small increase in EBL could result in the path diversity system using a larger code than before, while still allowing the no diversity system to use the same code as before. This will result in G_{ER} to drop.

channel diversity to yield a 25% improvement when two statistically identical channels are used. Note that the number 25% has been selected as a benchmark improvement. The left part of Figure 3 plots the EBL needed for channels whose LTER varies between 1% and 10%. We notice that as expected the smaller the LTER, the more bursty the channels should be in order for channel diversity to provide substantial gains⁷. When the LTER of the channels is 1%, an EBL of only 5.2 is enough in order for the use of channel diversity to be meaningful⁸.

Next, we assess the benefits of channel diversity when channels are not statistically identical, i.e., they have different EBL and/or LTER. We observe that the use of a second channel, even a relatively bad one (with higher long term error rate and/or burstiness), can result in significant performance benefits. This second bad channel will obviously not be used as much as the first channel, but it will nonetheless allow the sender to “break up” extended periods of error bursts. We consider the case where one of the two channels is the GSM channel, and the other is not a GSM channel. For different values of LTER we identify the maximum EBL of this second channel that still allows for a 25% improvement in Effective Rate by utilizing channel diversity as compared to the case of just using the GSM channel (the right part of Figure 3). Again, we notice that a 25% improvement can be achieved even for a very poor second channel, e.g., when the second channel has a LTER of 0.30 (with EBL=3.2), or when the second channel is very bursty (EBL=937 when LTER=0.055). Notice that independent of the EBL, all channels with LTER less than or equal to 0.05 are able to increase the performance of the GSM channel by at least 25%. Obviously, all channels with LTER greater than 0.05 whose EBL is smaller than the channels shown in the right part of Figure 3 will also increase the performance of the GSM channel by at least 25%.




Fig. 3.

Left: Two identical channels with these characteristics result in a 25% relative gain in ER.

Right: Channels that when used together with a GSM channel improve its performance by 25%.

We now investigate how the benefits of path diversity vary as a function of the frame size k . First observe that for Bernoulli channels the optimal strategy is to transmit all packets over the channel with the lower long-term error rate. Thus, path diversity does not provide any advantage in this case. When k , and as a result the code length N , is much higher than the EBL of the channels, the difference between bursty and Bernoulli loss processes becomes less significant. We therefore expect that for a large enough k , path diversity will also not offer significant advantage. For example, for two GSM channels and $P_{min} = 0.995$, for $k = 25, 50, 100, 250, 500$ the relative differences in ER over the benchmark policies with $p = 0$ and $p = 1$ are 24.41%, 16.85%, 10.1%, 5.4%, 3.5%, 2.3%, respectively. However, note that large codes are complex to implement and can introduce unacceptable decoding latency. Thus, an important advantage of path diversity is that for any given frame length k , it helps attain the requisite performance with a smaller code length N . Furthermore, even if it decreases, the advantage that path diversity offers remains present up to relatively large frame lengths. The conclusions remain similar for non-identical channels. However, the magnitude of the benefits depend now on the individual channel characteristics. Figure 4 shows the optimal policy as a function of the frame length k for two fairly bursty non-identical

⁷ Note that the smallest required EBL for a 25% performance improvement is not monotonically decreasing as LTER increases. This is again due to the discrete nature of the code lengths that can be used.

⁸ Note that the EBL of a Bernoulli channel with LTER of 1% is about 1.01.

channels. The first channel has a smaller long term error rate (1% instead of 9%), and as a result as k becomes large, the optimal policy eventually converges to $p^* = 1$. The relative difference in Effective Rate is about 11.1% when $k = 10$, about 6.5% when $k = 20$ and about 4.7% when $k = 30$. The benefits of path diversity are somewhat less than in the case of two GSM channels. This is because the second channel has worse characteristics than the first channel, and as a result it is not used as much. But, for $k < 287$, the optimal policy still uses the second channel, and attains a higher Effective Rate as compared to the benchmark policy that uses only the “best” channel. This is yet another instance where adding a relatively bad second channel delivers improved performance.

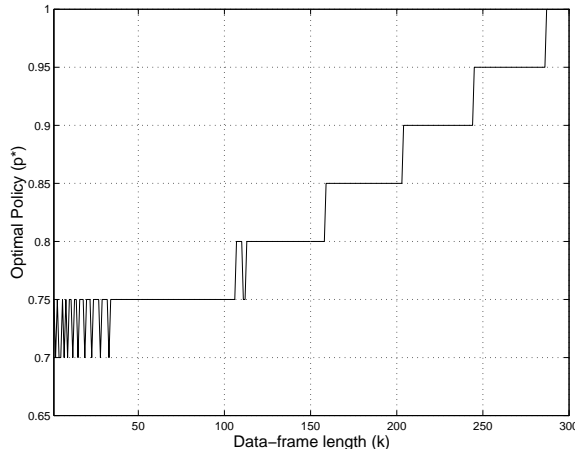


Fig. 4. The optimal policy (p^*) of the (N^*, k) code that achieved $P_{succ} \geq 0.995$ as a function of the frame length k . Both channels are bursty, with the second having a worse LTER (9% versus 1%).

The natural next question is to ask how the benefits of path diversity evolve as more channels become available. The work of other authors as well as our own intuition suggest that most of the benefits of path diversity are realized with just two channels. We are currently investigating this issue, and preliminary results appear to confirm that in most instances two channels deliver the bulk of the possible improvement achievable through path diversity. We also note that in all 55 cases of identical channels that we investigated, the average additional gain by using a third channel was about 5.13%, although in some cases it was as high as 20%⁹. We are currently in the process of systematically evaluating the benefits of using multiple channels, and pinpointing the specific channel characteristics where the use of more than two channels can yield meaningful further improvements.

7 Conclusion

This paper explored the potential benefits associated with channel (or path) diversity in improving performance in packet networks. The focus was on schemes that impose no or minimum added complexity, and do not increase the number of packet transmissions. Only long term channel characteristics were assumed known, and the channel selection strategies considered were limited to simple random policies that can be implemented with only minor changes to existing systems. We concentrated on scenarios that involve only two bursty channels, as bursty loss characteristics are common in many wireless and wireline settings, and intuitively represent conditions where diversity is most likely to help, i.e., by allowing the break-up of extended periods of consistently bad performance.

Our findings demonstrated that considerable improvements in the useful data throughput (Effective Rate) achieved by applications are commonly feasible. For some channel combinations, this improvement reached over 70%. These improvements were not the product of “exotic” channel combinations, e.g., they were present when using two standard GSM channels, and more often than not, adding a bad channel to a good channel was found to yield substantial gains. Our main contributions have been to provide computationally efficient procedures for computing policies that enable effective use of path diversity, and for broadly characterizing when it can be of benefit.

⁹ For the same set of 55 *identical* channels, the average gain by going from one to two channels was about 14.31%. The maximum gain was 70.35%.

Experimental validation using an IEEE 802.11 testbed is an area we are actively pursuing. Another extension we are working on involves generalizing the problem to consider how multiple users may optimally share multiple channels, when each one of them wants to leverage path diversity to maximize performance.

References

1. Rabin, M.: Efficient dispersal of information for security, load balancing, and fault tolerance. *Journal of the Association for Computing Machinery* **36** (1989) 335–348
2. Ayanoglu, E., I, C.L., Gitlin, R., Mazo, J.: Diversity coding for transparent self-healing and fault-tolerant communication networks. *IEEE Transactions on Communications* **41** (1993)
3. Herzel, F., Fischer, G., Gustat, H.: An integrated CMOS RF synthesizer for 802.11a wireless LAN. *IEEE Journal of Solid-state Circuits* **38** (2003)
4. Bolot, J.: End-to-end packet delay and loss behavior in the Internet. In: *Proc. ACM SIGCOMM'93*, San Francisco, CA (1993)
5. Salamatian, K., Vaton, S.: Hidden Markov modeling for network communication channels. In: *Proc. ACM SIGMETRICS'01*, Cambridge, MA (2001)
6. Kanal, L., Sastry, A.: Models for channels with memory and their applications to error control. *Proceedings of the IEEE* **66** (1978)
7. Konrad, A., Zhao, B., Joseph, A., Ludwig, R.: A Markov-based channel model algorithm for wireless networks. *Wireless Networks* **9** (2003)
8. Jaiswal, S., Iannaccone, G., Diot, C., Kurose, J., Towsley, D.: Measurement and classification of out-of-sequence packets in a tier-1 IP backbone. In: *Proc. IEEE INFOCOM'03*, San Francisco, CA (2003)
9. Paxson, V.: End-to-end internet packet dynamics. In: *Proc. SIGCOMM'97*, Cannes, France (1997)
10. Wang, F., Zhang, Y.: Improving TCP performance over mobile ad-hoc networks with out-of-order detection and response. In: *Proc. ACM MobiHoc'02*, Lausanne, Switzerland (2002)
11. Ji, P., Liu, B., Towsley, D., Ge, Z., Kurose, J.: Modeling frame-level errors in GSM wireless channels. *Performance Evaluation (Internet Performance Symposium - IPS 2002)* **55** (2004)
12. Diggavi, S., Al-Dhahir, N., Stamoulis, A., Calderbank, A.: Great expectations: The value of spatial diversity in wireless networks. *Proceedings of the IEEE* **92** (2004)
13. Laneman, J., Martinian, E., Wornell, G., Apostolopoulos, J., Wee, S.: Comparing application- and physical-layer approaches to diversity on wireless channels. In: *Proc. of IEEE International Conference on Communications (ICC)*, Anchorage, AK (2003)
14. Wang, J., Zhai, H., Liu, W., Fang, Y.: Reliable and efficient packet forwarding by utilizing path diversity in wireless ad hoc networks. In: *Proc. Milcom '04*, Monterey, CA (2004)
15. Zimmermann, E., Herhold, P., Fettweis, G.: A novel protocol for cooperative diversity in wireless networks. In: *Proc. of the Fifth European Wireless Conference*, Barcelona, Spain (2004)
16. Laneman, J., Tse, D., Wornell, G.: Cooperative diversity in wireless networks: Efficient protocols and outage behavior. *IEEE Trans. Inform. Theory* **50** (2004)
17. Pradhan, S., Puri, R., Ramchandran, K.: n-Channel symmetric multiple descriptions – Part I: (n,k) source-channel erasure codes. *IEEE Transactions on Information Theory* **50** (2004)
18. Puri, R., Pradhan, S., Ramchandran, K.: n-Channel symmetric multiple descriptions – Part II: An achievable rate-distortion region. *IEEE Transactions on Information Theory* **51** (2005)
19. Chakareski, J., Han, S., Girod, B.: Layered coding vs. multiple descriptions for video streaming over multiple paths. In: *Proc. ACM Int'l. Conf. on Multimedia*, Berkeley, CA (2003)
20. Mao, S., Lin, S., Panwar, S., Wang, Y., Celebi, E.: 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** (2003) 1721–1737
21. Miu, A., Apostolopoulos, J., Tan, W., Trott, M.: Low-latency wireless video over 802.11 networks using path diversity. In: *Proc. IEEE Int'l. Conf. on Multimedia and Expo (ICME)*, Baltimore, MD (2003)
22. Nguyen, T., Zakhori, A.: Path diversity with forward error correction (PDF) system for packet switched networks. In: *Proc. IEEE INFOCOM'03*, San Francisco, CA (2003)
23. Golubchik, L., Lui, J.C., Tung, T., Chow, A.L., Lee, W.J., Franceschinis, G., Anglano, C.: Multi-path continuous media streaming: What are the benefits? *Performance Evaluation Journal* **39** (2002)
24. Abdouni, B., Cheng, W., Chow, A.L., Golubchik, L., Lee, W.J., Lui, J.C.: Multi-path streaming: Optimization and evaluation. In: *12th Annual Multimedia Computing and Networking (MMCN)*, San Jose, CA (2005)
25. Tsigas, A., Haas, Z.: Analysis of multipath routing-Part I: The effect on the packet delivery ratio. *IEEE Transactions on Wireless Communications* **3** (2004)
26. Elliott, E.: Estimates of error rates for codes on burst-noise channels. *Bell System Technical Journal* **42** (1963)
27. Gilbert, E.: Capacity of a burst-noise channel. *Bell System Technical Journal* **39** (1960)
28. Ross, S.: *Stochastic Processes*, 2nd ed. Wiley (1996)
29. Yee, J., Weldon, E.: Evaluation of the performance of error-correcting codes on a Gilbert channel. *IEEE Transactions on Communications* **43** (1995)
30. Sharma, G., Hassan, A., Dholakia, A.: Performance evaluation of burst-error-correcting codes on a Gilbert-Elliott channel. *IEEE Transactions on Communications* **46** (1998)