

Design and Implementation of an “Approximate” Communication System for Wireless Media Applications

Sayandeep Sen[#], Syed Gilani[§], Shreesha Srinath[§], Stephen Schmitt[§] and Suman Banerjee[#]
University of Wisconsin-Madison
{sdsen,suman}@cs.wisc.edu[#], {gilani,ssrinath,sschmitt1}@wisc.edu[§]

ABSTRACT

All practical wireless communication systems are prone to errors. At the symbol level such wireless errors have a well-defined structure: when a receiver decodes a symbol erroneously, it is more likely that the decoded symbol is a good “approximation” of the transmitted symbol than a randomly chosen symbol among all possible transmitted symbols. Based on this property, we define *approximate communication*, a method that exploits this error structure to natively provide unequal error protection to data bits. Unlike traditional (FEC-based) mechanisms of unequal error protection that consumes additional network and spectrum resources to encode redundant data, the approximate communication technique achieves this property at the PHY layer without consuming any additional network or spectrum resources (apart from a minimal signaling overhead). Approximate communication is particularly useful to media delivery applications that can benefit significantly from unequal error protection of data bits. We show the usefulness of this method to such applications by designing and implementing an end-to-end media delivery system, called Apex. Our Software Defined Radio (SDR)-based experiments reveal that Apex can improve video quality by 5 to 20 dB (PSNR) across a diverse set of wireless conditions, when compared to traditional approaches. We believe that mechanisms such as Apex can be a cornerstone in designing future wireless media delivery systems under any error-prone channel condition.

Categories and Subject Descriptors

C.2.1 [Computer Communication Networks]: Network Architecture and Design—*Wireless Communication*

General Terms

Algorithms, Design, Experimentation, Measurement, Performance

Keywords

Cross layer, Media Delivery, Wireless PHY

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

SIGCOMM’10, August 30–September 3, 2010, New Delhi, India.
Copyright 2010 ACM 978-1-4503-0201-2/10/08 ...\$10.00.

1. INTRODUCTION

Any communication system, whether wired or wireless, is impacted by errors. Typically, a transmitter encodes application data in a frame as a set of symbols, where each symbol encodes one or more data bits. When errors happen, one or more of these symbols are incorrectly decoded at the receiver, leading to loss in performance.

Conventional design of networked systems have assumed that communication channels are error-prone and handle them in various ways. Often a checksum field is added to each data unit. A receiver verifies the received checksum against the expected checksum for the data bits to validate correct reception. Higher layers deal with these errors in different ways.

At the MAC layer, some protocols (such as 802.11) typically discard packets received with checksum errors. A newer class of MAC-PHY mechanisms attempt to recover correct bits from packets with partial errors. Examples include partial packet recovery (PPR) [18], SOFT [36], ZipTx [21], and Maranello [13]. Similarly, at the transport layer TCP would re-transmit erroneous or lost segments. UDP simply discards them silently. UDP-lite delivers erroneous packets to applications, and allows applications to recover correct portions of these packets using partial checksums.

In all the above techniques, it is the job of the receiver to determine and discard erroneous bits within a packet, and the transmitter has no a-priori knowledge of which bits (or bit positions) are likely to be in error. In this paper, we demonstrate an alternate possibility — *At the PHY layer, there is a systematic structure to wireless errors, and a wireless transmitter can leverage this structure to ensure that certain bit positions are less likely to be in error than other bit positions.* If the transmitter is aware of relative importance of data bits, it can place more important data bits in more protected positions and the less important data bits in less important positions. In other words, a transmitter can ensure an unequal error protection (UEP) of its data by simply placing data bits in specific bit positions. Thus, unlike prior approaches of UEP that explicitly need to add redundant bits to higher priority data [24, 35, 15], this proposed approach would not require any redundant data bits to be transmitted to achieve UEP, and is available *natively* from the wireless channel (apart from minimal signaling overhead).

Further, adaptation of the relative degree of UEP available to different bit-positions can be done based on application needs and ambient channel condition.

The relevant error structure across the wireless channel is a consequence of the following phenomenon: When a wireless symbol is decoded in error, this erroneous symbol is still a good “approximation” of the original transmitted symbol. (We will make this notion more precise in the next section.) Based on this approximation property, we design a wireless communication method to

provide UEP that we call, *approximate communication*. Approximate communication can provide significant performance gains to all applications in which different data bits have different levels of priority, e.g., the I-, P-, and B-frame structure in MPEG-4-AVC [3] encoded video. In particular, we design and implement a specific approximate communication system called Apex (an Approximate communication system for media exchange) that combines knowledge of relative bit priorities of the application with the structure of wireless errors and client feedback on channel conditions to achieve improved performance for media applications.

Apex requires a few, albeit marginal, modifications to the radio transmit-receive paths. These modifications include a few additional buffers, multiplexers, and a little amount of additional addressing logic. The latencies due to such additional data processing is negligible. Through detailed experiments on the WARP SDR platform [6] under various interference and multi-path scenarios, we demonstrate how Apex can improve the quality of video delivered over a wireless link, using a mechanism that is complementary to existing approaches.

Key contributions: We summarize the above discussion by identifying the main contributions of our work:

- *Apex, an approximate communication system to improve applications:* We show that wireless errors have a well-defined structure at the PHY layer. This structure allows the wireless PHY layer to natively provide UEP to data bits, which form the foundation of approximate communication. Further, we show that by suitably adapting system parameters, based on current channel conditions and application requirements Apex significantly improves performance of applications that prioritize different data bits differently.
- *Design of modifications to transmit-receive paths of an 802.11 radio to support approximate communication in Apex:* Apex requires minor modifications to the transmit-receive paths of a wireless radio. We describe how the transmit-receive chains of present day 802.11 hardware should be modified to implement approximate communication. We demonstrate such modifications can co-exist with other complementary mechanisms implemented in the lower layers, such as a scrambler system, convolution coding, data interleaving, data modulation, and PHY transmission rate selection.
- *Implementation of the Apex-based video delivery system:* We demonstrate that Apex can be practically implemented by building a prototype system using the WARP Software Defined Radio (SDR) hardware and running H.264 video streams through them over the air.
- *Experimentation over range of scenarios:* We show that Apex can, indeed, provide significant performance gains for media applications over a range of wireless communication scenarios. In different experiments, we varied the transmit power levels, the position of nodes, and the degree of external interference to demonstrate the robustness of our schemes. Over various experiment scenarios, it provided a video quality improvement of 5-20 dB(measured in terms of PSNR).

The fact that wireless symbol errors of popular modulation schemes have a specific error structure has both been analytically studied by researchers [23, 2, 38, 8] and experimentally demonstrated [30]. However, many popular wireless communication standards, e.g., IEEE 802.11, consider such a structure to be a nuisance. In fact, standards such as IEEE 802.11 typically employ a set of randomization and data protection mechanisms, (convolution coding, data




Figure 1: High level overview of wireless communication. (Some MAC layer operations, e.g., coding, scrambling, and interleaving are not shown.)

scrambling, and data interleaving) to ensure that all application bits are equally likely to be in error. These data randomization mechanisms have a number of other advantages. Hence, their wide adoption and specific placement in the protocol stack ensured that opportunities of natively providing UEP on the wireless channel was lost.

In this work we explain how the wireless symbol error structure can be leveraged to implement UEP in approximate communication (Sections 2), how it can be practically implemented in Apex (Section 3) to co-exist with the popular data randomization and protection techniques (convolution coding, data scrambling, and data interleaving) and how relevant parameters like modulation scheme and constellation mappings can be dynamically adapted (Section 3) to improve the quality of media delivered.

We believe that approximate communication can offer a useful design alternative for wireless media delivery systems and facilitate greater interaction between characteristics observable at the PHY layer and expectations of the applications.

2. INTUITION AND APPROACH FOR APPROXIMATE COMMUNICATION

Figure 1 shows a high level (and somewhat simplified) schematic of data transfer across a wireless link as it happens for common technologies today. Let us assume that the content is encoded into a data bit sequence using a popular format, e.g., MPEG-4 [3]. After packetization, the transmitter of the wireless link maps these bits into symbols for transmission across the wireless channel. In our example, each symbol represents a set of three data bits. The receiver attempts to infer the transmitted symbols, but sometimes makes decoding errors. Hence, when the received symbols are mapped back into a bit sequence, bit errors might result (erroneous bits and symbols are shaded). An important thing to note is that when a symbol error occurs, not all its constituent bits are actually in error. Symbol errors in common wireless technologies, e.g., those based on the popular I/Q modulation schemes, have a very well defined structure. When a receiver decodes a symbol erroneously, it is more likely that the decoded symbol is a good “approximation” of the transmitted symbol, than a randomly chosen symbol among all possible symbols. In this section, we explain this phenomenon through an example based on Quadrature Amplitude Modulation (QAM) ¹.

In QAM, data elements are encoded into amplitude values of two sinusoidal waves that are 90 degrees out-of-phase with each other. A QAM modulation scheme is usually represented by a I/Q constellation diagram, as shown in Figure 2. Each constellation point (or symbol) is mapped to the amplitude of the in-phased and the quadrature-phased signals, and corresponds to a certain bit sequence to be transmitted. In a 256-QAM scheme (shown in Fig-

¹Some of the different mechanisms manipulated in the paper might appear similar in meaning. Hence, for the sake of clarity, we define the use of our terms and their meaning in Table 1.

Term	Meaning / Examples
Modulation scheme	16-QAM, 64-QAM, etc.
Constellation map	Mapping from bit-sequence to symbols (Gray, Block, etc.)
Data bit placement	Mapping of data bits to bit positions in symbols

Table 1: Table summarizing some terminology used for different mechanisms used in the paper.




Figure 2: Quadrature Amplitude Modulation (QAM) constellations.

ure 2(a)), there are 256 distinct symbols. Each such symbol encodes a 8-bit sequence of data to be transmitted. In contrast, in a 64-QAM scheme, there are 64 distinct symbols, and each symbol encodes a 6-bit data sequence (Figure 2(b)). When a transmitter wants to communicate a specific bit-sequence, it emits the corresponding symbol. The goal of the receiver is to identify which symbol was sent.

Let us consider the case of a transmitter-receiver pair using 64-QAM and the transmitter emits a symbol, say C (Figure 2(b)). Depending on the conditions of the channel, there is a reasonable chance that the receiver will correctly decode this symbol. However, there is also some possibility for the receiver to make an error in decoding this symbol. It turns out that if an error does occur, such errors are most likely confined to the near neighborhood of the transmitted symbol C , i.e., the receiver is more likely to decode this symbol to be one of A , B , D , or E , which are the nearest neighbors of C (indicated by the inner circle of Figure 2(b)), as compared to a far away symbol, X . In other words, within the 2-dimensional I/Q space, an erroneously decoded symbol tends to be a reasonable approximation of the transmitted symbol.

As noise in the environment increases, the likelihood of making incorrect decoding decisions to “faraway” symbols can increase to some extent (as indicated by the larger circle in Figure 2(b)). However, likelihood of making errors within close proximity also increases, and continues to significantly dominate such faraway errors. We illustrate this behavior in Figure 3 in which 10 million 64-QAM symbols were transmitted between a transmitter-receiver pair under two different conditions, ambient noise, and explicitly generated interference in the background. The Y-axis represents the probability that the transmitted symbol is decoded erroneously to be another symbol at a distance given in the X-axis. The logarithmic scale of the Y-axis demonstrates the dramatic decay in this probability with increasing distance between the transmitted and decoded symbols.

Further, if the noise conditions continue to increase in this manner, a reasonable communication system will adjust to a lower rate modulation scheme, such as 16-QAM (Figure 2(b)), which separates constellation points further apart. This allows the system to revert back to the original scenario where erroneous symbols are mostly the nearest neighbors of the transmitted symbol.

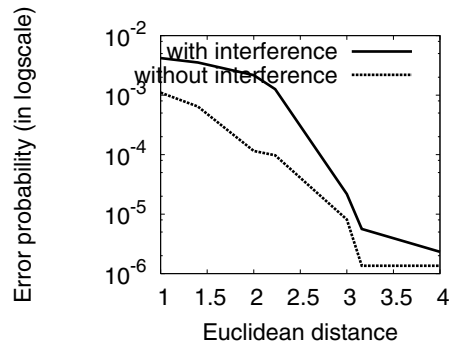


Figure 3: Fast decay in probability of erroneous decoding as the distance between the correct symbol and the erroneously decoded symbol increases. Data based on 10 million known 64-QAM symbols transmitted. In the ambient noise scenario, the overall BER was 3×10^{-4} . In the explicit interference scenario, the overall BER was 6×10^{-3} .

2.1 UEP and approximate communication

The wireless error structure demonstrated in Figure 3 leads to a construction of UEP for data bits. This is achieved by appropriately choosing a constellation map, i.e., the map between bit-sequences and symbols in a constellation diagram.

For the sake of simplicity, let us assume that when symbol errors occur, the erroneous symbol is just one unit away from the actual transmitted symbol. In the 64-QAM example in Figure 2(b), this assumption implies that when C is transmitted, a decoding error would result in the received symbol to be one among A , B , D , and E .

Now imagine if a constellation map was to map bit-sequences randomly to symbol positions. Then, when an error occurs, the likelihood of a correct value in any given bit position is purely by chance. The probability of such an event is 0.5, i.e., the probability that a neighboring symbol has the same most significant bit (MSB) value as the MSB of the actually transmitted symbol is 0.5, and the same is true for the least significant bit (LSB) and all other bit positions in between. Such a constellation map does not help in achieving our desired goal in approximate communication.

Fortunately, various common constellation maps, e.g., the Gray code (used in 802.11) map bit sequences to symbols in a way that increases the resilience of certain bits, even when the symbols are in error. As an example, the Gray code corresponding to a 16-QAM constellation is shown in Figure 4. If we examine the MSB of the different symbols to the left of the Q-axis, we can observe that all of them have a value of 0. Similarly, the MSB of different symbols to the right of Q-axis all have a value of 1. In such a case, if a decoding error occurs for any symbol within the shaded region, there are no errors in the MSB (under our assumption that symbol errors do not exceed 1 unit). For the remaining symbols, we expect an error probability of either 1/4 or 1/3 for the MSB, depending on its position in the constellation. If the likelihood of transmitting each symbol is identical, then the probability of error in the MSB, given a symbol error has occurred, is 1/6. For the LSB, on the other hand, the probability of making an error, given a symbol error has occurred, is 1/3. This gives rise to an intriguing possibility. If an application identifies different priority levels for its various data bits, then instead of performing data bit placement (mapping these data bits to bit positions in symbols) in an agnostic manner, we can achieve the desired impact of differential data protection by placing the higher priority data bits to the MSB positions of symbols, and the lower priority data bits to the LSB (and other) positions.

Note that such differential protection of data values is possible by

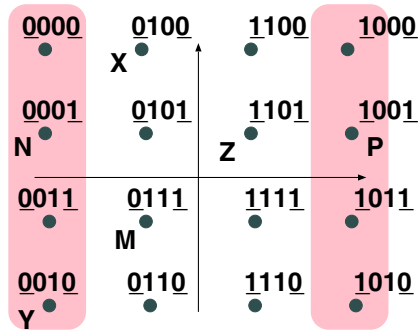


Figure 4: 16-QAM with Gray code constellation map. MSBs and LSBs are underlined.

simply placing the data bits to appropriate bit positions, and without adding any form of redundancy into the system. In a way, such data protection is available natively through the wireless channel, and can be better exploited by applications, without imposing any overheads. This is in contrast to the traditional FEC-based methods for providing UEP that require communication of additional bits that redundantly encode more important data.

Also, while we explained the phenomenon of differential error resilience of bit positions with the example of QAM modulation the same would hold for other modulation schemes such as Phase Shift Keying (PSK), Pulse Position Modulation (PPM), Frequency Shift Keying (FSK) etc. as well.

Based on the above intuition, we build an approximate communication system that can provide differential protection to different data bits of an application. We use streaming of MPEG-4 encoded video as an example of such an application. Such a video stream consists of frames with different levels of importance, e.g., I-frames are more important than P-frames, which in turn are more important than B-frames.

Akin to regular communication systems, the application creates and continuously sends data bits corresponding to different frames (I, P, and B) down the network protocol stack. Bits from each video frame will be packetized in smaller chunks at the network and MAC layers, and then will be handed into the digital component of the PHY layer. This is shown in Figure 5, with four video frames — two I-frames with data 0100 and 0010, a P-frame with data 1101, and B-frame with data 1001. In traditional communication systems, the 4-bit sequences of each frame gets placed together in a single wireless frame, e.g., the first I-frame is mapped to the first wireless frame (consisting of a single symbol X). The key difference in approximate communication is that data bits of different priority levels will be placed together and combined into a single symbol, such that the most protected bit positions are occupied by the higher priority data bits (say, I bits), and the least protected bit positions are occupied by the lower priority data bits (say, P and B bits). For example, the first two bits of the first I-frame, i.e., 01, are mapped to MSBs of the first symbol of a wireless frame, while the second two bits of the same I-frame, i.e., 00, are mapped to the MSBs of the symbol of the next frame². The bits of P- and B-frames are placed into the LSB positions of different symbols.

²In our actual implementation, the I-, P-, and B-frame bits would first go through the usual 802.11-style digital PHY processes, including the scrambler, convolution coder, and the interleaver for added resilience, before being mapped to the symbols.

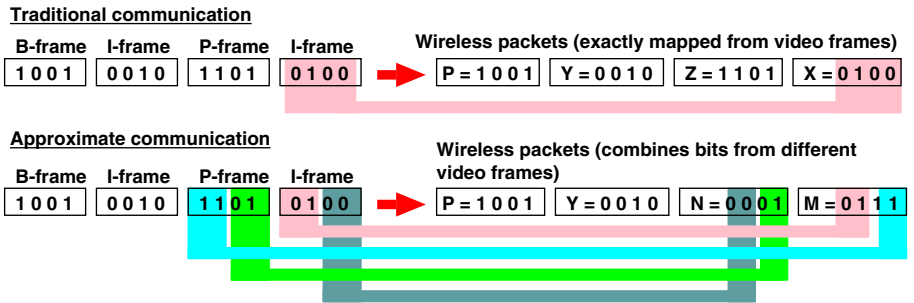


Figure 5: Wireless frames in approximate communication combines bits from video frames of different priority levels. Example uses 16-QAM and assumes that the two MSBs can be better protected than the two LSBs. The data bit placement strategy maps I-bits MSBs and P- and B-bits to LSBs. Symbols M , N , P , X , Y , and Z are as shown in Figure 4 that uses a Gray code constellation map. (Only subset of data bit placements shown.)

Overall, this means the I-frame bits will be better protected than the P- and B-frame bits.

When symbols are decoded (possibly in error), the receiver simply passes them along to the higher layers, even if in error. When the different application bits are extracted out of these erroneous symbols, the highest priority bits are most likely to be correct and are of great value to the application. Lower priority ones maybe in error and will be suitably handled. The key observation is that the higher priority bits are, thus, successfully extracted out of the symbols, even if in error, through this mechanism, without requiring re-transmissions.

This differential protection is in contrast to the traditional communication method used in 802.11a/b/g, where I-frame bits occupy both MSB and LSB positions of a symbol, the same as P- and B-frame bits, and there is no difference in the level of protection for bits of different video frames.

To validate that our approach of approximate communication is, indeed, possible, we need to identify one or more constellation maps that provide unequal error protection among different bit positions, *across a range of wireless conditions* — *different transmit power levels, different degrees of interference at the receiver, and with and without PHY layer convolutional codes*. Such confirmation of unequal protection will provide further evidence that our approximation property holds true for a range of common scenarios, where approaches such as Apex can be quite worthwhile.

Given a constellation diagram of k points, each symbol will represent a bit-sequence of length $\log_2(k)$ bits. Overall, there are no more than $k!$ different constellation maps possible, although some of them can be identical due to rotation and reflection based symmetries. Each constellation map is likely to protect individual bit positions differently. The above property can be leveraged by altering the constellation map during an ongoing transmission, based on the channel conditions and the relative priority of the application data-bits.

In this section, we present results for two constellation maps, each of which offer varying amount of error protection for different bit positions.

Both are, therefore, suited for approximate communication. We note that our contribution is not in designing these constellation maps, but in identifying suitable ones and in exploiting them to implement approximate communication for media applications. We first briefly describe the schemes, and then evaluate their error protection properties in the rest of this section.

Modulation schemes	16 and 64 QAM
Encoding schemes	Gray, Block
Frequency	2.4 GHz
Convolutional codes (when used)	Generator polynomials: $g_0(133)$ $g_1(171)$, Rate = 1/2, 2/3, 3/4
Base data rate	485 Kbps (at rate 1/2)

Table 2: Properties of our implemented system.




Figure 6: Setup for communication between the two WARP SDR radios. The relative position, transmit power levels, and degree of external interference varied during different experiments.

2.2 Example constellation maps

We evaluate the following two constellation maps for use in approximate communication.

Gray code: In Gray code [12], symbols that are immediate neighbors either along the I-axis or the Q-axis of the I/Q space differ in exactly one bit position; the rest of the bits are identical between the two neighbors. An example of a Gray code is shown in Figure 4, for a 16-QAM system. Gray codes can be constructed for any QAM scheme in a systematic manner based on the above observation. Gray code is widely used in many popular communication systems, including in 802.11 a/g/n and 802.16 based systems.

Block code: In block code [23], the constellation points on the same side of the I (Q) axis have the same value for the first (second) bit position. Hence, all points in a given I-Q quadrant have the same value for the first two bit positions. For each of these quadrants, the symbols are partitioned into four sub-quadrants, and the same process is repeated to assign bit values for the next two bit positions. The process is repeated iteratively for the remaining positions.

Both of these schemes can be implemented on different QAM based modulation schemes. While we have implemented and experimented with different QAM schemes in our SDR radio platform, for sake of uniform comparison, in this section we focus on the 64-QAM scheme. In 64-QAM, each symbol encodes 6-bit sequences, and in diverse experiments we observe how the different schemes differentially protect the two most significant bits (MSBs), the two middle bits (MID), and the two least significant bits (LSBs).

2.3 Experiment configurations

We use the following experiment parameters.

Hardware and software: We implemented our system on the WARP SDR radio platform [6]. In our implementation we used the WARP boards as RF front-end to transmit the packets over air and carried out the digital layer PHY signal processing activities in the connected laptop. The laptops also acted as the traffic source and sink. The setup for our experiments is depicted in Figure 6. We summarize various aspects of our implementation in Table 2. Frequency and phase synchronization between the sending and re-

ceiving side was achieved using a 11 symbol barker sequence sent as a preamble along the packet. We experimented by varying the transmit power levels of the transmitter, adding external source of interference and by changing the location of the nodes. The receiver sensitivity of WARP hardware in our implementation is -70 dBm for 64-QAM and at lower RSSI values (< -70 dBm)³ it becomes difficult to derive any meaningful results. Hence, for all of transmit power variations, we report on RSSI values that varied between -30 and -70 dBm.

Metrics: In this section, we are interested in the relative bit error rates (BERs) experienced by different bit positions of a symbol. Hence, we compare the two constellation maps by examining the BERs of the MSBs, MIDs, and LSBs across different conditions. In certain plots, we report *BER gain* which is the ratio of the BER of Block code to that of Gray code, for specific bit positions. BER gain can, thus, be a number greater than 1.

Experiment parameters: In each experiment we sent 10 million randomly generated symbols. For each symbol received, we decode its value using both maps, and calculate the error rates for different bit positions.

A typical 802.11 PHY layer (optionally) applies convolution code (a PHY layer FEC) to further protect the data-bits. In order to evaluate the effect of such PHY layer convolution codes, we have also experimented with (and without) different convolution codes (rates include 1/2, 2/3, and 3/4) as well. Our results show that unequal error protection property holds both with and without convolution codes. In this paper, we show the error resilience of different bit positions in presence of PHY layer convolution coding only for a convolution code of rate 1/2.

2.4 Experimental results

We describe our experiments and associated observations below.

With and without PHY convolution codes: We plot the error rate at the MSBs, MIDs, and LSBs for the two constellation maps at an intermediate transmit power level (RSSI -55 dBm) in Figure 7, both with and without PHY convolution codes. Each constellation map provides different levels of protection to these bit positions. In both cases, Block provides the greatest error protection to the MSBs at the cost of lower protection to the LSBs. The use of convolution codes clearly reduce the BERs suitably, but the effect of unequal error protection exists both with and without use of convolution codes. Hence, in the rest of this paper, we focus on results with PHY convolution codes applied.

Ambient noise, varied transmit power levels, with PHY convolution codes: We present the performance of the Gray and Block constellation maps for an entire range of received power (varied by changing the transmit power levels as well as the separation between nodes and their locations), when using a convolution rate of 1/2 in Figure 8. Instead of the BER, we plot the *BER gain*. With decreasing RSSI, symbol errors increase. As a consequence, the BER of Block shrink in comparison to BER of Gray for MSBs, i.e., Block protects MSBs even better compared to Gray as RSSI decreases. For the range of operating parameters the relatively greater protection of MSBs using Block changes from a factor of 2 (at -30 dBm) to a factor 15(at -70 dBm). The increased protection of MSBs comes at a cost of decreased protection for MIDs and LSBs in Block.

Varied external interference: We next examine the behavior of different encoding schemes in presence of an explicit source of external interference (Figure 9). We had used a laptop sending

³We calculate the RSSI values by reading the MAX2829 ADC output register value and mapping it to corresponding RSSI value from the Data-sheet [7].

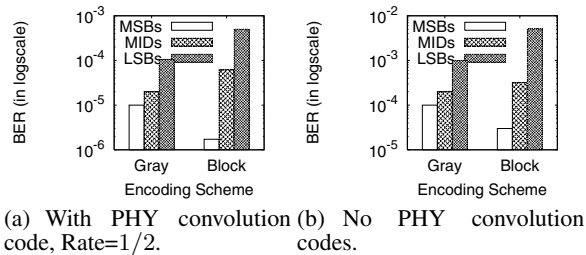


Figure 7: BER of various bit positions when using Gray, and Block data coding schemes. The overall uncoded BER of during the experiment was order of 10^{-3} for both schemes. 10 million QAM-64 symbols were transmitted.

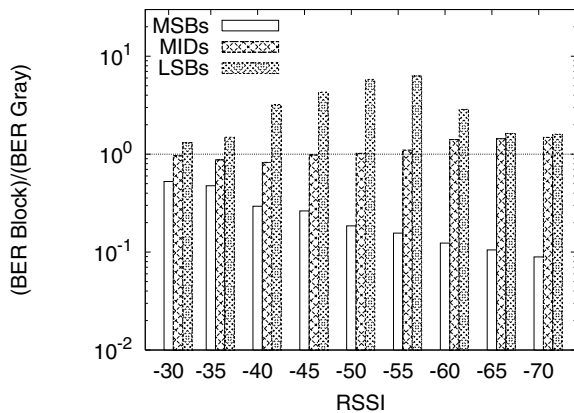


Figure 8: Variation in BER gains at different bit-positions with Block relative to Gray with varying RSSI, and with PHY convolution coding at rate 1/2. The horizontal line corresponds to the case where both schemes offer equal error protection for a given bit position.

WiFi traffic as the source of interference. As a nearby WiFi source, would back-off on sensing transmissions from our system, we had to place our interfering source with some care. To turn the WARP radio and the interfering laptop into a hidden nodes for each other. For this, we reduced the signal strength on the laptop and placed it at a far by location such that the achieved throughput for both systems at their highest data-rates (manually fixed) was same regardless of whether the other node was working or not. We then increased the power on the laptop to a level such that it started interfering with the Warp node (confirmed by increased errors when the laptop was transmitting).

We show the BER gain in the plot. The external interferer injects regular 802.11 traffic into the wireless medium, which causes significant symbol errors in the data transmitted between WARP nodes. Across all levels of interference Block provides significantly higher protection of MSBs than Gray.

3. DESIGN & IMPLEMENTATION OF APEX

The previous section validates that unequal error protection for different bit positions is consistently possible across very diverse wireless channel conditions. In this section, we describe our specific approximate communication scheme that can leverage these

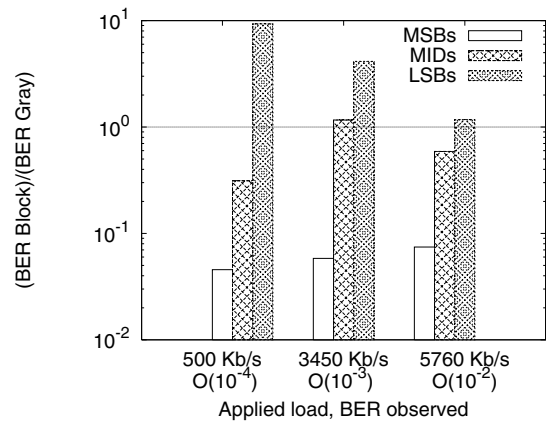


Figure 9: Variation in BER gain at various bit positions of Block relative to Gray in presence of an external interferer. Convolution coding at rate 1/2. The horizontal line corresponds to the case where both schemes offer equal error protection for a given bit position.

specific constellation maps to improve the performance of media applications in a natural way.

In particular, we focus on video streams in which content is partitioned into multiple classes, and each class has a different levels of importance in correct decoding of the stream. This is commonly observed in all MPEG-based video streams in which a “Group of Pictures” (GOP) consist of a sequence of I-, P-, and B-frames that have been coded together.

The application (or an application helper) provides a simple estimate of the relative priority value of these different data bits and passes it down the stack. There are many effective ways to compute priority of data bits to an application [22, 32], and in this paper we use a relatively simple one based on decoding dependencies. More specifically, in Apex the *value* of any data unit i is proportional to the number of data bytes, that require i for successful decoding at the receiver. In our case, we use frames within a GOP as a data unit. Hence, all bytes of a given frame type in a GOP has the same value, making it easy to perform the computation once for each GOP. By design, I-frame data is necessary to decode the highest number of dependent frames, and hence they have the highest priority. B-frames have no other dependent except themselves, so they have the least priority. P-frame has a priority which is intermediate of the two other frame types. The relative priorities of the three frame types can change between different GOPs. As mentioned before, better ways of defining the value of different bytes are, indeed, possible, we choose to use our simple approach, as it is sufficient to illustrate the performance of approximate communication.

We also note, that our approach generalizes to any other media encoding, where the content is structured in layers, and there is different levels of priority (value) for each layer.

The MAC-PHY layer of Apex is provided with a few different constellation maps. In our case, we use two alternatives — Block and Gray. Based on the application-specified value of different data units, the MAC and PHY layers of Apex makes two simple decisions: (i) given a constellation map, how to place application data bits to bit positions for the desired levels of protection, and (ii) which specific constellation map to use. We discuss how each of these components is implemented by making small changes to an 802.11-style PHY-MAC subsystem, while retaining all its features.

Placing application data bits to bit positions: Each constel-

lation map provides varying degree of protection for different bit positions of a symbol. However, once a constellation map has been chosen, there is only one optimal way to place application data bits to various protection levels to maximize value of protected data — *greedily*. Given x data bits across all priority levels, and m bits per symbol, the total number of symbols available is x/m . We start by placing the highest priority bits to the most protected bit positions of each of these symbols. We move over to the next protected bit positions, if we exhaust the most protected bit positions. Once we complete assigning all highest priority bits, we move to bits at the next priority level. We continue until we exhaust all bits. Unlike a traditional communication system, where each wireless frame carries data bits from a single video frame, *in approximate communication, each typical wireless frame carries data bits from an I-frame, a P-frame, and a B-frame*. The I-frame bits within the wireless frame will be most protected, while the B-frame bits will be least protected. Note that the above greedy approach does not require that the number of protection levels available in different bit positions be equal to the number of desired priority levels in application data, and it just provides a relative ordering in the level of protection for these bits.

Also, the above approach ensures that we do not need a fixed ratio of I-, P- and B- frame data. As the excess data bits belonging to one priority level can be sent by using all bit positions in a symbol.

Choosing a constellation map: As Section 2.1 shows, constellation maps differ in the level of protection available at different bit positions. Hence, the choice of a constellation map depends on the relative *utility* of protecting different application bits differently.

In order to evaluate the utility of a constellation map, we need to estimate the BERs at different bit positions for different constellation maps. For this, we add a small number of well-known pilot symbols into each wireless frame (1 pilot byte for every 100 bytes). The receiver will decode these symbols with the different constellation maps in consideration, e.g., Gray and Block. Since the pilot symbols are known, the receiver can estimate the BERs for different positions and provide this as a feedback in the 802.11-style ACK frame. Note that such BER computation occurs at transmission speeds, as such computation already occurs for wireless frame decoding and ACK generation. We use a single byte to carry each BER value of the MSBs, MIDs, and LSBs of 64-QAM, for a total overhead of 6 additional bytes for both schemes in the ACK frame. (16-QAM and 256-QAM has two and four protection levels respectively, and the overheads in the ACK frame are 4 and 8 bytes respectively.)

With this information, the transmitter can calculate the utility, $u(\mathcal{E})$, of a constellation map, \mathcal{E} , as follows. Let v_i be the *value* associated with application data units at the i th priority level. Let p_j indicate the protection level of the j th significant bit position group, and let x_{ij} be the number of application data units with priority level i be allocated to bit positions with protection level j (using the greedy approach). Then, $u(\mathcal{E}) = \sum_i v_i \sum_j p_j x_{ij}$, where v_i is presented as an input from the application, and is computed once for each video stream GOP. x_{ij} is obtained from our bit mapping strategy, and can also be computed per constellation map, once for each video stream GOP, and p_j is received as feedback in the ACK frame. We pick the constellation map that maximizes this value. We illustrate the benefits of switching constellation maps as part of our evaluation in Section 4. This computation itself is quite fast. Hence, in principle, the constellation map can be changed for every frame. However, in our implementation the decoding process for the different constellation maps happen, not in the FPGA of the WARP radios, but in the associated laptops. This adds latencies to the process such that our constellation map usually changes at a




Figure 11: Modifications to PHY PLCP header and payload for Apex.

rate of once every ten packets or more. Our evaluation illustrates the performance advantages of this dynamic choice in spite of this latency. The actual gains of this component in a real system is, therefore, likely to be even better.

Interaction with choice of modulation scheme and PHY rate of transmission: Approximate communication has a direct interaction with the choice of data modulation scheme and the PHY layer transmission rate of data units. In general, any good rate and modulation selection scheme will ensure that the approximation property of symbol errors hold. In our work, we implement a rate selection scheme that is based on the SoftRate algorithm [34], which is known to be fairly agile and accurate. Each change in the rate triggers an immediate re-evaluation of the specific constellation map, although the constellation map might sometimes change at an even faster rate, if necessary due to change in channel conditions. We anticipate that the performance of the system can be improved even further if the rate selection decision is combined with the constellation map selection. In this paper we do not explore this joint problem further, and relegate it as part of our future work.

Modifications to PHY PLCP header and payload: A few modifications need to be made to the PHY layer PLCP header. First, we added a two-bit constellation map selector within the PLCP header, to inform the receiver which encoding scheme is used in the wireless frame. This limits the number of constellation map alternatives to four, which we believe should be sufficient for most applications. We also added information on how the data of different priority levels are placed into various bit positions using our greedy approach. This can be simply expressed by indicating the number of bytes (expressed in multiples of eight bytes) in each level (the greedy algorithm for placing data bits to bit positions can be used to partition the transmitted data). We limit ourselves to four priority levels for application data currently. We add a one-byte field in the PLCP header for each level, for a total of four additional bytes. This limits the maximum payload size to 2048 bytes.

Note that only the two-bit constellation map selector is placed in the early part of the PLCP header to be transmitted at the base data rate, e.g., 802.11a uses BPSK with 1/2 PHY convolution codes, and using a pre-defined constellation map (Gray). The four priority level size fields are in the latter part of the PLCP header can be transmitted at higher data rates like the rest of the frame. Hence, although the PLCP header increases from 40 bits to 72 bits, only two of these additional bits need to be transmitted at the base rate. Finally, we add the pilot symbols through the data field for BER estimation of different bit positions with different constellation maps. This is presented in Figure 11.

Modifications to the digital component of the PHY layer pipeline: We describe modifications needed to a standard 802.11 PHY layer pipeline to implement all features in Apex, and present it pictorially in Figure 10. All components in the picture are from the 802.11a

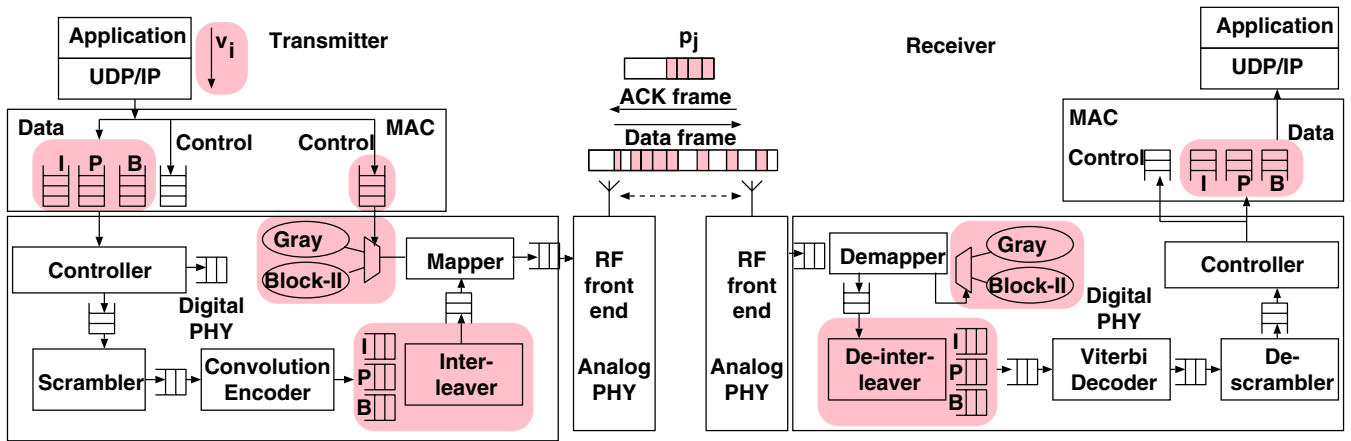


Figure 10: Changes to basic 802.11 layers for approximate communication in Apex. I, P, B labels are for illustration purpose only indicating three arbitrary priority layers. Shaded parts are modifications made.

reference pipeline and the shaded parts of the picture indicate locations where some changes are needed.

In the transmitter, data is available to the PHY pipeline in multiple queues, one for each priority level present. Our system allows for four priority levels and in the figure we show only three queues populated with data (corresponding to I, P, and B frame data). In a regular 802.11a pipeline, all data coming into the controller will be a single queue and for a regular sized Ethernet-style frame, may have a maximum payload of 1500 bytes. In our case, we expect the total bytes corresponding to a single wireless frame, still to be 1500 bytes, but this data may consist of a different number of bytes from I, P, and B queues. Each of these I, P, and B fragments will pass through the scrambler and the convolution encoder *independently* like in a regular 802.11 pipeline, and arrive at the interleaver. Note that, at the interleaver, the scrambled and the convolutionally encoded data are still separated according to their priority levels, and are placed in separate queues.

In the interleaver, the bits will be placed into bit positions based on the greedy algorithm. Each bit-sequence, thus generated, will correspond to a symbol and will likely consist of bits collected from the three different queues. Like a regular 802.11 interleaver, we will allow interleaving of bit values across different OFDM sub-carriers (if multiple OFDM sub-carriers are used). But unlike a regular interleaver, we will only limit interleaving of bits within a single sub-carrier to within the bit positions allocated to data from the same priority class.

Finally, based on the application data value and the BER feedback from the receiver, one of the data constellation maps will be selected. The bit-sequences will be mapped to the corresponding symbol values using a simple table lookup process as in a regular pipeline, and handed over the RF front-end.

A similar, but reverse process would occur in the receive chain.

4. EVALUATION

We now evaluate the performance benefits offered by Apex in detail. In addition, we also show that the advantages of approximate communication are complementary to some alternate mechanisms of adding unequal error protection, e.g., application-layer FECs, and data re-transmissions.

4.1 Experiment configurations

We describe the new aspects of our experimental configuration vis-a-vis Section 2.1 for evaluating performance of video streaming

MOS Rating of video quality	PSNR range
Excellent	> 37
Good	31-37
Fair	25-31
Poor	20-25
Bad	< 20

Table 3: Table mapping the MOS based user perception of video quality to the PSNR range [25]

in Apex. We use the same WARP SDR platform, with the signal processing functions implemented in laptops connected to these radios.

Video content: We use Claire and Foreman video clips, encoded to MPEG-4 format using *ffmpeg* tool. The rate of encoded video is 384 kbps, frame rate is 30 frames per second and each of its GOP consists of 30 frames. We induce a buffering of 1 second to collect all the frames in a GOP. The video is looped multiple times to get a playback length of 500 seconds. Each experiment is repeated 20 times unless otherwise stated. We use Evalvid tool [4] to stream the video. The tool has been modified to provide information about frame-type and priority, while streaming.

Metrics: We have compared the quality of the received video both visually and using a widely used metric — peak signal to noise ratio (PSNR). The PSNR of a video is well correlated with the perceived quality of video experienced by the user. The relationship between user perception expressed in Mean Opinion Score (MOS) and the PSNR range were detailed in [25, 17] and are summarized in Table 3.

Schemes compared: We compare the performance of an approximate communication (Apex) and a traditional communication (Trad.) system (which is unaware of the relative protection levels of the different bit positions), when both are allowed to transmit the same amount of application data. We allow partial recovery of bits not in error [18].

We compare the relative volume of re-transmissions needed for the Apex and traditional systems to achieve similar PSNR. We also measure the additive gains of Apex in presence of application layer FEC.

Apex is complementary to other data protection mechanisms that can be implemented at the different layers, e.g., MAC-layer mecha-

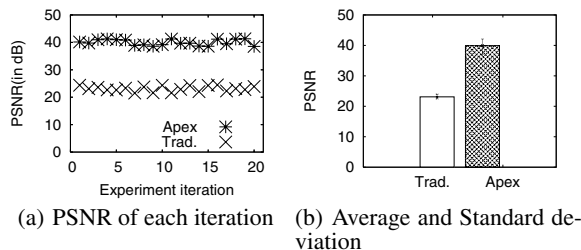


Figure 12: PSNR differences between traditional and approximate communication with rate adaptation at a representative wireless environment (Claire video).

nisms (ZipTx and MRD), PHY-MAC mechanisms (PPR and SOFT). It was challenging for us to implement all of these schemes within the experimental WARP SDR platform. In Section 5 present an intuitive explanation of why other schemes, PPR, SOFT, MRD, and ZipTx, are also expected to provide gains that are complementary to Apex.

Trace-based evaluation for dynamic rate adaptation scenarios: In our SDR platform (WARP radios) a high latency is associated with processing RF samples from a received wireless frame and sending feedback to the transmitter for efficient rate adaptation. To avoid this latency related inaccuracy in evaluation, we adopt the practice of trace-driven analysis as is commonly applied in evaluating rate adaptation algorithms with PHY layer processing needs [34].

Symbols carrying application data are sent at different rates in succession. We iteratively send small packets of 200 bytes at all the different rates for building the trace. Also, we ensure that the BER is monotonically increasing with increasing data rate for a iteration (discarding samples for which this criteria does not hold) signifying a coherent channel [34]. At the receiver we calculate the achieved throughput by different data-rates and only those symbols that correspond to the rate selected by the rate adaptation algorithm are considered to be part of the media flow. As the selected rate varies over time, the symbols selected to be part of the flow at different time instants are picked accordingly. The actual performance of the media delivery process can then be evaluated by identifying loss behaviors due to errors in an off-line manner.

4.2 Experimental results

In this section, we show how Apex helps improve the media quality under different interference scenarios, how the gains of Apex are significant when compared to FEC-based or re-transmission based data recovery schemes, and describe some parameter selections made in the system.

Apex vs traditional — a representative scenario: We present the relative performance for traditional media delivery and the Apex scheme with dynamic rate and constellation map adaptation enabled for the same channel conditions as above in Figure 12. The average uncoded BER for this scenario was 4.1×10^{-3} (corresponds to a coded BER of $O(10^{-5})$, which is normal operating condition for 802.11 based radios). The average PSNR improvement due to Apex across a set of 20 runs was about 16 dB. The better performance of Apex stems from the fact that a throughput optimal data delivery can be further improved upon by the importance-aware mapping of data bits to bit positions within symbols.

We next illustrate how the dynamic adaptation of constellation map is useful to improving the performance of Apex.

Impact of constellation map selection: The best choice of constellation map depends both on the channel conditions (p_j values)

Video	Trad.	Apex (Gray)	Apex (Block)
Claire	21.8	29.5	37.8
Foreman	21.0	34.5	31.0

Table 4: Table showing the PSNR for the two video clips in the representative wireless scenario with statically chosen constellation map.

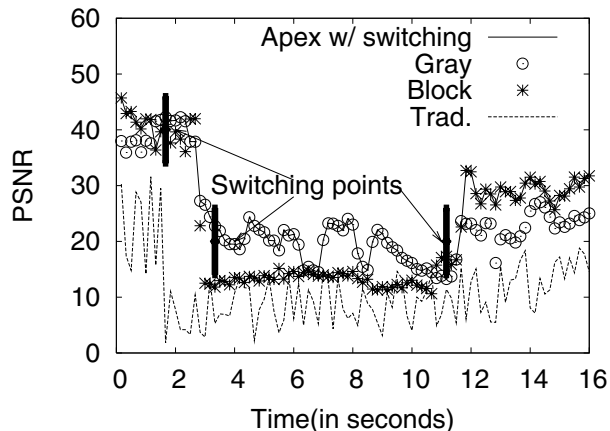


Figure 13: Dynamic switching of data encoding scheme computed based on the utility of each encoding alternative (Block and Gray).

and the differential value of application bit (v_i values). channel conditions. We demonstrate this in turn.

- *Dependence on application data:* For the same representative scenario as in Figure 12, we show the relative performance of traditional and Apex for two different video clips, Claire and Foreman, in Table 4, but with statically chosen constellation maps for the entire duration of the clip. Apex outperforms the traditional system in all cases. But the relative amount of improvement depends on the specific constellation map selection. Due to the distribution of bytes between I-, P-, and B-frames in these clips, the relative value of different data bits are different in these two schemes. As a consequence, Gray code leads to a better average PSNR performance for the Foreman video than Block code, while the reverse is true for the Claire video. By dynamically selecting the most appropriate constellation map, Apex would be able to deliver the best performance among various alternatives.

- *Dependence on channel conditions:* We next illustrate how the dynamic adaptation of the constellation map occurs in practice, also due to changing channel conditions (Figure 13). The figure shows how the dynamic choice of the data encoding scheme ensures that the best encoding scheme is picked as the quality of the channel changes. In particular, our algorithm decides to switch constellation maps around time 1.8, 3.3, and 10.9 seconds, soon after the PSNR quality using the constellation map falls below the other alternative. We also show the performance of traditional communication system in the figure for the sake of completeness. We find that Apex performs better than traditional communication at all times.

Currently, our feedback latencies are significantly higher since the decoding operation with different encoding schemes are performed in the associated laptop, incurring high latencies, although they are adequate to provide performance gains. A future version

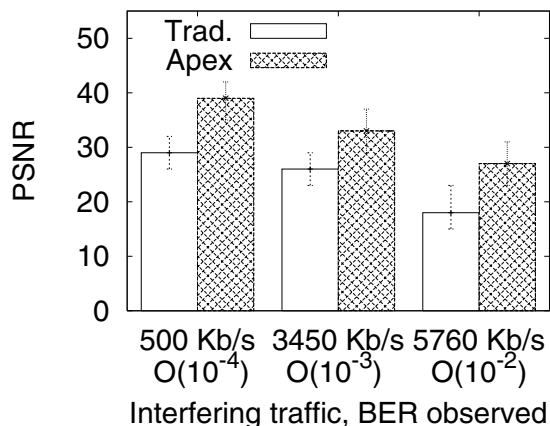


Figure 14: PSNR for traditional and approximate communication under varying external interference (Claire video). X-axis marks both the imposed (802.11) interference load on the channel and the corresponding average uncoded BER.

of the system will implement the feedback process in an FPGA, leading to a more agile response and potential performance gain.

With external interference: We next experimented by changing the level of external interference generated by a third (laptop-based) node in the vicinity, as described in Section 2.1. As expected, both schemes suffer with increasing interference. However as the channel condition worsens, due to increasing interference, the relative gains of Apex over traditional communication increases (Figure 14). At the higher interference levels, the PSNR in Apex is about 1.8 times that of traditional communication. The higher gains of Apex can be attributed to informed placement of application bits to bit-positions of symbols.

Complementarity to application-level FEC: We show that gains of Apex communication are complementary to application-level FECs through evaluation of traditional and approximate communication, both with and without such FEC. We used a Reed-Solomon RS(255,233) code for application-layer FEC, where the amount of coding was determined based on ambient channel conditions as suggested in related efforts [9, 21]. In Figure 15, we show the performance comparison for a scenario where traditional communication achieved the highest possible PSNR. The plot shows that Apex achieves better performance than traditional communication both in presence and absence of application-layer FECs, and the gap in PSNR between the two approaches is somewhat similar in both cases.

Comparison with a MAC layer scheme: Consider a MAC-layer variant of Apex as follows. Given an optimal PHY rate discovered by a rate adaptation algorithm, in this variant we assume that I-frames are sent at a rate which is one lower (to protect it better from errors), and the P- and B-frames are sent at some higher rate (reflecting their lower priority). A recently proposed work for video streaming [29] uses a similar MAC-layer rate adaptation strategy, albeit in context of multicast communication. To show that the PHY layer implementation of Apex is the most efficient version of this idea, we compared Apex to this variant of MAC layer rate adaptation. In our experiments, we found that Apex outperformed this MAC layer variant by at least 4 dB or more in different experiments, primarily owing to the fact that the former was able to achieve differential data prioritization at a finer granularity.

We next augmented the MAC layer scheme even further by allowing it to send an increasing number of re-transmissions until it was able to match the PSNR of Apex (without any re-transmissions). We found that the number of symbols required by the MAC layer

Pilot symbols	Error in estimating BER		
	in MSBs	in MIDs	in LSBs
1 in 25	0.05 (0.03)	0.05 (0.03)	0.04 (0.02)
1 in 50	0.06 (0.04)	0.05 (0.04)	0.05 (0.03)
1 in 100	0.08 (0.06)	0.07 (0.05)	0.06 (0.05)
1 in 150	0.12 (0.08)	0.11 (0.13)	0.09 (0.07)
1 in 200	0.14 (0.11)	0.16 (0.07)	0.10 (0.08)

Table 5: Error of BER estimation for different bit positions in 64-QAM, Block coding, with varying number of pilot symbols. We use 1 pilot symbol for each 100 symbols in Apex. Experiments used 1500 byte packets. Mean (standard deviation) of prediction error is presented.

scheme is $1.85\times$ that of Apex, to achieve the same PSNR performance (when the uncoded BER was 4.1×10^{-3}). As the channel condition became more error-prone, the gains of Apex increased even further. This result prompted us to study how Apex compares to various re-transmission strategies next.

Apex vs re-transmission strategies: One other way to compute the gains of Apex is to compute how much re-transmission traffic it can save. To characterize this, we let the traditional system to re-transmit each lost symbol multiple times. We bound the re-transmission limit of the traditional system to upto 7 re-transmission attempts (as is common in 802.11 system). In Figure 16, we show the PSNR of the two schemes (Apex with no re-transmissions, and traditional with 1, 2, or 3 re-transmissions). At the limit of 3 re-transmissions, the traditional scheme achieved the same PSNR as Apex. It is more instructive to see the relative bandwidth consumed by Apex with no re-transmissions, when compared to traditional with 3 re-transmission limit for this wireless scenario. Figure 17 shows the increased volume of traffic (including re-transmissions) for traditional communication, compared to the number of transmissions in Apex. Note that both schemes now achieve the same PSNR for the video, but as the channel deteriorates from time to time, the bandwidth requirements of traditional communication sharply increases.

Audio on Apex: We have characterized the performance improvements offered by Apex for audio as well. For this we generated a compressed encoding of an audio clip using WavPack [5]. The WavPack tool generates two output files, one corresponding to the audio frequencies which are important for human perception (the frequencies in the middle of the auditory range) and another with information about frequencies not perceived properly (frequencies at the fringe of auditory range) by human ears. We map the bits corresponding to the first set on more error resilient bit positions while, the bits corresponding to the second set to more error prone bit positions. We find that Apex can improve the Mean Opinion score calculated using Perceptual Evaluation of Speech Quality (PESQ) of the audio received audio by about 1 (on a scale of 0-5) in certain conditions. Research has shown that a MOS reduction of only 0.1 is noticeable by the human ear [26]. We omit details pertaining to rest of the experiments for the sake of brevity.

Parameter selection — number of pilot symbols needed: One of the parameters for the Apex system is the number of pilot symbols used per frame (1 in 100). We made this choice through a sequence of experiments, where we varied the number of pilot symbols in a 1500 byte packets and observed the accuracy of predicting the BERs in different bit positions. Greater the number of pilot symbols, greater is the overhead of the scheme, but higher is the likelihood of accurate estimation. Table 5 shows this variation, where error is defined as $(Actual\ BER - Estimated\ BER)/Actual$

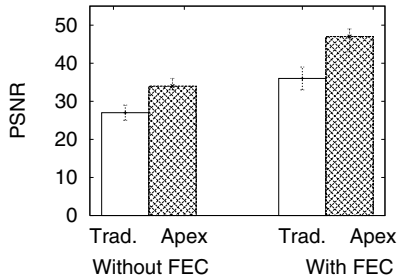


Figure 15: PSNR of video with and without application-level FEC (Claire video).




Figure 16: Variation in PSNR for traditional communication (with different number of retransmissions) and Apex (without retransmissions).




Figure 17: Saved retransmissions in Apex. Y-axis shows the normalized number of bytes transmitted by traditional communication (3 retransmissions) with respect to Apex.

BER. To balance the trade off between overheads and accuracy, we decided to use 1 pilot symbol for every 100 symbols of data, i.e., an overhead of 1% due to Apex BER estimation process.

5. RELATED WORK

We partition prior related efforts into the following categories. **PHY-MAC approaches:** In recent years, a growing number of efforts have designed and studied wireless communication techniques that utilize mechanisms at the PHY and MAC layers to recover more bits out of erroneous packets, e.g., PPR [18], ZipTx [21], SOFT [36], and Maranello [13]. While these schemes apply generally to all data, the transmitters in such schemes do not exploit the unique approximation properties of wireless errors. In contrast, transmitters in Apex learn about and utilize the structure in wireless symbol errors to improve the performance of media delivery applications.

SoftCast [17] is a recently proposed, related scheme that improves wireless media delivery performance in multicast settings by using properties of wireless errors in a manner similar to Apex. SoftCast represents data values in a “raw” numerical (analog) format and maps them directly to wireless symbols with specific transmit power levels. However, there are three main differences between the two schemes: (i) SoftCast requires that media content be represented and transmitted in the raw numerical format, different from what is used in popular standards, e.g., H.264. Media represented in these popular formats have to be converted into the SoftCast format (through a computationally-intensive process) to achieve its performance gains. In contrast, Apex does not attempt to define any new media representation format and can utilize any media format, including the one suggested by SoftCast. (ii) SoftCast uses a static data representation format that is agnostic of wireless channel conditions. In Apex, our mappings between bit sequences to symbols, and between data bits to bit positions are dynamically altered based on channel conditions. (iii) SoftCast is designed specifically for multicast traffic in which no receiver feedback is assumed. In contrast, Apex focuses on unicast traffic, where immediate MAC layer acknowledgments are used to dynamically fine-tune the mappings of bit-sequences to symbols, and data bits to bit positions for improved performance. Such dynamic adjustments are precluded in SoftCast. Hence, SoftCast and Apex are applicable to video delivery in vastly different scenarios.

Network-layer approaches: Some recent research efforts, e.g., MIXIT, have shown knowledge of correct or incorrect symbol de-

coding from the PHY layer can be leveraged in multi-hop wireless settings, to better recover data on opportunistic end-to-end paths [19]. Approaches such as MIXIT operate on a network-wide scale by leveraging diversity across multiple nodes. However, such a scheme also discards symbols received in error, unlike our approach. Further, our scheme is designed for a single wireless hop.

Transport-layer approaches: Transport layer protocols such as UDPlite [20], have been designed to allow applications to accept corrupted data values to recover some content from them. Several notable efforts [31] have shown that under certain network conditions (e.g., high latency), this leads to a better media streaming performance. Again, such schemes complement our approach.

Application-layer approaches: There have been several research efforts to improve the quality of streaming content, mainly at the application layer through effective data prioritization. The primary technique in this regard is addition of FEC to different data components to make them more resilient than others. A good survey of such schemes is present in [35]. Other researchers have also studied hybrid of ARQ-based and FEC-based mechanisms for loss recovery [11, 28]. Authors in [14] and [37] propose specific architectures for such cross layered systems. These techniques do not take advantage of the natural properties of the wireless channel, which we effectively exploit in Apex.

Prior theoretical studies [23, 16, 10, 33, 2, 27] have explored the existence of differential error properties at bit positions for individual constellation maps. Our work augments such theoretical findings with design of constellation map switching, data bit placement strategies, and radio RF pipeline modifications, and evaluates various nuances through practical implementation on real wireless hardware and experimentation, and for real media applications.

Finally, other recent work [8, 38] highlight benefits of exploiting differential error resilience of bit positions for video delivery in unicast and broadcast settings based on numerical simulations. In particular, authors in [38] propose some modifications to DVB-H [1] standard to leverage differential error resilience for terrestrial broadcast of video. In contrast, Apex explores the notion of approximate communication through a full system implementation and extensive experimentation.

6. CONCLUSIONS AND FUTURE WORK

Our efforts in this paper has focused on demonstrating that approximate communication is, indeed, a promising idea. Our results on the WARP SDR platform shows that our approximate commu-

nication system, provides significant improvement in video quality (ranging from a 5 dB to 20 dB in different scenarios). However, this is a small first step in realizing the full capability of this system.

We believe this work leaves open a few optimization problems that should lead to further performance gains.

For example, a joint construction of data modulation schemes and constellation map selection might lead to further performance gains. In particular, since our goal is unequal protection of bit positions, modulation schemes in which constellation points are unequally spaced, might be useful to explore.

Similarly, the effect of the buffering induced by Apex, on the performance needs to be better understood.

A FPGA based implementation of Apex scheme would be necessary for satisfactorily addressing the above. We are currently building such a system, using which, in future we will explore these and multiple other issues.

7. ACKNOWLEDGMENTS

We would like to thank Sateesh Addepalli for providing us with various comments and opinions that helped us improve the quality of the paper submitted to this conference. We are also grateful to our shepherd Srikanth Krishnamurthy whose feedback helped bring the paper to its final form. Sayandeep Sen and Suman Banerjee were supported in part by a Cisco University Research Program (URP) award and through multiple awards from the US National Science Foundation including CNS-0916955, CNS-0855201, CNS-0751127, CNS-0627589, CNS-0627102, and CNS-0747177.

8. REFERENCES

- [1] Hierarchical modulation in DVB-T. http://www.dvb.org/technology/fact_sheets/WP01_DVB-T%20Hierarchical.pdf.
- [2] Hierarchical QAM modulation. http://en.wikipedia.org/wiki/Hierarchical_modulation.
- [3] MPEG-4, Moving Picture Experts Group-4 Standard. <http://www.chiariglione.org/mpeg/standards/mpeg-4/mpeg-4.htm>.
- [4] Video Quality Evaluation Tool-set. <http://www.tkn.tu-berlin.de/research/evalvid/>.
- [5] WavPack Audio Compression. <http://www.wavpack.com/>.
- [6] Wireless Open-Access Research Platform (WARP). <http://warp.rice.edu/>.
- [7] Wireless Open-Access Research Platform (WARP) Public Forum. <http://warp.rice.edu/forums>.
- [8] A. Abdurrahman and M. Woodward. Unequal Error protection of H.264/AVC video Using Hierarchical QAM. In *Telekomunikacioni forum TELFOR*, 2008.
- [9] J. Byers, M. Luby, M. Mitzenmacher, and A. Rege. A digital fountain approach to reliable distribution of bulk data. In *SIGCOMM*. ACM, 1998.
- [10] A. R. Calderbank and N. Seshadri. Multilevel codes for unequal error protection. In *IEEE Transactions on Information Theory*, volume 39, pages 1234–1248, 1993.
- [11] N. Feamster and H. Balakrishnan. Packet Loss Recovery for Streaming Video. In *Packet Video Workshop*, 2002.
- [12] F. Gray. Pulse code communication. U.S. Patent 2,632,058, 1947.
- [13] B. Han, A. Schulman, F. Gringoli, N. Spring, B. Bhattacharjee, L. Nava, L. Ji, S. Lee, and R. Miller. Maranello: Practical Partial Packet Recovery for 802.11. In *NSDI*, 2010.
- [14] W. Heinzelman. Application Specific Protocol Architectures for Wireless Networks. Phd Thesis, 2000.
- [15] U. Horn, K. Stuhmuller, M. Link, and B. Girod. Robust Internet video transmission based on scalable coding and unequal error protection. In *Signal Processing: Image Communication*, 1999.
- [16] M. Isaka, M. P. C. Fossorier, R. H. Morelos-Zaragoza, S. Lin, and H. Imai. Multilevel Coded Modulation for Unequal Error Protection and Multistage Decoding - Part II: Asymmetric Constellations. *IEEE Trans. Commun*, 48:774–786, 1999.
- [17] S. Jakubczak, H. Rabul, and D. Katabi. SoftCast: One Video to Serve All Wireless Receivers. MIT-CSAIL-TR-2009-005.
- [18] K. Jamieson and H. Balakrishnan. PPR: partial packet recovery for wireless networks. In *SIGCOMM*, 2007.
- [19] S. Katti, D. Katabi, H. Balakrishnan, and M. Medard. Symbol-level Network Coding for Wireless Mesh Networks. In *ACM Sigcomm*, 2008.
- [20] L. Larzon and et. al. The udp-lite protocol. <http://www.ietf.org/rfc/rfc3828.txt>.
- [21] K. C.-J. Lin, N. Kushman, and D. Katabi. ZipTx: exploiting the gap between bit errors and packet loss. In *ACM Mobicom*, pages 351–362, New York, NY, USA, 2008. ACM.
- [22] A. Majumdar, D. Sachs, I. Kozintsev, K. Ramchandran, and M. Yeung. Multicast and unicast real-time video streaming over wireless LANs. *IEEE Transactions on Circuits and Systems for Video Technology*, 2002.
- [23] R. H. Morelos-Zaragoza, M. P. C. Fossorier, S. Lin, and H. Imai. Multilevel Coded Modulation for Unequal Error Protection and Multistage Decoding-Part I: Symmetric Constellations. In *IEEE Transactions on Communications*, 1999.
- [24] T. Nguyen and A. Zakhor. Distributed Video Streaming with Forward Error Correction. In *Proc. Packet Video Workshop*, 2002.
- [25] Z. Orlov. Network-driven adaptive video streaming in wireless environments. *PIMRC*, 2008.
- [26] A. Rix, J. Beerends, M. Hollier, and A. Hekstra. Perceptual Evaluation of Speech Quality (PESQ)-A New Method for Speech Quality Assessment of Telephone Networks and Codecs. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2001.
- [27] M. Sajadieh, F. Kschischang, and A. Leon-Garcia. Modulation-assisted unequal error protection over the fading channel. In *IEEE Transactions on Vehicular Technology*, 1998.
- [28] H. Seferoglu, Y. Altunbasak, O. Gurbuz, and O. Ercetin. Rate-Distortion Optimized Joint ARQ-FEC Scheme for Real-Time Wireless Multimedia. In *ICC*, 2005.
- [29] S. Sen, N. K. Madabushi, and S. Banerjee. Scalable WiFi Media Delivery through Adaptive Broadcasts. In *NSDI*, 2010.
- [30] S. Sen, S. Schmitt, M. Donahue, and S. Banerjee. Exploiting "Approximate Communication" for Mobile Media Applications. In *HotMobile*, 2009.
- [31] A. Singh, A. Konrad, and A. D. Joseph. Performance evaluation of UDPlite for cellular video. In *NOSSDAV*. ACM, 2001.
- [32] D. Tian, X. Li, G. Al-regib, Y. Altunbasak, and J. R. Jackson. Optimal packet scheduling for wireless video streaming with error-prone feedback. In *IEEE WCNC*, 2004.
- [33] G. Ungerboeck. Channel Coding with Multilevel/Phase Signals. In *IEEE Trans. on Information Theory*, 2006.
- [34] M. Vutukuru, H. Balakrishnan, and K. Jamieson. Cross-Layer Wireless Bit Rate Adaptation. In *ACM SIGCOMM*, Barcelona, Spain, August 2009.
- [35] Y. Wang and Q.-F. Zhu. Error control and concealment for video communication: a review. *Proceedings of the IEEE*, 86(5):974–997, May 1998.
- [36] G. R. Woo, P. Kheradpour, D. Shen, and D. Katabi. Beyond the bits: cooperative packet recovery using physical layer information. In *MOBICOM*, 2007.
- [37] S. Worrall, S. Fabri, A. Sadka, and A. Kondo. Prioritisation of Data Partitioned MPEG-4 Video over Mobile Networks. In *European Transactions on Telecommunications*, 2001.
- [38] Z. Wu, J. Boyce, and A. Stein. An Unequal Error Protection Framework for DVB-H and Its Application to Video Streaming. In *IEEE Globecom*, 2008.