Supporting Continuous Mobility through Multi-rate Wireless Packetization

(Extended Abstract)

Arunesh Mishra, Shravan Rayanchu, Dheeraj Agrawal, Suman Banerjee*
University of Wisconsin, Madison, WI 53706, USA
Email: {arunesh, shravan, dheeraj, suman}@cs.wisc.edu

ABSTRACT

Continuous mobility scenarios are those in which applications continue to use the radio interface while on the move. With the emergence of Voice-over-WiFi phones, WiFi-enabled music players, and many other such gadgets, continuous mobility is becoming a prevalent mode of operation for WiFi standards. We contend that the existing packetization structures employed in WiFi devices, is not the most suitable for these emerging class of continuous mobility applications. Therefore, in this paper, we suggest a new software-level, standards-compliant extension to the WiFi packetization techniques that provides greater agility and improved performance. In particular, we propose the notion of a multi-rate wireless packet, in which different segments of the same Protocol Data Unit (PDU) are modulated at different physical transmission rates. This is a departure from conventional modulation mechanisms in which the entire PDU is modulated using a single rate. In this paper, we (i) discuss some uses of such a packetization structure for continuous mobility applications, (ii) describe a practical approach to implementing multi-rate wireless packetization in the 802.11 context as a software-only modification that directly leverages current PHY and MAC layer implementations, and (iii) demonstrate the benefits of such an approach with some simple evaluation. We conclude by discussing some of the next steps needed to realize the full potential of this notion.

Categories and Subject Descriptors

C.2.m [**Computer-Communication Networks**]: Miscellaneous

General Terms

Algorithms, Design, Experimentation

Keywords

Multi-rate packetization, Wireless networks, 802.11

*All authors were supported in part by NSF awards CNS-0639434, CNS-0627589, CNS-0627102, CNS-0520152, and CNS-0747177.

1. INTRODUCTION

Over the last decade, usage patterns of WiFi devices have experienced some interesting transformations. The early widespread adoption of the WiFi standards was as a laptop radio interface. While this wireless standard allowed laptop users untethered access, the dominant usage pattern is, what we call, *discrete mobility*. A typical user will sit at a particular location, use his or her laptop for a period of time, suspend its usage while physically moving, and then resume usage in a new location. With respect to the 802.11-based WiFi protocols in use, the radio interface is pre-dominantly used while the user is static with respect to an Access Point (AP).

However, over the last few years, we have seen a great surge in small form factor, handheld WiFi devices. This includes Voice-over-WiFi phones (iPhone), WiFi-enabled music players (iTouch, Zune), and other WiFi-capable gadgets. There are two unique characteristics of this trend. First, it has facilitated the emergence of *continuous mobility* applications, where a particular user continues to use the WiFi radio interface while on the move. Second, the pre-dominant applications running on these devices are real-time and latency-sensitive in nature. Both these characteristics are placing new expectations of the existing adaptation mechanisms of the WiFi radio link. In particular, continuous mobility implies a faster changing wireless channel and the dominance of the new latency-sensitive applications imply the need for more aggressive adaptation of the channel. Our current experience reveals that the various adaptation mechanisms commonly used in the 802.11-based WiFi links are relatively slow for fast changing channel conditions. For example, our recent measurement efforts on a mobile user (at walking speeds) with a Netgear SPH101 VoWiFi phone running Skype using standard transmission rate adaptation and re-transmission mechanisms led to re-transmissions for about 80% of the voice traffic. This translates to wasted bandwidth, energy, and overall performance.

We believe that the root cause of this performance problem under higher mobility scenarios is the *reactive* nature of wireless 802.11 adaptation mechanisms. For example, in typical rate adaptation mechanisms, the transmission rate of wireless frames are reduced only after frames are lost. In contrast, a more *proactive* approach for link adaptation is likely to significantly benefit these application scenarios. For example, if one can continuously and quickly infer the performance of all possible transmission rates, then the rate adaptation mechanism can immediately select a better operating rate prior to packet losses occurring. In this paper, we focus on one specific building block — Multi-rate packetization — that can lead to such fast link adaptations for the emerging continuous mobility applications in the WiFi domain.

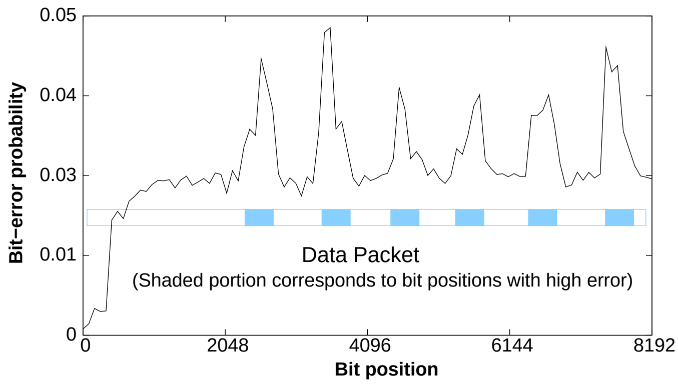

Figure 1: Distribution of bit-error probability within a packet computed over the set of packets received in error during the experiment. The error packets were modulated at 54 Mbps with a size of 1024 bytes.

Multi-rate packetization and uses: Traditionally, the entire Protocol Data Unit (PDU) carried in a single wireless frame has always been modulated at a single physical transmission rate. While such a choice is natural, we believe that such a packetization technique can limit the speed at which WiFi adaptations can be performed. Instead, in this paper, we propose a Multi-rate packet structure, where different segments of the same wireless packet is modulated at different transmission rates. While it might appear that Multi-rate packetization would require a fundamental change to the underlying MAC or PHY layers of the 802.11 system, in this paper we also discuss an implementation path of this mechanism, entirely in software, that can simply leverage the existing 802.11 PHY and MAC layer implementations. We now discuss two possible uses of Multi-rate packetization in fast link adaptation under continuous mobility scenarios, and why this structure can be an important building block for such emerging applications.

- *Unequal error protection within a packet:* Some bits within a single wireless packet are more vulnerable to errors than others. In Figure 1, we illustrate this by plotting the bit error probability of different bit positions, aggregated over a number 1 KB packets transmitted at 54 Mbps between a static AP and a mobile user. We can see that the first few bits of a packet are least likely to be in error. In addition, there is a periodic trend of high error probability approximately 1000 bit positions apart, starting around bit position 2200. Such a trend was observed across various different data rates and across many different experiments, and occurred due to the way these bits were mapped to the different OFDM sub-carriers. In particular, the sub-carriers at the outer edges of a particular channel experienced some interference from neighboring channels, leading to high bit error rates. Such phenomenon naturally suggests the following multi-rate transmission rate for the packet for maximum transmission efficiency. For bit positions with low error rates, we would like to use the highest possible transmission rate. For bit positions with higher error rates, we would like to reduce the physical transmission rate, so that the receiver has a greater chance of correctly decoding these bits.

- *Measuring performance of different transmission rates in a single packet:* Transmission rate adaptation is one of the important link adaptation mechanisms of the 802.11 WiFi links. Given that different 802.11 standards have a number of different transmission rates (e.g., 802.11g has 8 different rates ranging from 6 Mbps to 54 Mbps), as the channel quality degrades, a transmitter reverts to use of lower transmission rates. All rate adaptation algorithms, such as ARF [1], AARF [2], and SampleRate [3] use techniques to determine the feasibility of a certain transmission rate for a given channel. However, these techniques involve "experiments" in which different data packets are sent at different rates and loss patterns on these rates inform feasibility of these rates. However, given the large number of rates available, just understanding the performance of the 8 different rates would require 8 separate packets. In contrast, if multi-rate packets were available, then a transmitter can experiment with all the possible rates *within a single packet* leading to faster adaptation. In this case, different segments of a single packet can be encoded at the 8 different rates, and transmitted. By examining bit error patterns over these different segments, the rate adaptation algorithm can quickly estimate the rough performance at all of the different rates.

Main contributions: This work introduces the notion of Multi-rate packetization and suggests this notion to be a building block for emerging continuous mobility applications. The concept itself is fairly general in nature and can be implemented for any communication standard. In this paper, we focus on its realization for the IEEE 802.11 standard. A natural way to implement Multi-rate packets is to design new 802.11 hardware or to use existing Software-Defined Radio (SDR) platforms. However, to facilitate use of Multi-rate in existing off-the-shelf 802.11 hardware, we also present a *software-only design* of this technique that *can leverage 802.11 PHY and MAC layer implementations.*

2. MECHANICS OF Multi-rate

The concept of encoding a single packet at multiple rates offers new flexibilities and thus creates interesting applications, some of which we discussed above. One of our goals in designing Multi-rate mechanisms is to ensure that it can operate without any change to the PHY and MAC layers of an existing standard, such as 802.11. Therefore, in this section, we present a software-only implementation of Multi-rate that leverages currently implemented PHY layer mechanisms of the 802.11 standards. We achieve this goal by designing software-based bit spreading and de-spreading algorithms that allow us to implement different modulation rates within the same packet. These algorithms can be implemented at any point in the protocol stack prior to handing over the packet to the MAC layer. This software-only solution allows Multi-rate to be readily implementable over existing hardware. In addition, such a mechanism will also be compatible to almost any other wireless standards, including 802.11n and 802.16e.

Illustration through an example: We provide an overview of Multi-rate by illustrating its operation through an example of a phase-based modulation method, phase shift keying (PSK). PSK encodes data bits by changing the phase of the carrier signal. Figure 2 represents the outgoing signal in polar coordinates, with θ as the phase. The figure shows 8 different phase values $X_0 \ldots X_7$, in multiples of 45 degrees. These eight values can thus represent 3 bits of information. The amount of data that a modulation method can represent in a single unit of time is commonly referred to as a 'symbol.' A diagram such as Figure 2 which represents all the symbol values is also called constellation diagram and the values are themselves called constellation points. Each symbol in a constellation with n distinct values can, thus, represent $log(n)$ bits. Quadrature PSK (QPSK) and Binary PSK (BPSK) schemes are two other phase-based modulation schemes with lower modulation rates than

8-PSK that use four and two distinct constellation points respectively. QPSK is implemented by a constellation with phase-value points 0, 90, 180 and 270 degrees, while BPSK is implemented by a constellation with phase-value points 0 and 180 degrees.

While modulation schemes with larger constellation sizes can encode more bits per symbol, such modulation schemes are more prone to decoding errors. This is because the minimum 'distance' between constellation points of such a scheme is small and a receiver is more likely to confuse one symbol with another nearby symbol. Between 8-PSK, QPSK, and BPSK, we have three different modulation methods based on phase-modulation, which provide a tradeoff between rates and decoding accuracy.

Let us consider how we would implement Multi-rate for a wireless system using these three phase-based modulation schemes. In particular, we want to emulate the behavior of all of the three different modulation schemes within a single data packet, without actually changing the MAC and PHY implementation. Therefore, we keep the actual modulation scheme fixed to the one which offers the highest rate, i.e., 8-PSK. We consider a six bit data fragment within the packet, say '101110', and show how it may be encoded by each of the three modulation schemes. If these six bits are to be modulated using 8-PSK, the bit will be encoded as the two symbol sequence $\{X_5, X_6\}$ (Figure 2). We can emulate different lower modulation rates, in software, through a bit spreading operation, that is applied prior to the actual modulation of the data bits, as follows.

To encode these data bits using QPSK while keeping the hardware modulation method fixed at 8-PSK, our bit-spreading algorithm encodes this data by expanding it in such a manner that only those points from the constellation diagram of 8-PSK are used which correspond to the ones for QPSK. This is depicted in Figure 2. Our example bit sequence '101110' is, thus, transformed to '110 111 110' as shown. This transformation uses a mapping (shown in Figure 2) which converts every two bits into three, ensures that each three bit sequence corresponds to the four symbols that are exactly the ones used by QPSK, i.e., $\{X_0, X_2, X_4, X_6\}$. The transformed bit sequence '110 111 110' is now modulated at 8-PSK which results in the following sequence of symbols $\{X_6, X_4, X_6\}$. As shown in the figure, this is exact set of symbols which would have resulted from a hardware modulation of the original bit sequence '101110' using QPSK. We use the term Multi-rate based QPSK to denote the Multi-rate implementation of QPSK.

Similarly, it is also possible to realize BPSK modulation within the same packet (that is, Multi-rate based BPSK) by executing another bit-spreading operation. BPSK uses just two symbols which correspond to 0 and 180 degrees of phase. By spreading the data-bits to use only symbols X_0 and X_4, Multi-rate based BPSK can be implemented over 8-PSK. The spreading operation performs the following transformation : bit 0 is replaced by the sequence 000, and bit 1 by 111. As an example, data bits of '101' would be transformed to '111 000 111' resulting in the symbol sequence $\{X_4, X_0, X_4\}$ which would be exact set of phase-value points had the original data been modulated using BPSK in hardware. We can view the bit-spreading process as reducing the entropy of the symbols used.

A bit de-spreading operation using the same mapping can thus extract the encoded data bits at the receiver. Due to decoding error, it is possible that a receiver erroneously decodes a transmitted symbol to a different one that is not part of the lower rate constellation. For example, let us assume we are using Multi-rate based BPSK for the data bits '101'. The sender transmits the symbols $\{X_4, X_0, X_4\}$ encoded at 8-PSK, while the radio at the receiver decodes them as $\{X_5, X_7, X_4\}$. As discussed earlier, this can happen since the distance between the constellation points X_0 and X_7 might not be sufficient for the given SNR of the transmission thus causing an error at the receiver. We can recover from such errors, as described next.

Error resolution: Our Multi-rate de-spreading algorithm can recover from errors by performing what we call error resolution in software. The obtained constellation point can be resolved to the closest point which is a member of the set of constellation points representing the lower rate. For our example of transmitting the bit sequence '101' at Multi-rate based BPSK, the receiver would be able to map X_5 to X_4 since X_4 is the closest point which is a member of the set of constellation points representing BPSK ($\{X_0, X_4\}$). Similarly, the receiver resolves X_7 to X_0 and thus recovers the original data bits. It is possible that the error is significant enough to cause an incorrect resolution in software. Such errors can also occur when implementing the data rate in hardware. Excessive errors would trigger the rate adaptation algorithm resulting in the actual selection of the next lower rate.

Applying to modulation methods in 802.11: At a basic level, the mechanisms in Multi-rate are very simple and fully realizable in software. They map the given data through a bit-spreading operation such that only a subset of the available constellation points are utilized thus creating the exact effect of a lower rate modulation scheme. The simplicity of this technique makes it a powerful tool in that now different portions of a packet can be modulated at different rates. Also this whole mapping can be done dynamically – the *rate-map* of a packet, that is, which portions of a packet are modulated at which data rate, can now be dynamically programmed in software allowing for much greater flexibility at the MAC layer. We have successfully applied our software-only construction of Multi-rate packets to other 802.11 modulation schemes, such as different QAMs, but we omit the details of such a construction due to space constraints.

Ability to create new transmission rates: A unique and interesting application of Multi-rate lies in the possibility of designing new transmission rates which are not directly supported in hardware. For example, 802.11 standards only support 16-QAM and 64-QAM modulation schemes, while using our software-based implementation of Multi-rate packets, we can actually create a 5-QAM or a 6-QAM, as desired. Such flexibility allows a continuous mobility application to create and use arbitrary transmission rates for its communication, i.e., it is no longer restricted to the 8 discrete rates between 6 and 64 Mbps in the 802.11g standard, but can create almost any new transmission rate between this lower and upper bound. We illustrate this further in the next section. This ability opens up a new dimension for algorithmic work in rate adaptation.

2.1 Demonstrating Multi-rate

We demonstrate the usefulness of Multi-rate packets by measurements and trace-driven analysis that emulates our proposed Multi-rate system. The experiments were conducted using Linux laptops with Atheros based 802.11 b/g wireless NICs running the MadWiFi driver. We study Multi-rate packets relative to two metrics: (i) Transmission efficiency – defined as the channel time consumed by a transmission relative to 54 Mbps (which is maximum rate supported by the system), and (ii) Packet delivery ratio – defined as

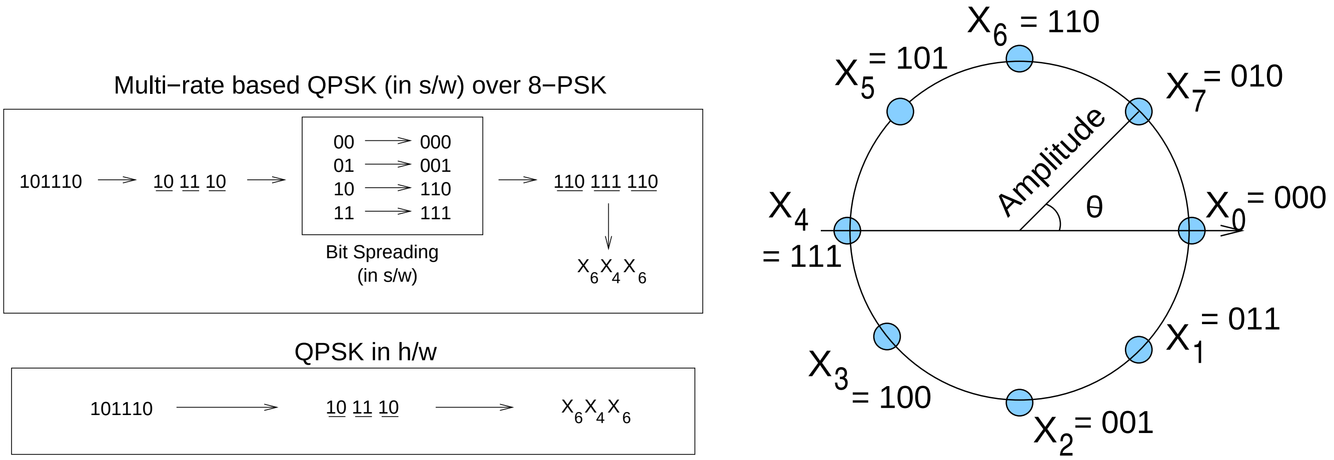

(a) Multi-QPSK over 8-PSK.

(b) Multi-rate over 8-PSK using X_0, X_2, X_4, X_6.

Figure 2: Multi-rate over 8-PSK.

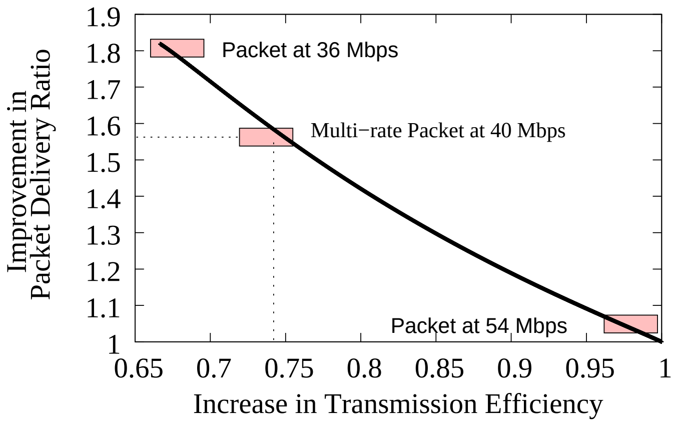

Figure 3: Tradeoff between transmission efficiency and packet delivery ratio for an 802.11 wireless channel. Multi-rate provides a continuum of data rates with different transmission efficiencies and packet delivery ratios to choose from.

the ratio of the number of packets successfully received to the total number of packets transmitted. By selecting lower rates, Multi-rate packets offer higher reliability in packet delivery at the cost of increased channel time (lower transmission efficiency). Multi-rate packets offer the flexibility of modulating variable amounts of data at lower rates and thus, offer a *fine-grained* control over the tradeoff between efficiency and reliability.

We analyze the performance of Multi-rate packets of size $m(= 8192)$ bits over two rates, 36 and 54 Mbps, in the following manner: Consider a parameter k which is the number of bits modulated at 36 Mbps; the remaining $m - k$ bits are modulated at 54 Mbps. We select the first k bits in decreasing order of their bit-error probabilities. These are modulated at the lower rate of 36 Mbps. The packet delivery probability for the newly constructed Multi-rate packet is then calculated based on the error distributions at 36 and 54 Mbps (as a function of the specific bit location). Figure 3 characterizes the specific Multi-rate packet with respect to transmission efficiency and delivery probability, as k is varied from $0 \ldots m$.

When $k = 0$, the entire packet is modulated at 54 Mbps, resulting in an efficiency of 1 (shown on the left end of Figure 3. Correspondingly, with $k = m$ the entire packet is modulated at 36 Mbps which yields a lower transmission efficiency of 0.66 but provides an 80% improvement in delivery ratio (shown on the right in Figure 3). For intermediate values of k (as marked in the figure), the corresponding Multi-rate packet offers a fine-grained control over the tradeoff between efficiency and delivery ratio. For example, depending on link conditions, the best rate to operate might be at 40 Mbps (as shown). A traditional wireless link, under such circumstances would operate at 36 Mbps. Multi-rate provides the flexibility of operating at the 'best' tradeoff point thus resulting in better utilization and throughput gains for the network as a whole.

3. RELATED WORK

The idea of modulating a single packet at different data rates itself has not been explored for any networking technology till date including the wireless media. Below, we present a summary of some of the prior research that focuses on enhancing the performance of wireless networks through three different mechanisms. These methods and their resulting performance are complementary to our approach.

Data rate adaptation: Wireless rate selection algorithms have been extensively studied in the literature [3, 1, 4, 5]. However, all the proposed rate adaptation mechanisms assume that a single packet is modulated at a particular data rate. Multi-rate allows the previously unexplored possibility of modulating a single packet with multiple data rates, thus creating vast space for design of new rate adaptation mechanisms.

Physical layer aware mechanisms: There is a growing interest in the wireless networking community to integrate hints from the physical layer to improve wireless network performance. Recent approaches such as PPR [6] and SOFT [7] use physical layer information about the received symbols to understand which parts of the packets are received in error. PPR uses a simple ARQ mechanism for selective retransmission of the incorrect portions of the packets, while SOFT uses information from multiple receivers to recover from errors. These approaches require significant changes to the current 802.11 MAC/PHY implementations. Multi-rate, on the other hand relies on reverse engineering the physical layer encoding process and uses software mechanisms which leverage *current* PHY and MAC layer implementations.

4. SUMMARY AND NEXT STEPS

By design, conventional communication systems modulate individual data packets at a single rate. In this paper, we demonstrated that in the context of 802.11 wireless networks, modulating a single packet at multiple different rates can have interesting new applications, especially in continuous mobility scenarios, apart from providing new solutions to existing problems. For example, we showed that communicating using Multi-rate packets can yield significant throughput gains by achieving a reduction in the bit-error probability for a selected set of error-prone bits within a packet. Also, Multi-rate packets can be used as a valuable tool to perform a quick sampling of the sustainability of multiple different rates without the need for explicitly sending packets at each rate. This can act as a valuable aid to rate-adaptation algorithms. Therefore, our next steps in this work are (i) a full implementation of the Multi-rate system including a protocol to manage, construct, and synchronize rate-maps between communicating mobile devices, and (ii) a reference proactive rate adaptation algorithm that leverages existence of Multi-rate packets. Our preliminary results in this paper indicate that such algorithms can lead to significant improvement in wireless packet delivery ratios in continuous mobility scenarios.

Based on the work presented in this paper, we believe that Multi-rate is a simple and powerful mechanism which provides fine-grained control over the transmission properties of bits within a packet by modifying the entropy of the symbol space. We believe that concept could lead to research and design of novel solution techniques and applications in emerging wireless mobility environments.

5. REFERENCES

[1] A. Kamerman and L. Monteban, “Wavelan ii: A high-performance wireless lan for the unlicensed band,” *Bell Labs Technical Journal*, 1997.

[2] Mathieu Lacage, Mohammad Hossein Manshaei, and Thierry Turletti, “Ieee 802.11 rate adaptation: A practical approach,” in *MSWiM*, 2004.

[3] J. Bicket, “Bit-rate selection in wireless networks,” MIT Master’s Thesis, 2005.

[4] S. Wong, S. Lu, H. Yang, and V. Bhargavan, “Robust rate adaptation for 802.11 wireless networks,” in *ACM Mobicom*, 2006.

[5] G. Holland, N. Vaidya, and Bahl. V., “A rate-adaptive mac protocol for multihop wireless networks,” in *ACM Mobicom*, 2001.

[6] K. Jamieson and H. Balakrishnan, “Ppr: Partial packet recovery for wireless networks,” in *ACM SIGCOMM*, 2007.

[7] G. Woo, P. Kheradpour, and D. Katabi, “Beyond the bits: Cooperative packet recovery using physical layer information,” in *ACM Mobicom*, 2007.