

# WhiteRate: A Context-Aware Approach to Wireless Rate Adaptation

Veljko Pejovic and Elizabeth M. Belding

**Abstract**—The increased demand for wireless connectivity emphasizes the necessity of efficient wireless communication as resources such as the available spectrum and energy reserves become limiting factors for network proliferation. Recent advancements in software-defined radio enable high flexibility of the physical layer allowing fine grained transmission adjustments. Although communication efficiency can greatly benefit from physical layer flexibility, modern wireless protocols can neither handle these new opportunities nor allocate resources according to the overlying application needs. In this work we develop WhiteRate, a method for physical layer parameter adaptation that efficiently utilizes available energy and spectrum resources, while maintaining the desired quality of communication. Our solution adjusts the modulation and coding scheme, and channel width to achieve a communication profile that matches application requirements. We implement WhiteRate in GNUradio and evaluate it in both indoor and outdoor environments. We demonstrate improvements on two important fronts: spectrum utilization and energy efficiency. Moreover, we show that by using WhiteRate, both benefits can be achieved simultaneously.

**Index Terms**—Wireless communication, network protocols, context awareness, energy efficiency, software-defined radio.

---

## 1 INTRODUCTION

WIRELESS connectivity is no longer limited to homes and offices of the western world; it is now reaching more than a half of the planet’s population, distributed across very diverse environments. In 2010 over one billion mobile phone handsets were shipped worldwide, 18% more than the previous year, and the predictions are that the number will continue to increase [4]. The rise is reflected in the wireless broadband trend, which currently grows five times faster than fixed connectivity [2]. The wireless revolution is happening worldwide, and the number of Internet users in the developing world has already surpassed the number of users in the developed world. The changes in wireless usage are not only quantitative, but also qualitative, as the demand pushes service provisioning to situations that were unforeseen when today’s wireless systems were designed. Always-online mobile computing, connectivity in remote rural areas and high quality voice and video communication are becoming common users’ expectations.

However, higher proliferation of wireless communication is hampered by the lack of physical resources: wireless spectrum and electrical energy. Spectrum scarcity is increasingly evident in highly populated areas where the sheer density of mobile wireless devices causes deteriorating quality of service. In another example, when brought to rural areas that are often characterized by

erratic power supply, traditional power-hungry WiFi equipment fails to deliver reliable connectivity.

While these problems are usually attributed to low infrastructure scalability, we argue that inefficient resource allocation significantly contributes to poor performance. As wireless connectivity spreads into new environments, we increasingly observe a wide variety of resource availability and user needs. Modern wireless protocols are unable to adapt in this large gamut of operating situations for two reasons: first, they can only adjust a restricted number of communication parameters over a fixed physical layer and, second, they react to changes in the wireless environment without considering the context in which the communication takes place.

To break out of the limitations of the fixed physical layer, we use software defined radio (SDR). We embrace an OFDM-based flexible physical layer, similar to the one presented in [32], which allows tight control over channel width, and modulation and coding scheme. We experimentally evaluate the potential of such a system for energy and spectrum efficient wireless operation. We then develop WhiteRate, a context-aware protocol that harnesses the physical layer flexibility to fulfill users’ communication expectations in the most resource efficient way. WhiteRate modifies channel width and modulation and coding scheme (MCS) according to the packet error rate (PER) threshold imposed by the upper protocol layers. The key observation that allows WhiteRate to operate with very little communication overhead is the fact that PER can be used to guide both energy efficiency and spectrum utilization, as well as to directly answer to the application needs. Thus, WhiteRate remains highly practical and, in terms of implementation complexity, comparable to standard rate adaptation protocols.

We implement WhiteRate in GNUradio, a software-

---

• Veljko Pejovic is affiliated with the School of Computer Science, University of Birmingham, UK. The work presented in this paper was done while he was a PhD candidate at the Computer Science Department, University of California, Santa Barbara. Elizabeth Belding is affiliated with the Computer Science Department, University of California, Santa Barbara.  
E-mail: v.pejovic@cs.bham.ac.uk

defined radio framework, and USRP2 hardware and experimentally evaluate the protocol in both indoor and outdoor environments. Our decision to use two different testbeds is motivated by the foreseen applications of WhiteRate. In one case, we see WhiteRate as a key tool for alleviating network congestion in urban areas where the growing appetite of mobile devices surpasses the infrastructure growth [3]. In our lab testbed, therefore, we investigate the communication quality under WhiteRate as the available spectrum and the number of active clients varies. On the other hand, we envision WhiteRate as an integral part of next-generation rural area networks operating in white spaces. These networks are expected to provide connectivity to remote regions where even basic voice and TV service is unavailable and where energy efficiency is of key importance. Thus, we evaluate WhiteRate’s potential for energy savings in a long distance outdoor wireless testbed in Pretoria, South Africa.

We show that WhiteRate operates along the tradeoff line at which the energy savings are balanced against the given application loss tolerance. We compare WhiteRate with a widely used rate adaptation solution and, where possible, intuitive alternatives. Through the experiments we demonstrate that by using WhiteRate we can support more clients without sacrificing end user application performance perception. We also show that, compared to application-agnostic solutions, WhiteRate saves up to a half of the transmission energy per bit. Finally, WhiteRate allows on-the-fly prioritization of resource utilization efficiency or packet delivery reliability and adaptability to varying spectrum availability.

The contributions of this paper are the following:

- Theoretical and practical evaluation of the flexible physical layer potential for context-aware operation.
- Development and full software defined radio implementation of WhiteRate, a channel width and MCS adaptation protocol that maximizes energy and spectrum efficiency while satisfying the application performance requirements.
- A comprehensive evaluation of WhiteRate’s performance in terms of achieved application quality, energy and spectrum efficiency with static and dynamic packet delivery requirements and spectrum availability, and in both indoor and a long-distance outdoor testbed.

In this paper we build on our previous work [19], however, we revise the original work and extend it on a number of important fronts. Specifically:

- We implement a fully functioning version of WhiteRate that works in real time<sup>1</sup>. On the contrary, in the earlier version we ran the algorithm in a simulator over a wireless link trace.
- In our previous testbed we had a 500m long line-of-sight link. For this work we deployed a 3 km long outdoor non line-of-sight white space link between

SDR nodes. This represents a more realistic rural-area network setup and is, to the best of our knowledge, the first such link reported in the literature.

- We analyze packet delivery performance with different channel widths and MCSs over a frequency selective long distance link in our testbed.
- We conduct a full experimental evaluation of WhiteRate, encompassing all aspects measured in our previous paper, and augment it with examination of dynamic protocol behavior.

## 2 RELATED WORK

Physical parameter adaptation to changing wireless conditions has been extensively researched and a number of bit rate adaptation protocols have been proposed for WiFi networks [5], [7], [31]. In addition, work has been done on adjusting the transmission power to achieve the most energy efficient operation [33], and even adjusting the channel width of the transmission [10]. The main goal of these algorithms is high throughput. In our work we concentrate on multiple, possibly conflicting goals. For example, when energy efficiency is important, we sacrifice some of the delivery reliability, if we are still able to satisfy the application performance requirements.

Since traditional WiFi cards offer little flexibility when it comes to rate adaptation, more flexible software defined radio platforms have been considered in recent rate adaptation work. Thus, AccuRate [26] extracts symbol distortion levels at the receiver to guide modulation adaptation, while SoftRate uses decoder confidence levels for the same purpose [30]. Strider [12] introduces rateless coding, where a linear combination of symbols is transmitted, effectively increasing the distance between constellation symbols until the throughput optimal transmission is achieved. In FARA [22] transmission bitrate is adjusted so that each of the OFDM subcarriers gets assigned a modulation level that corresponds to SNR observed on that particular subcarrier. While these schemes allow a high level of sophistication, they exhibit complexity which prevents easy implementation. SNR based adaptation requires in-situ training and receiver triggered rate selection. A large amount of feedback from the receiver needs significant time and inefficiently utilizes communication bandwidth. Our protocol, WhiteRate, builds upon WiFi-based solutions driven by a binary “packet delivered or not” outcome [7]. Thus, our solution remains highly practical and easy to implement despite the fine grained adaptation it provides.

Wireless networks are deployed in a variety of scenarios: personal and mobile communication, monitoring, disaster relief, to name a few. As a consequence, context, or circumstances that surround the communication, can be considered in order to better match wireless transmission with the end user needs. Context-aware wireless rate adaptation was proposed in CARS [27], a vehicular network rate adaptation protocol that accounts for vehicle speed and distance from its neighbours. ASTRA [15]

1. WhiteRate code is freely available at <https://github.com/vpejovic/whiterate>




Fig. 1. **Spectrum analyzer output for a single link occupying non-adjacent OFDM subcarriers.** In this example nine groups of subcarriers are active, while two are inactive. We activate groups in descending order of the observed channel gain, so that narrower channels always encompass “stronger” groups and perform better than wider channels.

relies on machine learning techniques in order to identify the relationship between the context, which is in ASTRA defined by the channel type, SNR and node velocity, and wireless parameters. Our work takes a different, holistic, view on the context. We embrace energy and spectrum resources as well as the application delivery requirements to guide our rate adaptation.

### 3 FLEXIBLE RADIO COMMUNICATION

Our primary objective is to efficiently use the available resources for wireless communication. A flexible physical layer is the key enabler of efficiency. Through sophisticated adaptation of PHY parameters, such as modulation level or channel width, we can impact energy and spectrum usage, and packet delivery performance.

We embrace orthogonal frequency-division multiplexing (OFDM) as a foundation of the PHY layer in our work. OFDM has emerged as a basis for a number of recent research efforts [22], [32] and also is the technique of choice for many wireless standards such as IEEE 802.11a/g (WiFi), 802.16 (WiMax) and 802.22 (WRAN in white spaces) due to its low sensitivity to harsh channel conditions. An OFDM channel consists of multiple narrow-band subcarriers, a high level of flexibility can be achieved if we manipulate the subcarriers individually. This allows us to explore the solution space defined by modulation and coding schemes, variable channel widths and the available spectrum distribution. Our approach to channel width change is similar to Jello [32] and 802.22 [11]: with OFDM subcarrier (de)activation we change the width and bind the available spectrum into a single OFDM channel, even if the available spectrum resides in non-contiguous bands (Figure 1).<sup>2</sup>

As we do not advocate any specific standard protocol stack in this paper, our work also remains operating-frequency agnostic. Our practical implementation, however, targets white spaces due to the vast number of potential channel widths and central frequencies.

2. We define a “channel” as a band over which a single link can operate. This bandwidth is limited by transmitter’s analog-digital converter speed. Within the channel, OFDM subcarriers can be activated in various, even non-contiguous, formations. Thus, we use “channel width” and “number of active subcarriers” interchangeably.

## 4 PHY PARAMETERS AND SYSTEM PERFORMANCE

In this section we examine the impact of PHY parameters on communication quality and energy and spectrum utilization. The findings are used to guide the design of our rate adaptation protocol, WhiteRate (Section 6).

### 4.1 Communication Performance

Packet delivery depends on both modulation and coding scheme, and the channel width. In an OFDM system we can consider each of the subcarriers individually, and for each of them Shannon’s capacity formula defines the maximum sustainable bitrate:

$$(1) \quad R_i = W \log_2 (1 + SNR_i),$$

where  $W$  represents the bandwidth occupied by a single subcarrier, and  $SNR_i$  represents the signal-to-noise ratio at the  $i^{th}$  subcarrier. This bitrate used in the calculations represents an upper bound. In physical systems, the choice of MCS, which is guided by the desired bit error rate (BER), determines the actual bitrate:

$$(2) \quad R_i = W \log_2 M,$$

where  $M \leq 1 + SNR_i$  has to be a practically feasible modulation level. Naturally, the higher the modulation level, the closer the system is to its full capacity. However, as the number of bits packed in a single transmitted symbol  $M$  increases, a more subtle difference among possible signal values results in easier misinterpretation of the signal. The robustness of MCSs is a well studied topic [21], and rate adaptation protocols usually select the MCS that maximizes throughput, i.e. the amount of information that reaches the receiver without errors. Since modern coding schemes rely on the Viterbi algorithm and soft decoding, in practical implementations all subcarriers use the same MCS to simplify the decoding<sup>3</sup>.

Wider channels always result in a higher capacity; the relationship between channel width and transmission errors, however, is not as clear. In [10] Chandra et al. report lower packet loss over narrower channels. They explain this phenomenon by increased power-per-Hz and better resilience to delay spread of narrow channels. The same reasoning does not hold in our case. As shown in Figure 1, we keep the same power per subcarrier irrespective of the number of active subcarriers and, unlike in [10], we do not change the subcarrier width. In our problem setting, frequency-selective fading, is the most probable cause of channel width-dependent delivery errors. In the case of frequency-selective fading, subcarriers at different frequencies experience mutually different channel gains. If a data packet is sent over an OFDM link where the subcarriers experience different fading, some bits are more likely to be corrupted than others. With subcarrier interleaving, the errors are distributed evenly throughout the OFDM symbol [14].

3. Viterbi decoder occupies a significant number of logical gates on a wireless NIC [14]. The operation is also the most processing intensive part of the demodulation in case of the SDR.

The impact of channel width on the packet error rate depends on the way the width change is envisioned. We discuss two cases: (i) when the subcarriers are activated according to the channel state they observe in the *best-first* fashion, and (ii) when the subcarriers are activated so that they always remain *contiguous*. The first case is optimal in a sense that it does not waste transmission power to overcome the impact of poorly performing subcarriers when better ones are available. However, from a practical point of view *best-first* comes with two major drawbacks. First, non-adjacent subcarriers demand precise narrow filters that require substantial computing power. Second, subcarrier selection based on the channel quality requires sweeping through the whole frequency range so that the well performing frequencies can be isolated. This increases the protocol delay and the communication overhead. Therefore, we devise a practical solution shown in Figure 1: adjacent subcarriers are assembled in a small number of groups and each group is individually (de)activated according to the average channel gain observed within a group. Consequently, channel widths are ordered in our system – a lower width always results in a better delivery rate. Note that due to multipath fading and imperfect filtering the relationship between achieved rate and the number of active subcarriers is not perfectly linear. Thus, in WhiteRate we rely on explicit channel probing in order to identify the optimal channel width.

## 4.2 Energy Efficiency

A good understanding of the energy consumption of wireless a network interface card (NIC) operating over flexible OFDM helps us to identify the space for improvement of the energy efficiency of the whole system.

A bit of information over an OFDM wireless channel with  $k$  active subcarriers is transmitted with energy:

$$(3) \quad E_{Tx} = \frac{\sum_{i=1}^k P_{Tx,i} + P_{TC}}{\sum_{i=1}^k R_i}$$

where  $P_{Tx,i}$  represents the transmission power over the  $i$ -th subcarrier and  $P_{TC}$  the transceiver circuit power of a wireless NIC [16]. The transceiver circuit power is constant irrespective of the transmission parameters and represents the power needed to keep the digital circuits powered whenever the NIC is in the transmission mode. The energy efficiency thus depends on the cumulative transmission power  $\sum P_{Tx,i}$  and transmission bitrate  $\sum R_i$  and their relationship as we change the channel width and/or MCS. Changing modulation levels and coding rate does not result in a noticeable power consumption change in modern WiFi NICs [6], [29], and same is true for USRP2 power consumption (Table 1). It follows that higher MCSs are more energy efficient as they reduce the transmission time without any power consumption penalty. Channel width change through

TABLE 1

**Power consumption breakdown for USRP2.** Fixed factors such as  $P_{TC}$  and  $P_{base}$  dominate the total consumption when the device is in transmission mode. Transmission power grows linearly with channel width, while it remains unaffected if the MCS changes.

Total power $P_{total}(tx\_amplitude = 0)$	12.77 W
Total power when <i>idle</i> , $P_{total}(idle) = P_{base}$	11.76 W
Transceiver circuit power $P_{TC} = P_{total}(0) - P_{total}(idle)$	1.01 W
The above values are independent on the number of active subcarriers, modulation and coding scheme and the USRP interpolation rate.	
<b>Changing MCS, <math>tx\_amplitude = 0.5</math></b>	
$P_{Tx} = P_{total}(max\_width, any\_mcs) - P_{base} - P_{TC}$	0.45 W
<b>Changing width, <math>tx\_amplitude = 0.5</math></b>	
$P_{Tx} = P_{total}(all\_widths, any\_mcs) - P_{base} - P_{TC}$	[0.22 - 0.45W]

subcarrier activation leads to a linear increase in the transmission power  $P_{TX} = \sum_{i=1}^k P_{Tx,i}$  and the transmission rate  $R = \sum_{i=1}^k R_i$ . The total power consumption, however, is a sum of  $P_{TX}$  and a large constant factor  $P_{TC}$  (see Table 1). It is clear that if we want to optimize the energy consumption  $E_{TX}$ , we have to maximize the channel width so that the constant factor in the numerator of equation (3) is diminished.

## 4.3 Spectrum Utilization

Spectrum efficient transmission utilizes the medium as close to its physical capacity as possible. Modulation levels and an adjustable channel width are the primary PHY parameters that impact spectrum utilization.

Modulation constellation that utilizes the I-Q space so that the distance between constellation points is minimized leads to maximal spectrum utilization. Subcarrier activation impact spectrum utilization, as activating fewer subcarriers than what is allowed by the available bandwidth leads to suboptimal transmission times. In addition, a subcarrier activation that leads to high spectrum fragmentation can prevent concurrent transmissions on other links in a network.

In summary, spectrum utilization increases with higher MCSs and higher number of active subcarriers, i.e. wider channels. The approach in our protocol is thus, to push the modulation level to the highest supported value and to occupy as much bandwidth as possible. Since these parameters impact the delivery error, the limits on their values are imposed by the packet error rate (PER) threshold that user can tolerate.

It is important to note that in a multiuser system, spectrum utilization is tightly connected with fairness. Here, a few wide links can starve nodes that do not have any available spectrum to use. In addition, it has been shown that low transmission rates, even at only one of the competing transmissions, significantly lower the throughput of all the links in the interference domain [23]. Although PHY parameters change spectrum utilization in a relatively straightforward way, MAC schemes can result in a non-trivial relationship between the MCS, channel width

and spectrum utilization. In the evaluation section we rely on trivial MAC implementations that allow us to isolate the impact of PHY layer parameters.

## 5 CONTEXT AWARENESS

Wireless systems are deployed for different purposes and even the same system can have varying priorities depending on the operational situation. A rural area wireless network may need to support high quality voice and video communication in order to connect remote regions. However, the same network might favor energy efficiency over the communication quality once grid power is unavailable and the electricity is supplied through limited charge UPSs. We argue that the traditional one-size-fits-all solution for PHY layer parameter adaptation can be counter-productive in certain contexts. We describe three avenues where a holistic approach is greatly beneficial.

### 5.1 Loss Tolerant Applications

Multiple layers of a standard networking stack are geared towards reliable packet delivery. However, many applications do not require lossless communication. For example, voice and video communication are often encoded so that they can tolerate non-zero packet loss. Moreover, voice and video are usually used to transmit data framed in a certain situation (e.g. a human conversation about a known topic). From the endpoint perspective, communication success is not a binary value, but can be described as more or less satisfactory depending on user tolerance. In some cases extra effort put into ensuring flawless packet delivery can even hurt performance. Real-time traffic requires stringent packet delay and jitter guarantees, whereas reliability can introduce additional unpredictable delay. Finally, reliable delivery, either through transmission with more robust MCSs (thus more redundancy and lower bitrate) or through packet retransmissions, requires extra energy and increases channel busy time.

TABLE 2  
**Commonly used voice and video codecs and their packet loss tolerance.** Real-time multimedia often tolerates non-zero packet loss.

Voice					
Codec	G.711 PLC	G.726	G.729a	G.723.1	G.722
Loss Tolerance	10%	5%	2%	1%	5%
Video					
Codec	H.262 (MPEG-2)		H.264 (MPEG-4)		
Loss Tolerance	0.3% - 0.9%		0.2% - 2.2%		

In Table 2 we list commonly used voice and video codecs along with the packet error rates (PERs) below which the quality of the communication is considered acceptable [1], [8], [9]. For successful voice and video communication over a wireless link, we need to ensure

that the packet error rate is below a certain, codec-specific, threshold, which is often significantly above zero. Thus, if relaxing the packet delivery constraints increases protocol efficiency, we can allow PER to be non-zero but below the application packet loss tolerance.

### 5.2 Limited Energy Budget

Mobile communication devices are expected to be compact yet powerful and as a result have very limited battery capacity. Power-hungry high-bandwidth data transfers often leave these devices depleted. Rural area wireless networks deployed in remote areas seldom have access to a reliable power grid and have to be powered by alternative energy sources such as unpredictable wind and solar energy. These two use cases represent wireless systems that are very sensitive to communication protocol energy inefficiency. Fortunately, these two instances can highly benefit from context aware PHY layer adaptation.

In Section 4.2 we analyzed the impact of PHY parameters on energy consumption. The need for energy efficiency depends on the battery charge - a value that is often exposed by a mobile device drivers. In a network of self-powered wireless routers, on the other hand, battery charge is a complex function of local wind speed, solar irradiation and hardware properties. In this case, the information about the energy needs can be inferred through energy flow models [20]. Our protocol designed in Section 6 ensures that the energy efficiency of PHY layer operation corresponds to the energy needs of the system. The actual means of determining those needs, however, is beyond the scope of this work.

### 5.3 Spectrum Demanding Operation

Wireless spectrum is unevenly utilized: urban environments suffer a shortage of available bandwidth, while in rural areas there is an abundance of unutilized channels [13]. Unfortunately, current protocols usually operate on a fixed central frequency and with a fixed bandwidth, which prevents them from using communication resources efficiently, and high spatial and temporal correlation of the bandwidth demand among users often leads to intermittent connectivity with virtually no quality of service guarantee.

Our system adapts by adjusting the channel width and MCS according to spectrum availability. The strategy of adaptation depends on the medium access scheme. For example, with a time ordered scheme (e.g. TDMA) the protocol should allow a high level of multiplexing by utilizing the widest available channel so that the transmission time is minimized and more flows can be packed in time. In a frequency division regime our protocol should enable as many multiple flows to operate concurrently as possible, by increasing the MCS and, thus, enabling a narrower channel width per each flow. Note that the available spectrum is often fragmented; for efficient spectrum utilization our protocol allows for non-adjacent subcarrier activation (Figure 1).

## 6 WHITERATE

Physical layer parameters impact communication performance (section 4.1), energy efficiency (section 4.2) and spectrum utilization (section 4.3). Often the impact is beneficial from one angle and detrimental from another. The usage context defines the importance of each of the above aspects. To provide a practical solution for PHY adaptation under various packet delivery, spectrum utilization and energy efficiency requirements, we design WhiteRate, a holistic approach that adapts the bitrate according to the usage context.

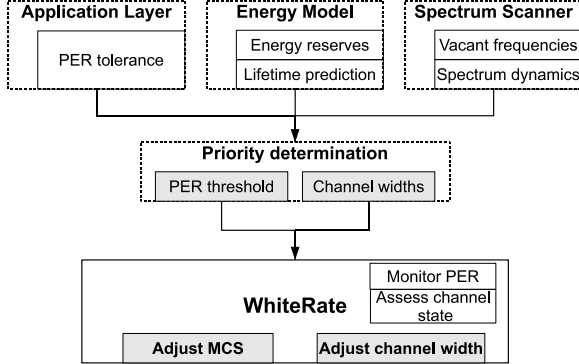


Fig. 2. **Relationship among the context-aware PHY adaptation system components:** The top three modules (application layer, energy model and spectrum scanner) determine the system requirements: application packet delivery, energy efficiency and spectrum utilization. In addition, the spectrum scanner marks the available spectrum. The information is passed to the priority determination module, which analyzes it and sets a single PER threshold and available channel widths that guide WhiteRate. WhiteRate adjusts MCS and channel width, while monitoring the PER, making sure that it stays within the limits given by the priority determination module.

### 6.1 Cross-layer Information Flow

We consider three types of constraints imposed by the context in which a device operates. The first one is the application packet delivery requirement. In our system we integrate “hooks” that allow the overlaying application to set the maximum tolerated PER. As seen in Section 5, even a single type of traffic can have different loss tolerances. The second constraint comes from the limited energy reserves of the system. In the case of mobile devices or self-powered routers, energy efficiency often takes precedence over lossless communication. Again, the “hooks” are provided for the platform to provide its energy reserves information. This vertical flow of information enables loosening of the PHY layer restrictions. Finally, modern SDRs can easily, and with a fine time and frequency granularity, monitor wireless spectrum. We let an independent scanning module deliver the available spectrum information to which WhiteRate abides.

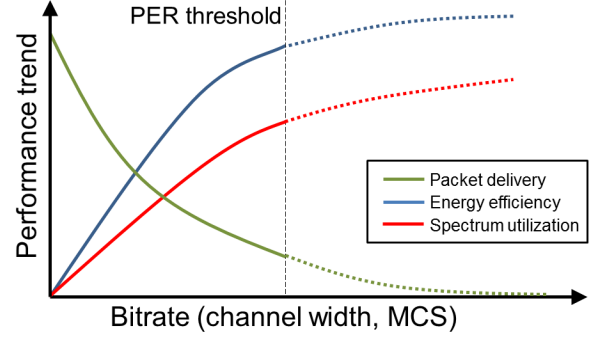


Fig. 3. **Packet delivery reliability, energy and spectrum efficiency trends with changing bitrate.** A restriction on the PER is sufficient to control rate adaptation as operation at the edge of tolerable PER ensures maximum energy and spectrum efficiency. Note: the lines in the graph are for illustrative purpose only, and do not represent actual values.

In Figure 2 we illustrate the relationship between different parts of the system and the information flow. *Application layer*, *energy model* and *spectrum scanner* are modules that rely on techniques from sections 5.1, 5.2, and 5.3, respectively. *Priority determination* enables WhiteRate to answer to complex and perhaps contradicting demands through a practical algorithm.

The trends in energy and spectrum efficiency and packet delivery reliability are illustrated in Figure 3. High bitrate positively impacts both energy savings and spectrum efficiency, yet can have an adverse effect on packet delivery. Therefore, the PER threshold is set by the priority determination module and WhiteRate ensures that the bitrate is adjusted so that the actual error rate is below the threshold. The threshold exact value depends on the context and external properties, such as the system hardware and user preferences. In the rest of the paper we use a stub implementation of the modules in the dotted boxes and concentrate on WhiteRate.

### 6.2 The Algorithm

We design WhiteRate as a packet-level MCS and channel width adaptation solution. WhiteRate keeps track of the PER and ensures that the MCS and channel width are set so that the PER stays within the specified tolerance (Algorithm 1). At the same time, it aggressively tries to find the highest achievable bitrate within this loss tolerance, as in that case the energy consumption is minimized and the spectrum utilization maximized (Algorithm 2). To identify the most efficient bitrate, WhiteRate periodically probes channel width and MCS combinations that have lower theoretical minimum bit transmission time than the current combination’s average bit transmission time. Packet delivery statistics are updated in the process, and if at any point WhiteRate finds a combination that results in higher energy and spectrum savings while keeping the PER below the threshold, it switches to it and continues probing as before.

In the event of a complete loss of wireless link, WhiteRate quickly adjusts to a lower bitrate setting. Once a reliable MCS is pinpointed, WhiteRate continues, with the help of channel gain estimations, to intelligently adjust the channel width as long as unsatisfactory end-to-end performance is observed.

---

**Algorithm 1** WhiteRate
 

---

**Require:**  $MCS = MCS_5, channel\_width = channel\_width_7$

**Ensure:**  $PER < PER_{thold}$

- 1: **if**  $c$  consecutive losses **then**
  - 2:    $MCS = \min(MCS - 1, MCS_1)$
  - 3: **else if**  $l^{th}$  packet **then**
  - 4:   Evaluate  $PER$
  - 5:   **if**  $PER > PER_{thold} - PER_{margin}$  **then**
  - 6:      $bin\_search\_reduce(channel\_width)$
  - 7:   **else if**  $PER < PER_{thold} - PER_{margin}$  **then**
  - 8:      $bin\_search\_expand(channel\_width)$
  - 9:   **end if**
  - 10: **else if**  $p^{th}$  packet **then**
  - 11:   set  $pick\_probe(channel\_width, MCS)$
  - 12: **else**
  - 13:    $pick\_best()$  {Average tx time table lookup.}
  - 14: **end if**
- 

---

**Algorithm 2** pick\_probe
 

---

- 1: Parameter:  $MCS_i$
  - 2: **if**  $\exists MCS_j \mid min\_tx\_time[MCS_j] < avg\_tx\_time[MCS_i]$  **then**
  - 3:    $MCS = MCS_j$
  - 4:   **if**  $PER(MCS_j) > PER_{thold} - PER_{margin}$  **then**
  - 5:      $bin\_search\_reduce(last\ channel\_width\ at\ MCS_j)$
  - 6:   **else if**  $PER(MCS_j) < PER_{thold} - PER_{margin}$  **then**
  - 7:      $bin\_search\_expand(last\ channel\_width\ at\ MCS_j)$
  - 8:   **end if**
  - 9: **end if**
- 

### 6.3 The Algorithm - Details

We give a fictional example of the algorithm flow for a single source-destination pair in Figure 4. The columns represent available MCSs and rows represent channel widths, which can be composed of non-adjacent sub-carriers if needed. Wider channels/higher MCSs are labelled with a higher subscript value (widths from 1 to 7, MCSs from 1 to 5). Cell shades represent average bit-transmission times, where a lighter shade means a lower transmission time. These times are initially unknown, and are computed as the algorithm visits each of the cells, while the minimum theoretically possible transmission times for MCSs (with the widest channel) are known in advance. Depicted is also (with a black line) an initially unknown region within which  $[MCS, channel\ width]$  combinations yield PER acceptable for the application.

**WhiteRate search dimension selection.** In an OFDM system that implements interleaving and operates over limited spectrum, the choice of MCS determines the order of the PER magnitude, while the channel width adjustment allows fine-grained tuning. WhiteRate implements mechanisms that enable fast recognition of the

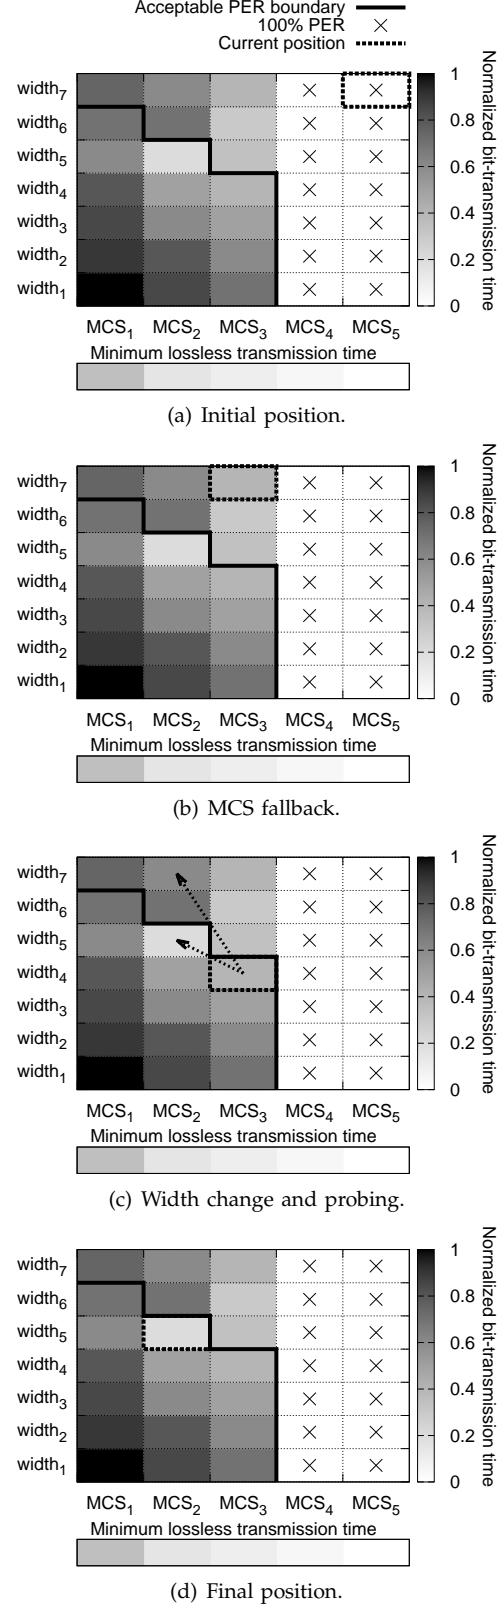


Fig. 4. **WhiteRate protocol example.** Each cell represents one  $[channel\ width, MCS]$  combination. The goal of WhiteRate is to identify the cell that results in minimum average bit-transmission time, while satisfying the PER threshold, i.e. the brightest cell left of the black line.

channel state, thus allowing the algorithm to adjust the appropriate parameter.

In the case of extremely bad link performance, recognized by  $c$  consecutive packet losses, WhiteRate adapts to a more robust MCS. In our example, the algorithm starts transmission over the highest possible MCS and the widest available channel,  $[MCS_5, width_7]$  (Figure 4(a)). The next lower combination ( $[MCS_4, width_7]$ ) might still be lossy enough to result in  $c$  consecutive packets dropped. Thus, the process is repeated until a link is established at  $[MCS_3, width_7]$  (Figure 4(b)).

**Well-performing MCS and channel width combination identification.** When the link is established, the PER is calculated and compared to the application provided threshold. If it is above the threshold the application can tolerate, the channel width has to be changed.

WhiteRate periodically estimates the channel gain by transmitting a known bit sequence. Afterwards, OFDM groups (Figure 1) are ordered based on the descending average gain. Thus, the narrowest channel ( $width_1$ ) in Figure 4 contains only the best performing groups of subcarriers, and we expect the channel conditions to deteriorate as we increase the width. This ordering allows WhiteRate to quickly, through binary search, identify a channel width that supports the PER requirements of the application, and prevent excessive loss. In the example in Figure 4(c) WhiteRate finds a well performing combination at  $[MCS_3, width_4]$ .

**Identifying the optimal solution.** The current  $[MCS, channel\ width]$  pair satisfies the PER requirement but might not be the optimal choice in terms of energy efficiency and spectrum usage. Therefore, the algorithm probes other MCSs, but only those that can potentially lower the transmission time and consequently improve the efficiency. Thus, for each of the MCSs we keep information on the so-far-observed average bit transmission time and the minimum possible bit transmission time.

WhiteRate searches for a better solution than the current one by transmitting every  $p^{th}$  packet on an MCS that has a minimum bit transmission time lower than the currently observed average bit transmission time and that has not experienced  $c$  consecutive losses recently. In Figure 4(c) these probes are depicted with arrows. First, WhiteRate picks a candidate MCS ( $MCS_2$ ) that could, ideally, result in better performance than the current combination ( $[MCS_3, width_4]$ ). If the candidate MCS is probed for the first time, the transmission takes place over the widest channel ( $[MCS_2, width_7]$ ). For the probed candidate  $[MCS, channel\ width]$  combination, WhiteRate calculates the average bit transmission time and the packet error rate. If the new bit transmission time is lower than the current one (corresponding to  $[MCS_3, width_4]$ ), and the PER satisfies the application requirements, WhiteRate switches to the candidate MCS and channel width. Otherwise, the next time the same MCS is selected for probing, WhiteRate readjusts the channel width through binary search, and probes a value that should result in lower PER (in our case

$[MCS_2, width_5]$ ). Once the combination that has the minimum of all average bit transmission times is found, WhiteRate selects it as a new ground point. In our example, as shown in Figure 4(d), WhiteRate settles at  $[MCS_2, width_5]$ , which is the combination that returns the lowest bit transmission time among all combinations that satisfy the application's packet loss requirements (the brightest cell under the PER boundary line).

## 6.4 Discussion

We presented an example WhiteRate run in the case when a link is established and WhiteRate has not yet explored any of the channel width, MCS combinations. Otherwise, in case that the PER is satisfied at the current combination, the adaptation process reverts to combination probing; in case the current PER is above the application threshold, the process reverts to channel width/MCS fallback. Note, however, that WhiteRate can adapt successfully only if the channel conditions are stable for sufficient enough time for the protocol to calculate PER and identify at least one well performing combination.

In the worst case protocol convergence requires MCS identification, which can be done in  $O(M)$ , where  $M$  is the number of MCSs, and channel width adjustment, which can be done in  $O(\log(W))$  where  $W$  is the number of available channel widths as the binary search is employed. The total complexity,  $O(M) + O(\log(W))$  is significantly lower than the complexity of exhaustive search ( $O(MW)$ ). To quickly bootstrap a link to a near optimal MCS, channel width combination a history of rates can be kept.

Compared to protocols that are oblivious to frequency selective fading, WhiteRate introduces very little overhead. Subcarrier quality evaluation is based on the PN sequence that is a standard part of the transmission, as it is needed for OFDM frame acquisition [24]. The receiver piggybacks the subcarrier group ordering to the sender via a packet acknowledgement. For  $N$  subcarrier groups that overhead is only  $N \log_2 N$  bits; for example, if we use seven groups we introduce only a three byte overhead in the acknowledgement.

The WhiteRate algorithm bears some resemblance to SampleRate [7] in its PER guided rate adaptation. Similar to the practical implementation of SampleRate from the MadWiFi driver, WhiteRate maintains a separate algorithm table for different packet sizes, and implicitly takes the relationship between BER and PER into account. Unlike SampleRate, WhiteRate changes both MCS and channel width and explicitly caters to application needs. Additionally, WhiteRate can ensure lossless communication if the PER threshold is set to zero. If the channel width is fixed and application hints not present, WhiteRate essentially behaves as SampleRate.

Finally, it is important to realize that WhiteRate aims for context-aware link adaptation. Network performance in case of multiple interfering links running WhiteRate depends on many other factors, such as the MAC



scheme, and is beyond the scope of this paper. Intuitively, however, WhiteRate’s tendency to narrow the channel in search for reliable transmission leads to lower contention among links.

## 7 GNURADIO IMPLEMENTATION

We implement WhiteRate in GNUradio - an open source software defined radio framework. The key components of our implementation are a MAC layer that takes care of transmission scheduling and delivery tracking, the rate selection algorithm, a flexible OFDM subcarrier activation system that allows on the fly channel width and modulation and coding scheme (MCS) selection, forward error correction coder, and channel gain estimator. We use the publicly available flexible OFDM system from Jello [32] and channel coding implementation from SoftRate [30], and integrate them with the other components we implemented.

Our initial implementation [19], did not allow runtime rate adaptation due to the following problems: a) GNUradio/USRP processing latency prevents the ACK-based stop-and-wait transmission needed to keep track of the packet error rate [17], and b) channel width/MCS adaptation requires that GNUradio flow graph is stopped, reconfigured and reconnected, incurring a non-negligible time delay. In our improved implementation we solve the first problem by using an Ethernet back-channel to transmit packet acknowledgements as well as control data, such as channel width change signals and channel gain updates. This cuts the minimum stop-and-wait time to about a half of its original value<sup>4</sup>. To solve the second problem, we move all GNUradio graph adaptation code from Python to C++ and cut the channel width/MCS adaptation time to nearly 100ms. This value, along with the halved ACK processing time, allows us to implement WhiteRate as a fully operational real-time protocol.

SampleRate is a widely accepted rate adaptation protocol often used as a benchmark for other rate adaptation protocols [5], [30]. We implement SampleRate and in our experiments compare WhiteRate and SampleRate. As SampleRate is width agnostic, we use it with the maximum channel width, unless otherwise stated.

## 8 EXPERIMENT SETUP

WhiteRate is geared towards improving two aspects of wireless communication: energy efficiency and spectrum utilization. We envision and focus on two scenarios where these aspects are crucial for overall network performance. The first scenario is a rural area wireless network connecting distant locations via white space links powered by alternative energy sources. Studies of existing rural area wireless networks show that critical battery resources, and voice and video communication

4. We found out that the actual value highly depends on the distance between the nodes, processing PC configuration and the over-the-air bandwidth. In our experiments we support a 500kHz channel with 200ms inter-packet time.

are the key descriptors of such networks [28]. The second scenario is a wireless network of numerous personal devices operating in the same wireless spectrum. This scenario is common in densely populated urban areas [3]. We build two testbeds so that we can examine WhiteRate’s performance in both scenarios.

Our first testbed is located in suburban Pretoria, South Africa. It consists of a single 3 km outdoor link between two PCs(Linux)/USRP2s with whitespace-enabled WBX cards and directional antennas. On the sender side we use a 1W amplifier, while on the receiver we do not have any amplification. A direct view between the nodes is obstructed by a hill, but transmission over white spaces allows a robust non-line of sight link. With its distance, outdoor surroundings, susceptibility to weather conditions and terrain configuration, this link encompasses all the phenomena we expect in a rural area network. Therefore, we use this testbed for experiments that provide insights about channel behavior. In addition, we use WhiteRate as a key protocol for enabling energy savings in self-powered rural area networks, thus we use this testbed for all experiments related to energy efficiency.

We deploy an indoor testbed of four PCs(Linux)/USRP2 nodes in a building at UC Santa Barbara campus. The nodes are located in the same room, and the distance between any two nodes is approximately three meters. This testbed allows us to examine WhiteRate spectrum saving potential in a spectrum-scarce setting. In addition, in this testbed we can achieve robust links over a much wider range of transmission amplitude values, thus we use it for a more complete analysis of WhiteRate’s correctness in terms of achieved delivery ratio, throughput and bitrate.

Our transmission uses white space frequencies (centered at 694MHz for the Santa Barbara testbed, and at 530MHz for the testbed in Pretoria) and a 512 FFT with 462 possibly active subcarriers, divided in 7 groups of 66 subcarriers. We use 1/2 convolution coding and BPSK, QPSK, 8-QAM, 16-QAM, 64-QAM MCSs. The default channel is 500kHz wide for the outdoor and 300kHz wide for the indoor testbed.

## 9 EXPERIMENTAL EVALUATION

We have two goals in the evaluation. In the first part (9.1) of the evaluation we experimentally justify two important protocol design decisions: prioritizing MCS over channel width for rate adaptation and relying on channel gain stability for width determination. In the second part (9.2), we thoroughly analyze WhiteRate’s ability to efficiently utilize resources while ensuring satisfaction of the application requirements.

Unlike the preliminary evaluation from our earlier work [19], the experiments presented in this paper are conducted with WhiteRate operating in real-time, over a longer NLOS outdoor link or a multi-link indoor setup. Moreover, in this work we evaluate WhiteRate under both static as well as dynamic application requirements.




Fig. 5. **Packet error rate under different channel widths and MCSs.** Narrower channels, composed of better performing subcarriers, deliver more packets. However, the difference between the PER at the widest and the narrowest channel is not as drastic as the difference among the PERs achieved at different MCSs.

### 9.1 WhiteRate design

WhiteRate needs to quickly identify a sustainable bitrate. To this end, our protocol switches to a lower MCS after  $c$  consecutively-dropped packets in a row (Section 6). An alternative approach is to prioritize the channel width and change it once consecutive losses are experienced. If MCS has more impact on packet loss than channel width, our approach leads to fast convergence and is the preferred alternative. The first goal of our experiments is to evaluate this claim.

In the South African testbed we transmit a sequence of 500 packets over each [channel width, MCS] combination and record the PER. We want to examine the change in combination performance while the channel is stable. Therefore we measure PER as we cycle through all the combinations starting a transmission immediately after a previous one ends. Subcarriers are allocated according to the gain averaged over a group of subcarriers, as explained in Section 4.1. In Figure 5 we show the error rate (y-axis) for different channel widths (x-axis) and MCSs (different point types). We observe that the choice of MCS impacts PER more than the choice of the channel width. A single MCS experiences no more than between 0% and 70% loss for all channel widths (the 16-QAM case), while channel widths can experience everything from 0% to 100% loss. Therefore, our decision to use MCS for coarse grain rate adaptation is experimentally validated.

WhiteRate uses channel gain measurements to rank subcarrier groups and activate them in descending order. If the channel gain ordering among groups varies over time, the protocol needs to repeat the measurements periodically. In order to examine the channel gain stability, and its role in width determination, we measure the gain change over time in our outdoor testbed. To estimate the channel gain, we prepend each packet with a known BPSK modulated preamble and compare the received copy with the original one. We then average the gain over all subcarriers in a group.




Fig. 6. **Channel gain observed at each of the seven subcarrier groups over time and with different total channel widths.** Each line represents gain of one group, different colors correspond to different widths. The gain does not change over time and the ordering of groups remains constant.

Figure 6 shows, with a single line, the gain for each of the seven subcarrier groups. We repeat the experiment over three different bandwidths: 500kHz, 2MHz and 5MHz. In Figure 6, lines of a different colour correspond to different bandwidths used in the measurements. Interpacket time is 200ms and a total of 600 packets is sent at each bandwidth. We observe that the channel gain is fairly constant over time. Even more important, the relative ordering of groups remains the same. The testbed with its long distance NLOS link may result in a non-negligible multipath, that would manifest as the frequency selective fading. What we observe though, is that fading varies over subcarrier groups, but remains constant over time. A closer examination reveals that groups closer to the edge of the channel are in a deeper fade. We suspect that rather than multipath, imperfect filtering and amplification cause this frequency selectivity. Irrespective of the cause, the results point out that once channel gain sampling is executed, there is no need for frequent gain re-estimation. In the rest of this section we report on experiments with flows that last fewer than two minutes, thus we perform gain measurements only once at the beginning of the flow.

### 9.2 WhiteRate performance

Because WhiteRate is designed to be a rate adaptation protocol for diverse environments, we evaluate it over a range of channel conditions: indoor, outdoor, long distance and short distance links. We quantify the energy and spectrum savings it delivers and examine its adaptability to dynamic system requirements.

#### *Packet loss, Throughput and Bitrate*

To understand how rate adaptation copes with changing link quality, we vary the transmission amplitude and measure the packet loss, throughput and bitrate. The results are shown in Figure 7. Each data point represents an average over five runs of an experiment that consists of 500 packets sent at channel width-MCS combinations




Fig. 7. **Rate adaptation under varying transmission amplitudes.** WhiteRate keeps packet errors below the given threshold (left), without sacrificing throughput (center). Transmission bitrate (right) is higher for SampleRate, but at the cost of higher PER. All data points are averages of five runs, with error bars marking two standard deviations.

picked by the protocol under study, either WhiteRate or SampleRate<sup>5</sup>. This set of experiments is run in the indoor testbed, as it allows a wider range of transmission amplitude values at which the link can be maintained, thus enabling comprehensive evaluation.

WhiteRate abides by the application PER requirements. In Figure 7(a) we see that the PER threshold remains satisfied with all but very low transmission amplitudes. To achieve the desired packet delivery WhiteRate lowers the bitrate (Figure 7(c)). SampleRate, on the other hand, does not explicitly take application requirements into account. The resulting PER is variable and above the threshold value.

Despite being constrained by the PER threshold, WhiteRate manages to keep the throughput on a par with SampleRate, a protocol specifically designed with maximum throughput in mind. We ran SampleRate with the highest channel width, thus the highest bitrates for every MCS. The highest throughput, however, is found at the [MCS, channel width] that balances the bitrate, and link conditions (PER). This provides a strong case for channel width adjustment even if we consider rate adaptation in a traditional sense of throughput maximization.

### Energy Savings

We quantify the energy savings that WhiteRate delivers under varying levels of PER tolerance. We calculate the energy per bit by multiplying the time needed to transmit a bit of information by the platform power consumption. Time to transmit a bit of information is calculated directly from the MCS and channel width used. The power consumption depends on the transmission device and differs drastically between an experimental SDR platform such as USRP2 and a commodity device such as a WiFi NIC. For system evaluation purposes we use power consumption figures of a USRP2 device from [18]. We discard the base power, since it is MCS and channel width independent, an order of magnitude higher than the base power of commodity devices and

5. In our previous, trace-based evaluation we analyzed channel width, MCS combinations that are optimal in terms of PER or result in zero loss [19]. Since we cannot guarantee the exact same channel in every run in the real-time experiment, here we omit Optimal selections.




Fig. 8. **Energy efficiency with varying PER threshold.**

incurs bias toward WhiteRate or any protocol that lowers the transmission time.

The results of experiments over our long-distance outdoor link are shown in Figure 8. At each error tolerance level we sent five hundred packets; each data point in the figure is an average value of five runs, with error bars representing two standard deviations. Delivery requirements have a strong impact on energy efficiency: the energy per bit can be cut to a third of its initial value if the delivery requirements are relaxed to the PER threshold of 20%. Even a relatively tight error rate requirement (5%) yields two-fold energy savings. As a consequence, the *priority determination* module described in Section 6 can strictly favor energy efficient operation by adjusting the required PER level.

### Energy Efficiency and Application Performance

We showed that WhiteRate increases energy efficiency when the PER threshold is loosened. However, higher packet error rate leads to lower end-to-end communication quality. We evaluate WhiteRate's ability to balance the energy savings and the communication quality. We are especially interested in voice-over-IP as it is the key application in rural areas where communication alternatives are non-existent [25].

In our South African testbed we transmit a stream of 500 packets that are modelled to resemble G.711 encoded speech and vary the application PER requirements from 1% to 20%. We measure the actual loss rate and the energy consumption; the results are averaged over five

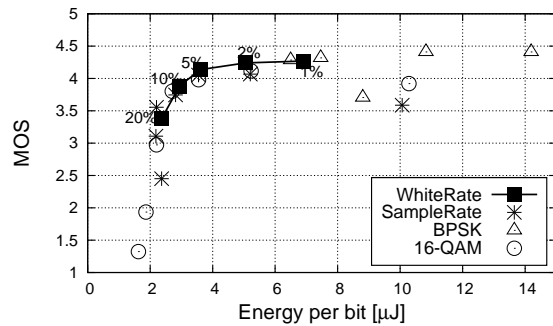


Fig. 9. **Tradeoff between voice quality and energy savings.** The plot shows WhiteRate runs with different PER requirements, and SampleRate, BPSK and 16-QAM with seven different channel widths. With a single parameter – PER threshold – WhiteRate allows fine grained selection between energy efficiency and VoIP call quality.

runs. Perceived quality of voice communication is related to loss rate through the e-model. The model reports the Mean Opinion Score (MOS) of voice quality ranging from 1 (poor) to 5 (excellent), with anything above 3.5 being considered satisfactory. In Figure 9 we show calculated MOS and energy per bit for WhiteRate as well as for SampleRate (tested with seven different channel widths) and fixed MCS, channel width combinations (BPSK and 16-QAM).

The results in Figure 9 show the energy saving benefits of loosening the packet delivery requirements. Relaxing the threshold from 1% to 5% PER yields 48% energy per bit savings with minimal impact on MOS. In the same figure we plot the performance of SampleRate that we run independently under seven different channel widths (from 66 to 462 active subcarriers). SampleRate is application-agnostic and results in unpredictable MOS values. We also plot the results of the operation on a fixed MCS (we plot 1/2 BPSK and 1/2 16-QAM) under changing channel width (each point represents results from a single channel width). Fixed MCS gives little freedom in picking the desired point on the tradeoff line. With a single MCS we can achieve either good voice quality or high energy savings, but cannot balance the two. Even with tight PER requirements (2%), WhiteRate halves the energy needs of BPSK, while significantly improving the sound quality delivered by 16-QAM, from poor (MOS 2) to good (MOS 4). An important advantage of WhiteRate over other protocols is that the savings can be adjusted to any point on the line depending on the priorities: application performance or energy savings.

### Spectrum Savings

WhiteRate is a cross layer protocol whose goal is to satisfy the application level PER requirement through physical layer adaptation. Spectrum efficiency – the information rate over a unit of bandwidth – should be analyzed from the application’s point of view. In practical terms, improvement in efficiency can be observed as




Fig. 10. **Number of supported VoIP calls with changing PER requirements.** We analyze two different medium access schemes: TDMA and FDMA. In both cases, WhiteRate harnesses the varying requirements to enable a higher number of simultaneous calls.

a larger number of simultaneous conversations that can be supported with the fixed amount of bandwidth over a fixed time interval.

To quantify WhiteRate’s spectrum efficiency under changing PER threshold we adopt two access schemes, TDMA and FDMA. In the indoor testbed we run a single 500-packet VoIP flow and measure its resource usage under each of the schemes. The result is then averaged over five runs and used to calculate the maximum number of identical G.711 VoIP flows that can be simultaneously supported. In the TDMA case, we use the full band for all transmissions and try to pack as many flows as possible in time, while in the FDMA case each flow has its own band of spectrum.

In Figure 10 we show the number of supported VoIP calls for PER thresholds ranging from 0.1% to 20%. For both the TDMA and FDMA schemes, the number of supported increases with the higher loss tolerance. As a side consequence, the call quality drops, yet, as long as the PER remains above 10%, the resulting MOS is above 3.5, i.e. satisfactory. In our example, three times more calls can be supported with 10% than with 0.1% loss tolerance. In highly populated urban areas spectrum availability is often critically low. With WhiteRate network administrators can easily balance between the individual user’s communication quality and the number of supported clients in the network.

### Dynamic Behavior

The PER threshold is set according to the system requirements, which can vary over time. For example, a solar powered wireless router does not need to be energy efficient during long stretches of sunny weather. Similarly, in a cognitive radio network presence of primary users dictates the spectrum availability. Thus, our protocol needs to adapt to changing operating conditions.

In the indoor testbed, we first examine how WhiteRate reacts to changing PER threshold. We start a packet flow and set the PER tolerance to 5%. After 400 packets are transmitted, we change the threshold to 20% and, 300 packets later we change it to 10%. In Figure 11 we plot




Fig. 11. **WhiteRate adapts to changing PER threshold.** The figure shows average packet loss ratio for every 20 packets. The first forty packets experience high loss until the algorithm finds a well performing [MCS, channel width]. The figure also shows a third order polynomial curve fitted according to the data points.

PER for each group of 20 transmitted packets over time. When threshold is raised WhiteRate’s aggressive probing quickly identifies a MCS, channel width combination that yields higher throughput. Once the threshold is tightened, WhiteRate immediately, at the subsequent transmission, adjusts MCS, channel width to values that result in acceptable loss. In the general case, WhiteRate adapts to a lower threshold faster than to a higher threshold. A rigorous initial threshold usually requires WhiteRate to explore a substantial part of the MCS, channel width space, thus the algorithm already knows what level of service is afforded by each of the combinations. In the reverse case, WhiteRate needs to explore the combinations before any decision can be made.

We next analyze how WhiteRate adapts to dynamic spectrum availability. A packet flow is started with 40kHz of available spectrum. After 300 transmitted packets the maximum width is extended to 300kHz, and after 300 additional packets it is decreased to 80kHz. It is important to note that we do not change any physical layer parameters of the link. Rather, we restrict the OFDM subcarrier groups that WhiteRate can activate.

We show the measured bitrate and throughput from a single experiment in Figure 12. WhiteRate immediately adapts to the restricted spectrum availability: if it already operated on a channel width that is lower than the new boundary, no change is necessary; otherwise, it momentarily deactivates the conflicting subcarrier groups. When more bandwidth becomes available, WhiteRate can gradually utilize it. Hence, the throughput in Figure 12(b) shows gradual increase. This is a consequence of WhiteRate’s conservative width occupation strategy: groups of subcarriers are activated one by one as long as the PER threshold is satisfied.

In the reported experiment WhiteRate performs a consecutive width occupation approach for the following reasons: it simplifies the protocol implementation; it satisfies the packet delivery requirement, which we see as the most important requirement of the system; and finally, we do not know channel gains over the whole

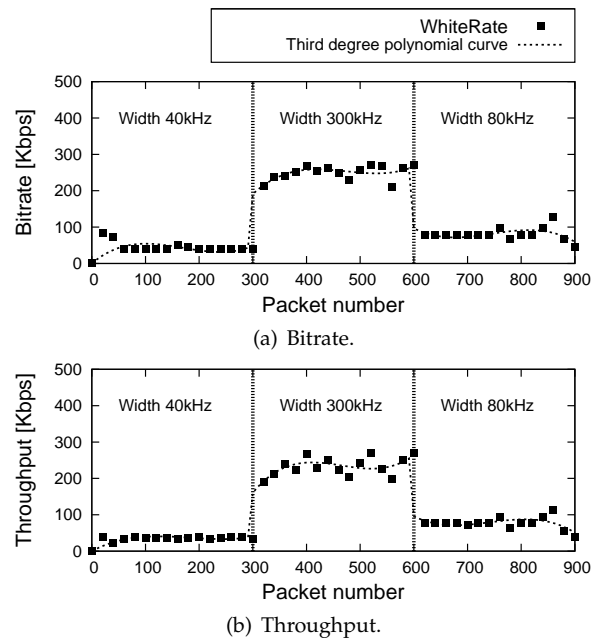


Fig. 12. **WhiteRate adapts to changing spectrum availability.** We modify available bandwidth from 40kHz to 300kHz and then to 80kHz.

new channel so it makes sense to be cautious about activating subcarriers. However, we also envision an alternative approach for dynamic and bandwidth critical applications: once the channel is extended a packet is sent over the whole width so that channel gain information is harvested; then, subcarrier groups are ordered and a binary search is performed to quickly identify the widest channel that supports the given PER threshold.

## 10 CONCLUSION

In this work we developed WhiteRate - a channel width and modulation and coding scheme adaptation protocol that adjusts a flexible physical layer according to the user needs for reliable communication and available resources such as energy and spectrum. WhiteRate is designed to answer to a wide variety of operating conditions, as demonstrated in extensive evaluation of the protocol on both indoor as well as outdoor white space links. The experimental results show that by abiding to relaxed delivery constraints, WhiteRate uses nearly a half of the energy and a third of the spectrum to successfully deliver a VoIP call. In addition, compared to a standard SampleRate adaptation protocol, WhiteRate explicitly takes communication reliability into account and enables high quality communication in various channel conditions. Even more importantly, WhiteRate enables operation that balances resource utilization and communication quality based on the resource availability and user needs. Through WhiteRate, bitrate adaptation goes beyond throughput maximization and becomes a valuable tool for wireless resource management.

In this work context awareness was considered in terms of energy reserves, spectrum availability and

packet error rate. Packet delay, device location and other aspects might be more important in certain scenarios. Additionally, we exploited two major physical layer parameters: MCS and channel width. However, the SDR-enabled flexibility allows virtually endless possibilities for PHY layer adaptation. From per-subcarrier transmission power to the length of OFDM cyclic prefix, a number of parameters that impact resource utilization can be modified. Exploration of this high dimensionality space promises to take us closer to a truly amorphous wireless system that serves applications in the most efficient, context-aware manner.

## 11 ACKNOWLEDGEMENTS

This work was funded in part through and NSF NetSE award CNS-1064821 and a gift from Google.

## REFERENCES

- [1] <http://www.voiptroubleshooter.com/diagnosis/emodel.html>.
- [2] *ITU World Telecommunication/ICT Indicators Database*, December 2010.
- [3] Mobile Broadband: The Benefits of Additional Spectrum. Tech. report, FCC, October 2010.
- [4] *IDC Worldwide Mobile Phone 2011/2015 Forecast and Analysis*, March 2011.
- [5] P. A. K. Acharya, A. Sharma, E. Belding, K. Almeroth, and K. Papagiannaki. Congestion-Aware Rate Adaptation in Wireless Networks: A Measurement-Driven Approach. In *IEEE SECON'08*, San Francisco, CA, June 2008.
- [6] Atheros. Power Consumption and Energy Efficiency Comparisons of WLAN Products, 2003.
- [7] J. Bicket. Bit-rate Selection in Wireless Networks. Master Thesis, MIT, 2005.
- [8] F. Boulos, D. S. Hands, B. Parrein, and P. Le Callet. Perceptual Effects of Packet Loss on H.264/AVC Encoded Videos. In *VPQM'09*, Scottsdale, AZ, January 2009.
- [9] L. N. Cai, D. Chui, M. McCutcheon, M. R. Ito, and G. W. Neufeld. Transport of MPEG-2 Video in a Routed IP Network. In *IDMS '99*, Toulouse, France, October 1999.
- [10] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl. A Case for Adapting Channel Width in Wireless Networks. In *SIGCOMM'08*, Oct. 2008.
- [11] C. Cordeiro, K. Challapali, D. Birru, B. Manor, and S. Diego. IEEE 802.22: An Introduction to the First Wireless Standard based on Cognitive Radios. *Journal of Communications*, 1(1):38–47, 2006.
- [12] A. Gudipati and S. Katti. Strider: Automatic Rate Adaptation and Collision Handling. In *SIGCOMM'11*, Toronto, ON, Canada, August 2011.
- [13] K. Harrison, S. M. Mishra, and A. Sahai. How Much White-Space Capacity is There? In *DySpan'10*, Singapore, April 2010.
- [14] J. Heiskala and J. Terry. *OFDM Wireless LANs: A Theoretical and Practical Guide*. SAMS, Indianapolis, IN, 2001.
- [15] H. Liu, J. He, P. Cui, J. Camp, and D. Rajan. ASTRA: Application of Sequential Training to Rate Adaptation. In *SECON'12*, Seoul, Korea, June 2012.
- [16] G. W. Miao, N. Himayat, and G. Y. Li. Energy-Efficient Link Adaptation in Frequency-Selective Channels. *IEEE Transactions on Communications*, 58(2):545–554, 2010.
- [17] G. Nychis, T. Hottelier, Z. Yang, S. Seshan, and P. Steenkiste. Enabling MAC Protocol Implementations on Software-Defined Radios. In *NSDI'09*, Boston, MA, April 2009.
- [18] V. Pejovic and E. Belding. Energy-Efficient Communication in Next Generation Rural-Area Wireless Networks. In *CoRoNet'10*, Chicago, IL, September 2010.
- [19] V. Pejovic and E. Belding. A Context-aware Approach to Wireless Transmission Adaptation. In *SECON'11*, Salt Lake City, UT, June 2011.
- [20] V. Pejovic, E. Belding, and M. Marina. An Energy-Flow Model for Self-Powered Routers and its Application for Energy-Aware Routing. In *NSDR'09*, Big Sky, MT, October 2009.

- [21] J. G. Proakis. *Digital Communications*. McGraw-Hill, 2000.
- [22] H. Rahul, F. Edalat, D. Katabi, and C. Sodini. Frequency-aware Rate Adaptation and MAC Protocols. In *MobiCom'09*, Beijing, China, Sept. 2009.
- [23] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan. Measurement-based Characterization of 802.11 in a Hotspot Setting. In *EWIND*, Philadelphia, PA, August 2005.
- [24] T. Schmidl and D. Cox. Robust Frequency and Timing Synchronization for OFDM. *IEEE Transactions on Communication*, 45(12):1613 – 1621, December 1997.
- [25] S. Sen, S. Kole, and B. Raman. Rural Telephony: A Socio-Economic Case Study. In *ICTD'06*, Bangalore, India, May 2006.
- [26] S. Sen, N. Santhapuri, R. R. Choudhury, and S. Nelakuditi. Accurate: Constellation aware rate estimation in wireless networks. In *NSDI'10*, San Jose, CA, April 2010.
- [27] P. Shankar, T. Nadeem, J. Rosca, and L. Iftode. CARS: Context Aware Rate Selection for Vehicular Networks. In *ICNP'08*, Orlando, FL, October 2008.
- [28] S. Surana, R. Patra, S. Nedeveschi, M. Ramos, L. Subramanian, Y. Ben-David, and E. Brewer. Beyond Pilots: Keeping Rural Wireless Networks Alive. In *NSDI'08*, San Francisco, CA, April 2008.
- [29] Ubiquiti Networks. *XR5 Power Consumption Study*. [www.ubnt.net](http://www.ubnt.net).
- [30] M. Vutukuru, H. Balakrishnan, and K. Jamieson. Cross-layer Wireless Bit Rate Adaptation. In *SIGCOMM*, Barcelona, Spain, August 2009.
- [31] S. Wong, H. Yang, S. Lu, and V. Bharghavan. Robust Rate Adaptation for 802.11 Wireless Networks. In *MobiCom'06*, Los Angeles, CA, September 2006.
- [32] L. Yang, W. Hou, L. Cao, B. Y. Zhao, and H. Zheng. Supporting Demanding Wireless Applications with Frequency-agile Radios. In *NSDI'10*, San Jose, CA, April 2010.
- [33] M. Zafer and E. Modiano. A Calculus Approach to Minimum Energy Transmission Policies with Quality of Service Guarantees. In *INFOCOM'05*, Miami, FL, April 2005.



**Veljko Pejovic** received a BS from the School of Electrical Engineering, Belgrade, Serbia and a PhD/MS from the Department of Computer Science, University of California Santa Barbara. Currently he is a postdoctoral Research Fellow at the School of Computer Science, University of Birmingham, United Kingdom. His research is focused on adaptive and resource-efficient wireless technologies and their impact on the society. He is broadly interested in wireless networks, mobile sensing systems, machine learning, technologies for development, technology and society. More details available at [www.cs.bham.ac.uk/~pejovicv](http://www.cs.bham.ac.uk/~pejovicv).



**Elizabeth M. Belding** is a Professor in the Department of Computer Science at the University of California, Santa Barbara, and is currently the departments Vice Chair. Elizabeths research focuses on mobile networking, specifically multimedia, monitoring, advanced service support, and solutions for developing and underdeveloped regions. Elizabeth is the author of over 100 technical papers and has served on over 60 program committees for networking conferences. She is currently on the steering committee of the ACM Networked Systems in Developing Regions (NSDR) Workshop, the editorial board of the IEEE Pervasive Magazine, and is serving as a guest editor for the IEEE Pervasive Magazine special issue on Pervasive Information and Communication Technologies for Developing Regions. Elizabeth is the recipient of an NSF CAREER Award, and a 2002 MIT Technology Review 100 award, awarded to the worlds top young investigators. She is an ACM Distinguished Scientist.